Oracle Private Cloud Appliance

Release Notes for Release 3.0.1

F49448-09
December 2022

Oracle Private Cloud Appliance Release Notes for Release 3.0.1,

F49448-09

Contents

Preface

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

1 Accessibility and Oracle Private Cloud Appliance

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

2 Feature Updates

Latest Features 2-1
Features Released in Software Version 3.0.1-b697160 (August 2022) 2-1

3 Service Limits

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

4 Known Issues and Workarounds

Platform Issues 4-1
- Compute Node Provisioning Takes a Long Time 4-1
- Not Authorized to Reconfigure Appliance Network Environment 4-1
- Unable to List Compute Images Available from Management Node 4-2

Grafana Service Statistics Remain at Zero 4-2
Terraform Provisioning Requires Fully Qualified Domain Name for Region 4-3
Synchronizing Hardware Data Causes Provisioning Node to Appear Ready to Provision 4-3
Rack Elevation for Storage Controller Not Displayed 4-3
Free-Form Tags Used for Extended Functionality 4-3
Imported Images Not Synchronized to High-Performance Pool 4-4
API Server Failure After Management Node Reboot 4-5
CLI Command Returns Status 500 Due To MySQL Connection Error 4-5
Microservice Pods Unresponsive After Storage Controller Failover 4-5
Administrators in Authorization Group Other Than SuperAdmin Must Use Service CLI to Change Password 4-6
ZFS Storage Appliance Password Must Not Exceed 16 Characters 4-6
Compute Nodes Ready to Provision But Not Detected Due to Hardware Inventory Error 4-7
DNS Entries Cannot Be Removed From Initial Appliance Configuration 4-7
User Interface Issues 4-7
Moving Resources Between Compartments Is Not Supported in the Compute Web UI 4-7
No Available Compute Web UI Operation to Update Instance Pool 4-8
Saving Resource Properties Without Modifications Briefly Changes Status to Provisioning 4-8
Resource Name Change Not Shown Until Manual Refresh 4-8
NFS Export Squash ID Not Displayed 4-8
Scrollbars Not Visible in Browser 4-9
Authorization Failure When Retrieving Compartment Data 4-9
Object List Is Not Updated Automatically 4-9
File Storage Mount Target Link Not Available 4-10
UDP Ports Not Displayed In Security List Rules Table 4-10
Not All Resources Shown in Drop-Down List 4-10
Identity Provider Description Not Displayed 4-11
Volume Group Can Be Created Without Name 4-11
File Systems and Mount Targets Not Displayed 4-11
Time Stamp Indicates Job Ended on New Year's Day 1970 4-12
Optional ICMP Security Rule Parameters Cannot Be Removed 4-12
Compartment Selector Not Available When Creating DHCP Options 4-12
Appliance Initial Configuration Wizard Hangs After Setting Up Primary Administrative Account 4-13
Uplink VLAN Restrictions Not Enforced By Service Web UI 4-13
Custom Search Domain Error Not Rolled Back When Operation Is Canceled 4-13
DHCP Options Error Message for Custom Search Domain Is Misleading 4-14
Compute Node Provisioning State Not Automatically Refreshed in Service Web UI 4-14
Creating a Flex Shape Instance in the Compute Web UI Allows Invalid Memory Values 4-15
Networking Issues 4-15
DNS Zone Scope Cannot Be Set 4-15

To Update a DNS Record the Command Must Include Existing Protected Records 4-15
Oracle Linux 8 Instance Host Name Resolution Fails 4-15
Create Route Table Fails With Confusing Error Message 4-16
VCN Creation Uses Deprecated Parameter 4-16
File Storage Traffic Blocked By Security Rules 4-16
Stateful and Stateless Security Rules Cannot Be Combined 4-18
VCN With Single Subnet of Same Size Not Supported 4-18
Routing Failure With Public IPs Configured as CIDR During System Initialization 4-19
Admin Network Cannot Be Used for SEUI Access 4-19
Compute Service Issues 4-19
No Consistent Device Paths for Connecting to Block Volumes 4-19
Instance Pools Cannot Be Terminated While Starting or Scaling 4-20
Network Interface on Windows Does Not Accept MTU Setting from DHCP Server 4-20
Concurrent Instance Migrations From the Same Compute Node Not Supported 4-20
Oracle Solaris Instance in Maintenance Mode After Restoring from Backup 4-21
Instance Disk Activity Not Shown in Compute Node Metrics 4-21
Attached Block Volumes Not Visible Inside Oracle Solaris Instance 4-21
Host Name Not Set In Successfully Launched Windows Instance 4-22
Oracle Solaris Instance Stuck in UEFI Interactive Shell 4-22
Instance Yum Configuration Is Overwritten During Instance Launch 4-23
Instance Provisioning Terminates with "Could not retrieve path for iscsi_device_id" 4-23
Storage Services Issues 4-24
Creating Image from Instance Takes a Long Time 4-24
Offset in Seconds Not Supported for Block Volume Backup Policy 4-24
Object Storage Commands Fail With Authentication Error 4-24
Object Storage Not Compatible with Terraform 4-25
Object Storage Pre-Authenticated Request REST Syntax Error 4-25
Large Object Transfers Fail After ZFS Controller Failover 4-25
Use Multipart Upload for Objects Larger than 100MiB 4-25
File System Export Created Despite Job Failure 4-26
Compute Image Import Fails After ZFS Controller Failover 4-26
Compute Image with UEFI Appears with BIOS Launch Mode After Import 4-26
File System Export Temporarily Inaccessible After Large Export Options Update 4-27
Block Volume Stuck in Detaching State 4-27
Scheduled Volume Backups Do Not Appear in Backup List 4-27
Compute Image Exported without Launch Mode 4-28
Permission to Manage Object Storage Objects Does Not Allow Listing Objects 4-28
OCI CLI Might Not Return the Correct Value for Object Storage Namespace 4-29
Serviceability Issues 4-29
DR Configurations Are Not Automatically Refreshed for Terminated Instances 4-29
Message Attribute Missing With Disaster Recovery Commands 4-30

Enabling DR Replication Fails on Second System 4-30
ULN Mirror Is Not a Required Parameter for Compute Node Patching 4-31
Patch Command Times Out for Network Controller 4-31
Disaster Recovery Configuration Hard-Codes Netmask of Replication IP Addresses 4-32
Order of Upgrading Components Has Changed 4-32

Preface

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

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

Audience

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

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

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

Feedback

Provide feedback about this documentation at https://www.oracle.com/goto/docfeedback.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.

Convention	Meaning
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, code in examples, text that appears on the screen, or text that you enter.
`$` prompt	The dollar sign (`$`) prompt indicates a command run as a non-root user.
`#` prompt	The pound sign (`#`) prompt indicates a command run as the `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/.

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

Access to Oracle Support for Accessibility

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

Diversity and Inclusion

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

1

Accessibility and Oracle Private Cloud Appliance

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

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

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

Oracle JET User Interface Accessibility Features

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

Oracle Server X9-2 Accessibility Features

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

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

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

Oracle Server X9-2 Hardware Accessibility

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

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

system status, indicator interpretation, and system configuration. For details, see Oracle Integrated Lights Out Manager Manager Accessibility.

Oracle Integrated Lights Out Manager Manager Accessibility

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

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

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

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

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

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

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

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

Oracle Hardware Management Pack Accessibility

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

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

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

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

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

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

BIOS Accessibility

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

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

2
Feature Updates

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

Latest Features

Flexible Compute Shapes

A flexible compute shape lets you customize the number of OCPUs and the amount of memory when launching your instance. This flexibility lets you create instances that meet your workload requirements, while optimizing performance and using resources efficiently. For details see Standard Shapes in the Oracle Private Cloud Appliance Concepts Guide.

GUI Support for Viewing CPU and Memory Metrics

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

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

Compute Instance Availability

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

Optimized NUMA Alignment

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

View CPU and Memory Metrics at the Fault Domain Level

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

Secondary Private IP Addresses

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

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

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

3
Service Limits

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

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

Tenancy Resource Configuration Limits

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

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

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

System Load and Concurrency Limits

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

Resource Type	Operation	Concurrency Limit
compute instance	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

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

Note:

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

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

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

4

Known Issues and Workarounds

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

Platform Issues

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

Compute Node Provisioning Takes a Long Time

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

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

Bug: 33519372

Version: 3.0.1

Not Authorized to Reconfigure Appliance Network Environment

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

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

Bug: 33535069

Version: 3.0.1

Unable to List Compute Images Available from Management Node

A set of compute images is provided with the appliance. They are stored on the shared storage of the three management nodes, and made available for download and import through a web server. To retrieve an image, you need the address of the web server and the exact name of the image. Unfortunately, only privileged appliance administrators have access to the Oracle Linux command line in order to list the files stored in the images directory. The workaround below is provided to help tenancy users retrieve the images from the management node web server. Note that the content of the images directory may change over time as the system is patched or upgraded.

Workaround: To retrieve a compute image from the web server running on the management node, use the URL syntax below. If you do not know, ask an appliance administrator to provide the host name (`<mgmt_vip_hostname>`) associated with the virtual IP of the management node cluster.

```
URL syntax:
https://<mgmt_vip_hostname>.<system_name>.<domain_name>:8079/images/
<image_file_name>

Example:
https://manager.myprivatecloud.example.com:8079/images/uln-pca-Oracle-
Linux-8-2022.01.12_0.oci
```

The compute images available for download are stored on the shared storage of the management nodes. These are the image file names you need to construct the download URL:

```
uln-pca-Oracle-Linux-7.9-2022.01.12_0.oci
uln-pca-Oracle-Linux-8-2022.01.12_0.oci
uln-pca-Oracle-Solaris-11.4.35-2021.09.20_0.oci
```

Bug: 33765086

Version: 3.0.1

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

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. 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"}'
  ```

- `PCA_blocksize`

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

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

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

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

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

Use of these tags counts against your tag limit.

Version: 3.0.1

Imported Images Not Synchronized to High-Performance Pool

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

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

Bug: 33660897

Version: 3.0.1

API Server Failure After Management Node Reboot

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

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

Bug: 33191011

Version: 3.0.1

CLI Command Returns Status 500 Due To MySQL Connection Error

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

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

Bug: n/a

Version: 3.0.1

Microservice Pods Unresponsive After Storage Controller Failover

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

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

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

Bug: 33495030

Version: 3.0.1

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

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

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

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

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

Bug: 33749967

Version: 3.0.1

ZFS Storage Appliance Password Must Not Exceed 16 Characters

From the Service Enclave of Oracle Private Cloud Appliance it is possible to send a command to the ZFS Storage Appliance in order to apply a password that violates the password policy of the storage appliance. More specifically, the ZFS Storage Appliance synchronizes the host root password to the ILOM, which does not accept a password length greater than 16 characters. As a result, attempts to set a longer password will fail but the corresponding Vault key is updated. This breaks the synchronization of component passwords in a way that cannot be restored by the administrator.

Workaround: Make sure that the password you set for the ZFS Storage Appliance complies with its ILOM password policy. Follow the password complexity rules for appliance infrastructure components, but do not exceed the maximum length of 16 characters for the storage appliance password.

The password rules for infrastructure components still apply. The minimum length is 8 characters. The password must contain at least 1 lower case letter (a-z), upper case letter (A-Z), digit (0-9), and symbol (@$!#%*?&).

Bug: 33915195

Version: 3.0.1

Compute Nodes Ready to Provision But Not Detected Due to Hardware Inventory Error

It may occur that compute node network interfaces are not registered correctly in the internal inventory database. As a result, the platform layer services fail to detect a compute node that is ready to be provisioned. Attempts to synchronize the hardware data from the Service CLI results in an error: "*Compute node data mac is empty.*".

Workaround: Please apply the latest patches available for your system. A code workaround was released to attempt to obtain the missing compute node data. It may take several minutes to retrieve the data, depending on the number of affected compute nodes installed. When the missing data has been added, run the `synchardwaredata` command from the Service CLI. Compute nodes should now be available for provisioning.

Bug: 33848769

Version: 3.0.1

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

DNS Entries Cannot Be Removed From Initial Appliance Configuration

During the initial configuration of the appliance, you set a number of essential network parameters, which include up to three DNS servers. If the DNS configuration in your data center changes, you can modify the DNS server details in the appliance network configuration. However, the DNS entries ***cannot*** be deleted, only updated.

Workaround: There is no workaround. Please select and apply DNS settings carefully.

> **Note:**
>
> If you specify multiple DNS servers, make sure that they are all configured to return identical responses. Queries are sent to the DNS servers in a round-robin pattern.

Bug: 33788348

Version: 3.0.1

User Interface Issues

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

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

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

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

Bug: 33038606

Version: 3.0.1

No Available Compute Web UI Operation to Update Instance Pool

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

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

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

Bug: 33393214

Version: 3.0.1

Saving Resource Properties Without Modifications Briefly Changes Status to Provisioning

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

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

Bug: 33445209

Version: 3.0.1

Resource Name Change Not Shown Until Manual Refresh

When you change the display name of a resource in the Compute Web UI, it may not be automatically refreshed on your screen. For some resources, a name change is only displayed after you click the manual refresh button.

Workaround: If a resource name change is not automatically reflected in the UI, click the Refresh button to manually refresh the window contents.

Bug: 33467585

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

Version: 3.0.1

Object List Is Not Updated Automatically

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

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

Bug: 33519215

Version: 3.0.1

File Storage Mount Target Link Not Available

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

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

Bug: 33571007

Version: 3.0.1

UDP Ports Not Displayed In Security List Rules Table

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

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

Bug: 33575269

Version: 3.0.1

Not All Resources Shown in Drop-Down List

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

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

Bug: 33583708

Version: 3.0.1

Identity Provider Description Not Displayed

In the Service Web UI, the description of an identity provider is not displayed in the Identity Provider detail page. However, a description can be provided while you create the identity provider, and is also displayed when you edit the identity provider.

Workaround: The Edit Identity Provider window does display all relevant settings, including the description field. To see the description, click Edit as you would to modify the identity provider settings.

Bug: 33585827

Version: 3.0.1

Volume Group Can Be Created Without Name

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

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

Bug: 33608462

Version: 3.0.1

File Systems and Mount Targets Not Displayed

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

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

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

Bug: 33666365

Version: 3.0.1

Time Stamp Indicates Job Ended on New Year's Day 1970

The Service Web UI allows you to display a list of jobs, and click the job name to display its detail page. In both the Jobs table and the job detail pages it may occur that the job end date is displayed as 31 December 1969 or 1 January 1970, depending on the time zone. For several job types the correct end date eventually appears when the job is successfully completed. However, in these specific cases the wrong job end time stamp is persisted, regardless of whether the job succeeded or failed: when live-migrating instances away from a compute node, and when deleting a tenancy.

Workaround: There is no workaround to retrieve the correct job end time stamp. If the job end is set to 31 December 1969 or 1 January 1970, please ignore the error. Job state – succeeded or failed – is displayed correctly.

Bug: 33582003

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.

Appliance Initial Configuration Wizard Hangs After Setting Up Primary Administrative Account

The Oracle Private Cloud Appliance Installation Guide provides step-by-step instructions to guide you through the initial setup of the appliance: see "Complete the Initial Setup" in the chapter Configuring Oracle Private Cloud Appliance. After creating the primary administrative account, you are instructed to refresh the browser window before signing in with the newly created account. However, refreshing the browser is not sufficient. When you proceed with the initial configuration wizard, errors occur that prevent you from saving your settings.

Workaround: Instead of refreshing the browser when the primary administrative account is set up, close the browser and open it again. Return to the initial configuration wizard and sign in with the administrator account you created. You can now apply the required settings and complete the wizard.

Bug: 33774118

Version: 3.0.1

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

Uplink VLAN Restrictions Not Enforced By Service Web UI

When performing the appliance initial setup procedure in the Service Web UI, you begin the network configuration by selecting the logical connection or routing type. When configuring the uplinks with static routing, you can choose to send this traffic over a VLAN. However, the Service Web UI does not verify that the VLAN ID you enter is a supported value. If it is not, the operation fails due to an "invalid parameter" and an error is returned.

Workaround: When you select static routing for the appliance uplinks to the data center, and you need to configure a VLAN, be sure to use an ID within the supported range of 2–3899. This is documented in the Initial Installation Checklist, as well as in the section "Network Infrastructure" in the chapter Hardware Overview of the Oracle Private Cloud Appliance Concepts Guide.

Bug: 33805854

Version: 3.0.1

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

Compute Node Provisioning State Not Automatically Refreshed in Service Web UI

The Service Web UI displays the *Provisioning State* of a compute node, both in the Rack Units table and each compute node detail page. However, when a compute node is ready for provisioning and you execute the provisioning command from the UI, the compute node detail page is not automatically refreshed to show the state change.

Workaround: When you provision a compute node from its detail page in the Service Web UI, refresh the browser page manually to track the state of the compute node and confirm when provisioning is complete.

Bug: 33817224

Version: 3.0.1

Creating a Flex Shape Instance in the Compute Web UI Allows Invalid Memory Values

If you create a flex shape instance in the Compute Web UI and enter an invalid value for memory, you see an error flagging the invalid memory value, but the instance is still created with the default memory value of 16GB.

Workaround: Select a valid memory value for the flex shape instance. See the valid ranges for flex shapes resources in the "Compute Shapes" section of the Compute Instance Concepts chapter of the Oracle Private Cloud Appliance Concepts Guide.

Bug: 34720069

Version: 3.0.1

Networking Issues

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

DNS Zone Scope Cannot Be Set

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

Bug: 32998565

Version: 3.0.1

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

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

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

Bug: 33089111

Version: 3.0.1

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

Oracle Linux 8 Instance Host Name Resolution Fails

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

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

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

Bug: 34734562

Version: 3.0.1

Create Route Table Fails With Confusing Error Message

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

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

Bug: 33570320

Version: 3.0.1

VCN Creation Uses Deprecated Parameter

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

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

Bug: 33620672

Version: 3.0.1

File Storage Traffic Blocked By Security Rules

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

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

```
+++ Ingress Rules ++++++++++++++++++++

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

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

Destination       Protocol     Source Ports            Destination Ports
-----------       --------     ------------            -----------------
<subnet CIDR>     TCP          111, 389, 445, 4045,    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

VCN With Single Subnet of Same Size Not Supported

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

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

Bug: 33758108

Version: 3.0.1

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

Routing Failure With Public IPs Configured as CIDR During System Initialization

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

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

Bug: 33765256

Version: 3.0.1

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

Admin Network Cannot Be Used for SEUI Access

The purpose of the (optional) Administration network is to provide system adminsistrators separate access to SEUI. 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

Compute Service Issues

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

No Consistent Device Paths for Connecting to Block Volumes

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

Workaround: No workaround is available.

Bug: 32561299

Version: 3.0.1

Instance Pools Cannot Be Terminated While Starting or Scaling

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

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

Bug: 33038853

Version: 3.0.1

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

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

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

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

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

   ```
   mtu_use_dhcp_config=true
   plugins=cloudbaseinit.plugins.common.mtu.MTUPlugin
   ```

2. Enter the command `Restart-Service cloudbase-init`.
3. Confirm that the MTU setting has changed. Use this command: `netsh interface ipv4 show subinterfaces`.

Bug: 33541806

Version: 3.0.1

Concurrent Instance Migrations From the Same Compute Node Not Supported

The system does not prevent you from starting the process of migrating all compute instances away from a given compute node multiple times. As as result, failures are reported for these migration jobs, even though it is likely that the instances are migrated successfully.

Workaround: After you started migrating the compute instances from a particular compute node, do not attempt to execute the same command again, unless you have confirmed that the previous migration job has completed.

Bug: 33582029

Version: 3.0.1

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

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 Yum Configuration Is Overwritten During Instance Launch

Oracle Linux compute images built for Oracle Cloud Infrastructure have a metadata value to set repository URLs in the yum configuration when an instance is launched. When using these Oracle Linux images with Oracle Private Cloud Appliance, the instances' yum configuration is overwritten with non-functional links to "*yum-null.oracle.com*".

Workaround: You must correct the configuration manually to be able to use yum inside your instances. To make sure that the yum configuration is preserved after a reboot, set the correct value in the file `/etc/yum/vars/ociregion`.

This issue has been fixed in a patch update. After instance launch, the yum configuration will point to the public yum repositories on *yum.oracle.com*. However, certain packages from the Oracle Cloud Infrastructure regional repositories – Ksplice, for example – are not available in the public repositories.

Bug: 33756673

Version: 3.0.1

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

Instance Provisioning Terminates with "Could not retrieve path for *iscsi_device_id*"

When provisioning an instance using either the Compute Web UI or the OCI CLI, the instance terminates with a work request error that states it cannot retrieve an iSCSI path.

```
$ oci work-requests work-request-error list --profile
PM_ADMIN --work-request-id
ocid1.workrequest.AK00661530.scasg01.compute-
qflfmywujn2enwiolj1ig4da1domp4k76
hv4hwbo5h1xdjbd67rl
{
  "data": [
    {
      "code": "Exception",
      "message": "Could not retrieve path for
600144f096933b920000635be5920390",
      "timestamp": "2022-10-28T14:23:23.829067+00:00"
    }
  ]
}
```

This error is the result of the failed devices on the compute node from previous errors, such as live migration, terminate instance, and detach volume.

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

Bug: 34747118

Version: 3.0.1

Storage Services Issues

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

Creating Image from Instance Takes a Long Time

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

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

Bug: 33392755

Version: 3.0.1

Offset in Seconds Not Supported for Block Volume Backup Policy

When creating a backup policy for your block volumes, you specify whether a backup should be made daily, weekly, monthly or yearly. Depending on the selected frequency you can also specify the month of the year, day of the month, day of the week, and hour of the day. The API also offers a schedule parameter to specify an offset in seconds, but it has not been implemented. Therefore this setting cannot be passed on to the ZFS storage appliance where the backups are made.

Workaround: To define the time of day in a volume backup policy, specify only the hour of the day; the `offsetSeconds` parameter is not supported.

Bug: 33529592

Version: 3.0.1

Object Storage Commands Fail With Authentication Error

Users in a tenancy can add up to 3 API keys to their profile, for authentication and authorization of various cloud operations. For each operation in the tenancy a command is sent to the API server, with an API key that the IAM service uses to determine whether the command may be executed. For performance optimization, the ZFS storage appliance caches the keys it received as part of object storage commands, along with the user OCIDs. If a user executes another object storage command with a different API key, and the previously cached key has not yet expired, the command results in an authentication error because a different key was expected.

Workaround: When working with the object storage service, users who have more than one API key in their profile must ensure they use the same key for all their object storage commands. When switching between API keys, do not perform any object storage operations until the cached key has expired. Each time a cached key is touched by a command, even if it fails, the lifetime of the key is extended.

Bug: 33560908

Version: 3.0.1

Object Storage Not Compatible with Terraform

The Object Storage service implemented in Private Cloud Appliance is currently not compatible with Terraform and the Oracle Cloud Infrastructure Go SDK.

Workaround: There is no workaround to use Terraform with object storage.

Bug: 33654986

Version: 3.0.1

Object Storage Pre-Authenticated Request REST Syntax Error

To create a pre-authenticated request, so users can access a bucket without having to provide an API key, you are instructed by the Oracle Cloud Infrastructure API reference to use the following syntax for REST API calls: `POST /n/{namespaceName}/b/{bucketName}/p/`. Unfortunately, using this syntax results in an authentication error. The Private Cloud Appliance API server does not accept the forward slash at the end, so you should remove it.

Workaround: Format the REST API command as follows: `POST /n/{namespaceName}/b/{bucketName}/p`.

Bug: 33540180

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 Created Despite Job Failure

When you create a file system export with a large number of export options, a timeout may occur because the system is busy and cannot process all instructions quickly enough. At this point, an error message is returned to indicate the "lock wait timeout" was exceeded, which suggests the operation has failed. In actual fact, the file system service continues to retry in the background, because the request was accepted but not yet completed due to system load. It is likely that the export is eventually created successfully despite the timeout error.

Workaround: To verify the status of the export, consider using the CLI command `oci fs export get`.

Bug: 33690163

Version: 3.0.1

Compute Image Import Fails After ZFS Controller Failover

If a ZFS controller failover or failback occurs, it is possible that for some time following the event, importing a compute image cannot be completed successfully. This behavior of the ZFS Storage Appliance is unpredictable, but it is expected that image import operations will succeed again, though several attempts may be required.

Workaround: If your image import fails shortly after a ZFS controller failover or failback occurred, retry the operation until it succeeds.

Bug: 33677366

Version: 3.0.1

Compute Image with UEFI Appears with BIOS Launch Mode After Import

When you import a compute image through the Compute Web UI, there is only one Launch Mode option available. If you select *Paravirtualized Mode*, the image will be set to "BIOS" during import, even if its metadata specifies UEFI. As a result, when you try to launch an instance from that image, the operation will fail. It is possible to override the launch options when using the CLI to launch an instance.

Workaround: If you imported a UEFI compute image with *Paravirtualized* launch mode, use the CLI to launch an instance from this image, and include the correct launch options as shown below:

```
oci compute instance launch [...] \
--launch-options '{"boot-volume-type":"PARAVIRTUALIZED","firmware":"UEFI_64",
"is-consistent-volume-naming-enabled":false,"is-pv-encryption-in-transit-
enabled": false,
"network-type":"PARAVIRTUALIZED","remote-data-volume-type": "PARAVIRTUALIZED"}'
```

Alternatively, it might be possible with some browsers to import the image through the Compute Web UI without selecting a Launch Mode option. If you cancel and close the Import Image window and re-open it by clicking the Import Image button, the Paravirtualized Mode might be deselected. Importing the image with those settings allows the UEFI setting from the image metadata to be applied.

Bug: 33736077

Version: 3.0.1

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

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.

Scheduled Volume Backups Do Not Appear in Backup List

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

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

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

Bug: 33785277

Version: 3.0.1

Compute Image Exported without Launch Mode

When custom images are exported from your tenancy to object storage, they are stored in `*.qcow2` format. The image launch mode and launch options are not registered in that format. When such an image is imported into a tenancy, default launch parameters are applied with the setting `"firmware": "BIOS"`. If that image was originally set to UEFI, the setting is lost during import. When you try to launch an instance from that image, the operation will fail.

Workaround: If you imported a UEFI compute image in `*.qcow2` format, you need to use the CLI to launch an instance from that image and override the launch options by explicitly entering the correct settings with the launch command:

```
oci compute instance launch [...] \
--launch-options '{"boot-volume-type":"PARAVIRTUALIZED","firmware":"UEFI_64",
"is-consistent-volume-naming-enabled":false,"is-pv-encryption-in-transit-
enabled": false,
"network-type":"PARAVIRTUALIZED","remote-data-volume-type": "PARAVIRTUALIZED"}'
```

Alternatively, you can create an image in `*.oci` format, add metadata, and specify the correct launch options. These will be recognized and applied during image import, so you can use the image to launch instances from within the Compute Web UI.

Bug: 33861247

Version: 3.0.1

Permission to Manage Object Storage Objects Does Not Allow Listing Objects

Access to resources in the tenancy is managed through policies applied to user groups. When a policy states that a group is allowed to *manage objects* in the tenancy, it grants permission – to the users who are members of the group in question – to execute all operations available for the *object* resource type. However, for users who belong to a group other than the default *Administrators*, the authorization request fails when an attempt is made to list the objects in a bucket. Other operations on object storage objects are executed as expected.

Workaround: There is no workaround available to list objects in a bucket. Note that users in the Administrators group do not appear to be affected by the issue.

Bug: 33914998

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

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

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

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

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

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

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

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

Bug: 34133183

Version: 3.0.1

Serviceability Issues

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

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

Message Attribute Missing With Disaster Recovery Commands

In general, for `get`, `list`, and `show` commands, a message attribute is included in the response. The implementation for disaster recovery is inconsistent: the message attribute is not returned for the API calls `dr-admin.computeinstance.get`, `dr-admin.mapping.list`, and `dr-admin.service.show`.

Workaround: There is no workaround. This inconsistent behavior should only affect developers or applications interacting directly with the API.

Bug: 33676119

Version: 3.0.1

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

Enabling DR Replication Fails on Second System

In the chapter Disaster Recoveryof the Oracle Private Cloud Appliance Administrator Guide, when you follow the instructions to set up peering between the ZFS Storage Appliances, the final step is to enable replication on the two systems. However, the final command fails when you run it on the second appliance. This happens because replication targets are already configured as a result of enabling replication on the first appliance, and using an existing replication target is not accepted.

Workaround: When you have enabled replication from the first appliance, do not run the same command on the second appliance. Instead, log on to the storage appliance and modify its replication targets manually to set the correct target *hostname* and *label* parameters.

Use the same remote IP addresses as in the `drSetupService` command, for the standard as well as the high-performance replication target. In the example below we use `remoteIp=10.50.7.33` and `remoteIpPerf=10.50.7.34`.

```
# ssh pcamn01
pcamn01 ~]# ssh 100.96.2.2

sn011742XC3024:> shares replication targets
sn011742XC3024:shares replication targets> ls
Targets:
TARGET       LABEL                        ACTIONS
target-000   mypca1sn01-dr1.example.com   0
target-001   mypca1sn02-dr1.example.com   0

sn011742XC3024:shares replication targets> select target-000
sn011742XC3024:shares replication target-000> set hostname=10.50.7.33
                      hostname = 10.50.7.33 (uncommitted)
sn011742XC3024:shares replication target-000> set label=target-zfssa-replication
                         label = target-zfssa-replication (uncommitted)
sn011742XC3024:shares replication target-000> commit
[...]

sn011742XC3024:shares replication target-000> up

sn011742XC3024:shares replication targets> select target-001
sn011742XC3024:shares replication target-001> set hostname=10.50.7.34
                      hostname = 10.50.7.34 (uncommitted)
sn011742XC3024:shares replication target-001> set label=target-zfssa-replication-
```

```
high
                          label = target-zfssa-replication-high (uncommitted)
sn011742XC3024:shares replication target-001> commit
[...]

sn011742XC3024:shares replication target-001> exit
```

When the ZFS storage appliance configuration has been updated, and you have logged out of the storage controller, run the following command on the management node:

```
kubectl exec -it `kubectl get pods |grep dr-admin |awk '{print $1}'` --python3 -c
"from pca_dr_admin.metadata import access; \
from pca_dr_admin.replication import replication; access.set_replication_enabled();
access.set_replication_high_enabled(); \
replication.create_repl_action('pca_dr_metadata', 'PCA_POOL');
replication.wait_repl_update_synced('pca_dr_metadata', 'PCA_POOL')"
```

Bug: 33760254

Version: 3.0.1

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

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

Disaster Recovery Configuration Hard-Codes Netmask of Replication IP Addresses

During the configuration of the disaster recovery service, when you set up peering between the two ZFS storage appliances, you specify a local and remote IP address for standard storage, and optionally for high-performance storage. You enter these IP addresses as part of the `drSetupService` command executed from the Service CLI. The command expects individual IP addresses; they are automatically hard-coded with a /23 netmask. If this conflicts with the data center network environment, the netmask should be modified manually before completing the disaster recovery service setup.

Workaround: Correct the IP address and netmask of the replication interfaces from the Oracle Linux command line of one of the management nodes. Apply this workaround on both systems before updating the CA certificate on the ZFS storage appliances.

For the standard storage replication interface, run this command:

```
# kubectl exec -it `kubectl get pods |grep dr-admin | awk '{print $1}'` --
python3 -c \
"from pca_dr_admin.metadata import access; from pca_dr_admin.replication import
replication; \
zfssa=replication.__get_zfs_session(replication.PCA_DEFAULT_POOL); \
interface=replication.__find_interface_name(zfssa,access.ZFS_REPL_INTERFACE_LABEL
); \
zfssa.set_interface_ips(interface,['<ip_cidr>'])"
```

Replace `<ip_cidr>` with the replication interface IP address in CIDR notation; for example: 10.128.1.21/24.

If you use high-performance storage, run this command as well:

```
# kubectl exec -it `kubectl get pods |grep dr-admin | awk '{print $1}'` --
python3 -c \
"from pca_dr_admin.metadata import access; from pca_dr_admin.replication import
replication; \
zfssa=replication.__get_zfs_session(replication.PCA_DEFAULT_POOL_HIGH); \
interface=replication.__find_interface_name(zfssa,access.ZFS_REPL_INTERFACE_HIGH_
LABEL); \
zfssa.set_interface_ips(interface,['<ip_cidr>'])"
```

Bug: 33965318

Version: 3.0.1

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