

# Oracle Private Cloud Appliance Administrator Guide



F74802-14  
August 2025



Oracle Private Cloud Appliance Administrator Guide,  
F74802-14

Copyright © 2022, 2025, Oracle and/or its affiliates.

# Contents

## Preface

---

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

## 1 Working in the Service Enclave

---

Using the Service Web UI	1
Logging In	1
Navigating the Dashboard	2
Using the Service CLI	6
Accessing the CLI	6
Command Syntax	6
Help and Command Completion	7
Base and Custom Commands	8

## 2 Hardware Administration

---

Displaying Rack Component Details	1
Viewing Appliance Details	1
Using the Rack Units List	2
Changing Passwords for Hardware Components	4
Checking Component Health	5
Performing Compute Node Operations	6
Provisioning a Compute Node	7
Disabling Compute Node Provisioning	8
Locking a Compute Node for Maintenance	9
Migrating Instances from a Compute Node	12
Manually Shut Down a Non Migratable Instance	14
Configuring the Compute Service for High Availability	15
Using Instance and Compute Service High Availability Configuration	15

Viewing and Setting Compute Service Configuration	17
Compute Service Configuration Commands	18
Configuring the Recovery State for a Stopped Instance	20
Enabling Strict Fault Domain Enforcement	20
Starting, Resetting or Stopping a Compute Node	20
Deprovisioning a Compute Node	22
Integrating GPU Expansion Nodes	23
Integrating a Compute Expansion Rack	25
Configuring the Active Directory Domain for File Storage	25
Reconfiguring the Network Environment	27
Editing Routing Information	28
Editing Management Node Information	28
Editing Data Center Uplink Information	29
Updating NTP Server Information	31
Editing Administration Network Information	32
Updating DNS Information	35
Updating Public IP Information	36
Configuring Appliance Proxy Settings	37
Creating and Managing Flex Networks	38
Taskmap for Creating a Flex Network	40
Creating a Flex Network	41
Enabling Flex Network Access	42
List Flex Networks	43
Get Flex Network Details	43
Editing Flex Networks	44
Disabling Flex Network Access	45
Deleting a Flex Network	45
Flex Network Examples	45
Configuring Flex Network Edge Mode to an Oracle Exadata	46
Accessing External Interfaces with Your CA Trust Chain	53
Create Certificate Signing Requests	54
Uploading Your CA Certificates	56

### 3 Administrator Account Management

---

Creating a New Administrator Account	1
Changing Administrator Credentials	2
Managing Administrator Privileges	3
Working with Authorization Groups	4
Working with Authorization Families	7
Changing Administrator Account Preferences	10
Deleting an Administrator Account	11

Federating with Microsoft Active Directory	12
Gathering Required Information from ADFS	12
Verifying Identity Provider Self-Signed Certificates	13
Managing Identity Providers	14
Adding Active Directory as an Identity Provider	14
Updating an Identity Provider	15
Viewing Identity Provider Details	16
Listing Identity Providers	16
Deleting an Identity Provider	16
Working with Group Mappings for an Identity Provider	17
Creating Group Mappings	17
Updating a Group Mapping	17
Viewing Group Mappings	18
Deleting a Group Mapping	18
Adding Private Cloud Appliance as a Trusted Relying Party in ADFS	18
Providing Federated Users Sign In Information	20

## 4 Tenancy Management

---

Creating a New Tenancy	1
Providing Platform Images	3
Modifying the Configuration of a Tenancy	4
Deleting a Tenancy	4

## 5 Viewing and Setting Resource Limits

---

Viewing Resource Limits	1
Listing and Showing Limit Definitions	1
Viewing Current Resource Limits	3
Working with Resource Limit Templates	4
Viewing Resource Limit Templates	4
Creating Resource Limit Templates	5
Assigning Resource Limit Templates	6
Unassigning Resource Limit Templates	7
Deleting Resource Limit Templates	8
Setting Resource Limit Overrides	8
Listing Resource Limit Overrides	8
Creating Resource Limit Overrides	9
Updating Resource Limit Overrides	10
Deleting Resource Limit Overrides	10
Planning Resource Limit Settings	11

## 6 Status and Health Monitoring

---

Using Grafana	1
Adding Grafana Users	2
Using Grafana Dashboards	4
Using Grafana Alerts	4
Checking the Health and Status of Hardware and Platform Components	6
Viewing and Interpreting Monitoring Data	7
Monitoring System Capacity	8
Viewing CPU and Memory Usage By Fault Domain	9
Viewing Disk Space Usage on the ZFS Storage Appliance	9
Viewing System Log Data	10
Using Grafana Explore Queries	10
Loki Logs	10
Audit Logs	12
LBaaS Logs	13
Using the Vector Service	13
Using Oracle Auto Service Request	15
Understanding Oracle Auto Service Request	15
Oracle Auto Service Request Prerequisites	16
Registering Private Cloud Appliance for Oracle Auto Service Request	16
Testing Oracle Auto Service Request Configuration	18
Unregistering Private Cloud Appliance for Oracle Auto Service Request	18
Disabling Oracle Auto Service Request	18
Enabling Oracle Auto Service Request	19
Viewing Admin Service Health Data	19
Compute Node CPU and Memory Utilization Faults	22
Storage Utilization Faults	24
Hardware Run State Faults	24
Health Checker Notification Faults	25
Manually Clearing Faults	25
Using Support Bundles	26
Using the asrInitiateBundle Command	27
Using the support-bundles Command	27
Uploading Support Bundles to Oracle Support	32

## 7 Backup and Restore

---

Activating Standard Daily Backup	1
Executing a Backup Operation	3
Identifying Converted Snapshots	4

Restoring the System from a Backup	5
------------------------------------	---

## 8 Disaster Recovery

---

Enabling Disaster Recovery on the Appliances	1
Collecting System Parameters for Disaster Recovery	2
Connecting the Components in the Disaster Recovery Setup	2
Setting Up Peering Between the ZFS Storage Appliances	3
Setting Up Peering Between the ZFS Storage Appliances Before 302-b892153	4
Setting Up Peering Between the ZFS Storage Appliances	8
Managing Disaster Recovery Configurations	14
Creating a DR Configuration	15
Adding Site Mappings to a DR Configuration	16
Removing Site Mappings from a DR Configuration	18
Adding Instances to a DR Configuration	19
Removing Instances from a DR Configuration	20
Refreshing a DR Configuration	21
Deleting a DR Configuration	22

## 9 Native Disaster Recovery

---

Migrating to the Native Disaster Recovery Service	1
Establishing a Peer Connection	3
Adding Cable Connections	3
Creating a Local Endpoint	4
Creating the Peer Connection	6
Setting Up the Disaster Recovery Service	8
Managing Disaster Recovery Configurations	10
Creating a DR Configuration	10
Maintaining Site Mappings	12
Adding and Removing Compute Instances	15
Refreshing a DR Configuration	19
Deleting a DR Configuration	20
Working With Disaster Recovery Plans	21
About DR Operations and Default Plans	21
Creating and Maintaining DR Plans	22
Customizing the Steps in a DR Plan	26
Executing a DR Plan	32
Tracking Disaster Recovery Metrics	38

# Preface

This publication is part of the customer documentation set for Oracle Private Cloud Appliance Release 3.0. Note that the documentation follows the release numbering scheme of the appliance software, not the hardware on which it is installed. All Oracle Private Cloud Appliance product documentation is available at <https://docs.oracle.com/en/engineered-systems/private-cloud-appliance/index.html>.

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

## Audience

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

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

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

## Feedback

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

## Conventions

The following text conventions are used in this document:

Convention	Meaning
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.

---

Convention	Meaning
monospace	Monospace type indicates commands within a paragraph, code in examples, text that appears on the screen, or text that you enter.
\$ prompt	The dollar sign (\$) prompt indicates a command run as a non-root user.
# prompt	The pound sign (#) prompt indicates a command run as the root user.

---

## Documentation Accessibility

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

## Access to Oracle Support for Accessibility

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

## Diversity and Inclusion

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

# 1

## Working in the Service Enclave

The appliance administrator's working environment is the Service Enclave. It is the part of the system where the appliance infrastructure is controlled. It provides tools for hardware and capacity management, tenancy control, and centralized monitoring of components at all system layers.

### Note

The CLI only accepts characters from the 7-bit ASCII table of letters, numerics, and other characters. Not all are valid for all fields, but A-Z, a-z, and 0-9 are generally accepted. The CLI ignores accented characters or those from the UTF table. Special care is needed to screen values provided in certificates or other sources.

More detailed information about the Service Enclave is provided in the Oracle Private Cloud Appliance Concepts Guide. Refer to the "Enclaves and Interfaces" section in the chapter "[Architecture and Design](#)".

This chapter describes the general usage principles of the graphical user interface and command line interface to the Service Enclave.

### Note

You access the Service Web UI using a web browser. For support information, please refer to the [Oracle software web browser support policy](#).

## Using the Service Web UI

The Service Web UI is the graphical interface to the Service Enclave. You can use the Service Web UI on its own or with the Service CLI to complete tasks. The Service Web UI provides the same core functionality as the Service CLI, however, the Service CLI does have some additional operations that do not have a UI equivalent.

This section provides general guidelines for using the Service Web UI. The actual commands and their functions are documented throughout the Oracle Private Cloud Appliance Administrator Guide as part of the step-by-step instructions.

## Logging In

To log into the Service Web UI, complete the following steps:

1. In a supported browser, enter the URL for your Oracle Private Cloud Appliance.

For example, `https://adminconsole.pcasys1.example.com` where `pcasys1` is the name of your Private Cloud Appliance and `example.com` is your domain.

The Sign In page is displayed.

2. Enter your Username and Password, and then click Sign In.

The Private Cloud Appliance dashboard displays with quick action tiles.

#### **Note**

If this is the first log in after a Private Cloud Appliance installation, the dashboard displays the ASR Phone Home page so you can register your system with My Oracle Support.

For more information, see [Registering Private Cloud Appliance for Oracle Auto Service Request](#).

## Navigating the Dashboard

When you log into the Service Enclave, the dashboard is displayed with a Quick Actions area containing clickable tiles for common tasks, such as viewing rack unit, tenancy, and appliance details and managing users and the network environment.

In the Observability & Management part of the dashboard, there is a quick action tile for Monitoring. When you click Monitoring, the Grafana console opens. For more information, see [Using Grafana](#).

In the top bar of the dashboard you can locate the realm and the system and domain names for your Private Cloud Appliance. You will see your user name in the top bar, as well, with links to your profile information, hardware data sync, oracle.com, and the ability to sign out.

#### **Note**

The dashboard is static and not configurable.

The navigation menu, which you can click on or tab to, lists appliance components and resources that you can manage within the Service Enclave of Private Cloud Appliance. When you click on an item in the navigation menu, a page is displayed that contains information about the component or resource. The following table provides details about what you can expect to find on these component and resource pages.

Component or Resource	Information Provided
Appliance Details	Read-only appliance configuration details and an option to edit rack name and description. For more information, see <a href="#">Displaying Rack Component Details</a> .

Component or Resource	Information Provided
Network Environment	<p>Read-only network configuration information and an Edit button that opens a Configure Network Params wizard where you can modify:</p> <ul style="list-style-type: none"> <li>• Routing uplink gateway, VLAN, and HSRP group, and spine virtual IP</li> <li>• Management nodes IPs and hostnames</li> <li>• Uplink port speed, count, port FEC, VLAN MTU, and netmask and spine IPs</li> <li>• NTP servers IP addresses</li> <li>• Admin network status</li> <li>• DNS servers IP addresses</li> <li>• Public IP ranges and object storage IP</li> </ul> <p>For more information, see <a href="#">Reconfiguring the Network Environment</a>.</p>
Rack Units	<p>Read-only list of all hardware components installed in the rack and detected by the appliance software and the following information for each:</p> <ul style="list-style-type: none"> <li>• Name</li> <li>• Rack unit type</li> <li>• State</li> <li>• Rack elevation</li> </ul> <p>Each component also has an Actions menu (three dots) with a View Details link to a component's detail page. For management nodes, switches, and storage controllers, the detail pages provide read-only rack unit and system information.</p> <p>For more information, see <a href="#">Displaying Rack Component Details</a>.</p> <p>For each compute node in the list, you can see additional information:</p> <ul style="list-style-type: none"> <li>• Provisioning state</li> <li>• Maintenance lock</li> <li>• Provisioned lock</li> </ul> <p>A compute node's detail page provides read-only compute node, rack unit, and system information. Additionally, from either its detail page or the Actions menu, you can perform several actions on a compute node, such as locking for maintenance, migrating all virtual machines, stopping, deprovisioning. For more information, see <a href="#">Performing Compute Node Operations</a>.</p>

Component or Resource	Information Provided
Tenancies	<p>Read-only list of all tenancies in the system and the following information for each:</p> <ul style="list-style-type: none"> <li>• Name</li> <li>• Description</li> <li>• Action menu</li> </ul> <p>Contains options to view a tenancy's details page, edit a tenancy's description, or delete a tenancy.</p> <p>You can also edit or delete a tenancy from its details page.</p> <p>A Create Tenancy button.</p> <p>For more information, see <a href="#">Tenancy Management</a>.</p>
Identity Providers	<p>Read-only list of identity providers and the following information for each:</p> <ul style="list-style-type: none"> <li>• Name</li> <li>• Force Authentication</li> <li>• Encrypt Assertion</li> <li>• Action menu</li> </ul> <p>Contains options to view an identity provider's details page and edit or delete the identity provider.</p> <p>You can also edit or delete an identity provider from its details page.</p> <p>A Create Identity Provider button.</p> <p>For more information, see <a href="#">Federating with Microsoft Active Directory</a>.</p>
IDP Group Mappings	<p>Read-only list of IDP group mappings in the system and the following information for each:</p> <ul style="list-style-type: none"> <li>• Name</li> <li>• IDP Group Name</li> <li>• Admin Group Name</li> <li>• Description</li> <li>• Action menu</li> </ul> <p>Contains options to view read-only information on an IDP group mapping details page. MORE...</p> <p>A Create Group Mapping button.</p> <p>For more information, see <a href="#">Federating with Microsoft Active Directory</a>.</p>

Component or Resource	Information Provided
Users	<p>Read only list of users in the system and the following information for each:</p> <ul style="list-style-type: none"> <li>• Name</li> <li>• Authorization Group</li> <li>• Default User</li> <li>• Action menu</li> </ul> <p>Contains options to view read-only information on a user's details page, change a user password, or delete a user. You can also change a user password or delete a user from its details page.</p> <p>A Create User button.</p> <p>For more information, see <a href="#">Administrator Account Management</a>.</p>
Jobs	<p>Read-only list of jobs that ran and the following information for each:</p> <ul style="list-style-type: none"> <li>• Object type</li> <li>• Start and end times</li> <li>• Run status - Active, Succeeded, Failed, or Aborted</li> <li>• Action menu</li> </ul> <p>Contains an option to view read-only information on a job's details page, which includes the user account that the job ran from.</p>
Upgrade & Patching	<p>Read-only list of upgrade and patching jobs that ran and the following information for each:</p> <ul style="list-style-type: none"> <li>• Job name</li> <li>• Request and job IDs</li> <li>• Start and end times</li> <li>• Command name</li> <li>• Result - Passed, Failed, Not Run, Canceled, or None</li> </ul> <p>A Create Upgrade or Patch button, where you can select:</p> <ul style="list-style-type: none"> <li>• Upgrade Request - includes several types of upgrades, such as compute node, host, ILOM, Kubernetes, and platform.</li> <li>• Patch Request - includes several types of patches, such as compute node, host, ILOM, Kubernetes, OCI Images, and platform.</li> </ul> <p>For more information, refer to the <a href="#">Oracle Private Cloud Appliance Upgrade Guide</a> and <a href="#">Oracle Private Cloud Appliance Patching Guide</a>.</p>
ASR Phone Home	<p>Read-only auto service request information and a Register button where you can register your Private Cloud Appliance.</p> <p>For more information, see <a href="#">Using Oracle Auto Service Request</a>.</p>

## Using the Service CLI

The command line interface to the Service Enclave, which we refer to as the *Service CLI* in the documentation, is available through the Oracle Linux shell on the management nodes. There is no additional installation or configuration required. The CLI provides access to all the functionality of the Service Web UI, as well as several additional operations that do not have a UI equivalent.

This section provides general guidelines for using the Service CLI. The actual commands and their functions are documented throughout the Oracle Private Cloud Appliance Administrator Guide as part of the step-by-step instructions in the chapters that follow.

## Accessing the CLI

To access the Service CLI, establish an SSH connection to TCP port 30006 on one of the following nodes and log in as an authorized administrator:

- On one of the management nodes.

```
$ ssh admin@pcamn02 -p 30006
Password authentication
Password:
PCA-ADMIN>
```

- On the Private Cloud Appliance.

```
$ ssh admin@admin.pca_hostname.example.com -p 30006
```

After successful authentication, you are in an interactive, closed shell environment where you perform administrative operations by entering commands at the PCA-ADMIN> prompt.

To terminate your CLI session, enter the `exit` command.

Command syntax and completion functions are described in the following sections.

## Command Syntax

In general, commands entered in the Service CLI have the following syntax:

```
PCA-ADMIN> command objectType <attributes> [options]
```

where:

- **command** is the command type to be initiated, for example: `list` or `create`.
- **objectType** is the target component or process affected by the command, for example: `list ComputeNode` or `create Tenant`.
- **attributes** are properties used to identify a specific object of the selected type to which the command must be applied, for example: `show ManagementNode name=pcamn02`.
- **options** are additional parameters that may be provided to affect the behavior of the command.

For example, you can add sorting and filtering options to the `list` command and select which data columns (fields) to display: `list RackUnit fields ipAddress,name,rackElevation,serialNumber,firmwareVersion where state eq running`.

The main elements of a command are separated by a space. Attributes are specified as "type=value". Lists are entered as a comma-separated series of values (such as fields *ipAddress,name,rackElevation,serialNumber,firmwareVersion*).

## Help and Command Completion

The Service CLI includes a help command. It shows how the most common types of commands are used, which helps you get familiar with the basics of the CLI.

```
PCA-ADMIN> help
For Most Object Types:
  create <objectType> [(attribute1)="value1"] ... [on <objectType> <instance>]
  delete <objectType> <instance>
  edit <objectType> <instance> (attribute1)="value1" ...
  list <objectType> [fields (attribute1,attribute2)]where [(filterableAttribute1) \
    <filterComparator> "value1" [AND|OR] [(filterableAttribute2) <filterComparator> "value2"]
  show <objectType> <instance>
For Most Object Types with Children:
  add <objectType> <instance> to <objectType> <instance>
  remove <objectType> <instance> from <objectType> <instance>
Other Commands:
  exit
  showallcustomcmds
  showcustomcmds <objectType>
  showobjtypes
```

The easiest way to learn which commands and object types are available, is to use the question mark ("?"). After logging in, you start by entering "?" at the CLI prompt, in order to display the set of base commands.

```
PCA-ADMIN> ?
  add
  clear
  count
  create
  delete
  edit
  [...]
```

You can drill down into the commands, object types and other elements by adding the "?" to see the available parameters at that cursor position.

### Note

Mind the position of the question mark: it is separated from the command by a space. If you omit the space, the CLI displays the parameters allowed at the level of that command, instead of the parameters that may follow *after* the command.

For example, if you want to see which object types you can list, type `list ?` and press Enter. Next, assume that you want to find compute nodes that have not yet been provisioned. To achieve this, you can display a list of compute nodes filtered by their provisioning state. The "?" allows you to navigate through the command parameters, as shown below. Each time you type "?" the CLI displays the parameters you can use at the cursor position. Press the Up arrow key to bring back the part of the command you already typed at the prompt, then add the next part of your command, and type "?" again to display the next set of parameters. When your command is complete, press Enter.

```

PCA-ADMIN> list ?
  AuthorizationGroup
  ComputeNode
  Event
  Fault
  [...]
PCA-ADMIN> list ComputeNode ?
  fields
  limit
  orderby
  where
PCA-ADMIN> list ComputeNode where ?
  id
  provisioningState
  provisioningStateLastChangedTime
  provisioningType
  faultDomain
  [...]
PCA-ADMIN> list ComputeNode where provisioningState ?
  EQ
  NE
  LIKE
  [...]
PCA-ADMIN> list ComputeNode where provisioningState EQ ?
  READYTOPROVISION
  PROVISIONED
PCA-ADMIN> list ComputeNode where provisioningState EQ READYTOPROVISION
Command: list ComputeNode where provisioningState EQ READYTOPROVISION
Status: Success
Time: 2021-06-25 14:04:16,837 UTC
Data:
id          name      provisioningState
--          -
bb940637-9825-4f7c-a5f4-1b49bcf6a5c9  pcacn005  Ready To Provision
76df44a9-6980-4242-a3a2-e1614b3d44d1  pcacn008  Ready To Provision
8fc0d06f-c64a-40ea-8a20-89680f03eb5e  pcacn011  Ready To Provision

```

The Service CLI also provides a form of tab completion. When you start to type a command and press the Tab key, the CLI auto-completes the part it can predict. If more than one possible value remains, you should add at least one more letter and press the Tab key again. The following examples illustrate how the CLI performs tab completion.

- Tab completion with one possible match
 

```

PCA-ADMIN> list Com<Tab>
PCA-ADMIN> list ComputeNode

```
- Tab completion with more than one possible match
 

```

PCA-ADMIN> list Ra<Tab>
PCA-ADMIN> list Rack

PCA-ADMIN> list RackU<Tab>
PCA-ADMIN> list RackUnit

```

## Base and Custom Commands

When you enter the help command or type the question mark ("?") at the PCA-ADMIN> prompt, the CLI returns information about its base commands, such as create, edit, add, remove, delete, list, show, and so on. However, there is another set of less commonly used *custom commands*. You

can display them all as a single list, or only those available for a particular object type. You can use the "?" to navigate through the commands.

```

PCA-ADMIN> showallcustomcmds
Operation Name: <Related Object(s)>
-----
asrClientDisable: ASRPhonehome
asrClientEnable: ASRPhonehome
asrClientRegister: ASRPhonehome
[...]
changeIloMPassword: ComputeNode, ManagementNode
changePassword: ComputeNode, LeafSwitch, ManagementNode, ManagementSwitch, SpineSwitch, User, ZFSAppliance
clearFirstBootError: NetworkConfig
configZFSAdDomain: ZfsAdDomain
configZFSAdWorkgroup: ZfsAdDomain
createAdminAccount:
createUserInGroup: User
deletePlatformImage: PlatformImage
deprovision: ComputeNode
disableVmHighAvailability: PcaSystem
drAddComputeInstance: ComputeInstance
drAddSiteMapping: DrSiteMapping
drConfigCleanupPrimary: DrConfig
[...]
maintenanceLock: ComputeNode
maintenanceUnlock: ComputeNode
migrateVm: ComputeNode
patchCN: PatchRequest
patchEtcD: PatchRequest
patchHost: PatchRequest
patchIloM: PatchRequest
patchKubernetes: PatchRequest
patchMySQL: PatchRequest
patchOCIImages: PatchRequest
patchPlatform: PatchRequest
patchSwitch: PatchRequest
patchVault: PatchRequest
patchZfssa: PatchRequest
[...]
start: CnUpdateManager, ComputeNode, Day0NetworkConfigManager, FaultManager, PurgeManager, ZfsPoolManager
stop: CnUpdateManager, ComputeNode, Day0NetworkConfigManager, FaultManager, PurgeManager, ZfsPoolManager
syncHardwareData:
syncUpstreamUlnMirror: PatchRequest
systemStateLock: PcaSystem
systemStateUnlock: PcaSystem
updateSauronCredentials:
upgradeCN: UpgradeRequest
upgradeEtcD: UpgradeRequest
upgradeFullMN: UpgradeRequest
upgradeHost: UpgradeRequest
upgradeIloM: UpgradeRequest
upgradeKubernetes: UpgradeRequest
upgradeMySQL: UpgradeRequest
upgradePlatform: UpgradeRequest
upgradePreConfig: UpgradeRequest
upgradeSwitch: UpgradeRequest
upgradeVault: UpgradeRequest
upgradeZfssa: UpgradeRequest

PCA-ADMIN> showcustomcmds ?
ASRBundle
ASRPhonehome

```

BackupJob  
CnUpdateManager  
ComputeInstance  
ComputeNode  
Day0NetworkConfigManager  
DrConfig  
DrJob  
DrSiteMapping  
Event  
ExadataNetwork  
FaultDomainInfo  
FaultManager  
Job  
LeafSwitch  
ManagementNode  
ManagementSwitch  
NetworkConfig  
PatchRequest  
PcaSystem  
PlatformImage  
PurgeManager  
SpineSwitch  
UpgradeJob  
UpgradeJobList  
UpgradeRequest  
User  
Vcn  
ZfsAdDomain  
ZFSAppliance  
ZfsPoolManager

PCA-ADMIN> showcustomcmds ComputeNode

provisioningLock  
provisioningUnlock  
maintenanceLock  
maintenanceUnlock  
provision  
deprovision  
migrateVm  
reset  
start  
stop  
changePassword  
changeIllumPassword  
getRunningInstances  
getRunningInstancesCount

# 2

## Hardware Administration

This chapter provides instructions for an administrator to verify the appliance hardware configuration, collect detailed information about the hardware components, and perform standard actions such as starting and stopping a component or provisioning a compute node.

### Displaying Rack Component Details

In the Service Enclave, administrators can obtain details about the appliance and its installed components. This can be done using either the Service Web UI or the Service CLI. The two interfaces display the results in a slightly different way.

### Viewing Appliance Details

The administrator can retrieve certain appliance properties, which may be required when communicating with Oracle, for troubleshooting purposes, or to configure or verify settings.

#### Using the Service Web UI

1. In the PCA Config navigation menu, click Appliance Details.

The detail page contains system properties such as realm, region and domain. The information is read-only, except for the name.

2. To change the rack name and add an optional description, click the Edit button.

The System Details window appears. Enter a Rack Name and Description. Click Save Changes.

The Service CLI provides additional information about hardware discovery and synchronization. Any faults are displayed at the end of the command output.

#### Using the Service CLI

1. Display system parameters and global status with a single command: `show PcaSystem`.

```
PCA-ADMIN> show PcaSystem
Command: show PcaSystem
Status: Success
Time: 2021-08-19 11:20:13,937 UTC
Data:
Id = 934732b6-9f08-4f44-a4fc-fddcdb9967e4
Type = PcaSystem
System Config State = Complete
Initial Hardware Discovery Time = 2021-07-31 00:37:49,763 UTC
Initial Hardware Discovery Status = Resync Success
Initial Hardware Discovery Details = Error retrieving hardware data from the hardware layer.
Resync Hardware Time = 2021-08-10 14:32:13,020 UTC
Resync Hardware Status = Success
Resync Hardware Details = Resync succeeded.
System Name = oraclepca
Domain Name = my.example.com
Availability Domain = AD-1
Realm = 1742XC3024
```

```

Region = oraclepca
ASR Reminder = true
Name = pca
Work State = Normal
FaultIds 1 = id:55f8de1e-ab25-4fc6-b6f4-a9ddd283605b type:Fault
name:PcaSystemInitialHwDiscoveryStatusStatusFault(pca)
FaultIds 2 = id:5c532489-6dad-45e1-a065-6c7649514ce1 type:Fault name:PcaSystemReSyncHwStatusStatusFault(pca)

```

2. Use the edit PcaSystem command to change these parameters:

- description
- name
- ASR reminder (whether or not to display the Oracle Auto Service Request configuration screen when an administrator logs in to the Service Web UI)

Note that the system name and domain name cannot be modified after the initial setup of the appliance.

```

PCA-ADMIN> edit PcaSystem name=myPca description="My Private Cloud" domainName=my.example.com
systemName=mycloud asrReminder=False
Command: edit PcaSystem name=myPca description="My Private Cloud" domainName=my.example.com
systemName=mycloud asrReminder=False
Status: Success
Time: 2021-08-19 11:58:50,442 UTC
JobId: 80cd1fb2-9328-42a0-89e2-7f3196246a28

```

Use the job ID to check the status of your edit command.

```
PCA-ADMIN> show Job id=80cd1fb2-9328-42a0-89e2-7f3196246a28
```

## Using the Rack Units List

The Rack Units list provides an overview of installed hardware components, and lets you drill down into more detailed component information.

### Using the Service Web UI

1. In the PCA Config navigation menu, click Rack Units.

The Rack Units table displays all hardware components installed in the rack and detected by the appliance software. For each component you see its host name, component type, global status information, and the rack unit number where the component is installed.

2. To view more detailed information about a component, click its host name in the table.

The detail pages for switches, storage controllers and management nodes are read-only. For compute nodes there are controls available to execute specific administrator tasks. For more information, see [Performing Compute Node Operations](#).

The Service CLI allows you to list rack units by component type or category. It also includes an option to display information about the rack as a component.

### Using the Service CLI

1. To display a list of all rack units, use the list RackUnit command.

```

PCA-ADMIN> list RackUnit
Command: list RackUnit
Status: Success
Time: 2021-08-19 12:23:55,300 UTC
Data:
id                objtype          name

```

```

--          -----
29f68a0e-4744-4a92-9545-7c48fa365d0a ComputeNode   pcacn001
7a0236f4-b00e-461d-93a0-b22673a18d9c ComputeNode   pcacn003
dc8ae567-b07f-48e0-89bd-e57069c20010 ComputeNode   pcacn002
6fb5ed14-b242-4dd5-842c-532d1c94d43f LeafSwitch    pcaswlf01
279fe518-0dff-40cb-aa3a-fa0966adc946 LeafSwitch    pcaswlf02
a13b5b83-0240-4014-b533-ef4a822e2a4b ManagementNode pcamn01
c24f0d26-8c22-4b2b-b8f5-be98cb25c06e ManagementNode pcamn03
c4e6bcc8-1e4c-44d5-8ca4-0ef9cd04d396 ManagementNode pcamn02
23c35224-d01e-4185-9ec6-22b538f5a5e1 ManagementSwitch pcaswmn01
8c4ecc55-7ac5-4704-bbd2-1023acf7c468 SpineSwitch    pcaswsp01
231276bd-be1f-454f-923f-ffc09f68c294 SpineSwitch    pcaswsp02
379690d6-4097-4637-9564-28ae890a20d2 ZFSAppliance   pcasn02
ca637f6f-5269-48be-81b9-ceda76a90daf ZFSAppliance   pcasn01

```

2. To display only rack units of a specific type, use one of these commands instead:

- list ManagementNode: displays a list of management nodes
- list LeafSwitch: displays a list of leaf switches
- list SpineSwitch: displays a list of spine switches
- list ManagementSwitch: displays a list of 1Gbit management switches
- list ZFSAppliance: displays a list of ZFS storage controllers
- list ComputeNode: displays a list of compute nodes
- list Rack: displays a list of racks that are part of the environment

Example:

```

PCA-ADMIN> list ManagementNode
Command: list ManagementNode
Status: Success
Time: 2021-08-19 12:34:09,429 UTC
Data:
  id          name
  --          ---
a13b5b83-0240-4014-b533-ef4a822e2a4b pcamn01
c24f0d26-8c22-4b2b-b8f5-be98cb25c06e pcamn03
c4e6bcc8-1e4c-44d5-8ca4-0ef9cd04d396 pcamn02

```

3. To view more detailed information about a component, use the show command with the component type and its name or ID.

4. Syntax (entered on a single line):

```

show
RackUnit|ComputeNode|LeafSwitch|ManagementNode|ManagementSwitch|Rack|RackUnit|SpineSwitch|ZFSAppliance
id=<component_id> OR name=<component_name>

```

Examples:

```

PCA-ADMIN> show SpineSwitch id=8c4ecc55-7ac5-4704-bbd2-1023acf7c468
Command: show SpineSwitch id=8c4ecc55-7ac5-4704-bbd2-1023acf7c468
Status: Success
Time: 2021-08-19 12:50:39,570 UTC
Data:
Id = 8c4ecc55-7ac5-4704-bbd2-1023acf7c468
Type = SpineSwitch
HW Id = FDO24290PQC
MAC Address = 3c:13:cc:bd:3a:7c
Ip Address = 100.96.2.20
Hostname = pcaswsp01

```

```

Firmware Version = 9.3(2)
Serial Number = FDO24290PQC
State = OK
Rack Elevation = 22
Validation State = Validated
RackId = id:dba2962d-c477-4a32-bdff-a3a256bf7972 type:Rack name:PCA X9-2 Base1
Name = pcaswsp01
Work State = Normal

```

```

PCA-ADMIN> show RackUnit name=pcamn02
Command: show RackUnit name=pcamn02
Status: Success
Time: 2021-08-19 12:48:51,852 UTC
Data:
Id = c4e6bcc8-1e4c-44d5-8ca4-0ef9cd04d396
Type = ManagementNode
HW Id = 1749XC302R
MAC Address = 00:10:e0:da:cb:7c
Ip Address = 100.96.2.34
Hostname = pcamn02
Firmware Version = 3.0.1
Serial Number = 1749XC302R
State = running
Rack Elevation = 6
Validation State = Validated
RackId = id:dba2962d-c477-4a32-bdff-a3a256bf7972 type:Rack name:PCA X9-2 Base1
Name = pcamn02
Work State = Normal

```

## Changing Passwords for Hardware Components

You can change the password for a compute node, leaf switch, management node, management switch, spine switch, or ZFS appliance component using the Service CLI. You can also change the ILOM password for a compute node or a management node.

### Important

The following password rules apply:

- Passwords for compute nodes, leaf switches, management nodes, management switches, or spine switches must contain at least 8 but no more than 20 characters.
- Passwords for ZFS appliance or ILOMs must contain at least 8 but no more than 16 characters.
- All passwords must contain at least 1 uppercase letter (A-Z), 1 lowercase letter (a-z), 1 digit (0-9), and 1 of the following symbols: @\$!%\*#&.
- Enclose the password in double quotes "*password*" when entering a password.

### Using the Service CLI

To view the components for which you can change passwords, use the `changepassword ?` or the `changeilompassword ?` command.

```

PCA-ADMIN> changepassword ?
ComputeNode
LeafSwitch

```

```
ManagementNode
ManagementSwitch
SpineSwitch
ZFSAppliance
```

```
PCA-ADMIN> changeIomPassword ?
ComputeNode
ManagementNode
```

To change the password for a hardware component, use the `changePassword` command.

Syntax (entered on a single line):

```
changePassword
ComputeNode|LeafSwitch|ManagementNode|ManagementSwitch|SpineSwitch|ZFSAppliance
id=<component_id> OR name=<component_name>
password="<new_password>" confirmPassword="<repeat_new_password>"
```

Example:

```
PCA-ADMIN> changePassword id=21ad5b60-d30d-4a95-b39f-5bf152005f0f
password=***** confirmPassword=*****
```

```
Status: Success
Time: 2022-08-16 17:13:22,674 UTC
JobId: fe772781-d0af-47cc-af87-2059f8e70b63
```

To change the ILOM password for a compute node or management node, use the `changeIomPassword` command.

Syntax (entered on a single line):

```
changeIomPassword ComputeNode|ManagementNode
id=<component_id> OR name=<component_name>
password="<new_password>" confirmPassword="<repeat_new_password>"
```

Example:

```
PCA-ADMIN> changeIomPassword
id=21ad5b60-d30d-4a95-b39f-5bf152005f0f password="*****" confirmPassword="*****"
```

```
Status: Success
Time: 2022-08-16 17:13:22,674 UTC
JobId: fe772781-d0af-47cc-af87-2059f8e70b63
```

## Checking Component Health

You can get a quick health check for compute nodes and management nodes by using the Service CLI `getComputeIomHealth` and `getMgmtIomHealth` commands. These commands return data from ILOM that shows, for example, the component health is OK, service is required, or faults need to be addressed.

See also [Viewing Admin Service Health Data](#) for information about the `list fault` and `show fault` commands.

### Using the Service CLI

To get basic health information from ILOM for compute nodes and management nodes, use the following commands.

Compute Nodes

```
PCA-ADMIN> getcomputeIomHealth
```

```
Status: Success
Time: 2022-08-16 11:24:42,961 EDT
Data:
Health Nodes 1 - macaddr = a8:69:8c:05:e8:c7
Health Nodes 1 - health = OK
Health Nodes 1 - time checked = 22-07-21T20:06:34
Health Nodes 2 - macaddr = a8:69:8c:05:e8:73
Health Nodes 2 - health = OK
Health Nodes 2 - time checked = 22-07-21T20:06:34
Health Nodes 3 - macaddr = 00:10:e0:fe:82:1b
Health Nodes 3 - health = OK
Health Nodes 3 - time checked = 22-07-21T20:06:34
```

### Management Nodes

```
PCA-ADMIN> getmgmtIomHealth
```

```
Status: Success
Time: 2022-08-16 11:25:19,486 EDT
Data:
Health Nodes 1 - macaddr = A8:69:8C:05:EC:C7
Health Nodes 1 - health = OK
Health Nodes 1 - time checked = 22-07-15T18:50:50
Health Nodes 2 - macaddr = A8:69:8C:05:EA:AB
Health Nodes 2 - health = OK
Health Nodes 2 - time checked = 22-07-15T18:50:50
Health Nodes 3 - macaddr = A8:69:8C:06:0F:A3
Health Nodes 3 - health = Service Required
Health Nodes 3 - time checked = 22-07-15T18:50:50
Health Nodes 3 - node Faults 1 - messageId = SPENV-8000-A7
Health Nodes 3 - node Faults 1 - fault type = fault
Health Nodes 3 - node Faults 1 - classId = fault.chassis.device.fan.fail
Health Nodes 3 - node Faults 1 - uuid = c6986589-07b5-ceb0-edfc-a8535eb2f442/115ed970-a382-668c-a50a-9e854dc8479f
Health Nodes 3 - node Faults 1 - time reported = 2022-07-14T22:24:36+0000
Health Nodes 3 - node Faults 1 - severity = Major
Health Nodes 3 - node Faults 1 - description = Fan module has a fan that is rotating too slowly.
Health Nodes 3 - node Faults 1 - action = Please refer to the associat
...
```

## Performing Compute Node Operations

From the Rack Units list of the Service Web UI, an administrator can execute certain operations on hardware components. These operations can be accessed from the Actions menu, which is the button with three vertical dots on the right hand side of each table row. In practice, only the View Details and Copy ID operations are available for all component types.

When compute nodes are in the discovery state or coming up, their status is 'Failed' until the hardware process transitions them to 'Ready to Provision'. This process typically takes under five minutes. If the failed state persists, use the Service CLI command `list ComputeNode` to determine the provisioning state of the compute nodes and take appropriate action.

For compute nodes, several other operations are available, either from the Actions menu or from the compute node detail page. Those operations are described in detail in this section, including the equivalent steps in the Service CLI.

## Provisioning a Compute Node

Before a compute node can be used to host your compute instances, it must be provisioned by an administrator. The appliance software detects the compute nodes that are installed in the rack and cabled to the switches, meaning they appear in the Rack Units list as *Ready to Provision*. You can provision them from the Service Web UI or Service CLI.

### Using the Service Web UI

1. In the navigation menu, click Rack Units.
2. In the Rack Units table, click the host name of the compute node you want to provision.  
The compute node detail page appears.
3. In the top-right corner of the page, click Controls and select the Provision command.

### Using the Service CLI

1. Display the list of compute nodes.

Copy the ID of the compute node you want to provision.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-20 08:53:56,681 UTC
Data:
id          name          provisioningState provisioningType
--          -
29f68a0e-4744-4a92-9545-7c48fa365d0a pcacn001 Ready to Provision Unspecified
7a0236f4-b00e-461d-93a0-b22673a18d9c pcacn003 Ready to Provision Unspecified
dc8ae567-b07f-48e0-89bd-e57069c20010 pcacn002 Ready to Provision Unspecified
```

2. Provision the compute node with this command:

```
PCA-ADMIN> provision id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Command: provision id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Status: Success
Time: 2021-08-20 11:35:40,152 UTC
JobId: ea93cac4-4430-4663-aafd-d70701593fb2
```

Use the job ID to check the status of your provision command.

```
PCA-ADMIN> show Job id=ea93cac4-4430-4663-aafd-d70701593fb2
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

3. Repeat the provision command for any other compute nodes you want to provision at this time.
4. Confirm that the compute nodes have been provisioned.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-20 11:38:29,509 UTC
Data:
id          name          provisioningState provisioningType
--          -
29f68a0e-4744-4a92-9545-7c48fa365d0a pcacn001 Provisioned      KVM
```

```
7a0236f4-b00e-461d-93a0-b22673a18d9c pcacn003 Provisioned KVM
dc8ae567-b07f-48e0-89bd-e57069c20010 pcacn002 Provisioned KVM
```

## Disabling Compute Node Provisioning

Several compute node operations can only be performed on condition that provisioning has been disabled. This section explains how to impose and release a provisioning lock.

### Using the Service Web UI

1. In the navigation menu, click Rack Units.
2. In the Rack Units table, click the host name of the compute node you want to make changes to.

The compute node detail page appears.

3. In the top-right corner of the page, click Controls and select the Provisioning Lock command.

When the confirmation window appears, click Lock to proceed.

After successful completion, the Compute Node Information tab shows Provisioning Locked = Yes.

4. To release the provisioning lock, click Controls and select the Provisioning Unlock command.

When the confirmation window appears, click Unlock to proceed.

After successful completion, the Compute Node Information tab shows Provisioning Locked = No.

### Using the Service CLI

1. Display the list of compute nodes.

Copy the ID of the compute node for which you want to disable provisioning operations.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-23 09:25:56,307 UTC
Data:
id          name      provisioningState provisioningType
--          -
3e62bf25-a26c-407e-ab8b-df01a4ad98b6 pcacn002 Provisioned     KVM
f7b8356b-052f-4911-babb-447e6ab9c78d pcacn003 Provisioned     KVM
4e06ebdf-faed-484e-996d-d77af786f123 pcacn001 Provisioned     KVM
```

2. Set a provisioning lock on the compute node.

```
PCA-ADMIN> provisioningLock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: provisioningLock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 09:29:46,568 UTC
JobId: 6ee78c8a-e227-4d31-a770-9b9c96085f3f
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=6ee78c8a-e227-4d31-a770-9b9c96085f3f
Command: show Job id=6ee78c8a-e227-4d31-a770-9b9c96085f3f
[...]
Done = true
```

```
Name = MODIFY_TYPE  
Run State = Succeeded
```

3. When the job has completed, confirm that the compute node is under provisioning lock.

```
PCA-ADMIN> show ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d  
[...]  
Provisioning State = Provisioned  
[...]  
Provisioning Locked = true  
Maintenance Locked = false
```

All provisioning operations are now disabled until the lock is released.

4. To release the provisioning lock, use this command:

```
PCA-ADMIN> provisioningUnlock id=f7b8356b-052f-4911-babb-447e6ab9c78d  
Command: provisioningUnlock id=f7b8356b-052f-4911-babb-447e6ab9c78d  
Status: Success  
Time: 2021-08-23 09:44:58,531 UTC  
JobId: 523892e8-c2d4-403c-9620-2f3e94015b46
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=523892e8-c2d4-403c-9620-2f3e94015b46  
[...]  
Done = true  
Name = MODIFY_TYPE  
Run State = Succeeded
```

5. When the job has completed, confirm that the provisioning lock has been released.

```
PCA-ADMIN> show ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d  
[...]  
Provisioning State = Provisioned  
[...]  
Provisioning Locked = false  
Maintenance Locked = false
```

## Locking a Compute Node for Maintenance

For maintenance operations, compute nodes must be placed in maintenance mode. This section explains how to impose and release a maintenance lock. Before you can lock a compute node for maintenance, you must disable provisioning first. Maintenance operations can only be performed if the compute node has no running compute instances.

**⚠ Caution**

Depending on the high-availability configuration of the Compute service, automatic instance migrations can prevent you from successfully locking a compute node. See [Configuring the Compute Service for High Availability](#). This situation is more likely to occur when available compute capacity is limited.

- Instance recovery or migration operations after a compute node outage can cause a maintenance lock to fail. Compute nodes involved in instance migrations will reject the maintenance lock until the migrations are complete.
- Displaced instances could be migrated back to their original fault domain when a compute node maintenance lock is released. A compute node from where a displaced instance is migrated back will reject the maintenance lock until the migration is complete.
- Migrating an instance typically takes no more than 30 seconds. However, large instances and heavy workloads increase the time required.
- In the event that an instance gets stuck in moving state and migration fails to complete, its host compute node cannot be locked for maintenance. Contact Oracle for assistance.

**Using the Service Web UI**

1. Ensure that provisioning has been disabled on the compute node.  
See [Disabling Compute Node Provisioning](#).
2. Ensure that the compute node has no active instances. They must be migrated or shut down.  
See [Migrating Instances from a Compute Node](#).
3. In the navigation menu, click Rack Units.
4. In the Rack Units table, click the host name of the compute node that requires maintenance.  
The compute node detail page appears.
5. In the top-right corner of the page, click Controls and select the Maintenance Lock command.  
When the confirmation window appears, click Lock to proceed.  
After successful completion, the Compute Node Information tab shows Maintenance Locked = Yes.
6. To release the maintenance lock, click Controls and select the Maintenance Unlock command.  
When the confirmation window appears, click Unlock to proceed.  
After successful completion, the Compute Node Information tab shows Maintenance Locked = No.

**Using the Service CLI**

1. Display the list of compute nodes.  
Copy the ID of the compute node that requires maintenance.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-23 09:25:56,307 UTC
Data:
  id                name      provisioningState provisioningType
  --                -
  3e62bf25-a26c-407e-ab8b-df01a4ad98b6 pcacn002 Provisioned     KVM
  f7b8356b-052f-4911-babb-447e6ab9c78d pcacn003 Provisioned     KVM
  4e06ebdf-faed-484e-996d-d77af786f123 pcacn001 Provisioned     KVM
```

2. Ensure that provisioning has been disabled on the compute node.

See [Disabling Compute Node Provisioning](#).

3. Lock the compute node for maintenance.

```
PCA-ADMIN> maintenanceLock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: maintenanceLock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 09:56:05,443 UTC
JobId: e46f6603-2af2-4df4-a0db-b15156491f88
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=e46f6603-2af2-4df4-a0db-b15156491f88
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

4. When the job has completed, confirm that the compute node has been locked for maintenance.

```
PCA-ADMIN> show ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
[...]
Provisioning State = Provisioned
[...]
Provisioning Locked = true
Maintenance Locked = true
```

The compute node is now ready for maintenance.

5. To release the maintenance lock, use this command:

```
PCA-ADMIN> maintenanceUnlock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: maintenanceUnlock id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 10:00:53,902 UTC
JobId: 625af20e-4b49-4201-879f-41d4405314c7
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=625af20e-4b49-4201-879f-41d4405314c7
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

6. When the job has completed, confirm that the provisioning lock has been released.

```
PCA-ADMIN> show ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
[...]
Provisioning State = Provisioned
[...]
```

Provisioning Locked = true  
Maintenance Locked = false

## Migrating Instances from a Compute Node

Some compute node operations, such as some maintenance operations, can only be performed if the compute node has no running compute instances. Administrators can migrate all running instances away from a compute node, also known as evacuating the compute node.

By default, if enough resources are available, running instances are live migrated to other compute nodes in the same fault domain.

### ! Important

Before you perform a compute node evacuation, check what the behavior will be for any instances that cannot be live migrated to another compute node in the same fault domain.

The remainder of this topic and [Configuring the Compute Service for High Availability](#) describe how to check settings and how instances are handled for different settings.

### i Note

Live migration between different types of compute nodes is not supported. For example, you can't migrate instances from an Oracle Server X10 to an Oracle Server X11.

See [Viewing and Setting Compute Service Configuration](#) to check whether strict fault domain enforcement is set.

When strict fault domain enforcement is disabled (Strict FD is set to Disabled in the Service Web UI or `Strict FD Enabled` is `false` in the Service CLI), instances that cannot be live migrated to another compute node in the same fault domain are live migrated to a different fault domain if enough resources are available in that fault domain.

When strict fault domain enforcement is enabled (Strict FD is set to Enabled in the Service Web UI or `Strict FD Enabled` is `true` in the Service CLI), instances that cannot be live migrated to another compute node in the same fault domain do not migrate; those instances are still running in the compute node that you want to evacuate.

Enable or disable strict fault domain enforcement to set whether instances that cannot live migrate to other compute nodes in the same fault domain will be live migrated to a different fault domain or still running in the same compute node after you attempt to evacuate the compute node.

If some instances cannot be live migrated, either because the current fault domain is not able to accommodate them and strict fault domain enforcement is enabled, or because strict fault domain enforcement is disabled but other fault domains also cannot accommodate the instances, then you can re-run the migrate operation with the `force` option specified. When the `force` option is specified, the Compute service soft stops any instances that fail to live migrate, allowing the evacuation to proceed.

**Restart stopped instances.** If instances were stopped by the Compute service (not manually stopped by an administrator) and you want them to be automatically restored to running when resources become available, check that the Auto Recovery property of the Compute service is

enabled and the instance availability recovery action is set to `RESTORE_INSTANCE`. See [Viewing and Setting Compute Service Configuration](#) and [Configuring the Recovery State for a Stopped Instance](#).

Instances can be stopped by the Compute service if the force option is used when an administrator evacuates a compute node, or in response to an unplanned compute node outage. You can change the Auto Recovery setting at any time before or after resources become available after an administrative maintenance or unplanned outage to restart instances that were stopped by the Compute service. If the instance availability recovery action is set to `STOP_INSTANCE`, the instance remains stopped even though the Auto Recovery property is enabled. If the instance availability recovery action is later changed to `RESTORE_INSTANCE`, a subsequent Auto Recovery pass will restart the instance.

**Return relocated instances.** If instances are live migrated to a different fault domain (displaced), and you want them returned to their selected fault domain (the fault domain that is specified in the instance configuration) when resources become available, check that the Auto Resolve property of the Compute service is enabled. See [Viewing and Setting Compute Service Configuration](#) and [Compute Service Configuration Commands](#). You can set the Auto Resolve property at any time before or after the compute node evacuation completes to relocate any displaced instances.

Use the following procedures to perform the migrate operation.

### Compute Node Evacuation: Before You Begin

1. Check fault domain and compute node resources. See [Viewing CPU and Memory Usage By Fault Domain](#). Based on this information, decide whether to do any of the following:
  - Terminate instances that are no longer needed.
  - Reconfigure some instances to use fewer resources. For example, specify a different shape.
  - Reconfigure some instances to specify a different fault domain.
  - Stop some instances while you perform the compute node evacuation.
  - Shutdown non migratable instances. See [Manually Shut Down a Non Migratable Instance](#).
  - Specify the force option on the migration operation to soft stop any instances that cannot be live migrated. See the discussion above of instance availability recovery action and Auto Recovery configuration.
2. Disable provisioning on the compute node. See [Disabling Compute Node Provisioning](#).

### Using the Service Web UI

1. In the navigation menu, click Rack Units.
2. In the Rack Units table, find the host name of the compute node that you want to evacuate. Click the Actions menu for that host, and click the Migrate All Vms option.

Alternatively, in the Rack Units table, click the host name of the compute node that you want to evacuate to display the details page for that compute node. Click the Controls menu, and click the Migrate All Vms option.

3. On the Confirm Migrating VMs dialog, choose whether to force stop any instances that cannot be migrated.

By default, the force stop option is not enabled, and instances that cannot be migrated will still be running on the node after the migrate operation completes. To force stop instances that cannot be migrated, enable the force stop option in the Confirm Migrating VMs dialog.

4. On the Confirm Migrating VMs dialog, click the Migrate button.

The Compute service live migrates the running instances to other compute nodes if enough resources are available and High Availability settings are configured to allow it. If the Force option was specified, any instances that could not be migrated are soft stopped. If any instances could not be migrated and Force was not specified, those instances remain running in the compute node that you are attempting to evacuate.

### Using the Service CLI

1. Display the list of compute nodes.

Copy the ID of the compute node that you that you want to evacuate.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-23 09:25:56,307 UTC
Data:
  id                name                provisioningState  provisioningType
  ---                -
  3e62bf25-a26c-407e-ab8b-df01a4ad98b6  pcacn002  Provisioned      KVM
  f7b8356b-052f-4911-babb-447e6ab9c78d  pcacn003  Provisioned      KVM
  4e06ebdf-faed-484e-996d-d77af786f123  pcacn001  Provisioned      KVM
```

2. Use the `migrateVm` command to live migrate all running compute instances off the compute node. To soft stop any instances that fail to migrate, set the `force` option:

```
PCA-ADMIN> migrateVm id=7a0236f4-b00e-461d-93a0-b22673a18d9c force=true
Command: migrateVm id=7a0236f4-b00e-461d-93a0-b22673a18d9c force=true
Status: Running
Time: 2021-08-20 10:37:05,781 UTC
JobId: 6f1e94bc-7d5b-4002-ada9-7d4b504a2599
```

The Compute service live migrates the running instances to other compute nodes if enough resources are available and High Availability settings are configured to allow it. If `force=true` was specified, any instances that could not be migrated are soft stopped. If any instances could not be migrated and `force=true` was not specified, those instances remain running in the compute node that you are attempting to evacuate.

Use the job ID to check the status of the `migrateVm` command.

```
PCA-ADMIN> show Job id=6f1e94bc-7d5b-4002-ada9-7d4b504a2599
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

## Manually Shut Down a Non Migratable Instance

While it's possible to specify the `force` option on the `vmMigrate` operation to soft stop any instances that cannot be live migrated, the best practice is to gracefully shut down non migratable instances before migration so any workloads running on the VM will be in a good state.

### Using the Service CLI

1. Display the list of non migratable instances.

Copy the ID of the running instances, so you can shut them down.

```

PCA-ADMIN> getNonMigratableInstances
Command: getNonMigratableInstances
Status: Success
Time: 2024-09-26 13:21:51,727 UTC
Data:
  id          Display Name  Compute Node Id  Domain State
  --          -
  ocid1.instance.unique_ID  instance202  CN_ID    running
  ocid1.instance.unique_ID  kqh027      CN_ID    shut off

```

2. Shut down the running instances.

See the *Working with Instances* section in [Compute Instance Deployment](#).

## Configuring the Compute Service for High Availability

[Migrating Instances from a Compute Node](#) describes how to evacuate a compute node for planned maintenance. If possible, the Compute service live migrates running instances to other compute nodes in the same fault domain.

In the case of a compute node unplanned outage, the Compute service stops the instances, and if the outage persists, attempts to evacuate the compute node using reboot migration.

The following sections describe how you can set high availability configuration to control how the Compute service handles an unplanned outage.

## Using Instance and Compute Service High Availability Configuration

The following sections describe how to use high availability configuration to manage outcomes for different types of compute node outages. Instance availability recovery action is the only high availability configuration that is set for each instance. All other high availability configuration is set on the Compute service and affects all instances.

The *selected* fault domain is the fault domain that is specified in the instance configuration. A *displaced* instance is in a fault domain that is not its selected fault domain.

### Planned Maintenance Outage

See [Migrating Instances from a Compute Node](#) for information about how to evacuate a compute node. If possible, the Compute service live migrates running instances to other compute nodes in the same fault domain.

"Migrating Instances from a Compute Node" also describes how to use the instance availability recovery action (set on each instance), and the Auto Recovery and Auto Resolve properties of the Compute service when performing a compute node evacuation.

### Unplanned Outage

The Compute service attempts to stop instances and reboot migrate the instances under the following compute node outage conditions:

- Power down from HW status
- Inability to reach the compute node data network

A compute node could experience an outage where the Compute service cannot migrate the instances. For example, if the Compute service cannot reach the compute node at all, then the Compute service cannot stop and reboot migrate the instances.

### Unplanned Outage Less Than Five Minutes

In an unplanned outage, the Compute service stops the affected instances. If the outage lasts less than five minutes, by default the Compute service attempts to restart instances that were running before the outage. Actual behavior depends on how the instances and the Compute service are configured. The following decision flow describes how you can control this behavior.

Do you want the Compute service to attempt to restart instances that were running prior to the outage? This is the default.

- Yes. Check that Auto Recovery is enabled and the instance availability recovery action is set to `RESTORE_INSTANCE`. See [Configuring the Recovery State for a Stopped Instance](#).

If some instances can no longer be accommodated in their selected fault domain, Auto Recovery will continue to poll and attempt to restart the instances. See also `getForcedStoppedInstances`.

If the instance availability recovery action is set to `STOP_INSTANCE`, the instance will remain stopped, even if Auto Recovery is enabled.

- No. Disable Auto Recovery. Instances that had been running prior to the outage will remain stopped.

The instance availability recovery action setting and Auto Recovery setting can be changed at any time, and the changes will be effective at the next polling time.

### Unplanned Outage More Than Five Minutes

In an unplanned outage, the Compute service stops the affected instances. If the outage lasts more than five minutes, by default the Compute service attempts to reboot migrate (cold migrate) instances off the compute node. Instances that cannot be accommodated on other compute nodes in the same fault domain are reboot migrated to other fault domains. Actual behavior depends on how the Compute service is configured. The following decision flow describes how you can control this behavior.

Do you want running instances to be reboot migrated? Reboot migration is stopping and starting each running instance on a given compute node. See also "Compute Instance Availability" in "High Availability" in the [Architecture and Design](#) chapter of *Oracle Private Cloud Appliance Concepts Guide*.

- Yes. Check that VM High Availability is enabled.

If some instances cannot be accommodated on another compute node in the same fault domain, do you want those instances to be reboot migrated to a different fault domain?

- Yes. Check that Strict FD is disabled. Instances that cannot be accommodated in any fault domain remain stopped by the Compute service.

After reboot migration, do you want instances that are running in a fault domain that is not their selected fault domain to be automatically live migrated to their selected fault domain when resources become available?

\* Yes. Check that Auto Resolve is enabled. See also `getDisplacedInstances`.

\* No. Disable Auto Resolve.

- No. Enable Strict FD. Instances that were running prior to the outage and cannot be migrated to another compute node in the current fault domain remain stopped by the Compute service.

- No. Disable VM High Availability. Instances that were running prior to the outage are stopped by the Compute service.

Do you want instances that were stopped by the Compute service to be automatically restored to running in their selected fault domain? If yes, check that Auto Recovery is enabled and the instance availability recovery action is set to RESTORE\_INSTANCE. See [Configuring the Recovery State for a Stopped Instance](#).

## Viewing and Setting Compute Service Configuration

For information about how these configuration settings work, see [Compute Service Configuration Commands](#).

### Using the Service Web UI

On the navigation menu, click FD Instances and then click Compute Service Detail.

The Compute Service Information page shows the current settings for Auto Recovery, Auto Resolve Displaced Instances, VM High Availability, and Strict FD. All of these settings are enabled by default except for Strict FD, which is disabled by default. By default, fault domain placement is not strictly enforced when the Compute service migrates instances.

Use the Controls menu on the Compute Service Information page to change the values of these configuration settings between Enabled and Disabled.

### Using the Service CLI

Use the `show computeservice` command to show the current Compute service configuration settings. In the following example, the default values are set for the four high availability configuration settings: Auto Recovery Action Enabled, Auto-Resolve Displaced Instances Enabled, VM High Availability Enabled, and Strict FD Enabled. All of these settings are true by default except for Strict FD Enabled, which is false by default.

```
PCA-ADMIN> show computeservice
Command: show computeservice
Status: Success
Time: 2023-04-17 20:37:42,296 UTC
Data:
Id = unique_ID
Type = ComputeService
total CN cpu usage percent = 23.3
total CN memory usage percent = 16.2
Auto Recovery Action Enabled = true
Auto-Resolve Displaced Instances Enabled = true
VM High Availability Enabled = true
Strict FD Enabled = false
Name = Compute Service
Work State = Normal
```

To change these settings, use the commands in the following list. The `showcustomcmds computeservice` command lists all high availability configuration commands in the Compute service.

```
PCA-ADMIN> showcustomcmds computeservice
enableAutoRecoveryAction
disableAutoRecoveryAction
enableAutoResolveDisplacedInstances
disableAutoResolveDisplacedInstances
enableVmHighAvailability
disableVmHighAvailability
```

```
enableStrictFD
disableStrictFD
getForcedStoppedInstances
getDisplacedInstances
```

For example, to disable Auto Recovery Action Enabled, run the `disableAutoRecoveryAction` command. To enable strict fault domain enforcement, run the `enableStrictFD` command.

## Compute Service Configuration Commands

This section describes the behavior of the high availability configuration settings in the Compute service. The Service CLI commands are shown in the list in this section. To access the equivalent Service Web UI settings, click the navigation menu and click FD Instances. See [Viewing and Setting Compute Service Configuration](#).

In these descriptions, the *selected* fault domain is the fault domain that is specified in the instance configuration. A *displaced* instance is in a fault domain that is not its selected fault domain.

### **enableAutoRecoveryAction**

Enables the automatic restart of instances that were stopped by the Compute service. This is the default. If the instance availability recovery action is set to `RESTORE_INSTANCE`, this command causes instances that were stopped by the Compute service to be automatically restarted in their selected fault domain when resources are available. See also [Configuring the Recovery State for a Stopped Instance](#) and `getForcedStoppedInstances`.

Instances could have been stopped by the Compute service for the following reasons:

- As a result of specifying the force option on a migrate all operation and some instances were not able to be migrated.  
See [Migrating Instances from a Compute Node](#).
- As a result of an unplanned compute node outage.  
See [Unplanned Outage](#).

You can set this Auto Recovery property at any time before or after an administrative maintenance outage or an unplanned outage to restart instances that were stopped by the Compute service. If the instance availability recovery action is set to `STOP_INSTANCE`, the instance remains stopped even though the Auto Recovery property is enabled. If the instance availability recovery action is later changed to `RESTORE_INSTANCE`, a subsequent Auto Recovery pass will restart the instance.

### **disableAutoRecoveryAction**

Disables the automatic restart of stopped instances. Instances that were stopped by the Compute service are not automatically restarted when resources are available.

### **enableAutoResolveDisplacedInstances**

Enables the return of running instances to their selected fault domain. This is the default. If instances were moved to a different fault domain (displaced) during compute node evacuation, this command enables those instances to be automatically live migrated to their selected fault domain once sufficient resources are available in that fault domain. See also `getDisplacedInstances`.

You can set this Auto Resolve configuration at any time before or after an outage to relocate any displaced instances.

Instances that are stopped are not migrated.

**disableAutoResolveDisplacedInstances**

Disables the return of instances to their selected fault domain. Instances that were moved to a different fault domain during compute node evacuation remain in the fault domain to which they were moved.

**enableVmHighAvailability**

Enables High Availability (reboot migration) off of an unreachable compute node. This is the default.

**disableVmHighAvailability**

Disables reboot migration.

**enableStrictFD**

Enables strict fault domain enforcement. During compute node evacuation, any instance that cannot be moved to a different compute node in the same fault domain is stopped if the force option was specified. If the force option was not specified, the migrate operation fails.

**disableStrictFD**

Disables strict fault domain enforcement. This is the default. During compute node evacuation, any instance that cannot be moved to a different compute node in the same fault domain is moved to a different fault domain. This move to a different fault domain is temporary if the Auto Resolve property of the Compute service is enabled: If Auto Resolve is enabled, then when resources become available, the moved instances are live migrated back to their selected fault domain. See also `getDisplacedInstances`.

**getForcedStoppedInstances**

Lists all instances that were stopped via the use of the force option on the migrate operation or that were stopped by the Compute service in response to an unplanned outage.

```
PCA-ADMIN> getForcedStoppedInstances
Command: getForcedStoppedInstances
Status: Success
Time: 2023-04-17 20:53:51,410 UTC
Data:
id          displayName compartmentId
--          -
ocid1.instance.unique_ID inst-name  ocid1.compartment.unique_ID
```

In the Service Web UI, click the navigation menu, click FD Instances, and then click Forced Stopped Instances. Use the Actions menu to copy the OCIDs.

**getDisplacedInstances**

Lists instances that are currently running in a fault domain that is not their selected fault domain. Instances that are not running are not shown.

In the following example, running instances are being migrated away from fault domain 1. One instance has been placed in fault domain 2 and one has been placed in fault domain 3.

```
PCA-ADMIN> getDisplacedInstances
Command: getDisplacedInstances
Status: Success
Time: 2023-04-18 23:20:41,484 UTC
Data:
id          displayName compartmentId          faultDomain  faultDomainSelected
--          -
ocid1.instance.unique_ID inst-name  ocid1.compartment.unique_ID  FAULT-DOMAIN-3  FAULT-DOMAIN-1
ocid1.instance.unique_ID inst-name  ocid1.compartment.unique_ID  FAULT-DOMAIN-2  FAULT-DOMAIN-1
```

In the Service Web UI, click the navigation menu, click FD Instances, and then click Displaced Instances. Use the Actions menu to copy the OCIDs.

## Configuring the Recovery State for a Stopped Instance

If the Compute service stopped an instance, you can configure how that stopped instance will be treated when resources are again available by setting the instance availability recovery action and the Auto Recovery property of the Compute service.

See the description of the `enableAutoRecoveryAction` command in [Compute Service Configuration Commands](#) for reasons that an instance can be stopped by the Compute service. See also the descriptions of `disableAutoRecoveryAction` and `getForcedStoppedInstances`.

During instance launch or in a subsequent instance update, set the instance recovery action in the instance availability configuration.

In the Compute Web UI, see the "Availability configuration" section in the dialog to create or edit an instance or create or edit an instance configuration. To restart instances that were stopped by the Compute service, check the box labeled "Restore instance lifecycle state after infrastructure maintenance". This is the default. To keep stopped instances stopped, uncheck the "Restore instance" box.

In the OCI CLI, use the `--availability-config` option or the `availabilityConfig` property in the `compute instance launch` or `update` command or the instance configuration `create` or `update` command. Set the `recoveryAction` to `RESTORE_INSTANCE` or `STOP_INSTANCE`. The default behavior is `RESTORE_INSTANCE`.

```
"availabilityConfig": {"recoveryAction": "STOP_INSTANCE"}
```

## Enabling Strict Fault Domain Enforcement

To enable strict fault domain enforcement, do one of the following:

- In the Service Web UI, click the navigation menu, click FD Instances, and click Compute Service Detail. On the Compute Service Information page, click the Controls menu, and click Enable Strict FD.
- In the Service CLI, run the `enableStrictFD` command.

For more information about the effect of fault domain enforcement, see [Compute Service Configuration Commands](#).

In case the current fault domain does not have enough resources to accommodate all instances that need to be migrated, do the following:

- If you are performing a planned compute node evacuation, specify the `force` option on the migration operation to stop the instances in their current fault domain.
- Run the `enableAutoRecoveryAction` command or select Enable Auto Recovery in the Service Web UI.
- Ensure that the instance availability recovery action for each instance is set to `RESTORE_INSTANCE`, which is the default. See [Configuring the Recovery State for a Stopped Instance](#).

See the example in [Migrating Instances from a Compute Node](#).

## Starting, Resetting or Stopping a Compute Node

The Service Enclave allows administrators to send start, reboot and shutdown signals to the compute nodes.

## Using the Service Web UI

1. Make sure that the compute node is locked for maintenance.  
See [Locking a Compute Node for Maintenance](#).
2. In the navigation menu, click Rack Units.
3. In the Rack Units table, locate the compute node you want to start, reset or stop.
4. Click the Action menu (three vertical dots) and select the appropriate action: Start, Reset, or Stop.
5. When the confirmation window appears, click the appropriate action button to proceed.  
A pop-up window appears for a few seconds to confirm that the compute node is starting, stopping, or restarting.
6. When the compute node is up and running again, release the maintenance and provisioning locks.

## Using the Service CLI

1. Display the list of compute nodes.

Copy the ID of the compute node that you want to start, reset or stop.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-23 09:25:56,307 UTC
Data:
id          name      provisioningState provisioningType
--          -
3e62bf25-a26c-407e-ab8b-df01a4ad98b6 pcacn002 Provisioned     KVM
f7b8356b-052f-4911-babb-447e6ab9c78d pcacn003 Provisioned     KVM
4e06ebdf-faed-484e-996d-d77af786f123 pcacn001 Provisioned     KVM
```

2. Make sure that the compute node is locked for maintenance.  
See [Locking a Compute Node for Maintenance](#).
3. Start, reset or stop the compute node using the corresponding command:

```
PCA-ADMIN> start ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: start ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 09:26:06,446 UTC
Data:
Success
```

```
PCA-ADMIN> reset id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: reset id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 09:27:06,434 UTC
Data:
Success
```

```
PCA-ADMIN> stop ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
Command: stop ComputeNode id=f7b8356b-052f-4911-babb-447e6ab9c78d
Status: Success
Time: 2021-08-23 09:31:38,271 UTC
Data:
Success
```

4. When the compute node is up and running again, release the maintenance and provisioning locks.

## Deprovisioning a Compute Node

If you need to take a compute node out of service, for example to replace a defective one, you must deprovision it first, so that its data is removed cleanly from the system databases.

### Using the Service Web UI

1. In the navigation menu, click Rack Units.
2. In the Rack Units table, click the host name of the compute node you want to deprovision. The compute node detail page appears.

3. In the top-right corner of the page, click Controls and select the Provisioning Lock command.

When the confirmation window appears, click Lock to proceed.

After successful completion, the Compute Node Information tab shows Provisioning Locked = Yes.

4. Make sure that no more compute instances are running on the compute node.

Click Controls and select the Migrate All Vms command. The system migrates the instances to other compute nodes.

5. To deprovision the compute node, click Controls and select the Deprovision command.

When the confirmation window appears, click Deprovision to proceed.

After successful completion, the Compute Node Information tab shows Provisioning State = Ready to Provision.

### Using the Service CLI

1. Display the list of compute nodes.

Copy the ID of the compute node you want to deprovision.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-20 08:53:56,681 UTC
Data:
  id                name      provisioningState provisioningType
  --                -
  29f68a0e-4744-4a92-9545-7c48fa365d0a pcacn001 Provisioned     KVM
  7a0236f4-b00e-461d-93a0-b22673a18d9c pcacn003 Provisioned     KVM
  dc8ae567-b07f-48e0-89bd-e57069c20010 pcacn002 Provisioned     KVM
```

2. Set a provisioning lock on the compute node.

```
PCA-ADMIN> provisioningLock id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Command: provisioningLock id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Status: Success
Time: 2021-08-20 10:30:00,320 UTC
JobId: ed4a4646-6d73-41f9-9cb0-73ea35e0d766
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=ed4a4646-6d73-41f9-9cb0-73ea35e0d766
[...]
```

```
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

3. Confirm that the compute node is under provisioning lock.

```
PCA-ADMIN> show ComputeNode id=7a0236f4-b00e-461d-93a0-b22673a18d9c
[...]
Provisioning Locked = true
```

4. Migrate all running compute instances off the compute node you want to deprovision.

```
PCA-ADMIN> migrateVm id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Command: migrateVm id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Status: Running
Time: 2021-08-20 10:37:05,781 UTC
JobId: 6f1e94bc-7d5b-4002-ada9-7d4b504a2599
```

Use the job ID to check the status of your command.

```
PCA-ADMIN> show Job id=6f1e94bc-7d5b-4002-ada9-7d4b504a2599
Command: show Job id=6f1e94bc-7d5b-4002-ada9-7d4b504a2599
Status: Success
Time: 2021-08-20 10:39:59,025 UTC
Data:
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

5. Deprovision the compute node with this command:

```
PCA-ADMIN> deprovision id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Command: deprovision id=7a0236f4-b00e-461d-93a0-b22673a18d9c
Status: Success
Time: 2021-08-20 11:30:43,793 UTC
JobId: 9868fdac-ddb6-4260-9ce1-c018cf2ddc8d
```

Use the job ID to check the status of your deprovision command.

```
PCA-ADMIN> show Job id=9868fdac-ddb6-4260-9ce1-c018cf2ddc8d
[...]
Done = true
Name = MODIFY_TYPE
Run State = Succeeded
```

6. Confirm that the compute node has been deprovisioned.

```
PCA-ADMIN> list ComputeNode
Command: list ComputeNode
Status: Success
Time: 2021-08-20 08:53:56,681 UTC
Data:
id          name          provisioningState provisioningType
--          ---          -
29f68a0e-4744-4a92-9545-7c48fa365d0a pcacn001 Provisioned    KVM
7a0236f4-b00e-461d-93a0-b22673a18d9c pcacn003 Ready to Provision Unspecified
dc8ae567-b07f-48e0-89bd-e57069c20010 pcacn002 Provisioned    KVM
```

## Integrating GPU Expansion Nodes

GPU nodes are installed in an expansion rack. Its networking components must be connected to the base rack so the new hardware can be integrated into the hardware administration and data networks. For installation requirements, physical hardware installation information, and

cabling details, refer to [Optional GPU Expansion](#) in the "Oracle Private Cloud Appliance Installation Guide".

In this section we assume the GPU expansion rack has been installed and connected to the Private Cloud Appliance base rack. The GPU nodes must be discovered and provisioned before their hardware resources are available for use within compute instances. Unlike standard compute nodes, which are added to the base rack and automatically integrated and prepared for provisioning, GPU nodes in an expansion rack go through a more strictly controlled process.

The GPU expansion rack is activated by running a script from one of the management nodes. With precise timing and orchestration based on a static mapping, this script powers on and configures each component in the GPU expansion rack. The required ports on the switches are enabled so that all hardware can be discovered and registered in the component database. When the scripted operations are completed, the data and management networks are operational across the interconnected racks. The operating system and additional software are installed on the new nodes, after which they are ready to provision.

Installation and activation of the expansion rack and GPU nodes are performed by Oracle. From this point forward, the system treats GPU nodes the same way as all other compute nodes. After provisioning, appliance administrators can manage and monitor them from the Service Enclave UI or CLI. See [Performing Compute Node Operations](#).

#### Note

Live migration is not supported for GPU instances. This impacts some compute node operations.

- Evacuating a GPU node will fail. Instances must be stopped manually.
- The high availability configuration of the Compute Service applies to GPU instances, but is further restricted by limited hardware resources.

When a GPU node goes offline and returns to normal operation, the Compute Service restarts instances that were stopped during the outage. An instance might be restarted, through cold migration, on another GPU node with enough hardware resources are available.

#### Caution

For planned maintenance or upgrade, best practice is to issue a shut down command from the instance OS, then gracefully stop the instance from the Compute Web UI or OCI CLI.

GPU nodes are added to the 3 existing fault domains, which is consistent with the overall Oracle cloud architecture. The fault domains might become unbalanced because, unlike standard compute nodes, GPU nodes can be added one at a time. This has no functional impact on the fault domains because the server families operate separately from each other. The GPU nodes can only host compute instances based on a GPU shape, and migrations between different server families in the same fault domain are not supported.

In the Compute Enclave, consuming resources provided by a GPU node is straightforward. Users deploy compute instances with a dedicated shape to allocate 1-4 GPUs. Instances based on a GPU shape always run on a GPU node.

## Integrating a Compute Expansion Rack

To expand compute capabilities, you can buy a compute expansion rack. Its networking components must be connected to the base rack so the new hardware can be integrated into the hardware administration and data networks. For installation requirements, physical hardware installation information, and cabling details, refer to [Optional Compute Expansion Rack](#) in the "Oracle Private Cloud Appliance Installation Guide".

In this section we assume the compute expansion rack has been installed and connected to the Private Cloud Appliance base rack. The compute nodes must be discovered and provisioned before their hardware resources are available for use within compute instances. Unlike standard compute nodes, which are added to the base rack and automatically integrated and prepared for provisioning, the compute nodes in an expansion rack go through a more strictly controlled process.

The compute expansion rack is activated by running a script from one of the management nodes. With precise timing and orchestration based on a static mapping, this script powers on and configures each component in the compute expansion rack. The required ports on the switches are enabled so that all hardware can be discovered and registered in the component database. When the scripted operations are completed, the data and management networks are operational across the interconnected racks. The operating system and additional software are installed on the new nodes, after which they are ready to provision.

Installation and activation of the expansion rack and compute nodes are performed by Oracle. From this point forward, the system treats compute nodes the same way as all other compute nodes. After provisioning, appliance administrators can manage and monitor them from the Service Enclave UI or CLI. See [Performing Compute Node Operations](#).

### Caution

For planned maintenance or upgrade, best practice is to issue a shut down command from the instance OS, then gracefully stop the instance from the Compute Web UI or OCI CLI.

## Configuring the Active Directory Domain for File Storage

The file storage service in Oracle Private Cloud Appliance enables users of Microsoft Windows instances to map a network drive, or mount a network share. Both the NFS and SMB protocols are supported, but for SMB it is required that the Microsoft Windows instances and Private Cloud Appliance belong to the same Active Directory domain. This section provides instructions to set up the Active Directory domain in the Service Enclave.

### Using the Service Web UI

1. Verify that DNS is configured on the appliance.
  - a. In the navigation menu, click Network Environment.
  - b. In the Network Environment Information detail page, select the DNS Servers tab and make sure that DNS servers are configured.

DNS is required because, during domain configuration, the system searches for a matching SRV record in order to locate the controllers of the Active Directory domain.

2. In the navigation menu, click Active Directory Domain.

3. Verify that no Active Directory domain is currently configured. The configuration details should show "Status = disabled" and "Domain = Not Available".
4. Click Edit to change the Active Directory domain configuration.
5. In the Active Directory Domain Setting window, enter these parameters:
  - the name of the Active Directory domain the appliance is meant to join
  - a user name and password that enable the appliance to join the domain
  - optionally, an organizational unit
6. Click Submit to apply the new configuration.
7. Verify that the Active Directory is configured correctly. The configuration details should show "Status = online" and the newly configured domain name should appear in the Domain field.
8. To remove the ZFS Storage Appliance from the Active Directory domain again, you must use the Service CLIs as documented below. Refer to the final step in the Service CLI instructions.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - the name of the Active Directory domain the appliance is meant to join
  - an account (user name and password) with authorization to join the Active Directory domain
2. Verify that DNS is configured on the appliance. During domain configuration, the system searches for a matching SRV record in order to locate the controllers of the Active Directory domain.

```
PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2021-12-17 12:20:51,238 UTC
Data:
Uplink Port Speed = 100
Uplink Port Count = 2
Uplink Vlan Mtu = 9216
[...]
DNS Address1 = 192.0.2.201
DNS Address2 = 192.0.2.202
DNS Address3 = 10.25.0.101
Management Node1 Hostname = mypca-mn1
Management Node2 Hostname = mypca-mn2
Management Node3 Hostname = mypca-mn3
[...]
Network Config Lifecycle State = ACTIVE
```

3. Verify that no Active Directory domain is currently configured.

```
PCA-ADMIN> show ZFSAdDomain
Command: show ZFSAdDomain
Status: Success
Time: 2021-12-17 12:17:42,734 UTC
Data:
Status = disabled
Mode = workgroup
Service href = /api/service/v2/services/ad
Domain href = /api/service/v2/services/ad/domain
Workgroup href = /api/service/v2/services/ad/workgroup
```

```

PasswordSet = false
Preexist = false
Workgroup = WORKGROUP

```

4. Configure the Active Directory domain by entering the name of the domain, and a user name and password that enables the appliance to join the domain.

```

PCA-ADMIN> configZFSAdDomain domain=ad.example.com user=Administrator password=*****
Command: configZFSAdDomain domain=ad.example.com user=Administrator password=*****
Status: Success
Time: 2021-12-17 12:24:25,333 UTC
JobId: 7e6abf2d-9f6a-4c32-8f18-5142f6eda3c5

```

5. Use the job ID to check the status of your command.

When the job has completed successfully, verify the Active Directory zone configuration and status.

```

PCA-ADMIN> show ZFSAdDomain
Command: show ZFSAdDomain
Status: Success
Time: 2021-12-17 12:35:04,944 UTC
Data:
Status = online
Mode = domain
Service href = /api/service/v2/services/ad
Domain href = /api/service/v2/services/ad/domain
Workgroup href = /api/service/v2/services/ad/workgroup
PasswordSet = false
Preexist = false

```

6. To remove the ZFS Storage Appliance from the Active Directory domain again, set its configuration back to *workgroup* mode.

```

PCA-ADMIN> configZFSAdWorkgroup workgroupName=WORKGROUP
Command: configZFSAdWorkgroup workgroupName=WORKGROUP
Status: Success
Time: 2022-08-31 07:47:38,916 UTC
JobId: 1329e43a-3ed6-4588-b90b-a45506271df8

```

```

PCA-ADMIN> show zfsAdDomain
Command: show zfsAdDomain
Status: Success
Time: 2022-08-31 07:48:07,837 UTC
Data:
Status = disabled
Mode = workgroup
Service href = /api/service/v2/services/ad
Domain href = /api/service/v2/services/ad/domain
Workgroup href = /api/service/v2/services/ad/workgroup
PasswordSet = false
Preexist = false
Workgroup = WORKGROUP

```

## Reconfiguring the Network Environment

From the Network Environment list of the Service Web UI, an administrator can edit the network environment information provided during initial system setup. Carefully plan any changes you make in this area, as these parameters provide the connections from your data center to the Private Cloud Appliance and can potentially disrupt system operations.

**Note**

It is *not* supported to turn off BGP authentication by changing the BGP password to null:

```
PCA-ADMIN> edit networkConfig adminbgppassword=
```

In this case, the BGP authentication is still set to true and no change is made to the password, even when the command to change the BGP password to null succeeds. This is not an error condition, but a security feature. To disable BGP authentication, you must do so explicitly. In the CLI, the command is

```
PCA-ADMIN> edit networkConfig adminbgpauthentication=false
```

## Editing Routing Information

**Caution**

It is not supported to change your routing information for your dynamic or static network topology.

## Editing Management Node Information

This section explains how to edit IP and hostname information for your management nodes.

**Caution**

Changing management node parameters can cause system disruption.

### Using the Service Web UI

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the Management Nodes tab.  
The Management Nodes details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.  
For field descriptions, see the [Initial Installation Checklist](#) section in the [Oracle Private Cloud Appliance Installation Guide](#).
5. Click Save Changes.

### Using the Service CLI

1. Display the current network configuration information using the `show NetworkConfig` command.

```

PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2021-09-28 17:31:33,990 UTC
Data:
  Uplink Port Speed = 100
  Uplink Port Count = 2
  Uplink Vlan Mtu = 9216
  Spine1 Ip = 10.n.n.12
  Spine2 Ip = 10.n.n.13
  Uplink Netmask = 255.255.255.0
  Management VIP Hostname = ukpca01mn
  Management VIP 100g = 10.n.n.8
  NTP Server(s) = 100.n.n.254
  Uplink Port Fec = auto
  Public Ip range/list = 10.n.n.2/32,10.n.n.3/32,10.n.n.4/32,10.n.n.5/32,10.n.n.6/32,10.n.n.7/32
  DNS Address1 = 206.n.n.1
  DNS Address2 = 206.n.n.2
  DNS Address3 = 10.n.n.197
  Management Node1 Hostname = ukpca01-mn1
  Management Node2 Hostname = ukpca01-mn2
  Management Node3 Hostname = ukpca01-mn3
  100g Management Node1 Ip = 10.n.n.9
  100g Management Node2 Ip = 10.n.n.10
  100g Management Node3 Ip = 10.n.n.11
  Object Storage Ip = 10.n.n.1
  Enable Admin Network = false
  Static Routing = true
  Spine VIP = 10.n.n.14
  Uplink Gateway = 10.n.n.1
  Uplink VLAN = 799
  Uplink Hsrp Group = 61
  BGP Authentication = false

```

2. Use the `edit NetworkConfig` command to change any of these management node parameters:

- Management Node 1 IP
- Management Node 1 Hostname
- Management Node 2 IP
- Management Node 2 Hostname
- Management Node 3 IP
- Management Node 3 Hostname
- Management Node VIP
- Management Node VIP Hostname

```

PCA-ADMIN> edit NetworkConfig mgmt01Ip100g=172.n.n.190 mgmt02Ip100g=172.n.n.191
Command: edit NetworkConfig mgmt01Ip100g=172.n.n.190 mgmt02Ip100g=172.n.n.191
Status: Success
Time: 2021-09-27 14:25:00,603 UTC
JobId: 52f5177d-402a-4a52-98fe-1cff9c1f26be
PCA-ADMIN>

```

## Editing Data Center Uplink Information

This section explains how to edit uplink information for your configuration.

**⚠ Caution**

Reconfiguring the Private Cloud Appliance connection to the data center causes an interruption of all network connectivity to and from the appliance. No network traffic is possible while the physical connections are reconfigured. All connections are automatically restored when the configuration update is complete.

**Using the Service Web UI**

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the Uplink tab.  
The Uplink details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.  
For field descriptions, see the [Initial Installation Checklist](#) section in the [Oracle Private Cloud Appliance Installation Guide](#).
5. Click Save Changes.

**Using the Service CLI**

1. Display the current network configuration information using the show NetworkConfig command.

```
PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2021-09-28 17:31:33,990 UTC
Data:
Uplink Port Speed = 100
Uplink Port Count = 2
Uplink Vlan Mtu = 9216
Spine1 Ip = 10.n.n.12
Spine2 Ip = 10.n.n.13
Uplink Netmask = 255.255.255.0
Management VIP Hostname = ukpca01mn
Management VIP 100g = 10.n.n.8
NTP Server(s) = 100.n.n.254
Uplink Port Fec = auto
Public Ip range/list = 10.n.n.2/32,10.n.n.3/32,10.n.n.4/32,10.n.n.5/32,10.n.n.6/32,10.n.n.7/32
DNS Address1 = 206.n.n.1
DNS Address2 = 206.n.n.2
DNS Address3 = 10.n.n.197
Management Node1 Hostname = ukpca01-mn1
Management Node2 Hostname = ukpca01-mn2
Management Node3 Hostname = ukpca01-mn3
100g Management Node1 Ip = 10.n.n.9
100g Management Node2 Ip = 10.n.n.10
100g Management Node3 Ip = 10.n.n.11
Object Storage Ip = 10.n.n.1
Enable Admin Network = false
Static Routing = true
Spine VIP = 10.n.n.14
Uplink Gateway = 10.n.n.1
Uplink VLAN = 799
```

```
Uplink Hsrp Group = 61  
BGP Authentication = false
```

2. Use the `edit NetworkConfig` command to change any of these data center uplink parameters:

- Uplink Port Speed
- Uplink Port Count
- Uplink VLAN MTU
- Uplink Port FEC

```
PCA-ADMIN> edit NetworkConfig uplinkPortCount=2  
Command: edit NetworkConfig uplinkPortCount=2  
Time: 2021-09-27 14:27:00,605 UTC  
JobId: 42f5137f-122a-4a52-98fe-1cfv9c1f26ve  
PCA-ADMIN>
```

## Updating NTP Server Information

This section explains how to edit or add NTP server IP addresses.

### Using the Service Web UI

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the NTP tab.  
The NTP details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.

For field descriptions, see the [Initial Installation Checklist](#) section in the [Oracle Private Cloud Appliance Installation Guide](#).

5. Click Save Changes.

### Using the Service CLI

1. Display the current network configuration information using the `show NetworkConfig` command.

```
PCA-ADMIN> show NetworkConfig  
Command: show NetworkConfig  
Status: Success  
Time: 2021-09-28 17:31:33,990 UTC  
Data:  
Uplink Port Speed = 100  
Uplink Port Count = 2  
Uplink Vlan Mtu = 9216  
Spine1 Ip = 10.n.n.12  
Spine2 Ip = 10.n.n.13  
Uplink Netmask = 255.255.255.0  
Management VIP Hostname = ukpca01mn  
Management VIP 100g = 10.n.n.8  
NTP Server(s) = 100.n.n.254  
Uplink Port Fec = auto  
Public Ip range/list = 10.n.n.2/32,10.n.n.3/32,10.n.n.4/32,10.n.n.5/32,10.n.n.6/32,10.n.n.7/32  
DNS Address1 = 206.n.n.1  
DNS Address2 = 206.n.n.2  
DNS Address3 = 10.n.n.197  
Management Node1 Hostname = ukpca01-mn1
```

```

Management Node2 Hostname = ukpca01-mn2
Management Node3 Hostname = ukpca01-mn3
100g Management Node1 Ip = 10.n.n.9
100g Management Node2 Ip = 10.n.n.10
100g Management Node3 Ip = 10.n.n.11
Object Storage Ip = 10.n.n.1
Enable Admin Network = false
Static Routing = true
Spine VIP = 10.n.n.14
Uplink Gateway = 10.n.n.1
Uplink VLAN = 799
Uplink Hsrp Group = 61
BGP Authentication = false

```

2. Use the `edit NetworkConfig` command to change the NTP servers. Enter multiple IP addresses in a comma-separated list:

```

PCA-ADMIN> edit NetworkConfig ntpIps=100.n.n.254,100.n.n.253
Command: edit NetworkConfig ntpIps=100.n.n.254,100.n.n.253
Time: 2021-09-27 14:31:00.605 UTC
JobId: 42f5137f-122a-4a52-98fe-1cfv9c1f26ve
PCA-ADMIN>

```

## Editing Administration Network Information

If you use the optional Administration Network, you can update the parameters using these procedures.

### Caution

If you are not currently using a separate Administration Network, the Network Environment Information page in the Service Web UI will not display an Admin Network tab or any of the related configuration parameters. You must first enable the Administration Network.

Once an Administration Network is configured, it cannot be disabled again.

To edit Administration Network information, see the following resources in the *Oracle Private Cloud Appliance Installation Guide*:

- For general configuration information, see [Administration Network Configuration Notes](#).
- For descriptions of Administration Network parameters, see [Initial Installation Checklist](#).

### Using the Service Web UI

#### Scenario 1: Administration Network Disabled

If you need to enable and configure a separate Administration Network, proceed as follows:

1. In the navigation menu, click Network Environment.
2. In the top-right corner of the page, click Edit.
3. In the wizard, navigate to the Admin Network tab and set Admin Networking to *Enable*.
4. Enter all the required parameters in the respective fields on the form.
5. Click Save Changes.

#### Scenario 2: Administration Network Enabled

If you already configured a separate Administration Network and need to edit its settings, proceed as follows:

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the Admin Network tab.  
The Admin Network details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.
5. Click Save Changes.

### Using the Service CLI

#### Caution

If you are not currently using a separate Administration Network, the Service CLI output will not display any Admin Network parameters. You must first enable the Administration Network.

1. Display the current network configuration information using the `show NetworkConfig` command.

```
PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2022-10-11 07:13:12,186 UTC
Data:
Uplink Port Speed = 100
Uplink Port Count = 4
Uplink Vlan Mtu = 9216
Spine1 Ip = 10.10.10.97,10.10.10.101
Spine2 Ip = 10.10.10.99,10.10.10.103
Uplink Netmask = 255.255.255.254,255.255.255.254
Management VIP Hostname = mypca
Management VIP = 10.10.10.107
NTP Server(s) = 10.80.211.105,10.211.17.1,10.68.48.1
Uplink Port Fec = auto
Public Ip range/list =
10.10.10.114/31,10.10.10.116/31,10.10.10.118/31,10.10.10.120/31,10.10.10.122/31,10.10.10.124/31,10.10.10.126/32
Management Node1 Hostname = pcamn01
Management Node2 Hostname = pcamn02
Management Node3 Hostname = pcamn03
Management Node1 Ip = 10.10.10.108
Management Node2 Ip = 10.10.10.109
Management Node3 Ip = 10.10.10.110
Object Storage Ip = 10.10.10.113
Enable Admin Network = true
Admin Port Speed = 100
Admin Port Count = 1
Admin Vlan Mtu = 9216
Admin Port Fec = auto
Admin VLAN = 3915
Admin Spine1 Ip = 10.25.0.111
Admin Spine2 Ip = 10.25.0.112
Admin Spine VIP = 10.25.0.110
Admin Netmask = 255.255.255.0
```

```

Admin Hsrp Group = 152
Static Routing = false
Uplink VLAN = 3911
Peer1 Asn = 50000
Peer1 Ip = 10.10.10.96,10.10.10.98
Oracle Asn = 136025
Bgp Topology = mesh
Peer2 Asn = 50000
Peer2 Ip = 10.10.10.100,10.10.10.102
BGP Authentication = false
BGP KeepAlive Timer = 60
BGP Holddown Timer = 180
Network Config Lifecycle State = ACTIVE
admin DNS Address1 = 10.25.0.1
admin Management Node1 Hostname = pcamn01admin.example.com
admin Management Node2 Hostname = pcamn02admin.example.com
admin Management Node3 Hostname = pcamn03admin.example.com
admin Management Node1 Ip = 10.25.0.101
admin Management Node2 Ip = 10.25.0.102
admin Management Node3 Ip = 10.25.0.103
admin Management VIP Hostname = mypcaadmin.example.com
admin Management VIP = 10.25.0.100

```

2. Use the `edit NetworkConfig` command to change any of these administration network parameters:

 **Tip**

Enter `edit networkConfig ?` to display the parameters available for editing.

- Admin Network enable
- Management node cluster Admin VIP and host name
- Management node 1-3 Admin IP and host name
- Admin DNS IP 1-3
- Admin Port count, speed, FEC
- Admin CIDR
- Admin VLAN and MTU
- Admin Gateway IP
- Admin Netmask
- Spine 1+2 Admin IP
- Spine Admin VIP

```

PCA-ADMIN> edit NetworkConfig adminPortSpeed=25
Command: edit NetworkConfig adminPortSpeed=25
Time: 2022-10-11 08:01:00,605 UTC
JobId: 62f8137f-772a-4a52-98f4-1cfv9c1f24te

```

```

PCA-ADMIN> edit NetworkConfig adminCidr=10.25.0.1/24
Command: edit NetworkConfig adminCidr=10.25.0.1/24
Status: Success
Time: 2022-10-11 08:10:02,640 UTC
JobId: 861381ae-cc63-44a2-a66e-8e095e4a99f9

```

## Updating DNS Information

This section explains how to edit or add DNS IP addresses.

### Using the Service Web UI

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the DNS tab.  
The DNS details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.

For field descriptions, see the [Initial Installation Checklist](#) section in the [Oracle Private Cloud Appliance Installation Guide](#).

5. Click Save Changes.

### Using the Service CLI

1. Display the current network configuration information using the `show NetworkConfig` command.

```
PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2021-09-28 17:31:33,990 UTC
Data:
Uplink Port Speed = 100
Uplink Port Count = 2
Uplink Vlan Mtu = 9216
Spine1 Ip = 10.n.n.12
Spine2 Ip = 10.n.n.13
Uplink Netmask = 255.255.255.0
Management VIP Hostname = ukpca01mn
Management VIP 100g = 10.n.n.8
NTP Server(s) = 100.n.n.254
Uplink Port Fec = auto
Public Ip range/list = 10.n.n.2/32,10.n.n.3/32,10.n.n.4/32,10.n.n.5/32,10.n.n.6/32,10.n.n.7/32
DNS Address1 = 206.n.n.1
DNS Address2 = 206.n.n.2
DNS Address3 = 10.n.n.197
Management Node1 Hostname = ukpca01-mn1
Management Node2 Hostname = ukpca01-mn2
Management Node3 Hostname = ukpca01-mn3
100g Management Node1 Ip = 10.n.n.9
100g Management Node2 Ip = 10.n.n.10
100g Management Node3 Ip = 10.n.n.11
Object Storage Ip = 10.n.n.1
Enable Admin Network = false
Static Routing = true
Spine VIP = 10.n.n.14
Uplink Gateway = 10.n.n.1
Uplink VLAN = 799
Uplink Hsrp Group = 61
BGP Authentication = false
```

2. Use the `edit NetworkConfig` command to change the DNS IP addresses:
  - DNS IP1

- DNS IP2
- DNS IP3

```
PCA-ADMIN> edit NetworkConfig DnsIp2=206.n.n.2
Command: edit NetworkConfig DnsIp2=206.n.n.2
Time: 2021-09-27 14:21:00,605 UTC
JobId: 42f5137f-122a-4a52-98fe-1cfv9c1f26ve
PCA-ADMIN>
```

## Updating Public IP Information

This section explains how to edit the public IP addresses for your appliance. You can add public IP addresses, or change the currently configured IP addresses.

### **Caution**

Changing public IP addresses that are in use can cause system disruption.

### Using the Service Web UI

1. In the navigation menu, click Network Environment.
2. In the Network Environment Information page, click the Uplink tab.  
The Uplink details appear.
3. In the top-right corner of the page, click Edit.
4. Click Next to navigate to the page you want to edit, then update the appropriate fields.  
For field descriptions, see the [Initial Installation Checklist](#) section in the [Oracle Private Cloud Appliance Installation Guide](#).
5. Click Save Changes.

### Using the Service CLI

1. Display the current network configuration information using the show NetworkConfig command.

```
PCA-ADMIN> show NetworkConfig
Command: show NetworkConfig
Status: Success
Time: 2021-09-28 17:31:33,990 UTC
Data:
Uplink Port Speed = 100
Uplink Port Count = 2
Uplink Vlan Mtu = 9216
Spine1 Ip = 10.n.n.12
Spine2 Ip = 10.n.n.13
Uplink Netmask = 255.255.255.0
Management VIP Hostname = ukpca01mn
Management VIP 100g = 10.n.n.8
NTP Server(s) = 100.n.n.254
Uplink Port Fec = auto
Public Ip range/list = 10.n.n.2/32,10.n.n.3/32,10.n.n.4/32,10.n.n.5/32,10.n.n.6/32,10.n.n.7/32
DNS Address1 = 206.n.n.1
DNS Address2 = 206.n.n.2
DNS Address3 = 10.n.n.197
Management Node1 Hostname = ukpca01-mn1
```

```

Management Node2 Hostname = ukpca01-mn2
Management Node3 Hostname = ukpca01-mn3
100g Management Node1 Ip = 10.n.n.9
100g Management Node2 Ip = 10.n.n.10
100g Management Node3 Ip = 10.n.n.11
Object Storage Ip = 10.n.n.1
Enable Admin Network = false
Static Routing = true
Spine VIP = 10.n.n.14
Uplink Gateway = 10.n.n.1
Uplink VLAN = 799
Uplink Hsrp Group = 61
BGP Authentication = false

```

2. Use the edit NetworkConfig command to change the public IP addresses or the object storage public IP address:

- Object Storage Public IP
- Public IP Range/List

```

PCA-ADMIN> edit NetworkConfig PublicIps= 10.n.n.17/32,10.n.n.18/32,10.n.n.19/32
Command: edit NetworkConfig PublicIps= 10.n.n.17/32,10.n.n.18/32,10.n.n.19/32
Time: 2021-09-27 14:21:00,605 UTC
JobId: 42f5137f-122a-4a52-98fe-1cfv9c1f26ve
PCA-ADMIN>

```

## Configuring Appliance Proxy Settings

The initial appliance setup procedure, as described in the [Oracle Private Cloud Appliance Installation Guide](#), provides no option to add a system-wide proxy configuration. However, some of the platform and infrastructure services require connectivity to endpoints outside the appliance environment. For example, for federation with an identity provider (IDP), the IAM and Admin services must acquire metadata from that external server (for users of the Compute Enclave and Service Enclave respectively.) If network traffic passes through a proxy server in the data center, requests between the appliance and the external server cannot be completed successfully.

To enable external network communication through a data center proxy server, add the proxy configuration to the appliance network setup. Ensure that the initial appliance setup procedure has been completed first.

### Using the Service Web UI

1. In the PCA Config navigation menu, click Appliance Details.

The Appliance Details page contains system properties such as realm, region and domain.

2. To configure a proxy at the appliance level, click the Set Rack-Wide Proxy button in the top-right corner.

A proxy configuration window appears.

3. Fill out the proxy configuration parameters:

- **Proxy Name:** Enter the fully qualified domain name of the proxy server.
- **Proxy Host:** Enter the proxy server IP address.
- **Proxy Port:** Enter the port number the proxy server uses for routing requests.
- **Proxy User Name:** If required, enter a user name for authentication with the proxy server.

- **Proxy Password:** If required, enter the password for the proxy user name.
- **Proxy Confirm Password:** If required, enter the proxy password again for confirmation.

The user name and password parameters are optional, in case the proxy server requires authentication. The details are stored in the Secret Service (Vault), where services can retrieve them securely to establish their external connection.

4. Click Set Rack-Wide Proxy to save the proxy configuration.

The proxy configuration is displayed in a separate tab on Appliance Details page.

5. To delete the proxy configuration from the appliance, go to the Appliance Details page and click Clear Rack-Wide Proxy in the top-right corner.
6. If you need to modify the stored proxy configuration, delete it and fill out the proxy configuration window again with the correct parameters.

### Using the Service CLI

1. Add the proxy configuration details using the `setProxy` command.

Syntax (entered on a single line):

```
PCA-ADMIN> setProxy
proxyName=<proxy_fqdn>
proxyHost=<proxy_ip>
proxyPort=<proxy_port>
proxyUsername=<proxy_user>
proxyPassword=<proxy_password>
proxyConfirmPassword=<proxy_password>
```

Example:

```
PCA-ADMIN> setProxy proxyName=myproxy.example.com
proxyHost=172.16.0.100
proxyPort=8080
proxyUsername=proxyuser
proxyPassword=*****
proxyConfirmPassword=*****
```

The user name and password parameters are optional, in case the proxy server requires authentication. The details are stored in the Secret Service (Vault), where services can retrieve them securely to establish their external connection.

2. To delete the proxy configuration from the appliance, enter the `clearProxy` command. No command parameters are required.
3. If you need to modify the stored proxy configuration, run the `setProxy` command again with the correct parameters.

## Creating and Managing Flex Networks

Oracle Private Cloud Appliance supports direct connectivity to external racks such as Oracle Exadata, Oracle Database Appliances, and external ZFS Appliances. In addition you can configure specific workloads to exit the rack through different physical ports on the spine switches to your data center. **This feature was formerly called Exadata Networks.**

You can create a Flex network in either Edge or Hub mode. The main difference between the modes is that you must configure a gateway for Hub mode.

This section describes creating and managing Flex networks from the Service Enclave. Before you can create a Flex network, you must physically connect your Private Cloud Appliance to an external rack. For instructions, see the "Optional Flex Network Connection" section in the chapter [Configuring Oracle Private Cloud Appliance](#) of the Oracle Private Cloud Appliance Installation Guide.

In order to use a Flex network, the VCNs containing compute instances that connect to the database nodes must have a dynamic routing gateway (DRG) configured. The enabled subnet needs a route rule with the external system CIDR as destination and the DRG as target.

**Note**

If a Flex network is in the **provisioning** or **updating** or **terminating** state, then a user cannot attach or detach a DRG or IGW, or create a NAT gateway. These operations need to be done once the Flex network is in the **available** or **terminated** state.

You can configure up to 128 Flex networks per rack, with a limit of 32 Flex networks per port.

**Note**

Exadata network commands are being deprecated and replaced by Flex network commands. The following table describes both the deprecated and new commands for this feature.

**Table 2-1 Flex Network Commands**

Deprecated Commands	New Commands
PCA-ADMIN> exaDataCreateNetwork ? *cidr *ports *spine1Ip *spine2Ip *spineVip advertiseNetwork exadataSpeed gatewayIp vlan	PCA-ADMIN> create FlexNetwork ? *cidr *ports *spine1Ip *spine2Ip *spineVip advertiseNetwork speed gatewayIp vlan
PCA-ADMIN> exaDataGetNetwork ? *exadataNetworkId	PCA-ADMIN> show FlexNetwork ? id=<object identifier>
PCA-ADMIN> exaDataListNetwork	PCA-ADMIN> list FlexNetwork
PCA-ADMIN> exaDataUpdateNetwork ? *exadataNetworkId *ports	PCA-ADMIN> edit FlexNetwork id=123 ? *ports
PCA-ADMIN> exaDataDeleteNetwork ? *exadataNetworkId	PCA-ADMIN> delete FlexNetwork ? id=<object identifier>

**Table 2-1 (Cont.) Flex Network Commands**

Deprecated Commands	New Commands
PCA-ADMIN> exaDataEnableAccess ? *exadataNetworkId *subnetId	PCA-ADMIN> flexNetworkEnableAccess ? *flexNetworkId *subnetId
PCA-ADMIN> exaDataDisableAccess ? *exadataNetworkId *subnetId	PCA-ADMIN> flexNetworkDisableAccess ? *flexNetworkId *subnetId

For more information about Flex Network Integration, see the "Network Infrastructure" section in the [Hardware Overview](#) chapter of the Oracle Private Cloud Appliance Concepts Guide.

## Taskmap for Creating a Flex Network

This task map describes the steps required to establish a Flex network between the Private Cloud Appliance and an external rack such as Oracle Exadata or Oracle Database Appliance.

No.	Task	Links
1.	Identify the physical ports on the Spine switch you plan to use for the external connection, then cable the hardware together.	<a href="#">Optional Connection to Flex Network</a>
2.	Create the Flex network from the Service enclave.	<a href="#">Creating a Flex Network</a>
3.	From the Compute enclave, create a DRG to provide a way for VMs to access the external system.	<a href="#">Create a Dynamic Routing Gateway</a>
4.	From the Compute enclave, create VCNs, Subnets, Route Tables and Internet Gateways, if needed. See <a href="#">Other Considerations</a> .	<a href="#">Creating a VCN</a> <a href="#">Managing VCNs and Subnets</a>
5.	From the Compute enclave, create DRG-attachments to enable VCN to use DRG.	<a href="#">Attach VCNs to a Dynamic Routing Gateway</a>
6.	From the Service enclave, enable communication between the Flex network and the VM subnets.	<a href="#">Enabling Flex Network Access</a>

### Other Considerations

When implementing a Flex network, consider the following:

- If a VM connected to a Flex network must also be accessed from a domain controller, you need to configure a second VNIC for that VM. See [Creating and Attaching a Secondary VNIC](#).
- Use an Internet Gateway (public subnet) for Domain Controller access using the primary VNIC.
- Use a Dynamic Routing Gateway (private subnet) to access the Flex network.

- Use separate Route Tables: one for the Internet Gateway with 0.0.0.0/0 and one for the DRG with a *specific* route rule for the Flex network.
- Update Security Lists as needed to enable ingress traffic.

## Creating a Flex Network

To set up a network connection between Private Cloud Appliance and an external system, you need this set of parameters:

Parameter	Example Value	Description
cidr	10.nn.nn.0/24	Choose a valid CIDR range that is within the CIDR range of the Oracle Exadata.
spine1Ip	10.nn.nn.2	A valid IP address in the CIDR specified.
spine2Ip	10.nn.nn.3	A valid IP address in the CIDR specified.
spineVip	10.nn.nn.1	A valid IP address in the CIDR specified.
vlan	3062	Choose a VLAN from 2 to 3899 that isn't in use by the uplink VLAN or other Oracle Exadata VLANs. This parameter can be unspecified for attaching a device not supporting VLAN tagging.
speed	10	Speed of the aggregated switch links under the port-channel must be 10, 20, 25, 40, 50, or 100 speed.
ports	7/1	7/1-4, 8/1-4, 9/1-4, or 10/1-4 are valid for 10G or 25G speeds. Ports 7, 8, 9, or 10 are valid for 40G or 100G speeds. For more detail, see the next table.
gateway IP	10.nn.nn.nn	Valid IP address of gateway. Default is null.
advertise Network	True	True or False - enables or disables the visibility of the Flex network to the customer's data center servers.

### Note

When the Flex network with, or without a gateway IP address is enabled, there is no access to the uplink using the Oracle Private Cloud Appliance DRG VRF or Oracle Exadata VRF. There needs to be a IGW or NAT on a separate interface in the VM on the Oracle Private Cloud Appliance for access to the uplink.

Valid speeds and valid port configurations are related. The following table shows the valid port configurations based on speed selected. Ports must be bonded on the external system side to match the Oracle Private Cloud Appliance configuration.

Speed	Valid Port Configurations
10	7/1-4, 8/1-4, 9/1-4, or 10/1-4
20	7/1-2, 8/1-2, 9/1-2, or 10/1-2 (20G bonds two 10G ports)
25	7/1-4, 8/1-4, 9/1-4, or 10/1-4
40	7, 8, 9, or 10

Speed	Valid Port Configurations
50	7/1-2, 8/1-2, 9/1-2, or 10/1-2 (50G bonds two 25G ports)
100	7, 8, 9, or 10

**Note**

For 25G flex networks forward error correction (FEC) is always set to off, with or without a gateway.

### Using the Service Web UI

1. Determine the Flex network parameters listed in the table above.
2. In the Dashboard, click the Racks quick action tile.
3. In the PCA Config navigation menu on the Racks page, click Flex Networks.
4. In the top-right corner above the table, click Create Flex Network.
5. Fill out the Flex Network form using the parameters you collected in advance.  
By default the network is not advertised to the data center network. You have to click the slider to set it to "on"/"true".
6. Click Submit to create the new network. It appears in the Flex Networks table and its Lifecycle State changes to Available when the configuration has been applied successfully.
7. Next, add a subnet to the Flex network. See [Enabling Flex Network Access](#).

### Using the Service CLI

1. Determine the Flex network parameters listed in the table above.
2. Create the Flex network by entering the parameters.  

```
PCA-ADMIN> create flexNetwork cidr=10.nn.nn.0/24 spine1Ip=10.nn.nn.1 spine2Ip=10.nn.nn.2 \
spinevip=10.nn.nn.3 vlan=900 gatewayIp=10.nn.nn.10 ports=7/1 advertiseNetwork=false
Command: create flexNetwork cidr=10.nn.nn.0/24 spine1Ip=10.nn.nn.1 spine2Ip=10.nn.nn.2 \
spinevip=10.nn.nn.3 vlan=900 gatewayIp=10.nn.nn.10 ports=7/1 advertiseNetwork=false
Status: Success
Time: 2025-03-05 18:07:12,546 UTC
JobId: unique_id
```
3. Next, add a subnet to the Flex network. See [Enabling Flex Network Access](#).

## Enabling Flex Network Access

Enable access from a subnet to the Flex network through the Service CLI. For for Flex network access from that subnet, ensure that the configured IP address ranges of Flex networks do not overlap.

Subnets that have been granted access, appear in the Flex network detail page under Access Lists, grouped by their parent VCN.

### Using the Service CLI

1. Get the OCID of the Flex network you want to enable, using the `list FlexNetwork` command.
2. Enable access to a configured Flex network.

```
PCA-ADMIN> flexNetworkEnableAccess flexNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Command: flexNetworkEnableAccess flexNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Status: Success
Time: 2024-11-17 18:56:45,251 UTC
Data:
id
--
ocid1.vcn.unique_id
```

3. If you are using a secondary NIC to access the Flex network, you must add a route to the Exadata CIDR address range for interface eth1 (the secondary NIC). Sign-in to the VM configured with the secondary NIC to add the route.

```
[root@hostname]# Flex-CIDR-address-range via gateway dev vlan-interface
```

For example, if the Flex address range is 192.168.0.0/24 and the gateway is 192.168.0.1 and the VLAN interface is bond0.900 :

```
[root@hostname]# 192.168.0.0/24 via 192.168.0.1 dev bond0.900
```

This entry appears as a second interface in the IP routing table:

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.168.1.0	192.168.1.1	255.255.255.0	0	0	0		eth0
192.168.0.0	192.168.0.1	255.255.255.0	0	0	0		<b>eth1</b>

A ping from the secondary NIC, eth1, now succeeds to the Flex network.

## List Flex Networks

### Using the Service Web UI

1. In the Dashboard, click the Rack Units quick action tile.
2. In the PCA Config navigation menu on the Rack Units page, click Flex Networks. The table contains all configured Flex networks.

### Using the Service CLI

1. Use the list FlexNetwork command to display configured Flex networks, including their OCIDs.

```
PCA-ADMIN> list FlexNetwork
Command: list FlexNetwork
Status: Success
Time: 2021-11-22 06:10:17,617 UTC
Data:
id          vlan cidr          spine1Ip    spine2Ip    spineVip    ports
--          -
ocid1.exadata.unique_id 2001 10.nn.nn.0/24 10.nn.nn.101 10.nn.nn.102 10.nn.nn.1 7/1,7/2
```

## Get Flex Network Details

### Using the Service Web UI

1. In the Dashboard, click the Rack Units quick action tile.
2. In the PCA Config navigation menu on the Rack Units page, click Flex Networks.

3. In the overview table, click the name (OCID) of the network for which you want to display details.

The Flex Network detail page shows the configuration parameters, the state of the network, and the subnets that have been granted access.

### Using the Service CLI

1. Get the OCID of the Flex network for which you want details, using the `exaDataListNetwork` command.
2. Use the `exaDataGetNetwork` command to display details about a specific Flex network, including the state of the network, subnet and VCN IDs.

```
PCA-ADMIN> show flexNetwork flexNetworkId=ocid1.exadata.unique_id
Command: show flexNetwork flexNetworkId=ocid1.exadata.unique_id
Status: Success
Time: 2024-11-22 19:34:56,917 UTC
Data:
  CIDR = 10.nn.nn.0/24
  Vlan = 2001
  Spine1Ip = 10.nn.nn.101
  Spine2Ip = 10.nn.nn.102
  SpineVip = 10.nn.nn.1
  Ports = 7/1,7/2
  advertiseNetwork = false
  Access List 1 - Vcn Id = ocid1.vcn.unique_id
  Access List 1 - Subnet Ids 1 = ocid1.subnet.unique_id
  Access List 1 - Subnet Ids 1 = ocid1.subnet.unique_id
  Access List 2 - Vcn Id = ocid1.vcn.unique_id
  Access List 2 - Subnet Ids 1 = ocid1.subnet.unique_id
  Lifecycle State = AVAILABLE
  gatewayIp = 10.nn.nn.21
  exaDataSpeed = 100
  name
```

## Editing Flex Networks

### Using the Service Web UI

1. In the Dashboard, click the Rack Units quick action tile.
2. In the PCA Config navigation menu on the Rack Units page, click Flex Networks. The table contains all configured Flex networks.
3. For the Flex network you want to edit, click the three dots in the Actions column, then click edit.
4. Enter the new ports and click Submit.

### Using the Service CLI

1. Use the `edit flexNetwork` command to add or remove Flex network ports.

```
PCA-ADMIN> edit flexNetwork id=ocid1.exadata.unique_id ports=7/1-4
Command: edit flexNetwork
Status: Success
Time: 2025-2-22 06:10:17,617 UTC
Data:
  id          vlan cidr          spine1Ip  spine2Ip  spineVip  ports
  --          --- ---          -
  ocid1.exadata.unique_id 2001 10.nn.nn.0/24 10.nn.nn.101 10.nn.nn.102 10.nn.nn.1 7/1-4
```

## Disabling Flex Network Access

Disabling access from a subnet to the Flex network must be done through the Service CLI.

Subnets that have been granted access, appear in the Flex network detail page under Access Lists, grouped by their parent VCN. When you disable access for a given subnet, it is removed from the Access Lists.

### Using the Service CLI

1. Get the OCID of the Flex network you want to disable, using the list FlexNetwork command.
2. Get the OCID of the subnet ID for the Flex network using the list FlexNetwork command.
3. Disable access to a configured Flex network.

```
PCA-ADMIN> flexNetworkDisableAccess flexNetworkId=ocid1.exadata.unique_id \  
subnetId=ocid1.subnet.unique_id  
Command: flexNetworkDisableAccess flexNetworkId=ocid1.exadata.unique_id \  
subnetId=ocid1.subnet.unique_id  
Status: Success  
Time: 2021-11-02 11:29:49,873 UTC
```

## Deleting a Flex Network

### Using the Service Web UI

1. Make sure that, for the Flex Exadata network you intend to delete, access has been disabled first.
2. Navigate to the Flex Network page.
3. Choose one of these options to delete the Flex network:
  - In the overview table, open the Actions menu on the right hand side of the row and select Delete. When prompted, click Confirm.
  - Open the Flex network detail page, then click the Delete button in the top-right corner.

### Using the Service CLI

1. Ensure that, for the Flex network you intend to delete, access has been disabled first.
2. Get the OCID of the Flex network you want to delete, using the exaDataListNetwork command.
3. Delete the Flex network.

```
PCA-ADMIN> delete FlexNetwork flexNetworkId=ocid1.exadata.unique_id  
Command: delete FlexNetwork flexNetworkId=ocid1.exadata.unique_id  
Status: Success  
Time: 2024-11-16 05:59:54,177 UTC
```

## Flex Network Examples

You can configure a Flex network in two different modes: Edge mode or Hub mode. The following sections contain examples for both modes.

## Configuring Flex Network Edge Mode to an Oracle Exadata

This example describes how to create a Flex network in Edge mode and then connect a virtual machine on Oracle Private Cloud Appliance to an Oracle Exadata within your data center.

This example, as shown in the diagram, creates a Flex network with a VCN that contains 3 virtual machines each connected to a private subnet, and also connected to one public subnet. The private subnets are routed through a dynamic routing gateway attached to the VCN, out to Oracle Exadata, which provides the VMs access to Oracle Exadata. This example also includes a public subnet, accessible by the VMs. This public subnet can be configured with an internet gateway to reach the data center ToR switches.



### Before you Begin

- Identify which physical ports on the Oracle Private Cloud Appliance spine switches will connect to the external system. See Valid Port Configurations in [Creating a Flex Network](#).
  - Identify the On-premises network subnet and reserve three IP addresses for the spine switches.
1. Create the Flex network from the Service Enclave. See [Creating a Flex Network](#).

To create a Flex network, at a minimum, you need the following parameters:

Parameter	Example Value
cidr	10.nn.nn.0/24
spine1Ip	10.nn.nn.101
spine2Ip	10.nn.nn.102
spineVip	10.nn.nn.1
vlan	2100
ports	7/1,7/2
advertiseNetwork	True

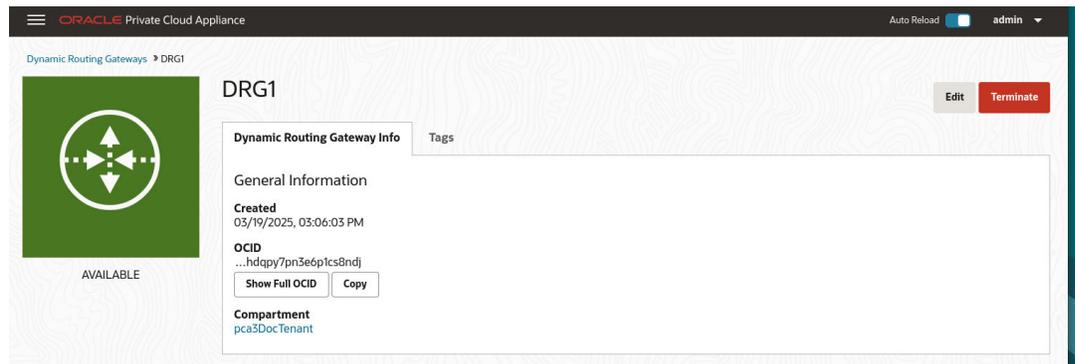
### Example:

```
PCA-ADMIN> create flexNetwork cidr=10.nn.nn.0/24 spine1Ip=10.nn.nn.101 spine2Ip=10.nn.nn.102
spinevip=10.nn.nn.1 \
vlan=2100 ports=7/1 advertiseNetwork=true
Command: create flexNetwork cidr=10.nn.nn.0/24 spine1Ip=10.nn.nn.101 spine2Ip=10.nn.nn.102
spinevip=10.nn.nn.1 \
vlan=2100 ports=7/1 advertiseNetwork=true
Status: Success
Time: 2025-03-05 18:07:12,546 UTC
JobId: 165f366-64c0-495e-sab1-34s8824b0da
PCA-ADMIN> list flexNetwork
Command: list flexNetwork
Status: Success
Time: 2025-03-05 18:07:21,480 UTC
Data:
  id          Vlan  CIDR          Spine1Ip  Spine2Ip  SpineVip  Ports
  --          ---  -
```

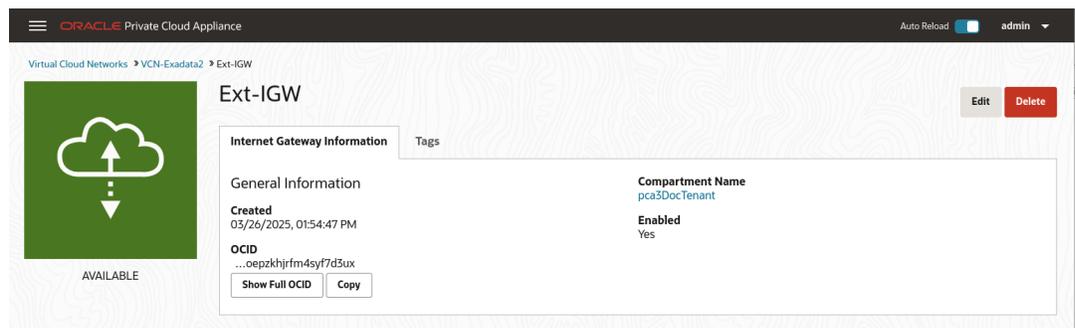
```
ocid1.cccexadata2.oc1.<unique_id> 2100 10.nn.nn.0/24 10.nn.nn.101 10.nn.nn.102 10.nn.nn.1 7/1,7/2
PCA-ADMIN>
```

Note the OCID of the Flex network, you need this OCID to enable the Flex network in step 5.

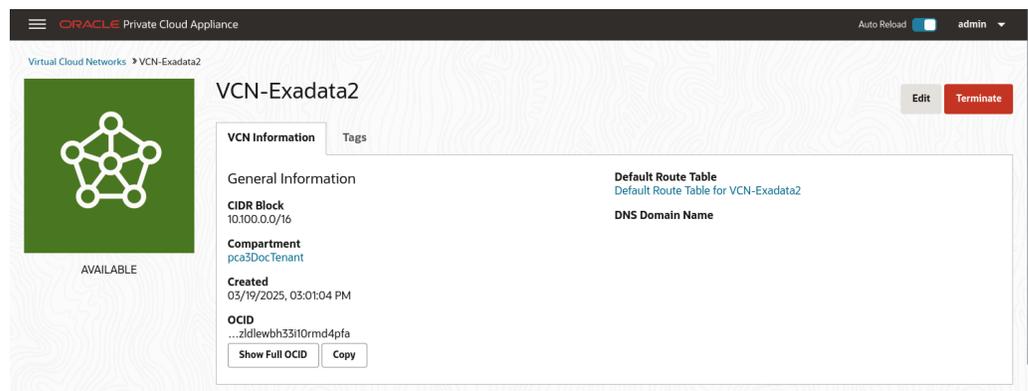
- From the Compute Enclave, create a DRG to provide a way for VMs to access the external system. See [Create a Dynamic Routing Gateway](#).



- From the Compute Enclave, create an Internet Gateway to provide a way for VMs to access the data center switches. See [Providing Public Access through an Internet Gateway](#).



- From the Compute Enclave, create VCNs, and Subnets. See [Managing VCNs and Subnets](#).
  - Create a VCN for use by the Flex network. In the Compute Web UI, create a VCN. Choose a compartment, a name for the VCN, then assign a CIDR block that meets your needs.



- Create subnets within the VCN. Navigate to the VCN in the Compute Web UI, and click Create Subnet. Name the subnet, choose the compartment, and assign a CIDR block within the CIDR block range of the VCN. Next select private or public subnet, then click Create Subnet.

The screenshot shows the Oracle Private Cloud Appliance web interface. The breadcrumb trail is "Virtual Cloud Networks > VCN-Exadata2 > Exa-Priv-SN1". The main content area features a green square icon with a white network diagram and the word "AVAILABLE" below it. To the right, the "Subnet Information" tab is active, displaying the following details:

<b>Subnet Information</b>	Tags
<b>General Information</b>	
<b>Created</b>	03/21/2025, 04:11:35 PM
<b>CIDR Block</b>	10.100.101.0/24
<b>Virtual Router Mac Address</b>	00:13:97:f1:f6:9c
<b>OCID</b>	...7whwrom5ovqdr1z4oq
<input type="button" value="Show Full OCID"/> <input type="button" value="Copy"/>	

The screenshot shows the Oracle Private Cloud Appliance web interface. The breadcrumb trail is "Virtual Cloud Networks > VCN-Exadata2 > Exa-Priv-SN2". The main content area features a green square icon with a white network diagram and the word "AVAILABLE" below it. To the right, the "Subnet Information" tab is active, displaying the following details:

<b>Subnet Information</b>	Tags
<b>General Information</b>	
<b>Created</b>	03/21/2025, 04:12:09 PM
<b>CIDR Block</b>	10.100.102.0/24
<b>Virtual Router Mac Address</b>	00:13:97:bb:7f:77
<b>OCID</b>	...fj3cqsm6v2dzrb3sytky
<input type="button" value="Show Full OCID"/> <input type="button" value="Copy"/>	

Virtual Cloud Networks > VCN-Exadata2 > Ext-Priv-SN3

**Ext-Priv-SN3**

Subnet Information Tags

General Information

**Created**  
03/21/2025, 04:12:39 PM

**CIDR Block**  
10.100.103.0/24

**Virtual Router Mac Address**  
00:13:97:3c:af:54

**OCID**  
...63m8qicz52enwdfk8q

Show Full OCID Copy

Virtual Cloud Networks > VCN-Exadata2 > Exa-DRG-RT

**Ext-Pub-SN4**

Subnet Information Tags

General Information

**Created**  
03/19/2025, 03:12:57 PM

**CIDR Block**  
10.100.104.0/24

**Virtual Router Mac Address**  
00:13:97:b4:3c:0a

**OCID**  
...a1xgi598b3tyckj6khkt

Show Full OCID Copy

- Note the OCID of each subnet, you need these OCIDs to enable communication between the VMs and the Flex network in step 5 of this tutorial.
- Create Route Tables in the VCN. Route tables are required to send traffic outside the VCN.

Navigate to the VCN in the Compute Web UI, select Route Tables in the Resources menu. Click Create Route Table, type a name for the Route Table and click Create Route Table.

Then create any associated route rules by clicking Add Route Rules and entering a target and the destination CIDR block.

This example shows a route table that directs network traffic for the DRG, and a route table that directs network traffic for the internet gateway.

Virtual Cloud Networks > VCN-Exadata2 > Exa-DRG-RT

### Exa-DRG-RT

Route Table Information Tags

General Information

Created  
03/19/2025, 03:16:57 PM

Compartment  
[pca3DocTenant](#)

OCID  
...oh5m8n5hfuj6kmhqp  
[Show Full OCID](#) [Copy](#)

Resources

Route Rules (1)

Route Rules

Destination	Target Type	Destination Type	Target	Description	Actions
10.80.65.0/24	Dynamic Routing Gateway	CIDR Block	DRG1		⋮

Virtual Cloud Networks > VCN-Exadata2 > Ext-IGW-RT

### Ext-IGW-RT

Route Table Information Tags

General Information

Created  
03/21/2025, 04:18:32 PM

Compartment  
[pca3DocTenant](#)

OCID  
...p9wjua59blrnup6uj5n  
[Show Full OCID](#) [Copy](#)

Resources

Route Rules (1)

Route Rules

Destination	Target Type	Destination Type	Target	Description	Actions
0.0.0.0/0	Internet Gateway	CIDR Block	Ext-IGW		⋮

- Create Security Lists to enable ingress traffic. See [Controlling Traffic with Security Lists](#).

Virtual Cloud Networks > VCN-Exadata2 > Default Security List for VCN-Exadata2

### Default Security List for VCN-Exadata2

Security List Information Tags

General Information

Created  
03/19/2025, 03:01:04 PM

Compartment  
pca3DocTenant

OCID  
...8j0x3Zzrq9mz0zymr

Show Full OCID Copy

Resources

Ingress Rules (4)

Egress Rules (1)

Ingress Rules

Stateless	Source	Source Type	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Actions
No	10.100.0.0/16	CIDR Block	ICMP			3: Destination Unreachable	⋮
No	0.0.0.0/0	CIDR Block	ICMP			3: Destination Unreachable, 4: Fragmentation Needed and Don't Fragment was Set	⋮
No	0.0.0.0/0	CIDR Block	TCP		22 - 22		⋮
No	0.0.0.0/0	CIDR Block	All				⋮

About Oracle | Contact Us | Copyright Information | Terms of Use | Your Privacy Rights

COPYRIGHT © 2022, 2025, ORACLE AND/OR ITS AFFILIATES.

- From the Compute Enclave, create DRG-attachments to enable the VCN to use the DRG. See [Attach VCNs to a Dynamic Routing Gateway](#)

From the VCN page, select Dynamic Routing Gateway from the left menu, then click Attach Dynamic Routing Gateway. Select the appropriate tenancy, then choose the DRG you want to attach from the dropdown list and click Attach Dynamic Routing Gateway.

Virtual Cloud Networks > VCN-Exadata2

### VCN-Exadata2

VCN Information Tags

General Information

CIDR Block  
10.100.0.0/16

Compartment  
pca3DocTenant

Created  
03/19/2025, 03:01:04 PM

OCID  
...zldlewbh33it0md4pfa

Show Full OCID Copy

Default Route Table  
[Default Route Table for VCN-Exadata2](#)

DNS Domain Name

Resources

Dynamic Routing Gateways

Name	Dynamic Routing Gateway	State	Route Table	Created	Actions
drgattachment20250319150712	DRG1	Attached		03/19/2025, 03:07:12 PM	⋮

- From the Service CLI, enable communication between the Flex network and the VM subnets.

```
PCA-ADMIN> flexNetworkEnableAccess flexNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Command: flexNetworkEnableAccess flexNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
```

Status: Success  
Time: 2024-11-17 18:56:45,251 UTC

7. Create VMs in the subnets and configure their access. See [Compute Instance Deployment](#)

For this example, each private subnet is configured with 2 VNICs: a primary and a secondary. Configure primary VNICs to attach to the public subnet, and secondary VNICs to attach to the DRG.

When creating an instance, choose the VCN then the subnet. For the public VM, assign a public IP address.

The screenshot shows the Oracle Private Cloud Appliance console for a VM named 'Exa-VM-1'. The interface includes a breadcrumb trail: Virtual Cloud Networks > VCN-Exadata2 > Exa-Priv-SN1 > Exa-VM-1. A green icon representing a network port is labeled 'AVAILABLE'. The main configuration area is divided into sections: 'VNIC Information' (with a 'Tags' tab), 'General Information', 'Primary IP Information', and 'IP Addresses'. The 'General Information' section includes fields for Created (03/26/2025, 03:46:54 PM), OCID (...yD8te291r13ov3rwoavt), Compartment (pca3DocTenant), Subnet (Exa-Priv-SN1), Skip Source/Destination Check (No), MAC Address (00:15:97:6a:0a:a6), and VLAN Tag (Not Available). The 'Primary IP Information' section includes Private IP Address (10.100.101.2), Private IP OCID (...xu12fdrrd95f4nidphhk), Private IP Assigned (05/26/2025, 03:46:54 PM), Fully Qualified Domain Name (exa-vm-1), Public IP Address (Not Set), and Public IP OCID (Not Set). The 'IP Addresses' table shows one entry for 'exa-vm-1' with a Private IP Address of 10.100.101.2 and a Public IP Address of 'Not Assigned'. A button 'Assign Secondary Private IP Address' is visible in the top right of the IP Addresses section.

Hostname	Private IP Address	Public IP Address	Created	Actions
exa-vm-1	10.100.101.2	Not Assigned	03/26/2025, 03:46:54 PM	⋮

From the Compute Instance page, select the instance. From the Resources menu select Attached VNICs. You will see the primary VNIC. To create the secondary VNIC, click Create VNIC Attachment. Choose the VCN and a subnet. and assign See [Creating and Attaching a Secondary VNIC](#).

Create primary and secondary VNICs for each instance such that the primary VINC attaches to the public subnet, and the secondary VNIC attaches to the DRG.

8. Verify connectivity between the VMs and the external system.

## Accessing External Interfaces with Your CA Trust Chain

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

You need the following three different CA certificates to access all external interfaces.

### ! Important

Update your "Regular uplink" CA certificate whenever a new service is added to Private Cloud Appliance.

- Admin-accessible
  - admin.<domain\_name>
  - adminconsole.<domain\_name>
  - alertmanager.<domain\_name>
  - api.<domain\_name>
  - grafana.<domain\_name>
  - prometheus.<domain\_name>
  - prometheus-gw.<domain\_name>
  - rps.<domain\_name>
- Regular uplink
  - auth.<domain\_name>
  - autoscaling.<domain\_name>

- backup-provider.<domain\_name>
- compute-containers.<domain\_name>
- console.<domain\_name>
- containerengine.<domain\_name>
- dns.<domain\_name>
- filestorage.<domain\_name>
- iaas.<domain\_name>
- identity.<domain\_name>
- limits.<domain\_name>
- network-load-balancer-api.<domain\_name>
- regionregistry.<domain\_name>
- regionrepository.<domain\_name>
- Object storage
  - objectstorage.<domain\_name>

The process to use your own CA trusted certificates is simple:

1. Create certificate signing requests (CSRs) on your Private Cloud Appliance.
2. Use these CSRs to generate certificates signed by your own CA.
3. Upload these CA certificates and your CA trust chain to Private Cloud Appliance.

## Create Certificate Signing Requests

To use your own CA, you must generate CSRs on Private Cloud Appliance and then use the CSRs to generate the certificates signed by your CA.

### Note

OpenSSH clients must be at least version openssh-clients-7.4p1 or later.

### Using the Service CLI

To generate the CSRs, use the `generateCustomerCsr` command.

1. Log into the Service CLI.
2. Run the `generateCustomerCsr` command:

```
PCA-ADMIN> generateCustomerCsr
Command: generateCustomerCsr
Status: Success
Time: 2023-05-17 18:43:55,904 UTC
Data:
status = success
message = Successfully generated customer csr:
Please download all CSR files from: /nfs/shared_storage/certs/customer_csr/
```

3. You can add Distinguished Names to the `generateCustomerCsr` command if needed:

```
PCA-ADMIN> generatecustomerCsr country=IN state=KA locality=blr \
organization=example organizationunit=adminexample.pca email=test@example.com
```

```
Command: generatecustomerCsr country=IN state=KA locality=blr \  
organization=example organizationunit=adminexample.pca email=test@example.com  
Status: Success  
Time: 2023-10-11 22:48:16,718 UTC  
Data:  
status = success  
message = Successfully generated customer csr:  
Please download all CSR files from: /nfs/shared_storage/certs/customer_csr/
```

Allowable Distinguished Names include country, state, locality, organization unit, and email.

You can find the newly-generated CSR files in the `/nfs/shared_storage/certs/customer_csr/` directory on the management node:

- `external_tls_term.csr.pem`
  - `external_admin_tls_term.csr.pem`
  - `zfssa.csr.pem`
4. Download the CSRs.
  5. Create certificates signed by your CA that are based on the CSRs.

### ! Important

When you generate your certificates you must add the FDQNs (and no IP addresses) from the SAN information in the CSRs.

If you supply outside certificates to establish a CA trust chain, you must add PTR records to the Data Center DNS. A PTR (Pointer record) in DNS maps an IP address to a hostname. This behavior is the reverse of the usual IP address lookup for a supplied hostname, which is provided by an A record in DNS.

You must create `ReverseIp` lookup zones for the two `ReplicationIps` used in disaster recovery. The DNS requests are forwarded to the Private Cloud Appliance in the same way as requests for the Private Cloud Appliance Service Zone are forwarded. If only the `zfsCapacityPoolReplicationEndpoint` is defined, then only a PTR record for that IP address is needed.

To create a `ReverseIp` lookup you need to create a DNS zone for the `ReverseIP` lookup. You create one or more reverse lookup zones depending on how the `Replication IPs` are configured. How to create these PTR records depends on the interface for the Data Center's DNS servers.

For example, if the rack domain is `myprivatecloud.example.com`, and the Capacity Pool IP is `10.170.123.98` and the Performance Pool IP is `10.170.123.99`, the Private Cloud Appliance requires two zones with the following mappings:

```
98.123.170.10.in-addr.arpa rtype PTR rdata sn01-dr1.myprivatecloud.example.com  
99.123.170.10.in-addr.arpa rtype PTR rdata sn02-dr1.myprivatecloud.example.com
```

For more information about DNS and PTR records, see the [Networking](#) section of the [Oracle Private Cloud Appliance User Guide](#).

You can proceed to the uploading process.

## Uploading Your CA Certificates

When you have the CA certificates, you must upload them along with the CA trust chain to the Private Cloud Appliance.

### Using the Service CLI

Use the `uploadCustomerCerts` command to upload the CA certificates. This command uses the following parameters to provide the full paths to the certificates and the CA trust chain:

- `caTrustChain`
- `externalAdminCert`
- `externalCert`
- `zfsCert`

1. Log into the Service CLI.
2. Copy the CA certificates you created in [Create Certificate Signing Requests](#) and your CA trust chain to the `/nfs/shared_storage` directory on the management node.
3. Run the `uploadCustomerCerts` command to upload all the CA certificates. For example:

```
PCA-ADMIN> uploadCustomerCerts externalcert=/nfs/shared_storage/external_tls_term.cert
zfsCert=/nfs/shared_storage/certs/zfssa.cert
caTrustChain=/nfs/shared_storage/CAPrivate.pem
Command: uploadCustomerCerts externalcert=/nfs/shared_storage/external_tls_term.cert
zfsCert=/nfs/shared_storage/certs/zfssa.cert
caTrustChain=/nfs/shared_storage/CAPrivate.pem
Status: Success
Time: 2023-05-17 18:43:55,904 UTC
Data:
status = success
message = Successfully uploaded customer CERTS
```

#### Important

Upload your CA trust chain with one of the CA certificate upload commands by using the `caTrustChain` parameter.

If your Private Cloud Appliance has the Admin networking feature enabled, the `uploadCustomerCerts` command requires the additional `externalAdminCert` parameter. For example:

```
PCA-ADMIN> uploadCustomerCerts externalcert=/nfs/shared_storage/external_tls_term.cert
zfsCert=/nfs/shared_storage/certs/zfssa.cert
caTrustChain=/nfs/shared_storage/CAPrivate.pem
externalAdminCert=/nfs/shared_storage/external_admin_tls_term.cert
Command: uploadCustomerCerts externalcert=/nfs/shared_storage/external_tls_term.cert
zfsCert=/nfs/shared_storage/certs/zfssa.cert
caTrustChain=/nfs/shared_storage/CAPrivate.pem
externalAdminCert=/nfs/shared_storage/external_admin_tls_term.cert
Status: Success
Time: 2023-05-17 18:49:03,904 UTC
Data:
status = success
message = Successfully uploaded customer CERTS
```

**Note**

If you need to backout your CA certificate and revert to an Oracle-supplied certificate, contact Oracle.

# 3

## Administrator Account Management

This chapter explains how the default administrator creates additional administrator accounts, and how the Service Enclave provides control over administrator account privileges, preferences and passwords.

Technical background information can be found in the [Oracle Private Cloud Appliance Concepts Guide](#). Refer to the section "Administrator Access" in the chapter "[Appliance Administration Overview](#)".

### Creating a New Administrator Account

During system initialization, a default administrator account is set up. This default account cannot be deleted. It provides access to the Service Enclave, from where additional administrator accounts can be created and managed.

#### Using the Service Web UI

1. Open the navigation menu and click Users.
2. Click Create User to open the Create User window.
3. Enter the following details:
  - **Name:** Enter a name for this administrator account. This is the name that will be used to log in.
  - **Authorization Group:** Select the authorization group to which the new administrator is added. This selection determines the access rights and privileges of the administrator account.
  - **Password:** Set a password for the new administrator account. Enter it a second time to confirm.
4. Click Create User. The new administrator account is displayed in the Users table.

#### Using the Service CLI

1. Display the list of authorization groups. Copy the ID of the authorization group in which you want to create the new administrator account.

```
PCA-ADMIN> list AuthorizationGroup
Command: list AuthorizationGroup
Status: Success
Time: 2021-08-25 08:38:58,632 UTC
Data:
  id Name
  ----
  9e6fef47-6ba7-4123-b25d-f9406173a609 OracleServiceAdmin
  2652ac1a-aa9e-4edf-bae7-d434efb23052 OCIAApp
  411ed79b-8f66-434b-862a-3c6e1b036fc4 SuperAdmin
  f5f9a82e-aa0a-4c31-a873-fae59fe20f38 Initial
PCA-ADMIN>
```

2. Create a new administrator account using the command `createUserInGroup`.

Required parameters are the user name, password and authorization group.

```
PCA-ADMIN> createUserInGroup name=testadmin password=***** confirmPassword=*****
authGroup=365ece7b-0a09-4a04-853c-7a0f6c4789f0
Command: createUserInGroup name=testadmin password=***** confirmPassword=*****
authGroup=365ece7b-0a09-4a04-853c-7a0f6c4789f0
Status: Success
Time: 2021-08-25 08:48:53,138 UTC
JobId: 6dd5a542-4399-4414-ac3b-636968744f79
```

3. Verify that the new administrator account was created correctly. Use the list and show commands to display the account information.

```
PCA-ADMIN> list User
Command: list User
Status: Success
Time: 2021-08-25 08:49:01,064 UTC
Data:
  id                name
  --                ----
  401fce73-5bee-48b1-b86d-fba1d85e049b  admin
  682ebc19-8493-4e9a-817c-148acea4b1d4  testadmin
```

```
PCA-ADMIN> show user name=testadmin
Command: show User name=testadmin
Status: Success
Time: 2021-08-25 08:50:04,245 UTC
Data:
  Id = 682ebc19-8493-4e9a-817c-148acea4b1d4
  Type = User
  Name = testadmin
  Default User = false
  AuthGroupIds 1 = id:365ece7b-0a09-4a04-853c-7a0f6c4789f0 type:AuthorizationGroup name:InternalGroup
  UserPreferenceId = id:1321249c-0651-49dc-938d-7764b9638ea9 type:UserPreference name:
```

## Changing Administrator Credentials

The administrator's password is set during account creation. You can always change the password for your own account. Depending on privileges, you may be authorized to change the password of another administrator as well.

### Using the Service Web UI

1. Open the navigation menu and click Users.
2. Click the administrator account for which you want to change the password. The user detail page is displayed.

Alternatively, to display your own user detail page, click your name in the top-right corner of the page and select My Profile.

3. Click Change Password to open the Change Password window.
4. Enter the new account password. Enter it a second time for confirmation. Click Save Changes to apply the new password.

### Using the Service CLI

1. Display the list of administrator accounts. Copy the ID of the account for which you want to change the password.

```
PCA-ADMIN> list User
Command: list User
```

```
Status: Success
Time: 2021-08-25 09:22:01,064 UTC
Data:
id                name
--                ----
401fce73-5bee-48b1-b86d-fba1d85e049b  admin
682ebc19-8493-4e9a-817c-148acea4b1d4  testadmin
```

2. Set a new password for the selected administrator account using the `changePassword` command.

```
PCA-ADMIN> changePassword id=682ebc19-8493-4e9a-817c-148acea4b1d4 password=*****
confirmPassword=*****
Command: changePassword id=682ebc19-8493-4e9a-817c-148acea4b1d4 password=***** confirmPassword=*****
Status: Success
Time: 2021-08-25 09:22:55,188 UTC
JobId: 35710cd9-26ac-4be9-8b73-c4cf634cc121
```

## Managing Administrator Privileges

An administrator is granted privileges through his membership in an authorization group or groups. When you create an administrator account, you select the authorization group to which the new administrator is added. However, you can change which authorization groups an administrator belongs to at any time.

For more information, see "Administrator Access" in the [Appliance Administration Overview](#) section of the [Oracle Private Cloud Appliance Concepts Guide](#).

### Using the Service Web UI

To add an administrator to an additional authorization group:

1. Open the navigation menu and click Authorization Groups.
2. Click the authorization group to which you want to add an administrator.
3. Under Resources, click Users and then click Add User to Group.
4. From the Add User to Group form, select an administrator and then click OK.

Before you can remove an administrator from an authorization group, you must make sure he belongs to at least one other group. To remove an administrator from an authorization group:

1. If the administrator belongs only to the authorization group you want to remove him from, add the administrator to another authorization group.
2. Open the navigation menu and click Authorization Groups.
3. Click the authorization group for which you want to remove an administrator.
4. Under Resources, click Users. The list of users in the authorization group is displayed.
5. From the list, click the Actions menu for the user you want to remove and then click Remove User from Group.

### Using the Service CLI

1. Gather the IDs of the administrator account you want to change, and the authorization groups involved in the configuration change.

```
PCA-ADMIN> list User
Command: list User
```

```
Status: Success
Time: 2021-08-25 09:22:01,064 UTC
Data:
id                name
--                ----
401fce73-5bee-48b1-b86d-fba1d85e049b  admin
682ebc19-8493-4e9a-817c-148acea4b1d4  testadmin
```

```
PCA-ADMIN> list AuthorizationGroup
Command: list AuthorizationGroup
Status: Success
Time: 2021-08-25 08:38:58,632 UTC
Data:
id                name
--                ----
587fc90d-3312-41d9-8be3-1ce21b8d9b41  MonitorGroup
c18cc6af-4ef8-4b1c-b85d-ee3b065f503e  DrAdminGroup
8f03faf2-c321-4455-af21-75cbffc269ef  AdminGroup
5ac65f5d-1f8c-42ea-a1de-95a1941f009f  Day0ConfigGroup
365ece7b-0a09-4a04-853c-7a0f6c4789f0  InitialGroup
7da8be67-758c-4cd6-8255-e9d2900c788e  SuperAdminGroup
```

2. To add an administrator to an authorization group, use the `add User` command.

```
PCA-ADMIN> add User id=682ebc19-8493-4e9a-817c-148acea4b1d4 to AuthorizationGroup
id=587fc90d-3312-41d9-8be3-1ce21b8d9b41
Command: add User id=682ebc19-8493-4e9a-817c-148acea4b1d4 to AuthorizationGroup
id=587fc90d-3312-41d9-8be3-1ce21b8d9b41
Status: Success
Time: 2021-08-25 08:49:54,062 UTC
JobId: 3facde6d-acb6-4fc4-84dc-93de88eea25c
```

3. Display the administrator account details to verify the changes you made.

```
PCA-ADMIN> show User name=testadmin
Command: show User name=testadmin
Status: Success
Time: 2021-08-25 08:50:04,245 UTC
Data:
Id = 682ebc19-8493-4e9a-817c-148acea4b1d4
Type = User
Name = testadmin
Default User = false
AuthGroupIds 1 = id:365ece7b-0a09-4a04-853c-7a0f6c4789f0 type:AuthorizationGroup name:InternalGroup
AuthGroupIds 2 = id:587fc90d-3312-41d9-8be3-1ce21b8d9b41 type:AuthorizationGroup name:MonitorGroup
UserPreferenceId = id:1321249c-0651-49dc-938d-7764b9638ea9 type:UserPreference name:
```

4. To remove an administrator from an authorization group, use the `remove User` command.

```
PCA-ADMIN> remove User name=testadmin from AuthorizationGroup id=587fc90d-3312-41d9-8be3-1ce21b8d9b41
Command: remove User name=testadmin from AuthorizationGroup id=587fc90d-3312-41d9-8be3-1ce21b8d9b41
Status: Success
Time: 2021-08-25 09:10:39,249 UTC
JobId: 44110d28-70af-4a42-8eb7-7d59a3bc8295
```

## Working with Authorization Groups

As an administrator, the specific functions you can perform is dependent on the *authorization group* to which you belong. Every authorization group must have at least one attached policy statement that allows users who belong to this group access to resources. An authorization group without a policy statement is valid, but its users would not have access to any resources.

You can create the policy statements immediately after you create the authorization group or you can add policy statements later. You can also list or delete policy statements using both the Service Web UI and Service CLI. Additionally, you can inactivate a policy statement using the Service CLI.

### Note

You cannot modify a policy statement. If you need to make changes to a policy statement, you must delete it and then recreate it.

For more information, see "Administrator Access" in the [Appliance Administration Overview](#) section of the [Oracle Private Cloud Appliance Concepts Guide](#).

### Using the Service Web UI

1. Open the navigation menu and click Authorization Group.
2. Click Create Group.
3. Enter a name using 1 to 255 characters, and then click Create Authorization Group.  
The new authorization group's details page displays.
4. Click Add Policy Statement. The Authorization Policy Statement Form window displays.
5. Enter a name using 1 to 255 characters.
6. Select an action: Inspect, Read, Use, or Manage.
7. Select a policy application:
  - Resources - Enter the resources you want the policy to apply to.
  - Function Family - Select one from the drop down.
  - Resource Family - Select one from the drop down.

### Note

For information on how to find the resource and function options, see the *Using the Service CLI* section.

8. Click Create Policy Statement.  
The new policy statement displays on the details page. Add up to 100 additional policy statements.

### Using the Service CLI

1. Create a new authorization group.

```
PCA-ADMIN> create AuthorizationGroup name=authors
Status: Success
Time: 2022-05-22 13:10:12,463 UTC
JobId: 14ea4d22-acf1-455d-a7a1-ec0a30f29671
Data:
id:c672d9c6-90ec-4776-bccb-caae128e86db name:authors
```

2. View the help for the create authpolicyStatement command.

```
PCA-ADMIN> create authpolicyStatement ?
*action
```

```
activeState
functionFamily
resourceFamily
resources
*on
```

3. Enter `showcustomcmds ?` to see options for resources, or enter `showallcustomcmds` to view options for functions, for example:

```
PCA-ADMIN> showcustomcmds ?
ASRBundle
ASRPhonehome
BackupJob
CnUpdateManager
ComputeInstance
ComputeNode
[...]
```

```
PCA-ADMIN> showallcustomcmds
Operation Name: <Related Object(s)>
-----
[...]
backup: BackupJob
changeIloMPassword: ComputeNode, ManagementNode
changePassword: ComputeNode, LeafSwitch, ManagementNode, ManagementSwitch, SpineSwitch, User, ZFSAppliance
clearFirstBootError: NetworkConfig
configZFSAdDomain: ZfsAdDomain
configZFSAdWorkgroup: ZfsAdDomain
createAdminAccount:
createUserInGroup: User
deletePlatformImage: PlatformImage
deprovision: ComputeNode
disableVmHighAvailability: PcaSystem
drAddComputeInstance: ComputeInstance
drAddSiteMapping: DrSiteMapping
[...]
```

### Note

For more information on resources and functions, see [Command Syntax](#) and [Base and Custom Commands](#).

4. Create a policy statement using `resources`, `functionFamily` or `resourceFamily`.

```
PCA-ADMIN> create authpolicyStatement action=manage resources=ComputeNode on authorizationGroup
id=c672d9c6-90ec-4776-bccb-caae128e86db
```

```
PCA-ADMIN> create authpolicyStatement action=manage authresourceFamily=rackops on authorizationGroup
id=c672d9c6-90ec-4776-bccb-caae128e86db
```

```
PCA-ADMIN> create authpolicyStatement action=manage authfunctionFamily=computeops on authorizationGroup
id=c672d9c6-90ec-4776-bccb-caae128e86db
```

5. View the details for the authorization group.

```
PCA-ADMIN> show authorizationGroup name=authors
Command: show authorizationGroup name=authors
Status: Success
Time: 2022-05-23 11:32:42,335 UTC
Data:
Id = c672d9c6-90ec-4776-bccb-caae128e86db
Type = AuthorizationGroup
```

```
Name = authors
Policy Statements 1 = dea601bf-9bfc-4b2c-a135-d98378e69c87(ACTIVE)-Allow authors to MANAGE ComputeNode
Is Predefined Authorization Group = false
AuthPolicyStatementIds 1 = id:4adde579-1f6a-49eb-a783-9478465f135e type:AuthPolicyStatement name:
AuthPolicyStatementIds 2 = id:be498a4e-3e0a-4cfa-9013-188542adb8e3 type:AuthPolicyStatement name:
```

To inactivate a policy statement:

1. View the help for the edit authpolicyStatement command.

```
PCA-ADMIN> edit authpolicyStatement ?
id=<object identifier>
```

2. Find the policy statement's ID using the show authorizationGroup name=*group-name* command.

```
PCA-ADMIN> show authorizationGroup name=authors
Command: show authorizationGroup name=authors
[...]
Policy Statements 1 = dea601bf-9bfc-4b2c-a135-d98378e69c87(ACTIVE)-Allow authors to MANAGE ComputeNode
Is Predefined Authorization Group = false
AuthPolicyStatementIds 1 = id:4adde579-1f6a-49eb-a783-9478465f135e type:AuthPolicyStatement name:
AuthPolicyStatementIds 2 = id:be498a4e-3e0a-4cfa-9013-188542adb8e3 type:AuthPolicyStatement name:
```

3. Using the ID of the policy statement (*AuthPolicyStatementIds Number = id:unique-identifier*) view the command to activate or inactivate the policy statement.

```
PCA-ADMIN> edit authpolicyStatement id=be498a4e-3e0a-4cfa-9013-188542adb8e3 ?
activeState
```

4. Inactivate the policy statement.

```
PCA-ADMIN> edit authpolicyStatement id=be498a4e-3e0a-4cfa-9013-188542adb8e3 activeState=inactive
Command: edit authpolicyStatement id=be498a4e-3e0a-4cfa-9013-188542adb8e3 activeState=inactive
Status: Success
Time: 2022-05-23 11:42:11,446 UTC
JobId: 842c444e-060d-461d-a4e0-c9cdd9f1d3c3
```

5. Verify the policy statement is inactive.

```
PCA-ADMIN> show authorizationGroup name=authors
Command: show authorizationGroup name=authors
Status: Success
Time: 2022-05-23 11:42:26,995 UTC
Data:
Id = c672d9c6-90ec-4776-bccb-caae128e86db
Type = AuthorizationGroup
Name = authors
Policy Statements 1 = 4adde579-1f6a-49eb-a783-9478465f135e(ACTIVE)-Allow authors to MANAGE ComputeNode
Policy Statements 2 = be498a4e-3e0a-4cfa-9013-188542adb8e3(INACTIVE)-Allow authors to MANAGE ComputeNode
Is Predefined Authorization Group = false
AuthPolicyStatementIds 1 = id:4adde579-1f6a-49eb-a783-9478465f135e type:AuthPolicyStatement name:
AuthPolicyStatementIds 2 = id:be498a4e-3e0a-4cfa-9013-188542adb8e3 type:AuthPolicyStatement name:
```

## Working with Authorization Families

Authorization families allow you to group resources and functions that make logical sense in the management of your appliance. There are two types of authorization families you can use in policy statements: Function Family and Resource Family.

For more information on resources and functions, see [Command Syntax](#) and [Base and Custom Commands](#).

For conceptual information on authorization groups, policies, and families, see "Administrator Access" in the [Oracle Private Cloud Appliance Concepts Guide](#).

### Using the Service Web UI

1. Open the navigation menu and click Authorization Families.
2. Click Create Authorization Family.
3. Select either authorization family type: Function Family or Resources Family.
4. Enter a name.
5. Enter the resources to include in the family.

#### Note

For information on how to find the resource and function options, see the *Using the Service CLI* section.

6. Click Create Family.

### Using the Service CLI

Create an authorization function family.

1. Display the options for the `create authfunctionFamily` command.

```
PCA-ADMIN> create authfunctionFamily ?
*name
*resources
```

2. Enter `showallcustomcmds` to view options for functions, for example:

```
PCA-ADMIN> showallcustomcmds
Operation Name: <Related Object(s)>
-----
[...]
backup: BackupJob
changeIloMPassword: ComputeNode, ManagementNode
changePassword: ComputeNode, LeafSwitch, ManagementNode, ManagementSwitch, SpineSwitch, User, ZFSAppliance
clearFirstBootError: NetworkConfig
configZFSAdDomain: ZfsAdDomain
configZFSAdWorkgroup: ZfsAdDomain
createAdminAccount:
createUserInGroup: User
deletePlatformImage: PlatformImage
deprovision: ComputeNode
disableVmHighAvailability: PcaSystem
drAddComputeInstance: ComputeInstance
drAddSiteMapping: DrSiteMapping
[...]
```

3. Create the authorization function family.

```
PCA-ADMIN> create authfunctionFamily name=cnops resources=ComputeNode.reset,ComputeNode.start,ComputeNode.stop
Command: create authfunctionFamily name=cnops resources=ComputeNode.reset,ComputeNode.start,ComputeNode.stop
Status: Success
Time: 2022-05-23 12:29:40,651 UTC
JobId: 4cd37ea7-161f-4b11-952f-ffa992a37d5f
Data:
id:ae0216da-20d1-4e03-bf65-c7898c6079b2 name:cnops
```

4. List the authorization function families.

```

PCA-ADMIN> list authfunctionFamily
Command: list authfunctionFamily
Status: Success
Time: 2022-05-23 12:29:57,164 UTC
Data:
id name
-----
7f1ac922-571a-4253-a120-e5d15a877a1e Initial
2185058a-3355-48be-851c-2fa0e5a896bd SuperAdmin
7f092ddd-1a51-4a17-b4e2-96c4ece005ec Day0
ae0216da-20d1-4e03-bf65-c7898c6079b2 cnops

```

Create an authorization resource family.

1. Display the options for the `create authresourceFamily` command.

```

PCA-ADMIN> create authresourceFamily ?
*name
*resources

```

2. Enter `showcustomcmds ?` to see options for resources, for example:

```

PCA-ADMIN> showcustomcmds ?
ASRBundle
ASRPhonehome
BackupJob
CnUpdateManager
ComputeInstance
ComputeNode
[...]

```

### Note

For more information on resources and functions, see [Command Syntax](#) and [Base and Custom Commands](#).

3. Create the authorization resource family.

```

PCA-ADMIN> create authresourceFamily name=rackops resources=ComputeNode,RackUnit
Command: create authresourceFamily name=rackops resources=ComputeNode,RackUnit
Status: Success
Time: 2022-05-23 11:52:37,751 UTC
JobId: eb49ac48-e3f3-4c2f-bf11-d5d18a066788
Data:
id:b54e4413-15bd-440e-b399-e2ab75f17c35 name:rackops

```

4. List the authorization resource families.

```

PCA-ADMIN> list authresourceFamily
Command: list authresourceFamily
Status: Success
Time: 2022-05-23 11:57:37,464 UTC
Data:
id name
-----
9aefc9c8-556d-42a4-9369-d7cdf0bf0c52 SuperAdmin
b591cc7b-b117-449e-af35-cb4fc6f0c213 Day0
87633db2-d724-45b6-97a5-30babb6c4869 cnops
b54e4413-15bd-440e-b399-e2ab75f17c35 rackops
a45c08b4-f895-4da8-87f4-c81ca0b2bf27 Initial

```

## Changing Administrator Account Preferences

When logged in to the Service CLI you can change certain settings for your own administrator account. Those changes take effect immediately and are persisted for all your future CLI connections.

However, you can also change settings temporarily for just your current CLI session. To do so, replace the object `UserPreference` with `CliSession` in the command examples below.

Setting	Options	Description
<code>alphabetizeMode</code>	YES, NO	Enable this setting to display any managed object's attributes in alphabetical order. The default setting is "No".
<code>attributeDisplay</code>	DISPLAYNAME, ATTRIBUTENAME	Use this setting to control whether the name of each object's attribute is displayed. The default setting is "displayName".
<code>endLineCharsDisplayValue</code>	CRLF, CR, LF	Specify the end-of-line character to be used when the CLI output consists of multiple lines. The default setting is "CRLF".
<code>outputMode</code>	VERBOSE, SPARSE, XML	Specify the CLI output format. The default setting is "Sparse".
<code>wsCallMode</code>	SYNCHRONOUS, ASYNCHRONOUS	Use this setting to determine whether the CLI output from a command is invoked synchronously or asynchronously. The default setting is "Asynchronous".
<code>wsTimeoutInSeconds</code>	<value>	When the CLI is set to "Synchronous" call mode, use this setting to determine how many seconds the CLI waits for a job returned by an operation to complete.

### Using the Service CLI

1. Display your current account preferences.

```
PCA-ADMIN> show UserPreference
Command: show UserPreference
Status: Success
Time: 2021-08-25 12:23:41,265 UTC
Data:
Id = ec433c0f-4208-4e92-859e-498218d0f5c9
Type = UserPreference
WS Call Mode = Asynchronous
Alphabetize Mode = No
Attribute Display = Display Name
End Line Characters Display Value = CRLF
Output Mode = Verbose
Command Wait Timeout In Seconds = 240
UserId = id:401fce73-5bee-48b1-b86d-fba1d85e049b type:User name:admin
```

2. Change the setting of your choice using the `edit userPreference` command.

```
PCA-ADMIN> edit UserPreference outputMode=XML
Command: edit UserPreference outputMode=XML
Status: Success
Time: 2021-08-25 12:32:02,102 UTC
JobId: 9d312d9b-6169-47cb-97d4-6a8984237fa0
```

3. Execute the same command for any other settings you wish to change.

4. Display your current account preferences again to verify the changes you made.

```
PCA-ADMIN> show UserPreference
Command: show UserPreference
Status: Success
Time: 2021-08-25 12:32:40,664 UTC
Data:
Id = ec433c0f-4208-4e92-859e-498218d0f5c9
Type = UserPreference
WS Call Mode = Asynchronous
Alphabetize Mode = No
Attribute Display = Display Name
End Line Characters Display Value = CRLF
Output Mode = Xml
Command Wait Timeout In Seconds = 180
UserId = id:401fce73-5bee-48b1-b86d-fba1d85e049b type:User name:admin
```

## Deleting an Administrator Account

This section describes how to delete an administrator account.

### Using the Service Web UI

1. Open the navigation menu and click Users.
2. Click the administrator account you want to delete. The user detail page is displayed.
3. Click Delete. Confirm the operation when prompted.

### Using the Service CLI

1. Look up the name and ID of the administrator account you want to delete.

```
PCA-ADMIN> list User
Command: list User
Status: Success
Time: 2021-08-25 08:49:01,064 UTC
Data:
id                name
--                ----
401fce73-5bee-48b1-b86d-fba1d85e049b  admin
682ebc19-8493-4e9a-817c-148acea4b1d4  testadmin
```

2. To delete the administrator account, use the delete User command followed by the account name or ID.

```
PCA-ADMIN> delete User name=testadmin
Command: delete user name=testadmin
Status: Success
Time: 2021-08-25 09:20:09,249 UTC
JobId: 56e9dfcb-6b64-4f9d-b137-171f538029d3
```

3. Verify that the deleted account is no longer displayed in the user list.

```
PCA-ADMIN> list User
Command: list User
Status: Success
Time: 2021-08-25 09:22:07,743 UTC
Data:
id                name
--                ----
401fce73-5bee-48b1-b86d-fba1d85e049b  admin
```

# Federating with Microsoft Active Directory

Many companies use an identity provider to manage user logins and passwords and to authenticate users for access to secure websites, services, and resources. To access the Oracle Private Cloud Appliance Service Web UI, users must also sign in with a user name and password. An administrator can *federate* with a supported identity provider so that each user can use their existing login and password, rather than having to create new credentials to access and use cloud resources.

Federation involves setting up a trust relationship between the identity provider and Private Cloud Appliance. When an administrator has established this relationship, federated users are prompted with a *single sign-on* when accessing the Service Web UI.

For more information, see "Federating with Identity Providers" in the chapter [Identity and Access Management Overview](#) of the Oracle Private Cloud Appliance Concepts Guide.

You can federate multiple Active Directory (AD) accounts with Private Cloud Appliance (for example, one for each division of the organization), but each federation trust that you set up must be for a *single* AD account. To set up a trust, you perform some tasks in the Private Cloud Appliance Service Web UI and some tasks in Active Directory Federation Services (ADFS).

Before you begin federating, make sure you already have:

- Installed and configured Microsoft Active Directory Federation Services for your organization.
- Set up groups in Active Directory that will map to groups in Private Cloud Appliance.
- Created users in Active Directory who will sign into the Private Cloud Appliance Service Web UI.

## Note

Consider naming Active Directory groups that you intend to map to Private Cloud Appliance groups with a common prefix to make it easy to apply a filter rule, for example, PCA\_Administrators, PCA\_NetworkAdmins, PCA\_InstanceLaunchers.

## Gathering Required Information from ADFS

To federate with Oracle Private Cloud Appliance you need to have the SAML metadata document and the names of the Active Directory (AD) groups that you want to map to Private Cloud Appliance groups.

1. Locate and download the SAML metadata document for your ADFS, which is by default at:

`https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml`

This is the document you will upload when you create the identity provider.

2. Make a note of all the AD groups that you want to map to Private Cloud Appliance groups.

**⚠ Caution**

Ensure that you have all the Private Cloud Appliance groups configured before you add AD as an identity provider.

## Verifying Identity Provider Self-Signed Certificates

**⚠ Caution**

You can skip this section if your ADFS certificate is signed by a known certificate authority because they should already exist in the Private Cloud Appliance certificate bundle.

The Oracle Private Cloud Appliance Certificate Authority (CA), is self-signed OpenSSL generated root and intermediate x.509 certificate. These CA certificates are used to issue x.509 server/client certificates allowing you to add outside Certificate Authority (CA) trust information to the rack. If you use a self-signed certificate for ADFS, you will need to add outside CA trust information from ADFS to the management nodes on the rack.

**📘 Note**

If you are using the `metadataUrl` property to create or update an identity provider, you will need to add the identity provider's web server's certificate chain to the Private Cloud Appliance outside CA bundle. See your identity provider's documentation on how to find the web server's certificate chain and then follow steps 3-8.

To add outside CA trust information, complete the following steps:

1. From a browser, enter the following URL and download the SAML metadata document for your ADFS, which is by default at:

```
https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml
```

2. Open the file in a text or XML editor and locate the signing certificate section, for example:

```
<KeyDescriptor use="signing">
<KeyInfo>
<X509Data>
<X509Certificate>
<!--CERTIFICATE IS HERE-->
</X509Certificate>
</X509Data>
</KeyInfo>
</KeyDescriptor>
```

3. Log on to management node 1 whose default name is `pcamn01`.
4. Navigate to `/etc/pca3.0/vault` and create a new directory named `customer_ca`.

**Note**

You can use this directory for multiple files. For example you can create a file for the identity provider certificate and one for the web server's certificate chain.

5. In the `customer_ca` directory, create a new file in PEM format.
6. Copy the certificate from the `FederationMetadata.xml` file, which is located within the `<X509Certificate>` tag, and paste into the new PEM file. Be sure to include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----`, for example:  

```
-----BEGIN CERTIFICATE-----  
<CERTIFICATE CONTENT>  
-----END CERTIFICATE-----
```
7. Save the file and close.
8. Run the following command to update the `ca_outside_bundle.crt` on all management nodes:  

```
python3 /usr/lib/python3.11/site-packages/pca_foundation/secret_service/cert_generator/cert_generator_app.py -copy_to_mns
```

## Managing Identity Providers

To federate with an identity provider in Oracle Private Cloud Appliance you create it in either the Service Web UI or the Service CLI and map account groups.

After you create your identity provider, you might have the need to make an update. For example, you will need to update your metadata XML file when it expires. You can also view all identity providers, view details of or delete an identity provider.

## Adding Active Directory as an Identity Provider

To federate with Active Directory (AD) in Oracle Private Cloud Appliance you must add it as an identity provider. At the same time, you can set up the group mappings or you can set them up later.

To add AD as an identity provider, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

1. Sign in with your Private Cloud Appliance login and password.
2. Open the navigation menu and click Identity Provider.
3. On the Identity Providers page, click Create Identity Provider.
4. On the Create an Identity Provider page, provide the following information:
  - a. **Display Name**  
The name that the federated users see when choosing which identity provider to use for signing in to the Service Web UI. This name must be unique across all identity providers and cannot be changed.
  - b. **Description**  
A friendly description of the identity provider.
  - c. **Authentication Contexts**  
Click Add Class Reference and select an authentication context from the list.

When one or more values are specified, Private Cloud Appliance (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the identity provider must contain an authentication statement with that authentication context class reference. If the SAML response authentication context does not match what is specified here, the Private Cloud Appliance authentication service rejects the SAML response with a 400.

**d. Encrypt Assertion (Optional)**

When enabled, the authorization service expects encrypted assertions from the identity provider. Only the authorization service can decrypt the assertion. When not enabled, the authorization service expects SAML tokens to be unencrypted, but protected, by SSL.

**e. Force Authentication (Optional)**

When enabled, users are always asked to authenticate at their identity provider when redirected by the authorization service. When not enabled, users are not asked to re-authenticate if they already have an active login session with the identity provider.

**f. Metadata URL**

Enter the URL for the FederationMetadata.xml document from the identity provider.

By default, the metadata file for ADFS is located at

`https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml`

**5. Click Create Identity Provider.**

Your new identity provider is assigned an OCID and is displayed on the Identity Providers page

After the identity provider is added, you must set up the group mappings between Private Cloud Appliance and Active Directory.

To set up group mappings, see [Creating Group Mappings](#).

## Updating an Identity Provider

To update an identity provider, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

**1. Open the navigation menu and click Identity Providers.**

A list of the identity providers is displayed.

**2. For the identity provider you want to update, click the Actions icon (three dots) and then click Edit.**

**3. Change any of the following information; however, be aware that changing this information can affect the federation:**

**a. Description**

**b. Authentication Contexts**

Add or delete a class reference.

**c. Encrypt Assertion**

Enable or disable encrypted assertions from the identity provider.

**d. Force Authentication**

Enable or disable redirect authentication from the identity provider.

**e. Metadata URL**

Enter the URL for a new FederationMetadata.xml document from the identity provider.

For more information, see step 4 in [Adding Active Directory as an Identity Provider](#).

4. Click Update Identity Provider.

## Viewing Identity Provider Details

The identity provider details page displays general information such as authentication contexts. It also provides the identity provider's settings, which include the redirect URL.

From this page, you can also edit the identity provider and manage the group mappings.

To view details for an identity provider, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

1. Open the navigation menu and click Identity Providers.

A list of the identity providers is displayed.

2. For the identity provider whose details you want to view, click the Actions icon (three dots) and then click View Details.

The identity provider details page is displayed.

## Listing Identity Providers

To list the identity providers, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

1. Open the navigation menu and click Identity Providers.

A list of the identity providers is displayed.

## Deleting an Identity Provider

If you want to remove the option for federated users to log into Private Cloud Appliance you must delete the identity provider, which also deletes all of the associated group mappings.

To delete an identity provider, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers is displayed.

2. For the identity provider you want to delete, click the Actions icon (three dots) and then click Delete.

3. At the Delete Identity Provider prompt, click Confirm.

## Working with Group Mappings for an Identity Provider

When working with group mappings, keep in mind the following:

- A given Active Directory group is mapped to a single Oracle Private Cloud Appliance group.
- Private Cloud Appliance group names cannot contain spaces and cannot be changed later. Allowed characters are letters, numerals, hyphens, periods, underscores, and plus signs (+).
- You can't update a group mapping, but you can delete the mapping and add a new one.

## Creating Group Mappings

After you have created an identity provider, you must create mappings from ADFS groups to Private Cloud Appliance groups.

To create a group mapping, follow the procedure for either the Service Web UI or the Service CLI. Repeat the steps for each identity provider group you want to map.

### Using the Service Web UI

1. Open the navigation menu and click IDP Group Mappings.  
A list of the identity provider group mappings is displayed.
2. Click Create Group Mapping.  
The IDP Group Mapping Form is displayed
3. In the Name field, enter a name for the IDP group mapping.
4. In the IDP Group Name field, enter the *exact* name of the identity provider group.
5. From the Admin Group Name list, select the Private Cloud Appliance group you want to map to the identity provider group.
6. Optionally, enter a Description of the group.
7. Click Create IDP Group Mapping.  
The new group mapping is displayed in the list.

## Updating a Group Mapping

To update a group mapping, follow the procedure for either the Service Web UI or the Service CLI. Repeat the steps for each group mapping you want to map.

### Using the Service Web UI

1. Open the navigation menu and click IDP Group Mappings.  
A list of the identity provider group mappings is displayed.
2. For the group mapping you want to update, click the Actions icon (three dots) and then click Edit.  
The IDP Group Mapping Form is displayed.
3. Modify any of the following fields; however, be aware that changing this information can affect the federation:
  - a. Name

- b. IDP Group Name
  - c. Admin Group Name
  - d. Description
4. Click Modify IDP Group Mapping.  
The updated group mapping is displayed in the list.

## Viewing Group Mappings

To view group mapping details, follow the procedure for either the Service Web UI or the Service CLI.

### Using the Service Web UI

1. Open the navigation menu and click IDP Group Mappings.  
A list of the identity provider group mappings is displayed.

## Deleting a Group Mapping

To delete a group mapping, follow the procedure for either the Service Web UI or the Service CLI. Repeat the steps for each identity provider group you want to delete.

### Using the Service Web UI

1. Open the navigation menu and click IDP Group Mappings.  
A list of the identity provider group mappings is displayed.
2. For the group mapping you want to delete, click the Actions icon (three dots) and then click Delete.
3. At the Deleting IDP Group Mapping prompt, click Confirm.

## Adding Private Cloud Appliance as a Trusted Relying Party in ADFS

To complete the federation process, you must add Private Cloud Appliance as a trusted relying party in ADFS and then add associated relying party claim rules.

### Add Relying Party in ADFS

1. In the Service Web UI on the Identity Providers page, view the following text block:  
**You need the Private Cloud Appliance Federation Metadata document when setting up a trust with Microsoft Active Directory Federation Services or with other SAML 2.0-compliant identity providers. This is an XML document that describes the Private Cloud Appliance endpoint and certificate information. Click Here**
2. Click "Click Here".  
A metadata XML file opens in the browser with a URL similar to:  
`https://adminconsole.system-name.domain-name/wsapi/rest/saml/metadata/`
3. Copy the metadata XML file URL.
4. From the system installed with ADFS, open a browser window and paste the URL.
5. Save the file, making sure to use the .xml extension, for example, *my-sp-metadata.xml*.
6. Go to the AD FS Management Console and sign in to the account you want to federate.

7. Add Private Cloud Appliance as a trusted relying party.
  - a. Under AD FS, right-click Relying Party Trusts and then select Add Relying Party Trust.
  - b. In the Add Relying Party Trust Wizard Welcome page, select Claims Aware and then click Start.
  - c. On the Select Data Source page, select "Import data about the relying party from a file".
  - d. Click Browse and navigate to your *my-sp-metadata.xml* and then click Open.
  - e. On the Specify Display Name page, enter a display name, add any optional notes for the relying party, and then click Next.
  - f. On the Choose Access Control Policy page, select the type of access you want to grant and then click Next.
  - g. On the Ready to Add Trust page, review the settings, and then click Next to save your relying party trust information.
  - h. On the Finish page, check "Configure claims issuance policy for this application" and then click Close.

The Edit Claim Issuance Policy dialog appears, which you can leave open for the next section.

### Adding Relying Party Claim Rules

After you add Private Cloud Appliance as a trusted relying party, you must add the claim rules so that the elements required (Name ID and groups) are added to the SAML authentication response.

To add a Name ID rule:

1. In the Edit Claim Issuance Policy dialog, click Add Rule.

The Select Rule Template dialog is displayed.
2. For Claim rule template, select Transform an Incoming Claim and then click Next.
3. Enter the following:
  - **Claim rule name:** Enter a name for this rule, for example, nameid.
  - **Incoming claim type:** Select Microsoft Windows account name.
  - **Outgoing claim type:** Select a claim type, for example, Name ID.
  - **Outgoing name ID format:** Select Persistent Identifier.
  - Select Pass through all claim values and then click Finish.

The rule is displayed in the rules list.

The Issuance Transform Rules dialog displays the new rule.

If your Active Directory users are in no more than 100 groups, you simply add the groups rule. However, if your Active Directory users are in more than 100 groups, those users cannot be authenticated to use the Private Cloud Appliance Service Web UI. For these groups, you must apply a filter to the groups rule.

### To add the groups rule:

1. In the Issuance Transform Rules dialog, click Add Rule.

The Select Rule Template dialog is displayed.

2. For Claim rule template, select Send Claims Using a Custom Rule and then click Next.
3. In the Add Transform Claim Rule Wizard, enter the following:
  - a. **Claim rule name:** Enter groups.
  - b. **Custom rule:** Enter the custom rule.

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"] => issue(store = "Active Directory", types = ("https://auth.oraclecloud.com/saml/claims/groupName"), query = ";tokenGroups;{0}", param = c.Value);
```

- c. Click Finish.

The Issuance Transform Rules dialog displays the new rule.

### Disable the Certificate Revocation Check

For ADFS to work with SAML, you must disable the Certificate Revocation List (CRL) checking.

1. Open Powershell on the ADFS system and enter the following command, where *TRUST\_NAME* is the name of the relying party trust:

```
Get-AdfsRelyingPartyTrust -Name '<TRUST_NAME>' | Set-AdfsRelyingPartyTrust -EncryptionCertificateRevocationCheck None -SigningCertificateRevocationCheck None
```

## Providing Federated Users Sign In Information

Before federated users can log in to the Private Cloud Appliance Service Web UI, you must provide them with the URL. You must also ensure that you have configured the groups mappings otherwise a federated user cannot do any work in Private Cloud Appliance.

# 4

## Tenancy Management

A tenancy is an environment where users create and manage cloud resources in order to build and configure virtualized workloads. At least one tenancy must be created. All the tenancies in the environment are collectively referred to as the Compute Enclave. However, tenancy management is a responsibility of the appliance administrator. Tenancies are created from the Service Enclave and subsequently handed over to the initial user in the tenancy: the primary tenancy administrator.

Technical background information about enclaves, tenancies and administrator roles can be found in the [Oracle Private Cloud Appliance Concepts Guide](#). Refer to the section "Enclaves and Interfaces" in the chapter [Architecture and Design](#).

### Creating a New Tenancy

An infrastructure administrator sets up a tenancy from the Service Enclave and provides access details to the primary tenancy administrator. Then the tenancy administrator can start configuring additional user accounts and cloud resources in the Compute Enclave.

#### Using the Service Web UI

1. In the navigation menu, click Tenancies.
2. In the top-right corner of the Tenancies page, click Create Tenancy.  
The Create Tenancy window appears.
3. Fill out the tenancy details:
  - **Name:** Enter a name for the new tenancy.
  - **Description:** Optionally, enter a description for the new tenancy.
  - **Service Namespace:** Set a unique namespace for all resources created within this tenancy.
  - **Authentication Credentials:** Set a user name and password for the primary tenancy administrator.

This account must be used to log in to the tenancy for the first time. The tenancy administrator sets up additional user accounts, defines compartments, policies and other resources, and generally configures the cloud environment so that users can start deploying their required resources.

4. Click Save Changes to create the new tenancy.  
The new tenancy is displayed in the Tenancies list.

#### Using the Service CLI

1. Create a new tenancy with the `create Tenant` command.

The name, namespace and admin account credentials are required parameters; a description is optional.

Syntax (entered on a single line):

```

create Tenant
name=<tenancy_name>
serviceNamespace=<tenancy_namespace>
description=<tenancy_description>
adminUserName=<tenancy_admin_user_name>
adminPassword=<tenancy_admin_password>
confirmPassword=<tenancy_admin_password>

```

**Example:**

```

PCA-ADMIN> create Tenant name=myTestTenancy serviceNamespace=test description="A tenancy for testing purposes" \
adminUserName=testadmin adminPassword=***** confirmPassword=*****
Command: create Tenant name=myTestTenancy serviceNamespace=test description="A tenancy for testing purposes"
adminUserName=testadmin adminPassword=***** confirmPassword=*****
Status: Success
Time: 2021-09-08 08:54:44,778 UTC
JobId: a0ee398f-5d44-4b3f-8b9c-e5a9692c36a4
Data:
  id:ocid1.tenancy.....<uniqueID> name:myTestTenancy

```

**2. Use the job ID to check the status of your command.**

```

PCA-ADMIN> show Job id=a0ee398f-5d44-4b3f-8b9c-e5a9692c36a4
Command: show Job id=a0ee398f-5d44-4b3f-8b9c-e5a9692c36a4
Status: Success
Time: 2021-09-08 08:55:11,125 UTC
Data:
  Id = a0ee398f-5d44-4b3f-8b9c-e5a9692c36a4
  Type = Job
  AssociatedObj = id:ocid1.tenancy.unique_ID type:Tenant name:myTestTenancy
  AssociatedObj Type = Tenant
  AssociatedObj Id = ocid1.tenancy.unique_ID
  Done = true
  Name = CREATE_TYPE
  Run State = Succeeded
  Transcript = null2021-09-08 08:54:44.753 : Created job CREATE_TYPE

  Username = admin

```

**3. Verify that the new tenancy was created correctly. Use the list and show commands to display the tenancy information.**

```

PCA-ADMIN> list Tenant
Command: list Tenant
Status: Success
Time: 2021-09-08 08:55:44,669 UTC
Data:
  id          name
  --          ---
  ocid1.tenancy.unique_ID myTenancy1
  ocid1.tenancy.unique_ID myTenancy2
  ocid1.tenancy.unique_ID myTenancy3
  ocid1.tenancy.unique_ID myTestTenancy

PCA-ADMIN> show Tenant name=myTestTenancy
Command: show Tenant name=myTestTenancy
Status: Success
Time: 2021-09-08 08:56:09,484 UTC
Data:
  Id = ocid1.tenancy.unique_ID
  Type = Tenant
  Name = myTestTenancy
  Description = A tenancy for testing purposes
  Service Namespace = test

```

4. Provide the Compute Web UI URL, tenancy name, user name and password to the primary tenancy administrator. The tenancy is now ready for use.

The tenancy administrator sets up additional user accounts, defines compartments, policies and other resources, and generally configures the cloud environment so that users can start deploying their required resources.

## Providing Platform Images

Platform images are provided during Private Cloud Appliance installation, and new platform images might be provided during appliance upgrade or patching operations.

During installation, upgrade, and patching, new platform images are placed on the management node in `/nfs/shared_storage/oci_compute_images`. During patching and upgrade, you can run commands to make these images available to Compute Enclave users. See the `patchOCIimages` command in "Patching Oracle Cloud Infrastructure Images" in the [Oracle Private Cloud Appliance Patching Guide](#), and the `upgradeOCIimages` command in "Upgrading Oracle Cloud Infrastructure Images" in the [Oracle Private Cloud Appliance Upgrade Guide](#).

The image import command described in [Importing Platform Images](#) also makes the images available to Compute Enclave users. Run this `importPlatformImages` command if images were not imported during patch or upgrade, or you need to re-import images. You can also use this command to make custom images available to all Compute Enclave users after you put the image in `/nfs/shared_storage/oci_compute_images` on the management node.

During upgrade and patching, new versions of an image do not replace existing versions on the management node. If more than three versions of an image are available on the management node, only the newest three versions are shown when images are listed in the Compute Enclave. Older platform images are still available to users by specifying the image OCID.

### Importing Platform Images

Run the `importPlatformImages` command to make all images that are in `/nfs/shared_storage/oci_compute_images` on the management node also available in all compartments in all tenancies in the Compute Enclave.

```
PCA-ADMIN> importPlatformImages
Command: importPlatformImages
Status: Running
Time: 2022-11-10 17:35:20,345 UTC
JobId: f21b9d86-ccf2-4bd3-bab9-04dc3adb2966
```

Use the `JobId` to get more detailed information about the job. In the following example, no new images have been delivered:

```
PCA-ADMIN> show job id=f21b9d86-ccf2-4bd3-bab9-04dc3adb2966
Command: show job id=f21b9d86-ccf2-4bd3-bab9-04dc3adb2966
Status: Success
Time: 2022-11-10 17:35:36,023 UTC
Data:
  Id = f21b9d86-ccf2-4bd3-bab9-04dc3adb2966
  Type = Job
  Done = true
  Name = OPERATION
  Progress Message = There are no new platform image files to import
  Run State = Succeeded
  Transcript = 2022-11-10 17:35:20.339 : Created job OPERATION
  Username = admin
```

## Listing Platform Images

Use the `listplatformImages` command to list all platform images that have been imported from the management node.

```
PCA-ADMIN> listplatformImages
Data:
id          displayName          lifecycleState
--          -
ocid1.image.unique_ID uln-pca-Oracle-Linux-7.9-2024.05.29_0.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux-8-2024.05.29_0.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux-9-2024.05.29_0.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux8-OKE-1.26.6-20240611.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux8-OKE-1.27.7-20240602.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux8-OKE-1.28.3-20240602.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Solaris-11-2024.05.07_0.oci  AVAILABLE
```

Compute Enclave users see the same `lifecycleState` that `listplatformImages` shows. Shortly after running `importPlatformImages`, both `listplatformImages` and the Compute Enclave might show new images with `lifecycleState` `IMPORTING`. When the `importPlatformImages` job is complete, both `listplatformImages` and the Compute Enclave show the images as `AVAILABLE`.

If you delete a platform image as shown in [Deleting Platform Images](#), both `listplatformImages` and the Compute Enclave show the image as `DELETING` or `DELETED`.

## Deleting Platform Images

Use the following command to delete the specified platform image. The image shows as `DELETING` and then `DELETED` in `listplatformImages` output and in the Compute Enclave, and eventually is not listed at all. However, the image file is not deleted from the management node, and running the `importPlatformImages` command re-imports the image so that the image is again available in all compartments.

```
PCA-ADMIN> deleteplatformImage imageId=ocid1.image.unique_ID_7
JobId: 401567c3-3662-46bb-89d2-b7ad1541fa2d

PCA-ADMIN> listplatformImages
Data:
id          displayName          lifecycleState
--          -
ocid1.image.unique_ID uln-pca-Oracle-Linux-7.9-2024.05.29_0.oci  AVAILABLE
ocid1.image.unique_ID uln-pca-Oracle-Linux-8-2024.05.29_0.oci  AVAILABLE
[...]
ocid1.image.unique_ID uln-pca-Oracle-Solaris-2024.05.07_0.oci  DELETED
```

# Modifying the Configuration of a Tenancy

The only tenancy property that an administrator can modify at this time is the description.

- **Service Web UI:** Open the tenancy detail page and click Edit.
- **Service CLI:** Use the command `edit Tenant name=<tenancy_name> description=<tenancy_description>`

## Deleting a Tenancy

Make sure that tenancy users have removed all their resources. The tenancy can only be deleted if it is empty.

## Using the Service Web UI

1. In the navigation menu, click Tenancies.
2. In the tenancies table, click the name of the tenancy you want to delete.  
The tenancy detail page is displayed.
3. In the top-right corner of the tenancy detail page, click Delete. Confirm the operation when prompted.

## Using the Service CLI

1. Look up the name and ID of the tenancy you want to delete.

```
PCA-ADMIN> list Tenant
Command: list Tenant
Status: Success
Time: 2021-09-08 11:08:17,042 UTC
Data:
id          name
--          ---
ocid1.tenancy.unique_ID myTenancy1
ocid1.tenancy.unique_ID myTenancy2
ocid1.tenancy.unique_ID myTenancy3
ocid1.tenancy.unique_ID myTestTenancy
```

2. To delete the tenancy, use the delete Tenant command followed by the tenancy name or ID.

```
PCA-ADMIN> delete Tenant name=myTestTenancy
Command: delete Tenant name=myTestTenancy
Status: Running
Time: 2021-09-08 11:10:00,288 UTC
JobId: 92b84ac2-1f2c-41d7-980e-d7549957ef93
```

3. Verify that the deleted tenancy is no longer displayed in the tenancy list.

```
PCA-ADMIN> list Tenant
Command: list Tenant
Status: Success
Time: 2021-09-08 11:11:20,358 UTC
Data:
id          name
--          ---
ocid1.tenancy.unique_ID myTenancy1
ocid1.tenancy.unique_ID myTenancy2
ocid1.tenancy.unique_ID myTenancy3
```

# 5

## Viewing and Setting Resource Limits

The Limit service enables you to view and change (override) limits that are currently set for Private Cloud Appliance resources.

You can create resource limit templates to easily apply a set of resource limits.

The following procedures describe how to list the services that are supported by the Limit Service. These are the services that allow you to override some limits and view the current limits.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Service.
2. Select the Configuration tab to see the list of services that are supported by the Limit Service.

### Using the Service CLI

Use the `show LimitService` command to list the services that support viewing and setting resource limits. The following example lists only two services that support viewing and setting resource limits: IAM and Load Balancer. More services are listed when you run this command.

```
PCA-ADMIN> show LimitService
Command: show LimitService
Status: Success
Time: 2025-05-06 16:51:31,049 UTC
Data:
Id = 175b1cc5-bfad-4990-940f-632d9263fa3d
Type = LimitService
Supported Services 1 = identity
Supported Services 2 = loadbalancer
Overall Communication State = Okay
Name = Limit Service
Work State = Normal
```

## Viewing Resource Limits

This section describes how to list the available resource limit definitions, show the full specified resource limit definition, and show the default and current limit values of the specified resources.

## Listing and Showing Limit Definitions

This topic describes how to list all available limit definitions for all supported services and show a description of each limit. This topic also describes how to show the details of a specified limit definition, including the default value of the limit, supported tags, and whether quotas are supported. See [Supported Tags and Quotas Supported](#).

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Definitions.

For each limit definition, the table shows the name of the limit definition, the name of the service, and the description of the limit.

- To show the details of a particular limit definition, either select the name of the limit definition in the list, or select the Actions menu for that limit definition and select View Details.

The Resources section of the details page shows overrides that have already been created for this resource limit. You can edit or delete an override in that list, view details of an override, or use the button above the list to create a new override for this resource limit. See [Creating Resource Limit Overrides](#).

## Using the Service CLI

To list all available limit definitions, use the list `LimitDefinition` command. The following example lists only two limit definitions. Many more limit definitions are listed when you run this command.

```
PCA-ADMIN> list LimitDefinition
Command: list LimitDefinition
Status: Success
Time: 2025-05-06 15:29:26,212 UTC
Data:
  id          Name      Service Name  Description
  --          -
identity~groups  groups  identity      Maximum number of user groups
loadbalancer~backends  backends  loadbalancer  Maximum number of backends per load balancer (system-wide)
```

To show the details of a specified limit definition, use the show `LimitDefinition` command. Provide the resource limit id from the list `LimitDefinition` output.

```
PCA-ADMIN> show LimitDefinition id=identity~groups
Command: show LimitDefinition id=identity~groups
Status: Success
Time: 2025-05-06 16:50:22,144 UTC
Data:
Id = identity~groups
Type = LimitDefinition
Name = groups
Service Name = identity
Description = Maximum number of user groups
Supported Tags 1 = {TENANCY}
Quotas Supported = false
Value = 100
PCA-ADMIN>
PCA-ADMIN> show LimitDefinition id=loadbalancer~backends
Command: show LimitDefinition id=loadbalancer~backends
Status: Success
Time: 2025-05-06 16:50:22,144 UTC
Data:
Id = loadbalancer~backends
Type = LimitDefinition
Name = backends
Service Name = loadbalancer
Description = Maximum number of backends per load balancer
Supported Tags 1 = SYSTEM
Value = 512
```

The "Value" is the default limit for this resource. For the current limit value (if an override has been set), see the example in [Viewing Current Resource Limits](#).

## Supported Tags and Quotas Supported

Supported Tags are tags that you can use to override values. Curly braces indicate variable tags. For example, {TENANCY} indicates that you can set different limit values on different tenancies. In the preceding Service CLI example, you can set different groups limits for different tenancies. You cannot set different backends limits for different load balancers; all load balancers have the same backends limit.

If Quotas Supported is Yes or true, a Compute Enclave user with sufficient authorization can create additional compartment-specific limits within the specified tenancy.

## Viewing Current Resource Limits

This topic describes how to show the current limit values of the specified resources. If an override has been set, the current limit value could be different from the default value.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Service Limits.
2. On the Service Name menu at the top of the page, select the service of the resource that you are interested in.
3. On the Tag menu, select the tag.

All resources that are supported in the Limits Service, that belong to the selected service, and that have the specified tag are listed.

The "Value" column shows the current value of the limit.

The "Assigned By" column shows whether the value is the default limit value from the limit definition or an override value set in a template or on a tag (see [Setting Resource Limit Overrides](#)). If the value is an override value set in a template, the name of the template is shown.

Any of the columns can be sorted by selecting the arrow keys in the column header. Select the Refresh button to re-load the original list sorted alphabetically by limit name.

### Using the Service CLI

The `getCurrentServiceLimits` command shows all current limits for the specified service and tag or shows the specified current limit.

The following example shows all current limits for the IAM service and the specified tenancy.

```
PCA-ADMIN> getCurrentServiceLimits serviceName=identity tag=ocid1.tenancy.unique_ID
Command: getCurrentServiceLimits serviceName=identity tag=ocid1.tenancy.unique_ID
Status: Success
Time: 2025-05-06 16:52:23,329 UTC
Data:
Limit Name  Service Name  Value  Assigned By  Tag  Template Name
-----
users      identity     200    Template    ocid1.tenancy.unique_ID  template_name
groups     identity     200    Tag Override  ocid1.tenancy.unique_ID  null
...
api-keys   identity     3      Default     ocid1.tenancy.unique_ID  null
```

The "Assigned By" column shows whether the "Value" is the default limit value from the limit definition or an override value set with `create LimitOverride` in a template or on a tag (see [Setting Resource Limit Overrides](#)). If the value is an override value set in a template, the name of the template is shown.

Provide the limit name to show only that particular limit for the specified service and tag.

```
PCA-ADMIN> getCurrentServiceLimits serviceName=identity tag=ocid1.tenancy.unique_ID limitName=users
Command: getCurrentServiceLimits serviceName=identity tag=ocid1.tenancy.unique_ID limitName=users
Status: Success
Time: 2025-05-06 16:52:23,329 UTC
Data:
Limit Name  Service Name  Value  Assigned By  Tag  Template Name
-----
users      identity     200    Template     ocid1.tenancy.unique_ID template_name
```

## Working with Resource Limit Templates

Resource limit templates enable you to easily apply a set of resource limits. For example, you can create a template to apply one set of resource limits to two different tenancies and create another template to apply a second set of resource limits to three other tenancies. You can apply no more than one template to a tag.

Templates only need to specify limit values that are different from the default. If a limit is not specified, then its value is the default value from the limit definition if the template's default value behavior is set to INHERIT.

The commands in this topic create a resource limit template, list the available templates, show more information about the specified template, assign a template to a tag, unassign a template, and delete the template.

## Viewing Resource Limit Templates

This topic describes how to list the available limit templates and show more information about the specified template.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Templates.

For each limit template, the table shows the name of the template, whether the template allows tag overrides, and the default value behavior: Inherit or Zero.

2. To show the details of a particular limit template, either select the name of the limit override in the list, or select the Actions menu for that limit override and select View Details.

The Resources section of the details page shows overrides that have already been created for this limit template. You can edit or delete an override in that list, view details of an override, or use the button above the list to create a new override for this template. See [Creating Resource Limit Overrides](#).

### Using the Service CLI

The list LimitTemplate command lists all available resource limit templates. Both templates shown in the following example allow you to override the limit values in the template.

```
PCA-ADMIN> list LimitTemplate
Command: list LimitTemplate
Status: Success
Time: 2025-05-06 16:57:24,796 UTC
Data:
id      Name      Allows Overrides  Default Value Behavior
--      -
template1  template1  true              INHERIT
template2  template2  true              ZERO
```

The `show LimitTemplate` command shows more information for the specified template, particularly overrides. Provide the template id from the output of the `list LimitTemplate` command.

In the following example, the default value of 100 for `identity~users` was overridden by using `create LimitOverride` (see [Setting Resource Limit Overrides](#)).

```
PCA-ADMIN> show LimitTemplate id=template1
Command: show LimitTemplate id=template1
Status: Success
Time: 2025-05-06 16:58:02,694 UTC
Data:
  Id = template1
  Type = LimitTemplate
  Name = template1
  Allows Overrides = true
  Default Value Behavior = INHERIT
  Overrides 1 - Limit Name = users
  Overrides 1 - Service Name = identity
  Overrides 1 - Value = 200
```

## Creating Resource Limit Templates

This topic describes how to create a resource limit template.

By default (Allow Overrides is True), if the same limit is set both in the template and directly on the tag, the limit that is set directly on the tag is used.

When Allow Overrides is false, limits set by the template are used; any limits set directly on the tag are ignored. If a limit for the tag to which the template is assigned is not explicitly set in the template, then the default value of that limit is used or 0 is used, depending on the value of Default Value Behavior.

By default (the value of Default Value Behavior is Inherit), any limit values that are not explicitly set in this template inherit their default value from their limit definition.

When the value of Default Value Behavior is Zero, any limit values that are not explicitly set in this template are set to 0: No new resources of that type can be created. Existing resources are not affected. Unassigning this template restores the ability to create new resources for the tag.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Templates.
2. Select the Create Limit Template button above the list.
3. On the Create Limit Template dialog, enter the following information:
  - Name: Maximum 40 characters.
  - Default Value Behavior: Select either Inherit or Zero. Inherit is selected by default.
  - Allow Overrides: Select either True or False. True is selected by default.
4. Select the Create Limit Template button in the dialog.

On the details page for this template, scroll to the Resources section to define limit overrides for this template. See [Creating Resource Limit Overrides](#).

### Using the Service CLI

Use the `create LimitTemplate` command to create a new template with the specified template name.

The following examples create resource limit templates. The `create LimitTemplate` command does not set any limit values. Use the `create LimitOverride` command, specifying a template instead of a tag, to set limit values for the specified template. See [Setting Resource Limit Overrides](#).

In the following example, `defaultValueBehavior` is set to INHERIT by default:

```
PCA-ADMIN> create LimitTemplate name=template1
Command: create LimitTemplate name=template1
Status: Success
Time: 2025-05-06 16:59:13,531 UTC
JobId: a2f6df33-e70e-408b-ae5-7583b24e53ce
```

In the following example, `defaultValueBehavior` is explicitly set to ZERO:

```
PCA-ADMIN> create LimitTemplate name=template2 defaultValueBehavior=ZERO
Command: create LimitTemplate name=template2
Status: Success
Time: 2025-05-06 16:59:13,531 UTC
JobId: a2f6df33-e70e-408b-ae5-7583b24e53ce
```

## Assigning Resource Limit Templates

This topic describes how to assign a limit template to a tag.

At most one limit template can be assigned to a given tag. If you assign a second template to a tag, the second template replaces the first template; you do not need to unassign the first template.

If you assign a template to a tag and a limit for that tag is directly applied to the tag, then the Allow Overrides setting determines which limit value is used:

- If Allow Overrides is false for the template, then the limit value specified in the template is used. If a limit value is not explicitly set in the template, then the default value of that limit is used or 0 is used, depending on the value of Default Value Behavior.
- If Allow Overrides is true for the template, then the limit value that was already specified on the tag is used.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Template Assignments.  
For each limit template assignment, the template name and the tag are shown.
2. To assign a limit template, select the Assign Limit Template button above the template assignments list.
3. On the Create Limit Template Assignment dialog, enter the following information:
  - Template Name: Select the template name from the drop-down menu.
  - Tag: Select the tag from the drop-down menu.
  - Allow Overrides: Select either True or False. True is selected by default.
4. Select the Create Limit Template Assignment button on the dialog.

### Using the Service CLI

Use the `assignLimitTemplate` command to assign the specified template to the specified tag.

```
PCA-ADMIN> assignLimitTemplate templateName=template1 tag=ocid1.tenancy.unique_IDI
Command: assignLimitTemplate templateName=template1 tag=ocid1.tenancy.unique_IDI
Status: Success
```

Time: 2025-05-06 11:17:53,557 UTC  
JobId: 4fe99f37-554d-4ce8-8084-639f6069cc44

Use the `list LimitTemplateAssignment` command to list all templates and the tags to which they are assigned. In the following example, `template1` is assigned to two different tags:

```
PCA-ADMIN> list LimitTemplateAssignment
Status: Success
Time: 2025-05-06 11:18:01,309 UTC
Data:
  Template Name  Tag
  -----
  template1     ocid1.tenancy.unique_ID1
  template1     ocid1.tenancy.unique_ID2
  template3     SYSTEM
```

## Unassigning Resource Limit Templates

This topic describes how to unassign a limit template from a tag.

When a template is unassigned, the limits set in that template are removed from the affected tags.

- If you specify only a template name, that template is unassigned from all tags to which it had been assigned.
- If you specify both a template name and a tag, the specified template is unassigned from the specified tag.
- If you specify only a tag, any template that had been assigned to that tag is unassigned from that tag.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Template Assignments.  
For each limit template assignment, the template name and the tag are shown.
2. To unassign a limit template, select the Actions menu for that template and then select the Unassign option, or select the Unassign Limit Templates button above the template assignments list.
3. On the Unassign Limit Templates dialog, enter the following information:
  - Template Name: Select the template name from the drop-down menu.
  - Tag: Select the tag from the drop-down menu.

You must specify at least one template name or tag. You can specify multiple template names or tags, or multiple template name and tag pairs. See the rules at the beginning of this topic.

4. Select the Unassign Limit Templates button on the dialog.

### Using the Service CLI

Use the `unassignLimitTemplate` command to unassign a limit template.

```
PCA-ADMIN> unassignLimitTemplate templateName=template1 tag=ocid1.tenancy.unique_ID2
Command: unassignLimitTemplate templateName=template1 tag=ocid1.tenancy.unique_ID2
Status: Success
Time: 2025-05-06 11:18:44,481 UTC
JobId: 48fa692d-a997-48da-8b60-309bc4be7ad1
```

## Deleting Resource Limit Templates

This topic describes how to delete a resource limit template. If the template overrides a limit that is set per tenancy, you might want to check which tenancies the template is assigned to. See [Assigning Resource Limit Templates](#).

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Templates.
2. For the limit template that you want to delete, select the Actions menu, and select the Delete option.
3. Confirm the delete operation.

### Delete a Resource Limit Template

Use the delete LimitTemplate command to delete the specified template.

```
PCA-ADMIN> delete LimitTemplate id=template2
Command: delete LimitTemplate id=template2
Status: Success
Time: 2025-05-06 17:01:36,806 UTC
JobId: 62ae5905-3150-4724-a921-727f61dcdcdc
```

## Setting Resource Limit Overrides

A resource limit override is a value that replaces the limit value that is currently set in the specified template or for the specified tag. Before you create a limit override, see [Viewing Resource Limits](#) to see the current limit value and the default limit value. See [Planning Resource Limit Settings](#) to see the allowable limit values.

This section describes how to list all resource limit overrides, set and edit overrides, and delete overrides.

## Listing Resource Limit Overrides

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Overrides.
 

For each limit override, the table shows the name of the limit, the name of the service, the limit value set by the override, the tag, and the template name if applicable.
2. To show the details of a particular limit override, either select the name of the limit override in the list, or select the Actions menu for that limit override and select View Details.

### Using the Service CLI

Use the list LimitOverride command to return all limit overrides, for both templates and tags.

```
PCA-ADMIN> list LimitOverride
Command: list LimitOverride
Status: Success
Time: 2025-05-06 16:54:04,609 UTC
Data:
id          Limit Name      Service Name Value Tag          Template Name
--          -
identity~users~TAG~ocid1.tenancy.unique_ID users          identity      200  ocid1.tenancy.unique_ID null
```

```
loadbalancer~load-balancers~TAG~SYSTEM load-balancers loadbalancer 10 SYSTEM null
loadbalancer~listeners~TEMPLATE~template3 listeners loadbalancer 10 null template3
```

## Creating Resource Limit Overrides

This topic describes how to create a resource limit override. The new limit value can be fewer than the current number of resources. Existing resources are not affected by setting a new limit. If the new limit is fewer than the current number, you will not be able to create a new one of those resources until their number drops below the new, lower, limit.

When you create an override, you specify either a tag or a template. If the same limit is set by specifying a tag name and also by specifying a template name, see [Assigning Resource Limit Templates](#) for the rule for which limit value will be used.

See [Planning Resource Limit Settings](#) for limit default values and allowable values. You might not be able to set a limit override to the maximum allowable value. The limit you can set depends on other configuration, including limit overrides that you have already set for related resources. See [Listing Resource Limit Overrides](#) for the list of limits that are already overridden.

If creating a limit override fails because you have attempted to set the limit too high, you will receive an error message with information about why the limit is too high and what the limit can be. Using this information, you might decide to override a different limit as well as, or instead of, the limit you were attempting to override.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Overrides.
2. Select the Create Limit Override button above the list.
3. On the Create Limit Override dialog, enter the following information:
  - Select the service name from the menu.
  - Select the limit name from the menu.
  - Enter the new value for the resource limit.
  - Select either a tag from the Tag menu or a template from the Template menu. If both are selected, the override defaults to the selected template.
4. Select the Create Limit Override button in the dialog.

### Using the Service CLI

Use the `create LimitOverride` command to create a resource limit override for either a template or a tag.

Provide the name of the limit, the name of the service, and the new resource limit value. Provide either a template name or a tag name.

The following example sets a limit override for a tenancy tag:

```
PCA-ADMIN> create LimitOverride limitName=users serviceName=identity value=200 tag=ocid1.tenancy.unique_ID
Command: create LimitOverride limitName=users serviceName=identity value=200 tag=ocid1.tenancy.unique_ID
Status: Success
Time: 2025-05-06 16:52:35,132 UTC
JobId: a5d75c66-8708-4355-b5ef-1a3213f9b690
```

The following example sets a limit override for a SYSTEM tag:

```
PCA-ADMIN> create LimitOverride limitName=load-balancers serviceName=loadbalancer value=10 tag=SYSTEM
Command: create LimitOverride limitName=load-balancers serviceName=loadbalancer value=10 tag=SYSTEM
Status: Success
Time: 2025-05-06 16:53:35,132 UTC
JobId: job_ID
```

The following example sets a limit override in a template:

```
PCA-ADMIN> create LimitOverride limitName=listeners serviceName=loadbalancer value=10 template=template3
Command: create LimitOverride limitName=listeners serviceName=loadbalancer value=10 template=template3
Status: Success
Time: 2025-05-06 16:54:53,132 UTC
JobId: job_ID
```

## Updating Resource Limit Overrides

This topic describes how to edit a limit override value. See [Planning Resource Limit Settings](#) for limit default values and allowable values. You might not be able to set a limit override to the maximum allowable value. The limit you can set depends on other configuration, including limit overrides that you have already set for related resources.

If the current limit is set by a template, see [Working with Resource Limit Templates](#) to see what other limits are set in that template and what tags the template is assigned to.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Overrides.
2. For the limit override that you want to edit, select the Actions menu, and select the Edit option.
3. The Update Limit Override dialog displays the current limit value in a text field. You can change that displayed value.
4. Select the Update Limit Override button in the dialog.

### Using the Service CLI

Use the `edit LimitOverride` command to reset the value of the specified resource limit. Use the `id` from the list `LimitOverride` command output. See [Listing Resource Limit Overrides](#).

```
PCA-ADMIN> edit LimitOverride id=loadbalancer~listeners~TEMPLATE~template3 value=8
Command: edit LimitOverride id=loadbalancer~listeners~TEMPLATE~template3 value=8
Status: Success
Time: 2025-05-06 16:55:31,903 UTC
JobId: 37baa077-6d75-4901-86c5-82568c514b26
```

## Deleting Resource Limit Overrides

This topic describes how to delete a resource limit override. Deleting the override deletes the limit value for the tag or template.

- The following rules describe what the value of the limit will be when more than one limit value is set and one is deleted:
- If you delete the limit value of a tag, and a template with an override for the same limit is assigned to that tag, then the limit value from the template is used.
- If you delete the limit value of a tag, and no template is assigned to that tag, then the default limit value from the limit definition is used.

- If you delete the limit value of a tag, and a template is assigned to that tag that does not include an override for the same limit, then the default limit value from the limit definition is used if `defaultValueBehavior` is `INHERIT` for the template.
- If you delete the limit value of a template that is assigned to a tag, and an override for the same limit is assigned to that tag, then the limit value assigned to the tag is used if `allowsOverrides` is `true` for the template.
- If you delete the limit value of a template that is assigned to a tag, and no other override for the same limit is assigned to that tag, then the default limit value from the limit definition is used if `defaultValueBehavior` is `INHERIT` for the template.

### Using the Service Web UI

1. On the navigation menu, select Limits, and then select Limit Overrides.
2. For the limit override that you want to delete, select the Actions menu, and select the Delete option.
3. Confirm the delete operation.

### Using the Service CLI

Use the `delete LimitOverride` command to delete the limit value for the tag or template.

Use the `id` from the `list LimitOverride` command output to delete the limit value.

The following example deletes the value for the load-balancers limit that is set for the SYSTEM tag:

```
PCA-ADMIN> delete LimitOverride id=loadbalancer~load-balancers~TAG~SYSTEM
Command: delete LimitOverride id=loadbalancer~load-balancers~TAG~SYSTEM
Status: Success
Time: 2025-04-12 16:57:39,847 UTC
JobId: 3bc59401-d8cd-41ba-ab75-140a8fcd087f
```

The following example deletes the value for the listeners limit that is set for `template3`:

```
PCA-ADMIN> delete LimitOverride id=loadbalancer~listeners~TEMPLATE~template3
Command: delete LimitOverride id=loadbalancer~listeners~TEMPLATE~template3
Status: Success
Time: 2025-04-12 16:58:39,847 UTC
JobId: job_ID
```

## Planning Resource Limit Settings

The tables in this topic show the maximum allowable limit for each resource and tag.

### Important

You might not be able to set a limit override to the maximum allowable value. The limit you can set depends on other configuration, including limit overrides that you have already set for related resources.

If creating a limit override fails because you have attempted to set the limit too high, you will receive an error message with information about why the limit is too high and what the limit can be. Using this information, you might decide to override a different limit as well as, or instead of, the limit you were attempting to override.

In the following example, the maximum value for load-balancers is 12, and the user attempts to set a new value of 14.

```
PCA-ADMIN> create LimitOverride serviceName=loadbalancer limitName=load-balancers value=14 tag=SYSTEM
Command: create LimitOverride serviceName=loadbalancer limitName=load-balancers value=14 tag=SYSTEM
Status: Failure
Time: 2025-05-06 18:08:33,621 UTC
Error Msg: PCA_GENERAL_000037: Error returned from Limit service. Code: 'InvalidParameter'.
Message: 'Rule 'lb' violated: {loadbalancer/load-balancers} <= 12'
```

Some of the limits in the following tables do not show a range in the Allowable Limit column. They show the same value in the Allowable Limit column that they show in the Default Limit column. These limit values cannot be overridden.

```
PCA-ADMIN> create LimitOverride serviceName=network limitName=subnet-count tag=SYSTEM value=30
Command: create LimitOverride serviceName=network limitName=subnet-count tag=SYSTEM value=30
Status: Failure
Time: 2025-05-06 18:34:16,008 UTC
Error Msg: PCA_GENERAL_000037: Error returned from Limit service. Code: 'NotSupportedError'.
Message: 'Limit 'network/subnet-count' does not support overrides.'
```

Use the procedures described in [Listing and Showing Limit Definitions](#) with the service names and limit names shown in the following tables to get additional information such as the limit description and supported tags.

### File System Limits

All Filesystem service limits have service name `file-system`.

The maximum allowable limit might not be possible in your environment. See the Important note at the beginning of this topic.

**Table 5-1 Filesystem Service Limits**

Limit Name	Default Limit	Allowable Limit
file-system-count	800	0 - 800
file-system-sys-count	800	0 - 800
mount-target-count	80	0 - 80
mount-target-sys-count	80	0 - 80

### IAM Service Limits

All IAM service limits have service name `identity`.

The maximum allowable limit might not be possible in your environment. See the Important note at the beginning of this topic.

**Table 5-2 IAM Service Limits**

Limit Name	Default Limit	Allowable Limit
users	100	0 - 800
groups	100	0 - 800
dynamic groups	50	0 - 400
compartments	50	0 - 400
policies	100	0 - 800
policy-statements	50	0 - 400

**Table 5-2 (Cont.) IAM Service Limits**

Limit Name	Default Limit	Allowable Limit
group-users	100	0 - 800
user-groups	100	0 - 800
idps	3	0 - 24
idp-group-mappings	100	0 - 800
tag-namespaces	100	0 - 800
tags	100	0 - 800
tag-defaults	5	0 - 40
cost-tracking-tags	10	0 - 80
api-keys	3	0 - 24

**Load Balancer Service Limits**

All Load Balancer service limits have service name `loadbalancer` and tag `SYSTEM`.

The maximum allowable limit might not be possible in your environment. See the Important note at the beginning of this topic.

**Table 5-3 Load Balancer Service Limits**

Limit Name	Default Limit	Maximum Allowable Limit
listeners	16	0 - 16
backend-sets	16	0 - 16
backends	512	0 - 512
hostnames	4	0 - 4
path-route-sets	16	0 - 16
network-security-groups	5	0 - 5
certificates	16	0 - 16
cipher-suites	16	0 - 16
load-balancers	12	0 - 12
load-balancers-min-total	36	36
load-balancers-total	144	0 - 144

Note that the `load-balancers-min-total` limit cannot be overridden. The limit value of `load-balancers-min-total` can only be 36. The only exception is if you assign a template to `SYSTEM` with no limit value set for `load-balancers-min-total` and `defaultValueBehavior` set to `ZERO`.

You can set `load-balancers-total` as high as 144, but the effective limit depends on how many compute nodes the system has. For example, if you have six compute nodes, the effective `load-balancers-total` is 72, even if you set the limit higher.

**Network Limits**

All Network service limits have service name `network`.

The maximum allowable limit might not be possible in your environment. See the Important note at the beginning of this topic.

**Table 5-4 Network Service Limits**

Limit Name	Default Limit	Allowable Limit
vcn-system-count	80	80
subnet-system-count	320	320
subnet-count	40	40
vcn-count	80	0 - vcn-system-count
nsg-count	100	0 - 100
nsg-rule-count	50	0 - 300

**Storage Controller Limits**

All Storage Controller service limits have service name `block-storage` and tag `SYSTEM`.

The maximum allowable limit might not be possible in your environment. See the Important note at the beginning of this topic.

**Table 5-5 Storage Controller Service Limits**

Limit Name	Default Limit	Allowable Limit
volume-count	32	0 - 32
vol-group-volume-count	32	0 - 32

# 6

## Status and Health Monitoring

The system health checks and monitoring data are the foundation of problem detection. All the necessary troubleshooting and debugging information is maintained in a single data store, and does not need to be collected from individual components when an issue needs to be investigated. The overall health of the system is captured in one central location: Grafana.

Oracle has built default dashboards and alerts into Grafana, as well as a mechanism to consult the logs stored in Loki. You should not change the dashboards that are delivered by Oracle because Oracle Support might need that information to resolve particular issues. You can create your own dashboards and alerts.

Implementation details and technical background information for this feature can be found in the [Oracle Private Cloud Appliance Concepts Guide](#). Refer to the section "Status and Health Monitoring" in the chapter [Appliance Administration Overview](#).

## Using Grafana

With Grafana, Oracle Private Cloud Appliance offers administrators a single, visual interface to the logs and metrics collected at all levels and across all components of the system.

This section provides basic guidelines to access Grafana and navigate through the logs and monitoring dashboards. For additional information about Grafana services and how to use Grafana, see the Oracle Systems blog [Oracle PCA X9-2 Monitoring and Alerting with Grafana](#).

### The Grafana Home Page

Do one of the following to access Grafana:

- Service Enclave admin user.

1. Log in to the Service Web UI.
2. On the right side of the dashboard, click the Monitoring tile.

The Grafana login page opens in a new browser tab. In the future, you can go directly to this Grafana login page as described in "Any Grafana user" below.

3. Enter your user name and password at the prompts.

You can create new users, and give the new users the direct URL to the Grafana login page. To create new Grafana users, see [Adding Grafana Users](#). For the direct URL to the Grafana login page, see "Any Grafana user" below.

- Any Grafana user.

You do not need to be a Service Enclave user.

1. Go to the Grafana login page.

The Grafana login page is [https://grafana.pca\\_name.your\\_domain/login](https://grafana.pca_name.your_domain/login), where *pca\_name* is the name of the Private Cloud Appliance.

2. Enter your user name and password at the prompts.

The Welcome panel on the Grafana home page contains many links to [grafana.com](https://grafana.com) for information about how to use Grafana, such as how to create your own dashboards, queries, and alerts. You can also find Grafana tutorials on Oracle Private Cloud Appliance 3.x on the [Oracle Learning YouTube channel](#) or search for Grafana on [Oracle Blogs](#).

On the left side of the home page is a vertical bar with icons that open the list of dashboards or the list of alerts, for example, or provide access to system logs as described in [Using Grafana Explore Queries](#). Your user icon near the bottom of the bar enables you to change your preferences settings or log out. The Grafana logo at the top of the bar takes you back to the Grafana home page.

### The Grafana Time Line

When logs and metrics are stored in Prometheus, they are given a time stamp based on the time and time zone settings of the Private Cloud Appliance. However, Grafana displays the time based on user preferences, which might result in an offset because you are in a different time zone.

Use the following instructions if you want to synchronize the time line in the Grafana visualizations with the time zone of the appliance:

1. Near the bottom of the vertical menu bar on the left side of the Grafana page, click your user account icon and click the Preferences option on the submenu that pops up.
2. In the Preferences section of the page, change the Timezone setting to the same time zone as the appliance.
3. Click the Save button at the bottom of that section to apply the change.

### Monitoring Multiple Private Cloud Appliance X9-2 Systems

If you need to deploy an external Grafana service with variable-driven dashboards to monitor multiple Oracle Private Cloud Appliance X9-2 systems, see the following resources:

- [Observability, Monitoring, and Alerting Across Multiple Oracle Private Cloud Appliance X9-2 Systems - Part 1](#).
- [Observability, Monitoring, and Alerting Across Multiple Oracle Private Cloud Appliance X9-2 Systems - Part 2](#).

## Adding Grafana Users

This section describes adding users and teams of users and granting permissions to use folders and dashboards.

To add a new user, perform the following procedure as the admin user:

1. In the vertical menu bar on the left side of the Grafana home page, click the Server Admin (shield) icon.
2. On the Server Admin drop down menu, click Users.
3. Click the New User button.
4. Enter the requested information, and click the Create User button.

By default, the new user has the Viewer role. You could modify the user to change the role. Another way to change a user's access is to add the user to a team that has the required access.

The following are the Grafana user roles:

**Admin:** Has access to all organization resources, including dashboards, users, and teams.

**Editor:** Can view and edit dashboards, folders, and playlists.

**Viewer:** Can view dashboards and playlists.

By default, an editor can edit all of the listed resources, and a viewer can view all of the listed resources. A user with the Admin role can grant or restrict permissions to specific resources for specific roles, teams, and users. For example, click the Permissions tab on a folder to change the permissions to that folder for the Editor or Viewer roles. Click the Add Permissions button on the Permissions tab to add permissions for specific users or teams.

To create a new team, perform the following procedure as the admin user:

1. In the vertical menu bar on the left side of the Grafana home page, click the Configuration (gear) icon.
2. On the Configuration drop down menu, click Teams.
3. Click the New Team button.
4. Enter the requested information, and click the Create button.
5. Click the Add Member button.
6. In the Add team member box, click the drop-down arrow and select the user you want to add to the team.
7. Click the Add to team button.
8. Click the Settings tab at the top of the page to modify team settings such as home dashboard and time zone.

Use folders to grant permissions to users and teams. Perform the following procedure as the admin user:

1. In the vertical menu bar on the left side of the Grafana home page, click the Dashboards (grid) icon.
2. On the Dashboards drop down menu, click Manage.
3. For the folder for which you want to grant teams and users permissions, click Go to folder.
4. At the top of the folder page, click the Permissions tab.
5. Click the Add Permission button.
6. In the Add Permission For box, select the team or user, and select the role for the user or for all users on the team.
7. Click the Save button.

You can also grant permissions for specific dashboards in a folder. Perform the following procedure as the admin user:

1. In the vertical menu bar on the left side of the Grafana home page, click the Dashboards (grid) icon.
2. On the Dashboards drop down menu, click Manage.
3. Click the name of the folder that contains the dashboard, and then click the dashboard.
4. At the top of the dashboard page, click the gear icon.
5. In the menu on the left side of the page, click Permissions.
6. Click the Add Permission button.
7. In the Add Permission For box, select the team or user, and select the role for the user or for all users on the team.

8. Click the Save button.
9. Click the Save Dashboard button.

## Using Grafana Dashboards

Oracle provides a number of predefined Grafana dashboards organized into folders. Use any of the following to display the list of folders of dashboards:

- The magnifying glass icon in the vertical menu bar on the left side the page
- The Dashboards > Manage option in the menu bar
- The dashboards Home button to the right of the Grafana logo at the top of the menu bar

Click a folder name or the arrow to the right of a folder name to show the dashboards in that folder.

Use the buttons at the top of the list to toggle between showing the list of folders and showing the list of all dashboards.

In the search field at the top of the page, enter text from the name of a folder or dashboard to show only those dashboards.

Click the name of a dashboard to show the content of that dashboard. On a dashboard, you can click the star to the right of the dashboard name at the top of the page to list this dashboard on the Grafana home page for faster access.

The dashboard shows information such as the query, graphs of the data collected over time, and alerts set for that data.

You are able to modify most dashboards, but note that Oracle Support might require that information. The Grafana home page contains links to information for how to create your own dashboards and queries rather than modify dashboards that were provided by Oracle. For your custom dashboards, first create one or more folders to keep these new dashboards separate from dashboards provided by Oracle.

## Using Grafana Alerts

Oracle provides a predefined a set of alerts. You can also add your own alerts. You can show only alerts that are in a specified state, such as Alerting. You can display detailed information about the alert, including the values that trigger the alert and that trigger a state change. In many cases, you can change these values.

If an alert is in the Alerting state, view the alert definition to determine what caused the alert to go to that state, and then use this information to evaluate the component that the alert is monitoring and determine what action might be needed.

### Browse Grafana Alerts

To view all alerts, click the bell icon in the vertical menu bar on the left side the page. An icon shows the status of each alert, and text below the alert name shows how long the alert has been in that status.

Enter text in the search field at the top of the list to show only alerts with that text in their names. Use the States list to show only alerts that are in the selected state: OK, Not OK, Alerting, No Data, Paused, Pending.

Use the "How to add an alert" button above the alerts list to create a new alert, or use information referenced on the Grafana home page to add or modify an alert, add a notification channel, and add a notification for a particular alert.

Click an alert name to see detailed information about the alert. This is the same page you see if you go to the dashboard, scroll to the metric, click the metric name, and select Edit.

Hover over the graph to list all data that is being monitored, for example each host, switch, device, or endpoint.

On the Alert tab below the graph, you can view and edit the rule. An alert rule consists of one or more queries and expressions, a condition, the frequency of evaluation, and optionally, the duration over which the condition is met. You can see how the alert state is set for various error conditions. You can send a notification message for this alert.

The state history button shows the last 50 state changes for this alert. Another button enables you to test the alert.

### Add or Configure Notification Channels

To add or configure notification channels, click the bell icon in the menu bar on the left side the page and then select the Notification channels option, or select the Notification channels tab at the top of the list of alerts.

To change the configuration of an existing notification channel, click the name of the channel. When you are finished making changes, click the Save button. Click the Test button to send a test notification.

To add a notification channel, go to the Notification channels tab, click the New channel button, and fill out the page. Click the Save button. Click the Test button to send a test notification. Click the Back button to cancel and not create a new notification channel.

### Configure Custom External Email Notifications

To configure email notification, open a service request (SR) for Oracle Support to do the initial configuration. When the initial configuration is complete, go to the Grafana alerts page, click the Notification channels tab, click the New channel button, and fill out the page, selecting Email in the Type field.

### Configure Custom External HTTP/HTTPS Notifications

To configure external HTTP or HTTPS based custom alerts, you must first configure the proxy for Grafana as shown in the following example.

Log in to the management node that owns the management virtual IP, and run the following command:

```
$ sudo curl -u admin_user_name -XPUT \  
'https://api.PCA_system_name.your_domain/v1/grafana/proxy/config?http-proxy=proxy_fqdn:proxy_port&https-  
proxy=proxy_fqdn:proxy_port'  
Enter host password for user 'admin_user_name':  
Grafana proxy config successfully updated!
```

Run the following command so that Grafana can still contact the internal Loki and Prometheus services:

```
$ sudo curl -u admin_user_name -XPUT \  
'https://api.PCA_system_name.your_domain/v1/grafana/proxy/config?no-proxy="sauron-sauron-prometheus,sauron-sauron-  
alertmanager,grafana-loki.loki.svc.cluster.local"'
```

The Grafana pod is restarted. Run the following command until you see that the Grafana pod (*sauron-sauron-grafana-unique\_ID*) is running:

```
$ kubectl get pods -n sauron
```

## Checking the Health and Status of Hardware and Platform Components

The hardware and platform layers form the foundations of the system architecture. Any unhealthy condition at this level is expected to have an adverse effect on operations in the infrastructure services. A number of predefined Grafana dashboards allow you to check the status of those essential low-level components, and see the real-time and historic details of the relevant metrics.

The dashboards described in this section provide a good starting point for basic system health checks, and for troubleshooting if issues are found. You might prefer to use different dashboards, metrics, and visualizations instead. The necessary data, collected across the entire system, is stored in Prometheus, and can be queried and presented through Grafana in many different ways.

Grafana Folder	Dashboard	Description
Service Monitoring	Server Stats	<p>This comprehensive dashboard displays telemetry data for the server nodes. It includes graphs for CPU and memory utilization, disk activity, network traffic, and so on.</p> <p>Some panels in this dashboard display a large number of time series in a single graph. Click to display a single time series, or hover over the graph to view detailed data at a specific time.</p>
PCA 3.0 Service Advisor	Platform Health Check	<p>This dashboard integrates the appliance health check mechanisms into the centralized approach that Grafana provides for logging and monitoring.</p> <p>By default, the Platform Health Check dashboard displays all health check services. Use the buttons above the Platform Health Check list to change the content of the list. Use the Platform Service list to select a single health checker. Use the Health Check Status list to display all results or only healthy results. Use the Filters list to select a filter and a value.</p> <p>Typically, if you see health check failures you want to start troubleshooting. For that purpose, each health check result contains a time stamp that serves as a direct link to the related Loki logs. To view the logs related to any health check result, click the time stamp.</p>
My Dashboards (Read Only)	Node Exporter Full	<p>This dashboard displays a large number of detailed metric panels for a single compute or management node. Use the Host button at the top of the page to display data for a different host.</p> <p>This dashboard could be considered a fine-grained extension of the Server Stats dashboard. The many different panels provide detailed coverage of the server node hardware status as well as the operating system services and processes. Information that you would typically collect at the command line of each physical node is combined into a single dashboard showing live data and its evolution over time.</p> <p>All dashboards in the My Dashboards folder provide data that would be critical in case a system-level failure needs to be resolved. Therefore, these dashboards cannot be modified or deleted.</p>

## Viewing and Interpreting Monitoring Data

The infrastructure services layer, which is built on top of the platform and enables all the cloud user and administrator functionality, can be monitored through an extensive collection of Grafana dashboards. These microservices are deployed across the three management nodes in Kubernetes containers, so their monitoring is largely based on Kubernetes node and pod metrics. The Kubernetes cluster also extends onto the compute nodes, where Kubernetes worker nodes collect vital additional data for system operation and monitoring.

The dashboards described in this section provide a good starting point for microservices health monitoring. You might prefer to use different dashboards, metrics and visualizations instead. The necessary data, collected across the entire system, is stored in Prometheus, and can be queried and presented through Grafana in many ways.

Grafana Folder	Dashboard	Description
Service Monitoring	ClusterLabs HA Cluster Details	<p>This dashboard uses a bespoke Prometheus exporter to display data for HA clusters based on Pacemaker. On each HTTP request it locally inspects the cluster status, by parsing preexisting distributed data provided by the cluster components' tools.</p> <p>The monitoring data includes Pacemaker cluster summary, nodes and resource stats, and Corosync ring errors and quorum votes.</p>
Service Monitoring	MySQL Cluster Exporter	<p>This dashboard displays performance details for the MySQL database cluster. Data includes database service metrics such as uptime, connection statistics, table lock counts, as well as more general information about MySQL objects, connections, network traffic, memory and CPU usage, etc.</p>
Service Monitoring	Service Level	<p>This dashboard displays detailed information about RabbitMQ requests that are received by the fundamental appliance services. It allows you to monitor the number of requests, request latency, and any requests that caused an error.</p>
Service Monitoring	VM Stats	<p>This comprehensive dashboard displays resource consumption information across the compute instances in your environment. It includes graphs for CPU and memory utilization, disk activity, network traffic, and so on.</p> <p>The panels in this dashboard display a large number of time series in a single graph. You can click to display a single time series, or hover over the graph to view detailed data at a specific point on the time axis.</p>
PCA 3.0 Service Advisor	Kube Endpoint	<p>This dashboard focuses specifically on the Kubernetes endpoints and provides endpoint alerts. These alerts can be sent to a notification channel of your choice.</p>
PCA 3.0 Service Advisor	Kube Ingress	<p>This dashboard provides data about ingress traffic to the Kubernetes services and their pods. Two alerts are built-in and can be sent to a notification channel of your choice.</p>

Grafana Folder	Dashboard	Description
PCA 3.0 Service Advisor	Kube Node	This dashboard displays metric data for all the server nodes, meaning management and compute nodes, that belong to the Kubernetes cluster and host microservices pods. You can monitor pod count, CPU and memory usage, and so on. The metric panels display information for all nodes. In the graph-based panels you can click to view information for just a single node.
PCA 3.0 Service Advisor	Kube Pod	This dashboard displays metric data at the level of the microservices pods, allowing you to view the total number of pods overall and how they are distributed across the nodes. You can monitor their status per namespace and per service, and check if they have triggered any alerts.
PCA 3.0 Service Advisor	Kube Service	This dashboard displays metric data at the Kubernetes service level. The data can be filtered for specific services, but displays all by default. Two alerts are built-in and can be sent to a notification channel of your choice.
Kubernetes Monitoring Kubernetes Monitoring Containers Kubernetes Monitoring Node	(all)	These folders contain a large and diverse collection of dashboards with a wide range of monitoring data that covers most of the operations of the Private Cloud Appliance system Kubernetes cluster. For example, these metrics provide information about deployment, ingress, and usage of CPU, disk, memory, and network resources.
OKE Monitoring	CAPOCI	This dashboard shows metrics from the Cluster API Provider for OCI (CAPOCI), which is a component of Oracle Private Cloud Appliance Kubernetes Engine (OKE). This dashboard monitors request status codes and response times for resources used by OKE such as compute instances and load balancers.  The information about controller reconciliation is for Oracle Support.
OKE Monitoring	Cluster Time Monitoring	This dashboard shows the time taken for operations such as create or update a particular OKE cluster or node pool. Average time for these operations across all clusters and node pools also is shown.
OKE Monitoring	Metrics Meter	This dashboard shows the health of various targets used by the OKE service such as the Cluster API Provider (CAPI), the Cluster API Provider for OCI (CAPOCI), OKE, and prometheus-k8s.
OKE Monitoring	OKE Service	This dashboard shows the service level metrics for OKE. Examples of metrics on this dashboard include counts of requests such as cluster and node pool create, update, and delete, and counts of exception codes for various requests. The exception code counts help expose any patterns in request failures.

## Monitoring System Capacity

It is important to track the key metrics that determine the system's capacity to host your compute instances and the storage they use. The detailed data for compute node load and

storage usage can be found in the Grafana dashboards. Administrators also have direct access to the current consumption of CPU and memory as well as storage space.

## Viewing CPU and Memory Usage By Fault Domain

These procedures display the number of compute nodes, the amount of total memory and free memory, and the number of total and free virtual CPUs for each fault domain.

The UNASSIGNED row refers to compute nodes that are not currently assigned to a fault domain. Because these compute nodes do not belong to a fault domain, their memory and CPU usage *in a fault domain* is zero.

To display this information and more for an individual compute node, select PCA Config > Rack Units from the navigation menu, or select the Rack Units tile on the Dashboard, and then click the name of a compute node in the list.

### Using the Service Web UI

1. In the navigation menu, select PCA Config > Fault Domains.
2. Click the name of a fault domain to see this information for only that fault domain.

### Using the Service CLI

Enter the `getFaultDomainInfo` command.

```
PCA-ADMIN> getFaultDomainInfo
Command: getFaultDomainInfo
Status: Success
Time: 2022-06-17 14:43:13,292 UTC
Data:
id      totalCNs totalMemory freeMemory totalvCPUs freevCPUs
--      -
UNASSIGNED 1      0.0      0.0      0      0
FD1      2      1072.0   976.0   176     164
FD2      1      984.0   984.0   120     120
FD3      1      984.0   984.0   120     120
```

The Notes column is omitted from the above example.

## Viewing Disk Space Usage on the ZFS Storage Appliance

The Service Enclave runs a storage monitoring tool called ZFS pool manager, which polls the ZFS Storage Appliance every 60 seconds. Using the Service CLI, you can display current information about the usage of available disk space in each ZFS pool. You can also set the usage threshold that triggers a fault when the threshold is exceeded.

### Check the Storage Status of ZFS Pools

List ZFS pools.

```
PCA-ADMIN> list ZfsPool
Command: list ZfsPool
Status: Success
Time: 2022-10-10 08:44:11,938 UTC
Data:
id      name
--      ---
e898b147-7cf0-4bd0-8b54-e32ec83d04cb PCA_POOL
c2f67943-df81-47a5-9713-06768318b623 PCA_POOL_HIGH
```

In a standard storage configuration, you only have one pool. If your system includes high-performance disk trays, then you can view usage information for each pool separately.

```
PCA-ADMIN> show ZfsPool id=e898b147-7cf0-4bd0-8b54-e32ec83d04cb
Command: show ZfsPool id=e898b147-7cf0-4bd0-8b54-e32ec83d04cb
Status: Success
Time: 2022-10-10 08:44:22,051 UTC
Data:
Id = e898b147-7cf0-4bd0-8b54-e32ec83d04cb
Type = ZfsPool
Pool Status = Online
Free Pool = 44879343128576
Total Pool = 70506183131136
Pool Usage Percent = 0.3634693989163486
Name = PCA_POOL
Work State = Normal
```

### Configure the Fault Threshold of the ZFS Pool Manager

By default, the fault threshold is set to 80 percent full: usage percentage 0.8.

```
PCA-ADMIN> show ZfsPoolManager
Command: show ZfsPoolManager
Status: Success
Time: 2022-10-10 08:58:11,231 UTC
Data:
Id = a6ca861b-f83a-4032-91c5-bc506394d0de
Type = ZfsPoolManager
LastRunTime = 2022-10-09 12:17:52,964 UTC
Poll Interval (sec) = 60
The minimum Zfs pool usage percentage to trigger a major fault = 0.8
Manager's run state = Running
```

The following example sets the fault threshold to 75 percent full: usageMajorFaultPercent=0.75.

```
PCA-ADMIN> edit ZfsPoolManager usageMajorFaultPercent=0.75
Command: edit ZfsPoolManager usageMajorFaultPercent=0.75
Status: Success
Time: 2022-10-10 08:58:27,657 UTC
JobId: 67cfe180-f2a2-4d59-a676-01b3d73cfae
```

## Viewing System Log Data

Logs are collected from all over the system and aggregated in Loki.

### Using Grafana Explore Queries

System log data can be queried, filtered, and displayed using Grafana Explore queries.

### Loki Logs

Loki uses labels to categorize log messages. A query specifies labels, and Loki displays the service and application log messages that match the query selections.

Labels are key-value pairs. Use the following procedure to select labels for your query.

1. Open the Grafana home page.
2. Open the Explore pane.

In the vertical menu bar on the left side of the page, click Explore (the compass icon).

3. To query Loki data, select Loki from the Explore data source menu at the top of the page to the right of the "Explore" title.

Loki query options are displayed. For example, a Log Browser menu is shown at the top of the page.

4. Query and filter the logs.

The following methods are similar. Both methods allow you to select labels and values from lists. The second method enables you to more easily select multiple labels and multiple values for one query.

- [Enter a Query in the Text Field](#)
- [More Easily Build a Complex Query](#)

Once you have created a query, you can select the same query again from the history list.

Additional query options:

- Add query. Click the Add query button to create another query and show the result of all separate queries together in the same timeline and message list.
- Query history. Run a query that was previously run, or copy or delete the query, add a comment to the query, or star the query so that you can use the Starred button to list only starred queries. At the top of the Query history list you can enter a search string to filter the list, and you can select how to order the list.
- Recurring run. Click the arrow on the Run query button, and select an interval from the menu. To stop the recurring runs, select Off at the top of the menu.

The timeline is displayed below the Log browser section of the Explore pane. Below the timeline, the log messages that match the query are displayed.

Messages are color-coded both in the timeline and in the message list to indicate whether the message is informational, a warning, error, or other.

Use the Query type button to choose to show the results over a range of time or at just one point in time. Use the range button at the top of the page (see the clock icon) to set the range.

Select a portion of the timeline to zoom in to focus on a smaller amount of data. To zoom out, use the magnifying glass button at the top of the page next to the range button.

In the message list, click the arrow on the left side of the time stamp of a message to display all labels that match that message. You can then click the plus + magnifying glass icon to add that label to your query results or click the minus - magnifying glass icon to remove that label from your query results. Notice that the query that you entered changes.

### Enter a Query in the Text Field

1. In the text field to the right of the Log browser button, enter the open brace { character.

The closed brace is automatically added, and a list of labels pops up.

2. Select a label from the list.

You might need to scroll the list to see all labels, or you can start typing a label name to filter the list.

The selected label is inserted into the query in the text field, an equals sign is added, and a list of values for that label pops up.

3. Select a value from the list.

You might need to scroll the list to see all values, or you can start typing a value name to filter the list.

The selected value is inserted inside quotation marks.

4. If you want to further filter the query result, enter a comma.

The list of labels pops up again, followed by the list of values after you select a label.

5. Run the query.

Type Shift+Enter, or click the Run query button in the upper right corner of the pane.

The timeline and log messages are displayed below the query building options.

### More Easily Build a Complex Query

Click the Log browser button so that the arrow on the button points down.

A query builder is displayed with the following steps:

1. Select labels.

Step 1 displays a row of buttons with a label name on each button. When you click one of these label buttons, a list pops up under Step 2 that shows the values for that label.

You can click more than one label button. If you click another label button, the list of values for the new label pops up with the first list of values under Step 2.

When you click a label button that is already selected, that label is removed from the query.

2. Choose values for the selected labels.

Step 2 shows the list of values for each label that is selected in Step 1. You might need to scroll the list to see all possible values, or you can start typing a value name in the search field to filter all value lists.

When you select a value from one list, some values might be removed from another list.

You can select more than one value from a particular list. Selecting a value that is already selected removes that value from the query.

As you select or deselect values, the query is built and displayed in Step 3.

3. Show the query result.

Click the Show Logs button in Step 3.

The timeline and log messages are displayed below the query building options.

The completed query is displayed in the field to the right of the Log browser button. You can edit the query in the Log browser field and click the Run query button to show a new result.

## Audit Logs

The audit logs can be consulted as separate categories. From Log browser lists, you can select the following audit labels. As described in [Loki Logs](#), either enter the queries shown in the following list in the text field, or select `job` or `log` from the Log labels list, and then select one of the values shown in the following list. See also the example custom query immediately following this list.

- `job="vault-audit"`

Use this log label to filter for the audit logs of the Vault cluster. Vault, a key component of the secret service, keeps a detailed log of all requests and responses. You can view every authenticated interaction with Vault, including errors. Because these logs contain sensitive information, many strings within requests and responses are hashed so that secrets are not shown in plain text in the audit logs.

- `job="kubernetes-audit"`  
Use this log label to filter for the audit logs of the Kubernetes cluster. The Kubernetes audit policy is configured to log request metadata: requesting user, time stamp, resource, verb, etc. Request body and response body are not included in the audit logs.
- `job="audit"`  
Use this log label to filter for the Oracle Linux kernel audit daemon logs. The kernel audit daemon (`auditd`) is the userspace component of the Linux Auditing System. It captures specific events such as system logins, account modifications, and `sudo` operations.
- `log="audit"`  
Use this log label to filter for the audit logs of the ZFS Storage Appliance.

In addition to using the log labels from the list, you can also build custom queries. For example, to filter for the audit logs of the admin service and API service, enter the following query into the Log browser text field:

```
{job=~"(admin|api-server)" | json tag="tag" | tag=~"(api-audit.log|audit.log)"}
```

To execute, either type Shift+Enter, or click the Run query button in the upper right corner of the Explore pane.

## LBaaS Logs

The Load Balancer as a Service (LBaaS) logs can be consulted as separate categories. From Log browser lists, you can select the following audit labels. As described in [Loki Logs](#), either enter the queries shown in the following list in the text field, or select `job` or `log` from the Log labels list, and then select one of the values shown in the following list.

- `job="pca-lbctl"`  
Use this log label to filter for the load balancer controller logs. You can view every client request that is being served. These logs contain API parameters and will contain error details when applicable.
- `job="pcalbmgr"`  
Use this log label to filter for the load balancer instances (manager) logs. You can view every request that is being served. These logs primarily contain the load balancer's configuration and management.

In addition to using the log labels from the list, you can also build custom queries. For example, you can view the controller and manager logs together:

```
{job=~"pca-lbctl|pca-lbmgr"}
```

To execute, either type Shift+Enter, or click the Run query button in the upper right corner of the Explore pane.

## Using the Vector Service

You can use the Vector service to send the information you want from Loki to an external location that you specify.

Beginning with Private Cloud Appliance Release 3.0.2-b1261765, Vector is installed, configured, and enabled on the appliance by default.

To specify which data you want and where you want the data sent, log in to the currently active management node as the root user and customize the Vector configuration file.

The Vector configuration file is at the following location on the management nodes:

```
/nfs/shared_storage/log_streaming/pca_vector.yaml
```

Edit the configuration file to customize the sinks section. See the Vector [Sinks reference](#). The following is a sample `pca_vector.yaml` file:

```
# Copyright (c) 2024, Oracle and/or its affiliates.

# DO NOT TURN API OFF
# otherwise livenessProbe will fail
api:
  enabled: true
  # Bind to 0.0.0.0. Otherwise the API will not be exposed outside the container.
  address: "0.0.0.0:8686"

sources:
  fluentd_source:
    type: fluent
    address: "0.0.0.0:8080"
    mode: tcp
    encoding:
      codec: json

transforms:
  log_event:
    type: remap
    inputs:
      - fluentd_source
    source: |
      log(.)

sinks:
  loki_sink:
    type: loki
    inputs:
      - fluentd_source
    endpoint: http://your_external_location:3100
    encoding:
      codec: json
    labels:
      job: "vector"
      namespace: "default"
      system: "pca_name.example.com"
      filename: "{{tag}}"
```

In the endpoint value, *your\_external\_location* can be an IP address or a domain name. At this location, you could install Grafana or use other tools to filter, manipulate, and display the data.

The value of the filename label that is shown in the example causes the name of the source log file to be shown in the Vector Loki Sink output. You can then use that file name as a label to search within Loki and Grafana.

The following is an example Splunk sink:

```
sinks:
  splunk_sink:
    type: splunkheclogs
    inputs:
      - source_id
    endpoint: https://splunk_endpoint
    token: splunkhec-token
```

```
encoding:  
  codec: json  
tls:  
  ca_file: "/path/to/ca.pem"
```

The `splunk-hec-token` is required to send logs to Splunk. The `ca_file` is optional if you are using HTTPS. For more information about HTTP Event Collector and how to configure and use Splunk, see the [Splunk documentation](#).

The following command reports the status of the log streaming pod (the Vector service):

```
# kubectl get pods -n log-streaming  
NAME                READY STATUS RESTARTS AGE  
log-streamer-bc4d65d78-ndrsk 1/1   Running 0      14d
```

The following command prints the logs from the `log-streamer-bc4d65d78-ndrsk` pod, which has only one container. For more information about the `kubectl logs` command, see [kubectl logs](#) on the [kubernetes.io](#) site.

```
# kubectl logs log-streamer-bc4d65d78-ndrsk -n log-streaming
```

## Using Oracle Auto Service Request

Oracle Private Cloud Appliance is qualified for Oracle Auto Service Request (ASR). ASR is integrated with My Oracle Support. When specific hardware failures occur, ASR automatically opens a service request and sends diagnostic information. The appliance administrator receives notification that a service request is open.

Using ASR is optional: the service must be registered and enabled for your appliance.

## Understanding Oracle Auto Service Request

ASR automatically opens service requests when specific Private Cloud Appliance hardware faults occur. To enable this feature, the Private Cloud Appliance must be configured to send hardware fault telemetry to Oracle directly at <https://transport.oracle.com/>, to a proxy host, or to a different endpoint. For example, you can use a different endpoint if you have the ASR Manager software installed in your data center as an aggregation point for multiple systems.

When a hardware problem is detected, ASR submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the administrator is even aware the problem exists.

ASR detects faults in the most common hardware components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR does not detect all possible hardware faults, and it is not a replacement for other monitoring mechanisms, such as SMTP alerts, within the customer data center. ASR is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

An email message is sent to both the My Oracle Support email account and the technical contact for Private Cloud Appliance to notify them of the creation of the service request. A service request might not be filed automatically in some cases, for example if a loss of connectivity to ASR occurs. Administrators should monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been filed automatically.

For more information about ASR, consult the following resources:

- Oracle Auto Service Request web page: <https://www.oracle.com/servers/technologies/auto-service-request.html>.
- Oracle Auto Service Request release notes on My Oracle Support: [Doc ID 2152198.1](#).
- Oracle Auto Service Request quick start guide on My Oracle Support: [Doc ID 2852505.1](#).
- Oracle Auto Service Request user documentation: [https://docs.oracle.com/cd/E37710\\_01/index.htm](https://docs.oracle.com/cd/E37710_01/index.htm).

## Oracle Auto Service Request Prerequisites

Before you register for the ASR service, ensure the following prerequisites are satisfied.

1. Ensure that you have a valid My Oracle Support account.  
If necessary, create an account at <https://support.oracle.com/portal/>.
2. Ensure that the following are set up correctly in My Oracle Support:
  - Technical contact person at the customer site who is responsible for Private Cloud Appliance.
  - Valid shipping address at the customer site where the Private Cloud Appliance is located, so that parts are delivered to the site where they must be installed.
3. Verify connectivity to the Internet using HTTPS.  
For example, try curl to test whether you can access <https://support.oracle.com/portal/>.

## Registering Private Cloud Appliance for Oracle Auto Service Request

To register a Private Cloud Appliance as an ASR client, the appliance must be configured to send hardware fault telemetry to Oracle in one of the following ways:

- Directly at <https://transport.oracle.com/>
- To a proxy host
- To a different endpoint

An example of when you would use a different endpoint is if you have the ASR Manager software installed in your data center as an aggregation point for multiple systems.

When you register your Private Cloud Appliance for ASR, the ASR service is automatically enabled.

### Using the Service Web UI

1. Open the navigation menu and click ASR Phone Home.
2. Click the Register button.
3. Fill in the user name and password, then complete the fields for the Phone Home configuration that you choose.
  - **Username:** Required. Enter your Oracle Single Sign On (SSO) credentials, which can be obtained from [My Oracle Support](#).
  - **Password:** Required. Enter the password for your SSO account.
  - **Proxy Username:** To use a proxy host, enter a user name to access that host.
  - **Proxy Password:** To use a proxy host, enter the password to access that host.
  - **Proxy Host:** To use a proxy host, enter the name of that host.

- **Proxy Port:** To use a proxy host, enter the port used to access the host.
- **Endpoint:** If you use an aggregation point, or other endpoint for ASR data consolidation, enter that endpoint in this format: `http://host[:port]/asr`

## Using the Service CLI

### Configure ASR directly to `https://transport.oracle.com/`

1. Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. Use the `asrClientRegister` custom command to register the appliance.

```
PCA-ADMIN> asrClientRegister ASRusername=asr-pca3_ca@example.com \  
password=***** confirmPassword=***** \  
endpoint=https://transport.oracle.com/ \  
Command: asrClientRegister username=asr-pca3_ca@example.com \  
password=***** confirmPassword=***** \  
endpoint=https://transport.oracle.com/ \  
Status: Success  
Time: 2021-07-12 18:47:14,630 UTC
```

3. Confirm the configuration.

```
PCA-ADMIN> show asrPhonehome  
Command: show asrPhonehome  
Status: Success  
Time: 2021-09-30 13:08:42,210 UTC  
Data:  
Is Registered = true  
Overall Enable Disable = true  
Username = asr.user@example.com Endpoint = https://transport.oracle.com/  
PCA-ADMIN>
```

### Configure ASR to a Proxy Host

1. Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. Use the `asrClientRegister` custom command to register the appliance.

```
PCA-ADMIN> asrClientRegister username=asr-pca3_ca@oracle.com \  
password=***** confirmPassword=***** \  
proxyHost=zeb proxyPort=80 \  
proxyUsername=support \  
proxyPassword=**** proxyConfirmPassword=**** \  

```

### Configure ASR to a Different Endpoint

1. Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. Use the `asrClientRegister` custom command to register the appliance.

```
PCA-ADMIN> asrClientRegister username=oracle_email@example.com \  
password=***** confirmPassword=***** \  
endpoint=https://transport.oracle.com/ \  
Command: asrClientRegister username=oracle_email@example.com \  
password=***** confirmPassword=***** \  
endpoint=https://transport.oracle.com/ \  
Status: Success  
Time: 2021-07-12 18:47:14,630 UTC
```

## Testing Oracle Auto Service Request Configuration

Once configured, test your Oracle Auto Service Request (ASR) configuration to ensure end-to-end communication is working properly.

### Using the Service Web UI

1. Open the navigation menu and click ASR Phone Home.
2. Select Test Registration in the Controls menu.
3. Click Test Registration. A dialog confirms whether the test is successful.
4. If the test is not successful, confirm your ASR configuration information and repeat the test.

### Using the Service CLI

1. Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. Use the `asrClientsendTestMsg` custom command to test the ASR configuration.

```
PCA-ADMIN> asrClientsendTestMsg  
Command: asrClientsendTestMsg  
Status: Success  
Time: 2021-12-08 18:43:30,093 UTC  
PCA-ADMIN>
```

## Unregistering Private Cloud Appliance for Oracle Auto Service Request

When you unregister your Private Cloud Appliance for ASR, the ASR service is automatically disabled; you do not need to perform a separate step.

### Using the Service Web UI

1. Open the navigation menu and click ASR Phone Home.
2. Click the Unregister button. Confirm the operation when prompted.

### Using the Service CLI

1. Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. Use the `asrClientUnregister` custom command to unregister the appliance.

```
PCA-ADMIN> asrClientUnregister  
Command: asrClientUnregister  
Status: Success  
Time: 2021-06-23 15:25:18,127 UTC  
PCA-ADMIN>
```

## Disabling Oracle Auto Service Request

You can disable ASR on an appliance to temporarily prevent fault messages from being sent and service requests created. For example, during system maintenance, components might be down but not failed or faulted. To restart the ASR service, see [Enabling Oracle Auto Service Request](#).

### Using the Service Web UI

1. Open the navigation menu and click ASR Phone Home.
2. Click the Disable button. Confirm the operation when prompted.

### Using the Service CLI

1. Using SSH, log into the management node VIP as admin.  

```
# ssh -l admin 100.96.2.32 -p 30006
```
2. Use the `asrClientDisable` custom command to halt the ASR service.

```
PCA-ADMIN> asrClientDisable  
Command: asrClientDisable  
Status: Success  
Time: 2021-06-23 15:26:17,753 UTC  
PCA-ADMIN>
```

## Enabling Oracle Auto Service Request

This section describes how to restart the ASR service if the ASR service is disabled.

### Using the Service Web UI

1. Open the navigation menu and click ASR Phone Home.
2. Click the Enable button. Confirm the operation when prompted.

### Using the Service CLI

1. Using SSH, log into the management node VIP as admin.  

```
# ssh -l admin 100.96.2.32 -p 30006
```
2. Use the `asrClientEnable` custom command to start the ASR service.

```
PCA-ADMIN> asrClientEnable  
Command: asrClientEnable  
Status: Success  
Time: 2021-06-23 15:26:47,632 UTC  
PCA-ADMIN>
```

## Viewing Admin Service Health Data

This section describes Private Cloud Appliance Admin service health metrics and the conditions that raise faults. This health information is not for hardware faults but is information about resource utilization (CPU, memory, and storage), hardware run state, and health checker notifications. The hardware faults listed at the bottom of [Table 6-1](#) are reported through ASR. For more information, see [\[PCA 3.x\] Private Cloud Appliance: Automatic Service Request \(ASR\) Event Coverage \(Doc ID 2833567.2\)](#).

### Admin Service Faults Summary

The threshold, run state, and health checker notification fault types that are listed in the following table are described in more detail in the following sections.

Table 6-1 Admin Service Fault Detection Configuration Summary

Fault Type	Fault Detection Frequency (seconds)	Fault Detection Delay (seconds)	Data Source	Method of Detection
<a href="#">Compute Node CPU and Memory Utilization Faults</a>	60	< 20	Admin calls ComputeNode service	Faults are raised by fault task based on the compute node object attributes stored in the database.
<a href="#">Storage Utilization Faults</a>	120	< 20	Admin calls Prometheus service	Faults are raised by fault task based on Prometheus ZFS pool usage and status data stored in the database.
<a href="#">Hardware Run State Faults</a>	150	< 20	Admin calls Hardware list REST API	Faults are raised by fault task based on hardware component node/ILOM run states stored in the database.
<a href="#">Health Checker Notification Faults</a>	Defined by the ZFS/Network health checker notification frequency	0	Various HealthChecker services send notifications	Faults are created based on the RabbitMQ notification fault results.
Platform ILOM Faults	150	0	Admin calls Hardware getMgmt and getCompute ILOM Health REST APIs	Faults are created based on the L1 API results for ILOM object data. See <a href="#">PCA X9-2 Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.1)</a> for a list of Private Cloud Appliance X9-2 events that are actionable by ASR. See <a href="#">[PCA 3.x] Private Cloud Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.2)</a> for Private Cloud Appliance X10 events.
Hardware Status Faults	On initialization, and when the syncHardwareData command runs	< 20	Admin calls Hardware list REST API	Faults are raised by fault task based on the PcaSystem object attribute. See <a href="#">PCA X9-2 Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.1)</a> for a list of Private Cloud Appliance X9-2 events that are actionable by ASR. See <a href="#">[PCA 3.x] Private Cloud Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.2)</a> for Private Cloud Appliance X10 events.

## Using the Service Web UI to View Admin Service Faults

1. Click the Active Faults link at the top of the Service Enclave Home page, or click Faults on the Navigation menu.

The Faults page is displayed.

2. At the top of the Faults page, you can toggle whether to list all faults or only active faults.
3. For more information about a fault, click the name of the fault, or click View Details on the Actions menu.

The details page shows the description, cause, and recommended action to take.

## Using the Service CLI to View Admin Service Faults

1. To view the list of Admin service faults, use the list fault command.

Both active and cleared faults are listed.

```
PCA-ADMIN> list fault
Command: list fault
Status: Success
Time: 2023-03-07 15:34:52,613 UTC
Data:
id                name                status severity
--                ---                -
33c61b8a-dcc7-4b8f-bc0f-56915ecc62f5 RackUnitIomRunStateFaultStatusFault(pcacn005) Cleared Critical
f7d22180-aeae-4159-b5c8-5e55a7906a78 RackUnitIomRunStateFaultStatusFault(pcacn004) Cleared Critical
a4fef907-8e54-4750-9fac-6829fbade90d ComputeNodeCpuFaultStatusFault(pcacn006) Cleared Minor
f8d93384-da30-43cd-9396-6e6671d240e2 RackUnitIomRunStateFaultStatusFault(pcacn010) Cleared Critical
8e61bb81-7a02-4c26-8ef4-c13b198f64da ComputeNodeCpuFaultStatusFault(pcacn007) Cleared Warning
3216b6f9-326b-4992-99a3-ab23cb18243b AK-8003-F9--PCIe 3 Active Minor
ef3fb25b-0573-4524-8d1c-fb704c814446 AK-8003-HF--vnic1 Active Major
f830cd46-21ff-4d74-ba81-c82fd6f52c67 ComputeNodeCpuFaultStatusFault(pcacn005) Cleared Minor
d2e71da0-ba63-4983-97da-24033d5c6447 ZfsPoolUsageFaultStatusFault(PCA_POOL) Cleared Major
eed5ef2-4a71-4137-be96-54c028212d2f ComputeNodeMemoryFaultStatusFault(pcacn004) Cleared Minor
cf68d2ee-e483-e573-b46e-c31bcb8e96f8 ISTOR-8000-1S--ORACLE SERVER E5-2L Cleared Major
0686c11d-b96b-e5aa-dfbc-a20154da4794 SPAMD-8002-FJ--ORACLE SERVER E5-2L Cleared Major
b488a45a-80df-46e3-b0b5-a35527eb9c0e AK-8003-F9--PCIe 10 Active Minor
ac48f88d-e181-4b03-b620-6bfbf4ad95ef RackUnitIomRunStateFaultStatusFault(pcacn007) Cleared Critical
b4c66a7c-def3-42c2-8842-d4763afc5184 RackUnitIomRunStateFaultStatusFault(pcacn006) Cleared Critical
9fc2e45a-1cff-4f95-828d-58742c8ce12f ComputeNodeMemoryFaultStatusFault(pcacn002) Active Minor
c0124122-a91c-4110-89cc-deebe54de7ba ComputeNodeMemoryFaultStatusFault(pcacn006) Cleared Critical
ca26ed46-4d1c-4ade-9e74-af27d94cf8f4 AK-8003-HF--vnic2 Active Major
58e9ab5d-d4e7-4d94-9ca6-e85a1c88b3b8 RackUnitRunStateFaultStatusFault(sn022147XLF014) Cleared Critical
474c269f-4018-45d7-97d5-da17c9c845f4 RackUnitIomRunStateFaultStatusFault(pcacn001) Cleared Critical
2b5ece1c-50fc-436a-81b3-da0c5b418fe3 RackUnitIomRunStateFaultStatusFault(pcacn003) Cleared Critical
1c164eb9-9a76-4592-8ab6-150edb8f7a75 ComputeNodeCpuFaultStatusFault(pcacn001) Cleared Warning
55ed1494-6aac-4248-91cb-9ac8295d668c AK-8003-HF--PCIe 6 Active Major
afb0c080-0b93-434b-8ead-fa673f302170 AK-8003-F9--PCIe 6 Active Minor
8b36c2db-a3b4-41c8-b416-8e733ace3aeb PcaSystemReSyncHwStatusStatusFault(null) Cleared Warning
28c5ba93-6b4e-42f3-ad61-90734b46bf30 SPENV-8000-RU--ORACLE SERVER E5-2L Cleared Critical
3d932188-0120-489f-a512-1a244ec01e49 RackUnitIomRunStateFaultStatusFault(pcacn009) Cleared Critical
21e6faa9-68e1-47ae-a298-e2cb14d2a406 ComputeNodeMemoryFaultStatusFault(pcacn007) Cleared Minor
db023304-fb7a-613b-ad9b-e277b7ce5675 SPENV-8000-A7--ORACLE SERVER E5-2L Cleared Major
63839bf5-335b-48ff-86a0-9e981e3e9902 RackUnitRunStateFaultStatusFault(sn012147XLF014) Cleared Critical
2e851c6e-aa29-4a25-846a-29b08967dd95 RackUnitValidationStateStatusFault(pcacn008) Cleared Major
76805c56-fcf6-48a2-b4fd-ffa77570e83c ComputeNodeCpuFaultStatusFault(pcacn002) Active Minor
9be74faf-df4d-ea20-cfc1-92b2a6a01b06 SPENV-8000-A7--ORACLE SERVER E5-2L Cleared Major
1624064f-d380-4ffc-9000-d293c185d7ac ComputeNodeCpuFaultStatusFault(pcacn003) Cleared Warning
7ca3f7af-f0bd-45d9-bad7-15794d49e7c6 RackUnitIomRunStateFaultStatusFault(pcacn008) Cleared Critical
3e7a3503-7a71-4ef1-a3ad-fba2162571ab ComputeNodeCpuFaultStatusFault(pcacn004) Cleared Warning
```

```
0922cd8e-297e-4356-b736-b09ac382b28b AK-8003-F9--PCIe 10 Active Minor
ab44ad2c-1105-417d-aa47-e8cb477ef0ec AK-8003-F9--PCIe 3 Active Minor
```

- To view the details of a specific fault, including description, cause, and recommended action to take, use the `show fault` command with the specific fault ID.

```
PCA-ADMIN> show fault id=ab44ad2c-1105-417d-aa47-e8cb477ef0ec
Command: show fault id=ab44ad2c-1105-417d-aa47-e8cb477ef0ec
Status: Success
Time: 2023-03-07 15:36:19,414 UTC
Data:
Id = ab44ad2c-1105-417d-aa47-e8cb477ef0ec
Type = Fault
Category = Internal
Severity = Minor
Status = Active
Last Update Time = 2023-03-06 20:04:11,668 UTC
Message Id = AK-8003-F9
Time Reported = Mon Mar 06 2023 16:50:24 GMT+0000 (UTC)
Action = Check the networking cable, switch port, and switch configuration. Contact your vendor for support
if the network port remains inexplicably down. Please refer to the associated reference document at
http://support.oracle.com/msg/AK-8003-F9 for the latest service procedures and policies regarding
this diagnosis.
Health Exporter = zfssa-analytics-exportersn022147XLF014
uuid = ab44ad2c-1105-417d-aa47-e8cb477ef0ec
Diagnosing Source = zfssa_analytics_exporter
FaultHistoryLogIds 1 = id:fdfaa42f-de8d-4622-a9df-ea229b7bad6f type:FaultHistoryLog name:
BaseManagedObjectId = id:2147XLF015/PCIe 3/465774J-2121701684 type:HardwareComponent name:
Description = Network connectivity via port mlxne4 has been lost.
Name = AK-8003-F9--PCIe 3
Work State = Normal
```

Additional examples of using the Service CLI to show Admin service faults are shown in [Compute Node CPU and Memory Utilization Faults](#).

## Compute Node CPU and Memory Utilization Faults

The Admin service raises faults for the percent of memory used and percent of CPU used for a `ComputeNode` object. More severe faults are raised as more memory and CPU are used. When the percent used drops below a certain percentage, any faults are cleared.

These are utilization faults (CPU and memory usage), not hardware faults. Problems with CPU and memory hardware are reported through ASR.

### CPU Usage

The following table shows the default percent of compute node CPU usage that raises different severities of faults.

CPU Percentage	Fault Severity	Fault State
< .75	Not applicable	Cleared
>= .75	Warning	Active
>= .80	Minor	Active
>= .90	Major	Active
>= .95	Critical	Active

## CPU Memory

The following table shows the default percent of compute node memory usage that raises different severities of faults.

Memory Percentage	Fault Severity	Fault State
< .75	Not applicable	Cleared
>= .75	Warning	Active
>= .80	Minor	Active
>= .90	Major	Active
>= .95	Critical	Active

## Using the Service CLI to View Compute Node Faults

To view the CPU and memory compute node usage default fault trigger settings using the Service CLI, use the `cnUpdateManager` command:

```
PCA-ADMIN> show cnUpdateManager
Command: show cnUpdateManager
Status: Success
Time: 2023-03-06 23:41:37,249 UTC
Data:
Id = caaaaaa1-a076-4e48-94b5-7bdcd4e0c42c
Type = CnUpdateManager
LastRunTime = 2023-03-06 23:41:33,676 UTC
Poll Interval (sec) = 60
The minimum CPU usage percentage to trigger a critical fault = 0.95
The minimum CPU usage percentage to trigger a major fault = 0.9
The minimum CPU usage percentage to trigger a minor fault = 0.8
The minimum CPU usage percentage to trigger a warning = 0.75
The minimum memory usage percentage to trigger a critical fault = 0.95
The minimum memory usage percentage to trigger a major fault = 0.9
The minimum memory usage percentage to trigger a minor fault = 0.8
The minimum memory usage percentage to trigger a warning = 0.75
```

To view the list of all faults and the details of a specific fault, see [Viewing Admin Service Health Data](#). The following example shows a specific compute node fault. Current usage is not shown except that it is at least the minor fault threshold but less than the major fault threshold. To see current usage, use the Service Web UI.

```
PCA-ADMIN> show fault id=76805c56-fcf6-48a2-b4fd-ffa77570e83c
Command: show fault id=76805c56-fcf6-48a2-b4fd-ffa77570e83c
Status: Success
Time: 2023-03-07 15:40:50,917 UTC
Data:
Id = 76805c56-fcf6-48a2-b4fd-ffa77570e83c
Type = Fault
Category = Status
Severity = Minor
Status = Active
Associated Attribute = cpuFault
Last Update Time = 2023-03-04 01:06:25,666 UTC
Cause = ComputeNode pcacn002 attribute cpuFault = MINOR.
FaultHistoryLogIds 1 = id:79b44c26-cb4e-4bec-a58c-6efc7fc63fed type:FaultHistoryLog name:
FaultHistoryLogIds 2 = id:fc90a99a-031b-457f-b585-5c905e61362e type:FaultHistoryLog name:
FaultHistoryLogIds 3 = id:48068f78-1328-447d-9506-efb6f22d154d type:FaultHistoryLog name:
FaultHistoryLogIds 4 = id:d97c5819-923c-480d-8f61-2341c8403182 type:FaultHistoryLog name:
FaultHistoryLogIds 5 = id:18cdd005-53c0-488c-a2df-28f2da3b1092 type:FaultHistoryLog name:
```

```

FaultHistoryLogIds 6 = id:bfe1ffcd-5899-4400-914c-b467d8671e0c type:FaultHistoryLog name:
FaultHistoryLogIds 7 = id:459fa55b-8654-4c07-8ae7-6d0ef011e3b1 type:FaultHistoryLog name:
FaultHistoryLogIds 8 = id:b9c8a909-f8ea-4de6-9bfe-2516e7addf73 type:FaultHistoryLog name:
FaultHistoryLogIds 9 = id:6ab5d1ca-3659-49a7-8e68-946bbbeccc9f type:FaultHistoryLog name:
FaultHistoryLogIds 10 = id:d04d06a1-1e2c-404c-ac67-680e0deb34c5 type:FaultHistoryLog name:
FaultHistoryLogIds 11 = id:22dd163e-528f-4346-b177-d62c7ceb9885 type:FaultHistoryLog name:
FaultHistoryLogIds 12 = id:cdb2dbf5-6999-43c2-bb5f-17192bfad3e2 type:FaultHistoryLog name:
FaultHistoryLogIds 13 = id:aa7b2e43-ab0b-4d78-bfe7-d4b0dd0fec4a type:FaultHistoryLog name:
BaseManagedObjectId = id:0dd96e90-de00-4fa0-82e3-16937e4601f8 type:ComputeNode name:
Description = ComputeNode pcacn002 attribute cpuFault = MINOR.
Name = ComputeNodeCpuFaultStatusFault(pcacn002)
Work State = Normal

```

## Storage Utilization Faults

The following table describes the two kinds of Oracle ZFS Storage Appliance faults raised in the Admin service.

These are utilization faults (ZFS pool usage), not hardware faults. Problems with ZFS hardware are reported through ASR.

Private Cloud Appliance uses Prometheus matrix data collected for ZFS Storage Appliance to report pool usage. Total pool size per pool (`zfssa_pool_total`) and free space per pool (`zfssa_pool_free`) are used to calculate pool usage percentage. The `zfssa_pool_status` metric reports the health of a pool.

Metric Name	Metric Value Description	Fault Condition
<code>zfssa_pool_total</code> <code>zfssa_pool_free</code>	Pool usage percentage is calculated using the following formula for each pool:  $\frac{(\text{zfssa\_pool\_total} - \text{zfssa\_pool\_free})}{\text{zfssa\_pool\_total}}$	If the pool usage percentage is above a pre-configured value, a major fault is raised. The default value is 80 percent.
<code>zfssa_pool_status</code>	The <code>zfssa_pool_status</code> metric can have the following values: <ul style="list-style-type: none"> <li>0 - exported</li> <li>1 - degraded</li> <li>2 - online</li> <li>-1 - offline</li> <li>-2 - faulted</li> <li>-3 - unavailable</li> <li>-4 - removed</li> </ul>	A major fault is raised for any pool/ <code>zfssa_node</code> combination that has any pool status value other than 0 or 2.

## Hardware Run State Faults

A critical or major fault is raised if a hardware unit on the rack such as a management node, compute node, storage node, or switch has an invalid run state.

The following table shows the severity of the fault that will be raised for the given run state. Any run state other than the listed run states results in clearing any fault.

Run State Value (case insensitive)	Fault Severity	Fault State
UNABLE TO CONNECT TO ILOM	Critical	Active
FAIL	Critical	Active

Run State Value (case insensitive)	Fault Severity	Fault State
SERVICE REQUIRED	Major	Active
<i>other</i>	Not applicable	Cleared

## Health Checker Notification Faults

Health Checker faults are raised from notifications from the ZFS Storage Appliance and Network Health Checker components. The Admin service raises a fault for every notification it receives.

Following are example attributes of the `faultedComponents` object in the Network Health Checker component fault data:

```
"class": "cisco.fan.fail",
"severity": "Major",
"description": "Fan module has failed and needs to be replaced. This can lead to overheating and temperature alarms.",
...
"class": "cisco.power.fail",
"severity": "Major",
"description": "Power Supply has failed or has been shutdown",
```

Following are example attributes of the `faultedComponents` object in the ZFS Storage Appliance Health Checker component fault data:

```
"severity": "Major",
"type": "Fault",
"description": "An internal power supply failure has been detected.",
```

Detailed information is provided about the part that has failed.

An action attribute contains a brief description of what to do to fix the problem and might include a link to the appropriate support document.

## Manually Clearing Faults

This section describes how to manually clear faults using the Service CLI. You cannot manually clear faults using the Service Web UI.

### Using the Service CLI

- Using SSH, log into the management node VIP as admin.

```
# ssh -l admin 100.96.2.32 -p 30006
```

- Use the `list fault` command to find the list of fault identifications.

```
PCA-ADMIN> list fault
Command: list fault
Status: Success
Time: 2024-01-31 21:38:05,472 UTC
Data:
id          Name          Status Severity
---          -
71671228...56a6a58947c6a6789 pcamn02-example Active Critical
524cb805...acc3458bb79t04295 RackUnit-example Active Major
PCA-ADMIN>
```

- Use the `clearFault` command with the fault identifier to clear the fault.

```
PCA-ADMIN> clearFault id=[524cb805...acc3458bb79t04295]
Command: clearFault
Status: Success
Time: 2024-01-31 21:39:30,094 UTC
PCA-ADMIN>
```

### Note

You can verify the clear fault result by using another `list fault` command.

```
PCA-ADMIN> list fault
Command: list fault
Status: Success
Time: 2024-01-31 21:40:02,685 UTC
Data:
id          Name          Status Severity
---          -
71671228...56a6a58947c6a6789 pcamn02-example Active Critical
PCA-ADMIN>
```

## Using Support Bundles

Support bundles are files of diagnostic data collected from the Private Cloud Appliance that are used to evaluate and fix problems.

Support bundles can be uploaded to Oracle Support automatically or manually. Support bundles are uploaded securely and contain the minimum required data: system identity (not IP addresses), problem symptoms, and diagnostic information such as logs and status.

Support bundles can be created and not uploaded. You might want to create a bundle for your own use. Creating a support bundle is a convenient way to collect related data.

Support bundles are created and uploaded in the following ways:

### Oracle Auto Service Request (ASR)

ASR automatically creates a service request and support bundle when certain hardware faults occur. The service request and support bundle are automatically sent to Oracle Support, and the Private Cloud Appliance administrator is notified. See [Using Oracle Auto Service Request](#).

### asrInitiateBundle

The `asrInitiateBundle` command is a PCA-ADMIN command that creates a support bundle, attaches the support bundle to an existing service request, and uploads the support bundle to Oracle Support. See [Using the asrInitiateBundle Command](#).

### support-bundles

The `support-bundles` command is a management node command that creates a support bundle of a specified type. Oracle Support might ask you to run this command to collect more data related to a service request, or you might want to collect this data for your own use. See [Using the support-bundles Command](#).

### Manual upload to Oracle Support

Several methods are available for uploading support bundles or other data to Oracle Support. See [Uploading Support Bundles to Oracle Support](#).

## Using the `asrInitiateBundle` Command

The `asrInitiateBundle` command takes three parameters, all required:

```
PCA-ADMIN> asrInitiateBundle mode=triage sr=SR_number bundleType=auto
```

A triage support bundle is collected and automatically attached to service request *SR\_number*. For more information about the triage support bundle, see [Triage Mode](#).

If the ASR service is enabled, `bundleType=auto` uploads the bundle to Oracle Support using the Phone Home service. For information about the Phone Home service, see [Registering Private Cloud Appliance for Oracle Auto Service Request](#). The bundle is saved on the management node for two days after successful upload. See [Using the support-bundles Command](#).

If you specify `mode=native` and do not specify any value for `nativeType`, then a ZFS\_BUNDLE is uploaded.

## Using the `support-bundles` Command

The `support-bundles` command collects various types of bundles, or modes, of diagnostic data such as health check status, command outputs, and logs. Depending on the options provided, these files might contain logs or status. All modes collect files into a bundle directory.

No more than one support bundle process is allowed at one time. A support bundle lock file is created at the beginning of bundle collection and removed when bundle collection is complete.

All `support-bundles` commands return immediately, and the bundle collection runs in the background. This is because bundle collections might take a long time, perhaps hours.

Bundles are stored for two days, then automatically deleted.

The following types of bundles are supported:

- [Triage Mode](#). Collects data about the current status of the Private Cloud Appliance.
- [Time Slice Mode](#). Collects data by time slots. These results can be further narrowed by specifying pod name, job, and `k8s_app` label.
- [Combo Mode](#). Collects a combination of triage and time slice data.
- [Native Mode](#). Collects data from management, compute, and ZFS nodes and from ILOM and Cisco hosts.

A good way to start to investigate an issue is to collect a `combo` bundle. Look for `NOT_HEALTHY` in the `trriage` mode results and compare that to what you see in the `time_slice` mode results.

The `support-bundles` command requires a mode option. All modes accept the service request number option. See the following table. Time slice and native modes have additional options.

Option	Description	Required
<code>-m mode</code>	The type of bundle.	yes
<code>-sr SR_number</code>	The service request number.	no
<code>--sr_number SR_number</code>		

The `support-bundles` command output is stored in the following directory on the management node, where *bundle-type* is the mode: `triage`, `time_slice`, `combo`, or `native`:

```
/nfs/shared_storage/support_bundles/SR_number_bundle-type-bundle_timestamp/
```

The *SR\_number* is used if you provided the `-sr` option. If you are creating the support bundle for a service request, specify the *SR\_number*.

This directory contains a bundle collection progress file and an archive file. The bundle collection progress file has the following name:

```
bundle-type_collection.log
```

The output archive file has the following name:

```
SR_number_bundle-type-bundle_timestamp.tar.gz
```

The archive file contains a `header.json` file with the following default components:

- `current-time` - the timestamp
- `create-support-bundle` - the command line that was used
- `sr-number` - the SR number associated with the archive file

### Log in to the Management Node

To use the `support-bundles` command, log in as `root` to the management node that is running Pacemaker resources. Collect data first from the management node that is running Pacemaker resources, then from other management nodes as needed.

If you do not know which management node is running Pacemaker resources, log in to any management node and check Pacemaker cluster status. The following command shows the Pacemaker cluster resources are running on `pcamn01`.

```
[root@pcamn01 ~]# pcs status
Cluster name: mncluster
Stack: corosync
Current DC: pcamn01
...
Full list of resources:

scsi_fencing (stonith:fence_scsi): Stopped (disabled)
Resource Group: mgmt-rg
vip-mgmt-int (ocf::heartbeat:IPaddr2): Started pcamn01
vip-mgmt-host (ocf::heartbeat:IPaddr2): Started pcamn01
vip-mgmt-ilom (ocf::heartbeat:IPaddr2): Started pcamn01
vip-mgmt-lb (ocf::heartbeat:IPaddr2): Started pcamn01
vip-mgmt-ext (ocf::heartbeat:IPaddr2): Started pcamn01
I1api (systemd:I1api): Started pcamn01
haproxy (ocf::heartbeat:haproxy): Started pcamn01
pca-node-state (systemd:pca_node_state): Started pcamn01
dhcp (ocf::heartbeat:dhcpd): Started pcamn01
hw-monitor (systemd:hw_monitor): Started pcamn01

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

## Triage Mode

In triage mode, Prometheus `platform_health_check` is queried for both HEALTHY and NOT\_HEALTHY status. If NOT\_HEALTHY is found, use `time_slice` mode to get more detail.

```
[root@pcamn01 ~]# support-bundles -m triage
```

The following files are in the output archive file.

File	Description
header.json	Time stamp and command line to generate this bundle.
compute_node_info.json	Pods running in the compute node.
hardware_info.json	Hardware component list retrieved from hms, all the ipmitool fru running at all the ready state management and compute nodes, all the zfssa heads information.
management_node_info.json	Pods running in the management node.
rack_info.json	Rack installation time and build version.
loki_search_results.log. <i>n</i>	Chunk files in json.

## Time Slice Mode

In time slice mode, data is collected by specifying start and end timestamps. Both of the following options are required:

- `-s start_date`
- `-e end_date`

Time slice mode has the following options in addition to the mode and service request number options. These options help narrow the data collection. If you do not specify either the `-j` or `--all` option, then data is collected from all health checker jobs.

- Only one of `--job_name`, `--all`, and `--k8s_app` can be specified.
- If none of `--job_name`, `--all`, or `--k8s_app` is specified, the pod filtering will occur on the default (`+.checker`).
- The `--all` option can collect a huge amount of data. You might want to limit your time slice to 48 hours.

Example:

```
[root@pcamn01 ~]# support-bundles -m time_slice -j flannel-checker -s 2021-05-29T22:40:00.000Z \
-e 2021-06-29T22:40:00.000Z -l INFO
```

See more examples below.

Option	Description	Required
<code>-s <i>timestamp</i></code>	Start date in format <code>yyyy-mmm-ddTHH:mm:ss</code>	yes
<code>--start_date <i>timestamp</i></code>	The minimum argument is <code>yyyy-mmm-dd</code>	
<code>-e <i>timestamp</i></code>	End date in format <code>yyyy-mmm-ddTHH:mm:ss</code>	yes
<code>--end_date <i>timestamp</i></code>	The minimum argument is <code>yyyy-mmm-dd</code>	

Option	Description	Required
<code>-j <i>job_name</i></code>	Loki job name. Default value: <code>.+checker</code>	no
<code>--job_name <i>job_name</i></code>	See Label List Query below.	
<code>--k8s_app <i>label</i></code>	The <code>k8s_app</code> label value to query within the <code>k8s-stdout-logs</code> job. See Label List Query below.	no
<code>--all</code>	Queries all job names except for jobs known for too much logging, such as <code>audit</code> , <code>kubernetes-audit</code> , and <code>vault-audit</code> and <code>k8s_app</code> label <code>pcacoredns</code> .	no
<code>-l <i>level</i></code>	Message level	no
<code>--levelname <i>level</i></code>		
<code>--pod_name <i>pod_name</i></code>	The pod name (such as <code>kube</code> or <code>network-checker</code> ) to filter output based on the pod. Only the starting letters are necessary.	no
<code>-t <i>timeout</i></code>	Timeout in seconds for a single Loki query. By default it is 180 seconds.	no
<code>--timeout <i>timeout</i></code>		

### Label List Query

Use the label list query to list the available job names and `k8s_app` label values.

```
[root@pcamn01 ~]# support-bundles -m label_list
2021-10-14T23:19:18.265 - support_bundles - INFO - Starting Support Bundles
2021-10-14T23:19:18.317 - support_bundles - INFO - Locating filter-logs Pod
2021-10-14T23:19:18.344 - support_bundles - INFO - Executing command - ['python3',
'/usr/lib/python3.6/site-packages/filter_logs/label_list.py']
2021-10-14T23:19:18.666 - support_bundles - INFO -
Label: job
Values: ['admin', 'api-server', 'asr-client', 'asrclient-checker', 'audit', 'cert-checker', 'ceui',
'compute', 'corosync', 'etcd', 'etcd-checker', 'filesystem', 'filter-logs', 'flannel-checker',
'his', 'hms', 'iam', 'k8s-stdout-logs', 'kubect', 'kubernetes-audit', 'kubernetes-checker',
'l0-cluster-services-checker', 'messages', 'mysql-cluster-checker', 'network-checker', 'ovm-agent',
'ovn-controller', 'ovs-vswitchd', 'ovsdb-server', 'pca-healthchecker', 'pca-nwctl', 'pca-platform-l0',
'pca-platform-l1api', 'pca-upgrader', 'pcsd', 'registry-checker', 'sauron-checker', 'secure',
'storagectl', 'uws', 'vault', 'vault-audit', 'vault-checker', 'zfssa-checker', 'zfssa-log-exporter']
```

```
Label: k8s_app
Values: ['admin', 'api', 'asr-client', 'asrclient-checker', 'brs', 'cert-checker', 'compute',
'default-http-backend', 'dr-admin', 'etcd', 'etcd-checker', 'filesystem', 'filter-logs',
'flannel-checker', 'fluentd', 'ha-cluster-exporter', 'has', 'his', 'hms', 'iam', 'ilom',
'kube-apiserver', 'kube-controller-manager', 'kube-proxy', 'kubernetes-checker',
'l0-cluster-services-checker', 'loki', 'loki-bnr', 'mysql-cluster-checker', 'mysqld-exporter',
'network-checker', 'pcacoredns', 'pcadnsmgr', 'pcanetwork', 'pcaswitchmgr', 'prometheus', 'rabbitmq',
'registry-checker', 'sauron-api', 'sauron-checker', 'sauron-grafana', 'sauron-ingress-controller',
'sauron-mandos', 'sauron-operator', 'sauron-prometheus', 'sauron-prometheus-gw',
'sauron-sauron-exporter', 'sauron.oraclelx.com', 'storagectl', 'switch-metric', 'uws', 'vault-checker',
'vmconsole', 'zfssa-analytics-exporter', 'zfssa-csi-nodeplugin', 'zfssa-csi-provisioner', 'zfssa-log-exporter']
```

Examples:

No job label, no `k8s_app` label, collect log from all health checkers.

```
[root@pcamn01 ~]# support-bundles -m time_slice -sr 3-xxxxxxxxxxx -s "2022-01-11T00:00:00" -e "2022-01-12T23:59:59"
```

One job `ceui`.

```
[root@pcamn01 ~]# support-bundles -m time_slice -sr 3-xxxxxxxxxxx -j ceui -s "2022-01-11T00:00:00" -e "2022-01-12T23:59:59"
```

One k8s\_app network-checker.

```
[root@pcamn01 ~]# support-bundles -m time_slice -sr 3-xxxxxxxxxxx --k8s_app network-checker -s "2022-01-11T00:00:00" -e "2022-01-12T23:59:59"
```

All jobs and date.

```
[root@pcamn01 ~]# support-bundles -m time_slice -sr 3-xxxxxxxxxxx -s `date -d "2 days ago" -u +%Y-%m-%dT%H:%M:%S.000Z` -e `date -d +u +%Y-%m-%dT%H:%M:%S.000Z`
```

All jobs.

```
[root@pcamn01 ~]# support-bundles -m time_slice -sr 3-xxxxxxxxxxx --all -s "2022-01-11T00:00:00" -e "2022-01-12T23:59:59"
```

The following files are in the output archive file.

File	Description
header.json	Time stamp and command line to generate this bundle.
loki_search_results.log. <i>n</i>	Chunk files in json. Time slice bundles have a limit of 500,000 logs per query, from start time.
rack_info.json	Rack installation time and build version.

### Combo Mode

The combo mode is a combination of a triage bundle and a time slice bundle. The output includes an archive file and two collection log files: triage\_collection.log and time\_slice\_collection.log.

The following files are in the output archive file.

File	Description
triage-bundle_ <i>timestamp</i> .tar.gz	The triage bundle archive file.
time_slice-bundle_ <i>timestamp</i> .tar.gz	The time slice bundle archive file. The time slice data collected is for --all jobs from one hour preceding the current time to the current time.

### Native Mode

The native\_collection.log file in the bundle directory provides collection progress information. Native bundles can take hours to collect.

The native mode has the following parameters in addition to mode and SR number.

Parameter	Description	Required
-t <i>nativetype</i> --type <i>nativetype</i>	<ul style="list-style-type: none"> <li>zfs_bundle</li> <li>sosreport</li> <li>ilom_snapshot</li> <li>cisco_bundle</li> </ul> Default value: zfs_bundle	no
-c <i>component</i> --component <i>component</i>	Component name, such as the name of a management, compute, or ZFS node, or an ILOM or Cisco host.	no

The following files are in the output archive file.

File	Description
header.json	Time stamp and command line to generate this bundle.
Native bundle files	These files are specific to the <i>nativetype</i> specified.
rack_info.json	Rack installation time and build version.

### ZFS Bundle

When *nativetype* is a ZFS support bundle, collection starts on both ZFS nodes and downloads the new ZFS support bundles into the bundle directory. When *nativetype* is not specified, *zfs\_bundle* is created by default.

```
[root@pcamn01 ~]# support-bundles -m native -t zfs_bundle
```

### SOS Report Bundle

When *nativetype* is an SOS report bundle, the report is collected from the management node or compute node specified by the `--component` parameter. If `--component` is not specified, the report is collected from all management and compute nodes.

```
[root@pcamn01 ~]# support-bundles -m native -t sosreport -c pcamn01
```

### ILOM Snapshot

When `nativeType=ilom_snapshot`, the value of the `--component` parameter is the ILOM host name of a management node or compute node. If the `--component` parameter is not specified, the report is collected from all ILOM hosts.

```
[root@pcamn01 ~]# support-bundles -m native -t ilom_snapshot -c ilom-pcacn007
```

### Cisco Bundle

When *nativetype* is `cisco-bundle`, the value of the `--component` parameter is an internal Cisco management, aggregation, or access switch management host name.

```
[root@pcamn01 ~]# support-bundles -m native -t cisco-bundle -c accsn01
```

To create a `cisco-bundle` type of collection, the following conditions must be met:

- The Cisco OBFL module must be enabled on all Private Cloud Appliance Cisco switches. The Cisco OBFL module is enabled by default on all Private Cloud Appliance Cisco switches.
- The Cisco EEM module must be enabled on all Private Cloud Appliance Cisco switches. The Cisco EEM module is enabled by default on all Private Cloud Appliance Cisco switches.
- EEM (Embedded Event Manager) policy

## Uploading Support Bundles to Oracle Support

After you create a support bundle using the `support-bundles` command as described in [Using the support-bundles Command](#), you can use the methods described in this topic to upload the support bundle to Oracle Support.

To use these methods, you must satisfy the following requirements:

- You must have a My Oracle Support user ID with Create and Update SR permissions granted by the appropriate Customer User Administrator (CUA) for each Support Identifier (SI) being used to upload files.
- For file uploads to existing service requests, the Support Identifier associated with the service request must be in your profile.
- To upload files larger than 2 GB, sending machines must have network access to connect to the My Oracle Support servers at [transport.oracle.com](https://transport.oracle.com) to use FTPS and HTTPS.

The Oracle FTPS service is a "passive" implementation. With an implicit configuration, the initial connection is from the client to the service on a control port of 990 and the connection is then switched to a high port to exchange data. Oracle defines a possible range of the data port of 32000-42000, and depending upon your network configuration you may need to enable outbound connections on both port 990 and 32000-42000. TLSv1.2 / ECDHE-RSA-AES128-GCM-SHA256 is the only encryption method enabled.

The Oracle HTTPS diagnostic upload service uses the standard HTTPS port of 443 and does not require any additional ports to be opened.

When using command line protocols, do not include your password in the command. Enter your password only when prompted.

- Oracle requires the use of TLS 1.2+ for all file transfers.
- Do not upload encrypted or password-protected files, standalone or within an archive. A Service Request update will note this as a corrupted file or reject the upload as disallowed file types were found. Files are encrypted when you use FTPS and HTTPS; additional protections are not required.
- Do not upload files with file type extensions `exe`, `bat`, `asp`, or `com`, either standalone or within an archive. A Service Request update will note that a disallowed file type was found.

### Uploading Files 2 GB or Smaller

Use the SR file upload utility on the [My Oracle Support](#) Portal.

1. Log in to [My Oracle Support](#) with your [My Oracle Support](#) user name and password.
2. Do one of the following:
  - Create a new service request and in the next step, select the Upload button.
  - Select and open an existing service request.
3. Click the Add Attachment button located at the top of the page.
4. Click the Choose File button.
5. Navigate and select the file to upload.
6. Click the Attach File button.

You can also use the methods described in the next section for larger files.

### Uploading Files Larger Than 2 GB

You cannot upload a file larger than 200 GB. See [Splitting Files](#).

The `curl` commands in this section show required options and arguments. You might want to add options such as `--verbose` and `--progress-bar` to get more information about your upload. The `--progress-meter` (more information than `--progress-bar`) should be on by default, but it is disabled when `curl` is writing other information to `stdout`. Note that some options might not be available or might behave differently on some operating systems or some versions of `curl`.

The following are the most common messages from uploading bundles to Oracle Support if you use the `--verbose` option with the `curl` command:

- **UPLOAD SUCCESSFUL.** The bundle is successfully uploaded to Oracle Support.
- **LOGIN FAILED.** The user has an authentication issue.
- **INVALID SR NUMBER.** The user does not have attach privilege to this Service Request.

### FTPS

Syntax:

Be sure to include the `/` character after the service request number.

```
$ curl -T path_and_filename -u MOS_user_ID ftps://transport.oracle.com/issue/SR_number/
```

Example:

```
$ curl -T /u02/files/bigfile.tar -u MOSUserID@example.com ftps://transport.oracle.com/issue/3-1234567890/
```

### HTTPS

Syntax:

Be sure to include the `/` character after the service request number.

```
$ curl -T path_and_filename -u MOS_user_ID https://transport.oracle.com/upload/issue/SR_number/
```

Example:

```
$ curl -T D:\data\bigfile.tar -u MOSUserID@example.com https://transport.oracle.com/upload/issue/3-1234567890/
```

### Renaming the file during send

```
$ curl -T D:\data\bigfile.tar -u MOSUserID@example.com https://transport.oracle.com/upload/issue/3-1234567890/NotSoBig.tar
```

### Using a proxy

```
$ curl -k -T D:\data\bigfile.tar -x proxy.example.com:80 -u MOSUserID@example.com https://transport.oracle.com/upload/issue/3-1234567890/
```

### Splitting Files

You can split a large file into multiple parts and upload the parts. Oracle Transport will concatenate the segments when you complete uploading all the parts.

Only HTTPS protocol can be used. Only the UNIX split utility can be used. The Microsoft Windows split utility produces an incompatible format.

To reduce upload times, compress the original file prior to splitting.

#### 1. Split the file.

The following command splits the file `file1.tar` into 2 GB parts named `file1.tar.partaa` and `file1.tar.partab`.

#### Important

Specify the `.part` extension exactly as shown below.

```
$ split -b 2048m file1.tar file1.tar.part
```

2. Upload the resulting file1.tar.partaa and file1.tar.partab files.

**! Important**

Do not rename these output part files.

```
$ curl -T file1.tar.partaa -u MOSuserID@example.com https://transport.oracle.com/upload/issue/SR_number/  
$ curl -T file1.tar.partab -u MOSuserID@example.com https://transport.oracle.com/upload/issue/SR_number/
```

3. Send the command to put the parts back together.

The spit files will not be attached to the service request. Only the final concatenated file will be attached to the service request.

```
$ curl -X PUT -H X-multipart-total-size:original_size -u MOSuserID@example.com https://transport.oracle.com/upload/issue/  
SR_number/file1.tar?multiPartComplete=true
```

In the preceding command, *original\_size* is the size of the original unsplit file as shown by a file listing.

4. Verify the size of the newly-attached file.

**i Note**

This verification command must be executed immediately after the concatenation command in Step 3. Otherwise, the file will have begun processing and will no longer be available for this command.

```
$ curl -I -u MOSuserID@example.com https://transport.oracle.com/upload/issue/SR_number/file1.tar  
X-existing-file-size: original_size
```

### Resuming an Interrupted HTTPS Upload

You can resume a file upload that terminated abnormally. Resuming can only be done by using HTTPS. Resuming does not work with FTPS. When an upload is interrupted by some event, the start with retrieving the file size of the interrupted file

1. Determine how much of the file has already been uploaded.

```
$ curl -I -u MOSuserID@example.com https://transport.oracle.com/upload/issue/SR_number/myinfo.tar  
HTTP/1.1 204 No Content  
Date: Tue, 15 Nov 2022 22:53:54 GMT  
Content-Type: text/plain  
X-existing-file-size: already_uploaded_size  
X-Powered-By: Servlet/3.0 JSP/2.2
```

2. Resume the file upload.

Note the file size returned in “X-existing-file-size” in Step 1. Use that file size after the -C switch and in the -H “X-resume-offset:” switch.

```
$ curl -Calready_uploaded_size -H "X-resume-offset: already_uploaded_size" -T myinfo.tar -u MOSuserID@example.com  
https://transport.oracle.com/upload/issue/SR_number/myinfo.tar
```

3. Verify the final file size.

```
$ curl -I -u MOSuserID@example.com https://transport.oracle.com/upload/issue/SR_number/myinfo.tar  
-H X-existing-file-size: original_size
```

In the preceding command, *original\_size* is the size of the original file as shown by a file listing.

# 7

## Backup and Restore

This chapter provides instructions for administrators who work with the integrated backup service. The purpose of this service is to store data that allows a crucial system service or component to be restored to its last known healthy state. It does not create backups of the environment created by users of the cloud resources in the Compute Enclave.

### **Caution**

#### **Backup Retention**

To optimize storage space consumption, the Backup and Restore Service applies a retention period of 14 days. When a backup operation runs, backups older than the retention period are deleted from shared storage on the ZFS Storage Appliance.

Automatic purging of backups – regardless of whether it is a standard daily backup or a manually triggered operation – is particularly critical for backups of the **MySQL database**. If a MySQL backup must be stored longer than the retention period, for example because it represents an important restore point, ensure that the data is copied to another location before the retention period expires. Contact your Oracle representative for assistance.

See [Identifying Converted Snapshots](#) for information about when the retention period does not apply and you need to manually delete snapshots.

### **Caution**

In appliance software version 3.0.2-b1081557 and newer, the monitoring data from Prometheus is not included in automated backups. To preserve your Prometheus data, create a backup and restore it manually. For more information, refer to the note with [Doc ID 3021643.1](#).

For implementation details and technical background information for this feature, see "Backup and Restore" in the [Appliance Administration Overview](#) chapter in the *Oracle Private Cloud Appliance Concepts Guide*.

## Activating Standard Daily Backup

System backups are not available by default. It is critical that you follow the instructions in this section to activate standard daily backups on your appliance.

**⚠ Caution**

Make sure that daily backups are activated after system initialization. If this procedure is omitted, you will not be able to restore a component or service from a last known good state.

To activate system backups, set up a Kubernetes CronJob by running the applicable script from the management node that owns the virtual IP of the cluster.

When the system initialization process is complete, execute the following procedure to activate system backups:

1. Log on to one of the management nodes.

```
# ssh root@pcamn01
```

2. Retrieve the name of the Kubernetes pod that runs the backup and restore service. Use the following command:

```
# kubectl get pods -A | grep brs
default brs-5bdc556546-gctx9 3/3 Running 0 17d
```

3. Execute the `default-backup` script as shown in the following example to set up the Kubernetes CronJob to make a daily backup.

```
kubectl exec brs-5bdc556546-gctx9 -c brs -- /usr/sbin/default-backup
```

This backup runs every day at 00:00 local appliance time and is retained for 14 days.

4. Verify that the CronJob has been added in the default namespace.

```
# kubectl get cronjobs -A
NAMESPACE NAME SCHEDULE SUSPEND ACTIVE LAST SCHEDULE AGE
default brs-cronjob-1629969790-backup 00 **** False 0 <none> 32s
health-check cert-checker */10 **** False 0 4m6s 17d
health-check etcd-checker */10 **** False 0 4m6s 17d
health-check flannel-checker */10 **** False 0 4m6s 17d
health-check kubernetes-checker */10 **** False 0 4m6s 17d
health-check 10-cluster-services-checker */10 **** False 0 4m6s 17d
health-check mysql-cluster-checker */10 **** False 0 4m6s 17d
health-check network-checker */10 **** False 0 4m6s 17d
health-check registry-checker */10 **** False 0 4m6s 17d
health-check sauron-checker */10 **** False 0 4m6s 17d
health-check vault-checker */10 **** False 0 4m6s 17d
sauron sauron-sauron-prometheus-gw-cj 30 19 *** False 0 18h 17d
```

When this `brs-cronjob-unique_ID-backup` CronJob runs, any backups that were previously created by a `brs` manual system backup job that are more than 14 days old are deleted. See [Executing a Backup Operation](#) for information about manual system backup.

Backups created by this CronJob are deleted regularly when they are more than 14 days old, as described in the previous step.

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

*projectname/filesystemname\_timestamp*

Backups are created on the ZFS Storage Appliance at the following location, as seen from the management node mount point:

```
/nfs/shared_storage/backups/
```

Each backup is identified by its unique path containing the job OCID and time stamp:

```
/nfs/shared_storage/backups/ocid1.backup_cronjob.unique_ID/backup_timestamp/
```

## Executing a Backup Operation

It is critical that the standard daily backups are activated on your appliance. Follow the procedure in [Activating Standard Daily Backup](#). In addition, you can initiate a system backup manually if necessary. Execute this procedure to perform a manual system backup.

### Using the Service CLI

1. Choose a strategy: create a full system backup, or back up individual components.
2. Run the backup command with the target parameter of your choice.

Target options are: layer0, zfs, vault, mysql, loki, sauron, all.

- To create a full system backup, select target=all.

```
PCA-ADMIN> backup target=all
Data:
Type = BackupJob
Job Id = ocid1.brs-job.2147XLD01D....<unique_ID>
Display Name = brs-job-1698401412-backup
Profile Id = ocid1.backup_profile.2147XLD01D....<unique_ID>
Time Created = 2023-10-27T10:10:12Z
Lifecycle State = CREATING
Retention = 14
```

- To create an individual component backup, select the appropriate target from the list: layer0 | zfs | vault | mysql | loki | sauron. For example:

```
PCA-ADMIN> backup target=layer0
Data:
Type = BackupJob
Job Id = ocid1.brs-job.2147XLD01D....<unique_ID>
Display Name = brs-job-1698401607-backup
Profile Id = ocid1.backup_profile.2147XLD01D....<unique_ID>
Time Created = 2023-10-27T10:13:27Z
Lifecycle State = CREATING
Retention = 14
```

3. Use the backup job ID to check the status of the backups.

```
PCA-ADMIN> getBackupJobs
Data:
id                                     Display Name                            Components
--                                     -
ocid1.brs-job.PCA3X62D9C1.mypca.wtczj8r3zl7wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nmrcuyo4vathc8rj brs-
job-1698401412-backup layer0,zfs,vault,mysql,loki,sauron
ocid1.brs-job.PCA3X62D9C1.mypca.089a7b8cuqz0cam2r7xexo4i4p7j7ia7sqhl9f8w89dyp9q3y10dnbaac6mu brs-
job-1698401607-backup layer0
```

```
PCA-ADMIN> getBackupJob backupJobId=ocid1.brs-
job.PCA3X62D9C1.mypca.wtczj8r3zl7wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nmrcuyo4vathc8rj
Data:
Type = BackupJob
Job Id = ocid1.brs-job.PCA3X62D9C1.mypca.wtczj8r3zl7wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nmrcuyo4vathc8rj
Display Name = brs-job-1698401412-backup
Time Created = 2023-10-27T10:10:12Z
Status = success
Components = layer0,zfs,vault,mysql,loki,sauron
```

4. Confirm that the backup operations have completed successfully.

Backups are created on the ZFS Storage Appliance at the following location, as seen from the management node mount point: `/nfs/shared_storage/backups/`. Each backup is identified by its unique path containing the job OCID and time stamp:

```
/nfs/shared_storage/backups/ocid1.brs-job.unique_ID/backup_timestamp/
```

## Identifying Converted Snapshots

Converted snapshots and any snapshots that you created manually are not governed by any retention policy and must be deleted manually.

### What Is a Converted Snapshot

If you change any settings in a backup policy other than the display name of the policy, snapshots that were created under that policy are detached and marked as converted. For example, if you update a backup policy to add a new schedule, change a schedule, or delete a schedule, all existing snapshots are detached and marked as converted on the ZFS Storage Appliance.

A converted snapshot has been converted from being managed by a backup policy to requiring manual management. Specifically, snapshots that are marked as converted will not be automatically deleted because they are no longer governed by any retention policy. When these snapshots are no longer needed, you must delete them manually.

### How to Identify Converted Snapshots

The following procedure describes how to identify converted snapshots. In addition to converted snapshots, any snapshots that you created manually must be deleted manually.

1. Log on to one of the management nodes.

```
# ssh root@pcamn01
```

2. Show storagectl content.

```
# helm history storagectl
REVISION UPDATED STATUS CHART APP VERSION DESCRIPTION
3 Mon Sep 23 14:01:17 2024 superseded storagectl-3.0.202-3.10.0.0.1.75.g94b57a2 3.0.202-3.10.0.0.1.75.g94b57a2 Upgrade
complete
4 Tue Sep 24 09:39:19 2024 superseded storagectl-3.0.202-3.10.0.0.0.79.g1f63430 3.0.202-3.10.0.0.0.79.g1f63430 Upgrade
complete
5 Wed Sep 25 03:20:06 2024 superseded storagectl-3.0.202-3.10.0.0.0.80.gbb09c40 3.0.202-3.10.0.0.0.80.gbb09c40 Upgrade
complete
6 Thu Sep 26 09:33:19 2024 superseded storagectl-3.0.202-3.10.0.0.1.66.g435af30 3.0.202-3.10.0.0.1.66.g435af30 Upgrade
complete
7 Tue Oct 1 19:37:24 2024 superseded storagectl-3.0.202-3.10.0.0.1.75.gf712a7e 3.0.202-3.10.0.0.1.75.gf712a7e Upgrade
complete
8 Tue Oct 1 19:46:22 2024 superseded storagectl-3.0.202-3.10.0.0.1.75.gf712a7e 3.0.202-3.10.0.0.1.75.gf712a7e Upgrade
complete
9 Tue Oct 1 19:51:08 2024 superseded storagectl-3.0.202-3.10.0.0.1.79.ged909dd 3.0.202-3.10.0.0.1.79.ged909dd Upgrade
complete
10 Tue Oct 1 19:53:25 2024 superseded storagectl-3.0.202-3.10.0.0.1.75.gf712a7e 3.0.202-3.10.0.0.1.75.gf712a7e Upgrade
complete
11 Tue Oct 1 20:00:18 2024 superseded storagectl-3.0.202-3.10.0.0.1.79.ged909dd 3.0.202-3.10.0.0.1.79.ged909dd Upgrade
complete
12 Tue Oct 1 21:45:03 2024 deployed storagectl-3.0.202-3.11.0.0.1.1.g1e7dc83 3.0.202-3.11.0.0.1.1.g1e7dc83 Upgrade
complete
```

3. Extract the storagectl chart from `/nfs/shared_storage/`.

In this example, use the chart from the latest revision (the last line) shown in Step 2.

```
# tar -xvf /nfs/shared_storage/charts/storagectl-3.0.202-3.11.0.0.1.1.g1e7dc83.tgz -C /target_directory
# cd /target_directory
```

4. Delete the file `job.batch/storagectl-list-converted-snapshots` if it exists.

```
# kubectl delete job.batch/storagectl-list-converted-snapshots
```

5. Rename the job definition file to remove the underscore prefix.

```
# cp storagectl/templates/_storagectl-list-converted-snapshots.yaml storagectl/templates/storagectl-list-converted-snapshots.yaml
```

6. (Optional) Set input parameters such as `volumeId` or `compartmentId` (the volume or compartment for which you want to list converted snapshots) or output file name.

If you do not set either a volume OCID or a compartment OCID, then the list will include all converted snapshots, which could be a large list.

```
# kubectl edit configmap storagectl-list-converted-snapshots-config
```

7. Deploy the Kubernetes job.

```
# helm template -s templates/storagectl-list-converted-snapshots.yaml storagectl/ | kubectl apply -f -
```

8. Check the status of the job.

```
# kubectl get all -A | grep storagectl
```

When the Kubernetes job is finished, the corresponding `pod/storagectl-list-converted-snapshots-*` is complete and the output file contains the list of converted snapshots.

9. View the list of converted snapshots in the output file.

```
# cat /nfs/shared_storage/pca-platform/kubernetes/storagectl-data/converted_snapshots_data.json
```

The snapshots in this list will not be deleted automatically.

## Restoring the System from a Backup

Restoring system data from a backup is a procedure that must be performed by Oracle support personnel. Contact your Oracle representative for assistance.

# 8

## Disaster Recovery

This chapter explains how to configure disaster recovery so that each of two Oracle Private Cloud Appliance systems in different physical locations operates as the fallback for the other system in case an outage occurs at one site.

It is important to understand what is covered under Oracle Private Cloud Appliance systems disaster recovery and what is not.

Disaster recovery supports:

- Compute instances
- The block volumes associated with these compute instances

The following limitations apply to the disaster recovery feature:

- File systems are not supported
- Object store is not supported
- OKE clusters are not supported
- Application and network load balancers are not supported
- SR-IOV instances are not supported

Implementation details and technical background information for this feature can be found in the [Oracle Private Cloud Appliance Concepts Guide](#). Refer to the section "Disaster Recovery" in the chapter [Appliance Administration Overview](#).

For additional information see:

- [Oracle Site Guard](#)
- [Oracle Application Disaster Recovery Using Site Guard](#)
- [Oracle Private Cloud Appliance: IMPLEMENTING ORACLE VM DR USING SITEGUARD](#)
- [Oracle VM DR with Oracle Site Guard: A switchover in action between two Oracle Private Cloud Appliances](#) (8 min video)

## Enabling Disaster Recovery on the Appliances

This section explains how to connect the systems that participate in the disaster recovery setup. It requires two Oracle Private Cloud Appliance systems installed at different sites, and a third system running an Oracle Enterprise Manager installation with Oracle Site Guard.

Oracle Private Cloud Appliance racks that have been factory reset to the 302-b892153, 302-b925538, or 302-b946415 versions need to have a common encryption key for the ZFS Storage Appliance storage pools at both the source and destination.

If you supply outside certificates to establish a CA trust chain for the Oracle Private Cloud Appliance, you must add two PTR records to the Data Center DNS when you set up disaster recovery. A PTR (Pointer record) in DNS maps an IP address to a hostname. This behavior is the reverse of the usual IP address lookup for a supplied hostname, which is provided by an A record in DNS.

You must create two ReverseIP lookup zones for the two ReplicationIps used in disaster recovery. The DNS requests are forwarded to the Private Cloud Appliance in the same way as requests for the Private Cloud Appliance Service Zone are forwarded. If only the zfsCapacityPoolReplicationEndpoint is defined, then only a PTR record for that IP address is needed.

To create a ReverseIP lookup you need to create a DNS zone for the ReverseIP lookup. You create one or more reverse lookup zones depending on how the Replication IPs are configured. How to create these PTR records depends on the interface for the Data Center's DNS servers.

For example, if the rack domain is myprivatecloud.example.com, and the Capacity Pool IP is 10.170.123.98 and the Performance Pool IP is 10.170.123.99, the Private Cloud Appliance requires two zones with the following mappings:

```
98.123.170.10.in-addr.arpa rtype PTR rdata sn01-dr1.myprivatecloud.example.com
99.123.170.10.in-addr.arpa rtype PTR rdata sn02-dr1.myprivatecloud.example.com
```

## Collecting System Parameters for Disaster Recovery

To set up disaster recovery for your environment, you need to collect certain information in advance. To be able to fill out the parameters required to run the setup commands, you need the following details:

- IP addresses in the data center network  
Each of the two ZFS Storage Appliances needs at least one IP address in the data center network. This IP address is assigned to the storage controller interface that is physically connected to the data center network. If your environment also contains optional high-performance storage, then two pairs of data center IP addresses are required.
- Fully Qualified Domain Names (FQDNs) in the data center network  
If you have upgraded your racks to 302-b892153, you need to use the FQDNs of the hosts and not their IP addresses. This FQDN is assigned to the storage controller interface that is physically connected to the data center network. If your environment also contains optional high-performance storage, then two pairs of data center FQDNs are required.
- Data center subnet and gateway  
The ZFS Storage Appliances need to be able to exchange data over the network. Their network interfaces connect them to a local subnet. For each interface included in the disaster recovery configuration, the subnet address and gateway address are required.

To complete the Oracle Site Guard configuration, you need the following details:

- The endpoints of both Private Cloud Appliance systems, where API calls are received. These are URIs, which are formatted as follows: `https://<myRegion>.<myDomain>`

For example:

```
https://myprivatecloud.example.com
```

- An administrative user name and password for authentication with the Private Cloud Appliance services and authorization of the disaster recovery API calls. These credentials are securely stored within Oracle Enterprise Manager.

## Connecting the Components in the Disaster Recovery Setup

The ZFS Storage Appliances installed in the two Oracle Private Cloud Appliance racks must be connected to each other, in order to replicate the data that must be protected by the disaster

recovery setup. This is a direct connection through the data center network; it does not use the uplinks from the spine switches to the data center.

To create the redundant replication connection, four cable connections are required at each of the two sites. The ZFS Storage Appliance has two controllers; you must connect both 25Gbit SFP28 interfaces of each controller's first dual-port Ethernet expansion card to the next-level data center switches. At the other site, the same four ports must also be cabled this way.

The replication connection must be used exclusively for data under the control of disaster recovery configurations. Any other data replicated over this connection might be automatically destroyed.

In the next phase, the network configuration is created on top of the interfaces you cabled into the data center network. On each storage controller the two interfaces are aggregated into a redundant 25Gbit connection. The aggregation interface is assigned an IP address: one controller owns the replication IP address for the standard performance storage pool; the other controller owns the replication IP for the high-performance storage pool, if one is present.

#### **Note**

Link aggregation needs to be configured on the data center switches as well. The MTU of the ZFS Storage Appliance data links is 9000 bytes; set the data center switch MTU to 9216 bytes.

The administrators at the two sites are not required to configure the replication network manually. The configuration of the ZFS Storage Appliance network interfaces is automated through the `drSetupService` command in the Service CLI. When executing the command, the administrator provides the IP addresses and other configuration settings as command parameters. Use of the `drSetupService` command is described in the next section.

Your Oracle Enterprise Manager does not require additional installations specific to Private Cloud Appliance in order to perform disaster recovery tasks. It only needs to be able to reach the two appliances over the network. Oracle Site Guard is available by default in the software library of Oracle Enterprise Manager.

To allow Oracle Site Guard to manage failover operations between the two Private Cloud Appliance systems, you must set up both appliances as *sites*. You identify the two sites by their endpoint URIs, which are used to configure the disaster recovery scripts in the failover operation plans. You also provide a user name and password to allow Oracle Site Guard to authenticate with the two appliances.

For additional information and instructions, please refer to the product documentation of Oracle Site Guard and Oracle Enterprise Manager.

## Setting Up Peering Between the ZFS Storage Appliances

After the physical connection between the ZFS Storage Appliances has been established, you set them up as peers using the `drSetupService` command in the Service CLI. You run this command from both systems so that each system operates as the replica of the other system.

The required replication parameters for standard storage are mandatory with the setup command. If Private Cloud Appliance systems also include high-performance storage, then add the replication parameters for the high-performance storage pool to the setup command.

However, only set up replication for high-performance storage if the high-performance storage pool is effectively available on the ZFS Storage Appliances. If not, run the setup command again to add the high-performance storage pool at a later time, after it has been configured on the ZFS Storage Appliances.

When you set up the replication interfaces for the disaster recovery service, the system assumes that the gateway is the first host address in the subnet of the local IP address you specify. This applies to the replication interface for standard storage and high-performance storage. For example, if you specify a local IP address as 10.50.7.31/23 and the gateway address is **not** 10.50.6.1 then you must add the gateway IP address to the `drSetupService` command using the `gatewayIp` and `gatewayIpPerf` parameters.

Optionally, you can also set a maximum number of DR configurations and a retention period for disaster recovery job details.

## Setting Up Peering Between the ZFS Storage Appliances Before 302-b892153

If Oracle Private Cloud Appliance racks are running a version of software before release 302-b892153, follow these Service Enclave API steps to set up peering between the racks and the ZFS Storage Appliance.

### Note

Both Private Cloud Appliance racks in the disaster recovery configuration must be running the same version of the system software.

Syntax (entered on a single line):

```
drSetupService
localIp=<primary_system_standard_replication_ip> (in CIDR notation)
remoteIp=<replica_system_standard_replication_ip>
localIpPerf=<primary_system_performance_replication_ip> (in CIDR notation)
remoteIpPerf=<replica_system_performance_replication_ip>
[Optional Parameters:]
gatewayIp=<local_subnet_gateway_ip> (default: first host IP in localIp subnet)
gatewayIpPerf=<local_subnet_gateway_ip> (default: first host IP in localIpPerf subnet)
maxConfig=<number_DR_configs> (default and maximum is 20)
jobRetentionHours=<hours> (default and minimum is 24)
```

Examples:

- With only standard storage configured:
 

```
system 1
PCA-ADMIN> drSetupService \
localIp=10.50.7.31/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.33

system 2
PCA-ADMIN> drSetupService \
localIp=10.50.7.33/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.31
```
- With both standard and high-performance storage configured:
 

```
system 1
```

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.31/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.33 \
localIpPerf=10.50.7.32/23 gatewayIpPerf=10.50.7.10 remoteIpPerf=10.50.7.34
```

system 2

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.33/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.31 \
localIpPerf=10.50.7.34/23 gatewayIpPerf=10.50.7.10 remoteIpPerf=10.50.7.32
```

### Important

When setting up disaster recovery, after you run `drSetupService` on the first system you must wait for the job to complete before running the command on the second system. You can monitor the job on the first system by running `drGetJob jobid=<unique-id>`.

The script configures both ZFS Storage Appliances.

After successful configuration of the replication interfaces, you must enable replication over the interfaces you just configured.

### Enabling Replication for Disaster Recovery

To enable replication between the two storage appliances, using the interfaces you configured earlier, re-run the same `drSetupService` command from the Service CLI, but this time followed by `enableReplication=True`. You must also provide the `remotePassword` to authenticate with the other storage appliance and complete the peering setup.

Examples:

- With only standard storage configured:

system 1

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.31/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.33 \
enableReplication=True remotePassword=*****
```

system 2

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.33/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.31 \
enableReplication=True remotePassword=*****
```

- With both standard and high-performance storage configured:

system 1

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.31/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.33 \
localIpPerf=10.50.7.32/23 gatewayIpPerf=10.50.7.10 remoteIpPerf=10.50.7.34 \
enableReplication=True remotePassword=*****
```

system 2

```
PCA-ADMIN> drSetupService \
localIp=10.50.7.33/23 gatewayIp=10.50.7.10 remoteIp=10.50.7.31 \
localIpPerf=10.50.7.34/23 gatewayIpPerf=10.50.7.10 remoteIpPerf=10.50.7.32 \
enableReplication=True remotePassword=*****
```

**! Important**

When enabling replication, after you run `drSetupService` on the first system you must wait for the job to complete before running the command on the second system. You can monitor the job on the first system by running `drGetJob` `jobid=<unique-id>`.

At this stage, the ZFS Storage Appliances in the disaster recovery setup have been successfully peered. The storage appliances are ready to perform scheduled data replication every 5 minutes. The data to be replicated is based on the DR configurations you create. See [Managing Disaster Recovery Configurations](#).

**Modifying the ZFS Storage Appliance Peering Setup**

After you set up the disaster recovery service and enabled replication between the systems, you can change the parameters of the peering configuration. You change the service using the `drUpdateService` command in the Service CLI.

Syntax (entered on a single line):

```
drUpdateService
localIp=<primary_system_standard_replication_ip> (in CIDR notation)
remoteIp=<replica_system_standard_replication_ip>
localIpPerf=<primary_system_performance_replication_ip> (in CIDR notation)
remoteIpPerf=<replica_system_performance_replication_ip>
gatewayIp=<local_subnet_gateway_ip> (default: first host IP in localIp subnet)
gatewayIpPerf=<local_subnet_gateway_ip> (default: first host IP in localIpPerf subnet)
maxConfig=<number_DR_configs> (default and maximum is 20)
jobRetentionHours=<hours> (default and minimum is 24)
```

**Example 1 – Simple parameter change**

This example shows how you change the job retention time from 24 to 48 hours and reduce the maximum number of DR configurations from 20 to 12.

```
PCA-ADMIN> drUpdateService jobRetentionHours=48 maxConfig=12
Command: drUpdateService jobRetentionHours=48 maxConfig=12
Status: Success
Time: 2022-08-11 09:20:48,570 UTC
Data:
Message = Successfully started job to update DR admin service
Job Id = ec64cef4-ba68-493d-89c8-22df51553cd8
```

Use the `drShowService` command to check the current configuration. Run the command to display the configuration parameters before you change them. Run it again afterward to confirm that the changes have been applied successfully.

```
PCA-ADMIN> drShowService
Command: drShowService
Status: Success
Time: 2022-08-11 09:23:54,951 UTC
Data:
Local Ip = 10.50.7.31/23
Remote Ip = 10.50.7.33
Replication = ENABLED
Replication High = DISABLED
Message = Successfully retrieved site configuration
maxConfig = 12
gateway IP = 10.50.7.10
Job Retention Hours = 48
```

### Example 2 – Replication IP change

There might be network changes in the data center that require you to use different subnets and IP addresses for the replication interfaces configured in the disaster recovery service. This configuration change must be applied in many commands on the two peer systems, and in a specific order. If the systems contain both standard and high-performance storage – as in the example that follows –, change the replication interface settings for both storage types in the same order.

1. Update the local IP and gateway parameters on system 1. Leave the remote IPs unchanged.

```
PCA-ADMIN> drUpdateService \  
localIp=10.100.33.83/28 gatewayIp=10.100.33.81 \  
localIpPerf=10.100.33.84/28 gatewayIpPerf=10.100.33.81
```

2. Update the local IP, gateway, and remote IP parameters on system 2.

```
PCA-ADMIN> drUpdateService \  
localIp=10.100.33.88/28 gatewayIp=10.100.33.81 remoteIp=10.100.33.83 \  
localIpPerf=10.100.33.89/28 gatewayIpPerf=10.100.33.81 remoteIpPerf=10.100.33.84
```

3. Update the remote IP parameters on system 1.

```
PCA-ADMIN> drUpdateService \  
remoteIp=10.100.33.88 remoteIpPerf=10.100.33.89
```

### Example 3 – Trusting a New ZFS Storage Appliance Certificate

The following example shows the command that must be run if the ZFS Storage Appliance certificate on the peer rack is updated. This command retrieves the new certificate from the remote host and adds it to the trust list,

```
PCA-ADMIN> drUpdateService \  
remoteIp=s10.100.33.88 remoteIpPerf=10.100.33.89
```

### Unconfiguring the ZFS Storage Appliance Peering Setup

If a reset has been performed on one or both of the systems in the disaster recovery solution, and you need to unconfigure the disaster recovery service to remove the entire peering setup between the ZFS Storage Appliances, use the `drDeleteService` command in the Service CLI.

#### **Caution**

This command requires no other parameters. Be careful when entering it at the PCA-ADMIN> prompt, to avoid executing it unintentionally.

You can't unconfigure the disaster recovery service while DR configurations still exist. Proceed as follows:

1. Remove all DR configurations from the two systems that have been configured as replicas for each other.
2. Sign in to the Service CLI on one of the systems and enter the `drDeleteService` command.
3. Sign in to the Service CLI on the second system and enter the `drDeleteService` command there as well.

When the disaster recovery service isn't configured, the `drShowService` command returns an error.

```
PCA-ADMIN> drShowService
Command: drShowService
Status: Failure
Time: 2022-08-11 12:31:22,840 UTC
Error Msg: PCA_GENERAL_000001: An exception occurred during processing: Operation failed.
[...]
Error processing dr-admin.service.show response: dr-admin.service.show failed. Service not set up.
```

## Setting Up Peering Between the ZFS Storage Appliances

If Oracle Private Cloud Appliance racks are running software release 302-b892153 or later, follow these Service Enclave API steps to set up peering between the racks and the ZFS Storage Appliance.

### Note

Both Private Cloud Appliance racks in the disaster recovery configuration must be running the either both earlier than build 302-b892153 or both build 302-b892153 or later.

Before beginning, the `show netNetworkConfig` output must have valid entries for the following:

- DNS IP addresses
- Management Node Hostnames
- Management Node IP Addresses
- Free Public IP Addresses
- A Valid IP address for the ZFS CapacityPool Replication Endpoint

You must add PTR entries for DNS:

- `sn01-dr1.<rack_name><domain_name>`
- `sn02-dr1.<rack_name><domain_name>` (if you use a Performance Pool)

For DNS mapping configured with the *zone delegation* option, these DNS mappings are managed by Private Cloud Appliance DNS.

To populate the rack core DNS, edit the network configuration:

- `system 1`  

```
PCA-ADMIN> edit networkConfig \  
zfsCapacityPoolReplicationEndpoint=10.0.7.31
```
- `system 2`  

```
PCA-ADMIN> edit networkConfig \  
zfsCapacityPoolReplicationEndpoint=10.0.7.32
```

For DNS mapping configured with the *manual* option, these DNS mappings are managed by the data center DNS.

For more information on creating Private Cloud Appliance DNS PTR entries, and DNS management in general, see "Working with Zone Records" in the [Networking](#) chapter of the *Oracle Private Cloud Appliance User Guide*.

Syntax (entered on a single line):

```
drSetupService
localIp=<primary_system_standard_replication_ip> (in CIDR notation)
remoteHost=<replica_system_standard_replication_fqdn_for_remoteHost>
localIpPerf=<primary_system_performance_replication_ip> (in CIDR notation)
remoteHostPerf=<replica_system_performance_replication_fqdn_for_remoteHostPerf>
[Optional Parameters:]
gatewayIp=<local_subnet_gateway_ip> (default: first host IP in localIp subnet)
gatewayIpPerf=<local_subnet_gateway_ip> (default: first host IP in localIpPerf subnet)
maxConfig=<number_DR_configs> (default and maximum is 20)
jobRetentionHours=<hours> (default and minimum is 24)
```

#### Examples:

- With only standard storage configured:

##### system 1

```
PCA-ADMIN> drSetupService \
localIp=10.0.7.31/23 gatewayIp=10.0.7.10 remoteHost=sn01-dr1.rack1.example.com
```

##### system 2

```
PCA-ADMIN> drSetupService \
localIp=10.0.7.33/23 gatewayIp=10.0.7.10 remoteHost=sn01-dr1.rack2.example.com
```

- With both standard and high-performance storage configured:

##### system 1

```
PCA-ADMIN> drSetupService \
localIp=10.0.7.31/23 gatewayIp=10.0.7.10 remoteHost=sn01-dr1.rack1.example.com \
localIp=10.0.7.32/23 gatewayIp=10.0.7.10 remoteHostPerf=sn02-dr1.rack1.example.com
```

##### system 2

```
PCA-ADMIN> drSetupService \
localIp=10.0.7.33/23 gatewayIp=10.50.7.10 remoteHost=sn01-dr1.rack2.example.com \
localIpPerf=10.0.7.34/23 gatewayIpPerf=10.0.7.10 remoteHostPerf=sn02-dr1.rack2.example.com
```

### Important

When setting up disaster recovery, after you run `drSetupService` on the first system you must wait for the job to complete before running the command on the second system. You can monitor the job on the first system by running `drgetjob jobid=<unique-id>`.

#### For example:

```
PCA-ADMIN> drgetjob jobid=<unique-id>
Command: drgetjob jobid=<unique-id>
Status: Success
Time: 2023-08-01 15:26:46,973 UTC
Data:
Type = setup service
Job Id = <unique-id>
Status = Success
Start Time = 2023-08-01 15:26:28.935479
Message = job successfully retrieved
```

**Note**

Ensure that the "Success" status message appears in the Data fields and not only the Command field.

The script configures both ZFS Storage Appliances.

After successful configuration of the replication interfaces, you must enable replication over the interfaces you configured.

**Enabling Replication for Disaster Recovery**

To enable replication between the two storage appliances, using the interfaces you configured earlier, run the same `drSetupService` command from the Service CLI, but this time followed by `enableReplication=True`. You must also provide the `remotePassword` to authenticate with the other storage appliance and complete the peering setup.

Examples:

- With only standard storage configured:

system 1

```
PCA-ADMIN> drSetupService \  
localIp=10.0.7.31/23 gatewayIp=10.0.7.10 \  
enableReplication=True remotePassword=***** remoteHost=sn01-dr1.rack2.example.com
```

system 2

```
PCA-ADMIN> drSetupService \  
localIp=10.0.7.33/23 gatewayIp=10.0.7.10 \  
enableReplication=True remotePassword=***** remoteHost=sn01-dr1.rack1.example.com
```

- With both standard and high-performance storage configured:

system 1

```
PCA-ADMIN> drSetupService \  
localIp=10.0.7.31/23 gatewayIp=10.0.7.10 remoteHost=sn01-dr1.rack2.example.com \  
localIpPerf=10.0.7.32/23 gatewayIpPerf=10.0.7.10 remoteHostPerf=sn02-dr1.rack2.example.com \  
enableReplication=True remotePassword=*****
```

system 2

```
PCA-ADMIN> drSetupService \  
localIp=10.0.7.33/23 gatewayIp=10.0.7.10 remoteHost=sn01-dr1.rack1.example.com \  
localIpPerf=10.0.7.34/23 gatewayIpPerf=10.0.7.10 remoteHostPerf=sn02-dr1.rack1.example.com \  
enableReplication=True remotePassword=*****
```

**Important**

When enabling replication, after you run `drSetupService` on the first system you must wait for the job to complete before running the command on the second system. You can monitor the job on the first system by running `drGetJob jobId=<unique-id>`.

At this stage, the ZFS Storage Appliances in the disaster recovery setup have been successfully peered. The storage appliances are ready to perform scheduled data replication every 5 minutes. The data to be replicated is based on the DR configurations you create. See [Managing Disaster Recovery Configurations](#).

## Modifying the ZFS Storage Appliance Peering Setup

After you set up the disaster recovery service and enabled replication between the systems, you can change the parameters of the peering configuration individually. You change the service using the `drUpdateService` command in the Service CLI.

Syntax (entered on a single line):

```
drUpdateService
localIp=<primary_system_standard_replication_ip> (in CIDR notation)
remoteHost=<replica_system_standard_replication_fqdn>
localIpPerf=<primary_system_performance_replication_ip> (in CIDR notation)
remoteHostPerf=<replica_system_performance_replication_fqdn>
gatewayIp=<local_subnet_gateway_ip> (default: first host IP in localIp subnet)
gatewayIpPerf=<local_subnet_gateway_ip> (default: first host IP in localIpPerf subnet)
maxConfig=<number_DR_configs> (default and maximum is 20)
jobRetentionHours=<hours> (default and minimum is 24)
```

### Example 1 – Simple parameter change

This example shows how you change the job retention time from 24 to 48 hours and reduce the maximum number of DR configurations from 20 to 12.

```
PCA-ADMIN> drUpdateService jobRetentionHours=48 maxConfig=12
Command: drUpdateService jobRetentionHours=48 maxConfig=12
Status: Success
Time: 2022-08-11 09:20:48,570 UTC
Data:
  Message = Successfully started job to update DR admin service
  Job Id = ec64cef4-ba68-493d-89c8-22df51553cd8
```

Use the `drShowService` command to check the current configuration. Run the command to display the configuration parameters before you modify them. Run it again afterward to confirm that your changes have been applied successfully.

```
PCA-ADMIN> drShowService
Command: drShowService
Status: Success
Time: 2022-08-11 09:23:54,951 UTC
Data:
  Local Ip = 10.0.7.31/23
  Remote Host = sn01-dr1.exmaple.com
  Replication = ENABLED
  Replication High = DISABLED
  Message = Successfully retrieved site configuration
  maxConfig = 12
  gateway IP = 10.0.7.10
  Job Retention Hours = 48
```

### Example 2 – Replication IP change

There might be network changes in the data center that require you to use different subnets and IP addresses for the replication interfaces configured in the disaster recovery service. This configuration change must be applied in several commands on the two peer systems, and in a specific order. If the systems contain both standard and high-performance storage – as in the example following – change the replication interface settings for both storage types in the same order.

1. Update the replication endpoint parameters on system 1.

```
PCA-ADMIN> edit networkConfig zfsCapacityPoolReplicationEndpoint=10.100.3.88 \
zfsPerfPoolReplicationEndpoint=10.100.3.89
```

2. Update the local IP and gateway parameters on system 1. Leave the remote IPs unchanged.

```
PCA-ADMIN> drUpdateService \
localIp=10.100.3.83/28 gatewayIp=10.100.3.81 \
localIpPerf=10.100.3.84/28 gatewayIpPerf=10.100.3.81
```

3. Update the replication endpoint parameters on system 2.

```
PCA-ADMIN> edit networkConfig zfsCapacityPoolReplicationEndpoint=10.100.3.88 \
zfsPerfPoolReplicationEndpoint=10.100.3.89
```

4. Update the local IP, gateway, and remote host parameters on system 2.

```
PCA-ADMIN> drUpdateService \
localIp=10.100.3.88/28 gatewayIp=10.100.3.81 remoteHost=sn01-dr1.rack1.example.com \
localIpPerf=10.100.3.89/28 gatewayIpPerf=10.100.3.81 remoteHostPerf=sn02-dr1.rack1.example.com
```

### Example 3 – Configuration Without Performance Pool

The following example applies these four commands to a configuration using only the basic pool and not the performance pool.

1. Update the replication endpoint parameters on system 1.

```
PCA-ADMIN> edit networkConfig zfsCapacityPoolReplicationEndpoint=10.16.9.43
Command: edit networkConfig zfsCapacityPoolReplicationEndpoint=10.16.9.43
Status: Success
Time: 2023-08-16 12:08:30,585 UTC
JobId: 175b1600-eabe-4a0f-aa45-xxxxxx65599c1
```

2. Update the local IP parameters on system 1. Leave the remote IPs unchanged. Check that job has finished successfully.

```
PCA-ADMIN> drUpdateService localIp=10.16.9.43/12
Command: drUpdateService localIp=10.16.9.43/12
Status: Success
Time: 2023-08-16 12:09:45,137 UTC
Data:
Message = Successfully started job to update DR admin service
Job Id = 2844b731-f53c-4d92-850d-xxxxx22b49e3
```

```
PCA-ADMIN> drgetJob jobId=2844b731-f53c-4d92-850d-xxxxx22b49e3
Command: drgetJob jobId=2844b731-f53c-4d92-850d-xxxxx22b49e3
Status: Success
Time: 2023-08-16 12:15:19,560 UTC
Data:
Type = update_service
Job Id = 2844b731-f53c-4d92-850d-xxxxx22b49e3
Status = finished
Start Time = 2023-08-16 12:09:45.017743
End Time = 2023-08-16 12:15:19.443415
Result = success
Message = job successfully retrieved
Response = Successfully updated DR service
```

**3. Update the replication endpoint parameters on system 2.**

```
PCA-ADMIN> edit networkConfig zfsCapacityPoolReplicationEndpoint=10.16.11.43
Command: edit networkConfig zfsCapacityPoolReplicationEndpoint=10.16.11.43
Status: Success
Time: 2023-08-16 12:22:36,218 UTC
JobId: b7bff723-0237-4a11-9d08-xxxxxd166e1d
```

**4. Update the local IP parameters on system 2. Leave the remote IPs unchanged. Check that job has finished successfully.**

```
PCA-ADMIN> drUpdateService localIp=10.16.11.43/12
Command: drUpdateService localIp=10.16.11.43/12
Status: Success
Time: 2023-08-16 12:24:54,882 UTC
Data:
  Message = Successfully started job to update DR admin service
  Job Id = 1d6826ac-04db-49f9-aa27-35996f69410a
```

```
PCA-ADMIN> drgetjob jobId=1d6826ac-04db-49f9-aa27-xxxxx69410a
Command: drgetjob jobId=1d6826ac-04db-49f9-aa27-xxxxx69410a
Status: Success
Time: 2023-08-16 12:31:55,828 UTC
Data:
  Type = update_service
  Job Id = 1d6826ac-04db-49f9-aa27-xxxxx69410a
  Status = finished
  Start Time = 2023-08-16 12:24:54.655686
  End Time = 2023-08-16 12:30:16.461914
  Result = success
  Message = job successfully retrieved
  Response = Successfully updated DR service
```

**Example 4 – Trusting a New ZFS Storage Appliance Certificate**

The following example shows the command that must be run if the ZFS Storage Appliance certificate on the peer rack is updated. This command retrieves the new certificate from the remote host and adds it to the trust list,

```
PCA-ADMIN> drUpdateService \
remoteHost=sn01-dr1.rack1.example.com remoteHostPerf=sn02-dr1.rack1.example.com
```

**Unconfiguring the ZFS Storage Appliance Peering Setup**

If a reset has been performed on one or both of the systems in a disaster recovery solution, and you need to unconfigure the disaster recovery service to remove the entire peering setup between the ZFS Storage Appliances, use the `drDeleteService` command in the Service CLI.

**⚠ Caution**

This command requires no other parameters. Be careful when entering it at the PCA-ADMIN> prompt, to avoid executing it unintentionally.

You cannot unconfigure the disaster recovery service while DR configurations still exist. Proceed as follows:

1. Remove all DR configurations from the two systems that have been configured as replicas for each other.

2. Log in to the Service CLI on one of the systems and enter the `drDeleteService` command.
3. Log in to the Service CLI on the second system and enter the `drDeleteService` command there as well.

When the disaster recovery service isn't configured, the `drShowService` command returns an error.

```
PCA-ADMIN> drShowService
Command: drShowService
Status: Failure
Time: 2022-08-11 12:31:22,840 UTC
Error Msg: PCA_GENERAL_000001: An exception occurred during processing: Operation failed.
[...]
Error processing dr-admin.service.show response: dr-admin.service.show failed. Service not set up.
```

## Managing Disaster Recovery Configurations

This section explains how to configure disaster recovery settings on two Oracle Private Cloud Appliance systems such that each system is the fallback for the other system.

### Rules and Conditions

When populating DR configurations, respect the following rules regarding compute and storage resources.

- A compute instance must be stopped before it can be added to a DR configuration. There is one exception: when all volumes attached to the instance are also attached to one or more instances already included in the same DR configuration.
- A compute instance must be stopped before it can be removed from a DR configuration. There is one exception: when all volumes attached to the instance are also attached to one or more instances still included in the same DR configuration.
- All compute instances in a DR configuration must be stopped before the DR configuration can be deleted.
- A volume attached to a compute instance might be created from another source volume or volume backup. Such an instance (instance *T*) can be added to a DR configuration on condition that the source volume is not attached to any instance in any DR configuration. Note that *source volume* also refers to the volume used for the volume backup, and its direct or indirect source.

Alternatively, the instance with the source volume attached can be added to a DR configuration on condition that *instance T* is not added to any DR configuration. Due to the volume source/target relationship, only one of the instances involved can be part of a DR configuration, not both.

- A DR configuration must be refreshed when configuration for any of its instances has changed, including changes for the instance and the attached storage or network resources. This is to ensure the instances on the standby rack after switchover or failover are started with needed attributes preserved.
- Refreshing a DR configuration results in a failure in case a volume **and** the source from which it was created, are both attached to one or more compute instances in any DR configuration.
- Instances in a DR configuration preserve the primary and secondary private IPs upon switchover or failover if the subnet for the IP address is assigned a freeform tag with the key `preserve_private_ips` and a value set to anything. Once the subnet is tagged, any instance

with a VNIC attached to this subnet has the subnet's primary and secondary IP addresses preserved after switchover or failover. The precheck on the standby rack checks to see if the IP address is already in use and fails if this is true.

## Creating a DR Configuration

A DR configuration is the parent object to which you add compute instances that you want to protect against system outages.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - a unique name for the DR configuration
  - a unique name for the associated ZFS storage project
2. Create an empty DR configuration with the `drCreateConfig` command.

Syntax (entered on a single line):

```
drCreateConfig
configName=<DR_configuration_name>
project=<ZFS_storage_project_name>
```

Example:

```
PCA-ADMIN> drCreateConfig configName=drConfig1 project=drProject1
Command: drCreateConfig configName=drConfig1 project=drProject1
Status: Success
Time: 2021-08-17 07:19:33,163 UTC
Data:
  Message = Successfully started job to create config drConfig1
  Job Id = 252041b1-ff44-4c8e-a3de-11c1e47d9217
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobid=252041b1-ff44-4c8e-a3de-11c1e47d9217
Command: drGetJob jobid=252041b1-ff44-4c8e-a3de-11c1e47d9217
Status: Success
Time: 2021-08-17 07:21:07,021 UTC
Data:
  Type = create_config
  Job Id = 252041b1-ff44-4c8e-a3de-11c1e47d9217
  Status = finished
  Start Time = 2021-08-17 07:19:33.507048
  End Time = 2021-08-17 07:20:16.783743
  Result = success
  Message = job successfully retrieved
  Response = Successfully created DR config drConfig1: 439ad078-7e6a-4908-ffa-ac89210d76ac
```

4. When the DR configuration is created, the storage project for data replication is set up on the ZFS Storage Appliances.

Note the DR configuration ID. You need it for all subsequent commands to modify the configuration.

5. To display a list of existing DR configurations, use the `drGetConfigs` command.

```
PCA-ADMIN> drGetConfigs
Command: drGetConfigs
Status: Success
Time: 2021-08-17 07:44:54,443 UTC
```

```
Data:
id configName
-----
439ad078-7e6a-4908-affa-ac89210d76ac drConfig1
e8291afa-a413-4932-880a-abb8ac22c85d drConfig2
7ad05d9f-731c-41b8-b477-35da4b999071 drConfig3
```

- To display the status and details of a DR configuration, use the `drGetConfig` command.

Syntax:

```
drGetConfig drConfigId=<DR_configuration_id>
```

Example:

```
PCA-ADMIN> drGetConfig drConfigId=439ad078-7e6a-4908-affa-ac89210d76ac
Command: drGetConfig drConfigId=439ad078-7e6a-4908-affa-ac89210d76ac
Status: Success
Time: 2021-08-17 07:47:53,401 UTC
Data:
Type = DrConfig
Config State = ENABLED
Config Name = drConfig1
Config Id = 439ad078-7e6a-4908-affa-ac89210d76ac
Project Id = drProject1
```

## Adding Site Mappings to a DR Configuration

Site mappings are added to determine how and where on the replica system the instances should be brought back up in case the primary system experiences an outage and a failover is triggered. Each site mapping contains a source object for the primary system and a corresponding target object for the replica system. Make sure that these resources exist on both the primary and replica system before you add the site mappings to the DR configuration.

These are the site mapping types you can add to a DR configuration:

- **Compartment:** specifies that, if a failover occurs, instances from the source compartment must be brought up in the target compartment on the replica system
- **Subnet:** specifies that, if a failover occurs, instances connected to the source subnet must be connected to the target subnet on the replica system
- **Network security group:** specifies that, if a failover occurs, instances that belong to the source network security group must be included in the target security group on the replica system

### Using the Service CLI

- Gather the information that you need to run the command:

- DR configuration ID (`drGetConfigs`)
- Mapping source and target object OCIDs

Use the Compute Enclave UI or CLI on the primary and replica system respectively.  
CLI commands:

- `oci iam compartment list`
- `oci network subnet list --compartment-id "ocid1.compartment.....uniqueID"`
- `oci network nsg list --compartment-id "ocid1.compartment.....uniqueID"`

- Add a site mapping to the DR configuration with the `drAddSiteMapping` command.

Syntax (entered on a single line):

```
drAddSiteMapping
drConfigId=<DR_configuration_id>
objType=[compartment | subnet | networksecuritygroup]
sourceId=<source_object_OCID>
targetId=<target_object_OCID>
```

Examples:

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=compartment \
sourceId="ocid1.compartment.....<region1>...uniqueID" \
targetId="ocid1.compartment.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=compartment
sourceId="ocid1.compartment.....<region1>...uniqueID" targetId="ocid1.compartment.....<region2>...uniqueID"
Status: Success
Time: 2021-08-17 09:07:24,957 UTC
Data:
9244634e-431f-43a1-89ab-5d25905d43f9
```

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=subnet \
sourceId="ocid1.subnet.....<region1>...uniqueID" \
targetId="ocid1.subnet.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=subnet
sourceId="ocid1.subnet.....<region1>...uniqueID" targetId="ocid1.subnet.....<region2>...uniqueID"
Status: Success
Time: 2021-08-17 09:07:24,957 UTC
Data:
d1bf2cf2-d8c7-4271-b8b6-cdf757648175
```

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=networksecuritygroup \
sourceId="ocid1.nsg.....<region1>...uniqueID" \
targetId="ocid1.nsg.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=networksecuritygroup
sourceId="ocid1.nsg.....<region1>...uniqueID" targetId="ocid1.nsg.....<region2>...uniqueID"
Status: Success
Time: 2021-08-17 09:07:24,957 UTC
Data:
422f8892-ba0a-4a89-bc37-61b5c0fbbbaa
```

- Repeat the command with the OCIDs of all the source and target objects that you want to include in the site mappings of the DR configuration.

#### Note

Mappings for compartments and subnets are always required in order to perform a failover or switchover. Missing mappings will be detected by the Oracle Site Guard scripts during a precheck on the replica system.

- To display the list of site mappings included in the DR configuration, use the `drGetSiteMappings` command. The DR configuration ID is a required parameter.

Syntax:

```
drGetSiteMappings drConfigId=<DR_configuration_id>
```

**Example:**

```
PCA-ADMIN> drGetSiteMappings drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Command: drGetSiteMappings drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Status: Success
Time: 2021-08-17 09:19:22,580 UTC
Data:
  id              name
  --            ---
  d1bf2cf2-d8c7-4271-b8b6-cdf757648175 null
  9244634e-431f-43a1-89ab-5d25905d43f9 null
  422f8892-ba0a-4a89-bc37-61b5c0fbbbaa null
```

5. To display the status and details of a site mapping included in the DR configuration, use the `drGetSiteMapping` command.

Syntax (entered on a single line):

```
drGetSiteMapping
drConfigId=<DR_configuration_id>
mappingId=<site_mapping_id>
```

**Example:**

```
PCA-ADMIN> drGetSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=d1bf2cf2-d8c7-4271-
b8b6-cdf757648175
Command: drGetSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=d1bf2cf2-d8c7-4271-b8b6-
cdf757648175
Status: Success
Time: 2021-08-17 09:25:53,148 UTC
Data:
  Type = DrSiteMapping
  Object Type = subnet
  Source Id = ocid1.nsg.....<region1>...uniqueID
  Target Id = ocid1.nsg.....<region2>...uniqueID
  Work State = Normal
```

## Removing Site Mappings from a DR Configuration

You can remove a site mapping from the DR configuration if it is no longer required.

**Using the Service CLI**

1. Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)
  - Site mapping ID (`drGetSiteMappings`)
2. Remove the selected site mapping from the DR configuration with the `drRemoveSiteMapping` command.

Syntax (entered on a single line):

```
drRemoveSiteMapping
drConfigId=<DR_configuration_id>
mappingId=<site_mapping_id>
```

**Example:**

```
PCA-ADMIN> drRemoveSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=422f8892-
ba0a-4a89-bc37-61b5c0fbbbaa
Command: drRemoveSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=422f8892-ba0a-4a89-
bc37-61b5c0fbbbaa
```

Status: Success  
Time: 2021-08-17 09:41:43,319 UTC

3. Repeat the command with the IDs of all the site mappings that you want to remove from the DR configuration.

## Adding Instances to a DR Configuration

Once a DR configuration has been created and the relevant site mappings have been set up, you add the required compute instances. Their data and disks are stored in the ZFS storage project associated with the DR configuration, and replicated over the network connection between the ZFS Storage Appliances of both Private Cloud Appliance systems.

If your system contains optional high-performance disk shelves, you must set up peering accordingly between the ZFS Storage Appliances. As a result, two ZFS projects are created for each DR configuration: one in the standard pool and one in the high-performance pool. When you add instances to the DR configuration that have disks running on standard as well as high-performance storage, those storage resources are automatically added to the ZFS project in the appropriate pool.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)
  - Instance OCIDs from the Compute Enclave UI or CLI (`oci compute instance list --compartment-id <compartment_OCID>`)
2. Add a compute instance to the DR configuration with the `drAddComputeInstance` command.

Syntax (entered on a single line):

```
drAddComputeInstance
drConfigId=<DR_configuration_id>
instanceId=<instance_OCID>
```

Example:

```
PCA-ADMIN> drAddComputeInstance \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
instanceId=ocid1.instance.....<region1>...uniqueID
```

```
Command: drAddComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
instanceId=ocid1.instance.....<region1>...uniqueID
```

Status: Success

Time: 2021-08-17 07:24:35,186 UTC

Data:

```
Message = Successfully started job to add instance ocid1.instance.....<region1>...uniqueID to DR config
63b36a80-7047-42bd-8b97-8235269e240d
Job Id = 8dcbd22d-69b0-4319-b09f-1a4df847e9df
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=8dcbd22d-69b0-4319-b09f-1a4df847e9df
```

```
Command: drGetJob jobId=8dcbd22d-69b0-4319-b09f-1a4df847e9df
```

Status: Success

Time: 2021-08-17 07:36:27,719 UTC

Data:

```
Type = add_computeinstance
Job Id = 8dcbd22d-69b0-4319-b09f-1a4df847e9df
Status = finished
Start Time = 2021-08-17 07:24:36.776193
```

```
End Time = 2021-08-17 07:26:59.406929
Result = success
Message = job successfully retrieved
Response = Successfully added instance [ocid1.instance.....<region1>...uniqueID] to DR config
[63b36a80-7047-42bd-8b97-8235269e240d]
```

- Repeat the `drAddComputeInstance` command with the OCIDs of all the compute instances that you want to add to the DR configuration.
- To display the list of instances included in the DR configuration, use the `drGetComputeInstances` command. The DR configuration ID is a required parameter.

Syntax:

```
drGetComputeInstances drConfigId=<DR_configuration_id>
```

Example:

```
PCA-ADMIN> drGetComputeInstances drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Command: drGetComputeInstances drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Status: Success
Time: 2021-08-17 08:33:39,586 UTC
Data:
id                               name
--                               ----
ocid1.instance.....<region1>...instance1_uniqueID    null
ocid1.instance.....<region1>...instance2_uniqueID    null
ocid1.instance.....<region1>...instance3_uniqueID    null
```

- To display the status and details of an instance included in the DR configuration, use the `drGetComputeInstance` command.

Syntax (entered on a single line):

```
drGetComputeInstance
drConfigId=<DR_configuration_id>
instanceId=<instance_OCID>
```

Example:

```
PCA-ADMIN> drGetComputeInstance \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
instanceId=ocid1.instance.....<region1>...instance1_uniqueID
Command: drGetComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
instanceId=ocid1.instance.....<region1>...instance1_uniqueID
Status: Success
Time: 2021-08-17 08:34:42,413 UTC
Data:
Type = ComputeInstance
Compartment Id = ocid1.compartment.....uniqueID
Boot Volume Id = ocid1.bootvolume.....uniqueID
Compute Instance Shape = VM.PCAStandard1.8
Work State = Normal
```

## Removing Instances from a DR Configuration

Instances can only be part of a single DR configuration. You can remove a compute instance from the DR configuration to which it was added.

### Using the Service CLI

- Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)

- Instance OCID (drGetComputeInstances)
2. Remove the selected compute instance from the DR configuration with the drRemoveComputeInstance command.

Syntax (entered on a single line):

```
drRemoveComputeInstance
drConfigId=<DR_configuration_id>
instanceId=<instance_OCID>
```

Example:

```
PCA-ADMIN> drRemoveComputeInstance \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
instanceId=ocid1.instance.....<region1>...instance3_uniqueID
Command: drRemoveComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
instanceId=ocid1.instance.....<region1>...instance3_uniqueID
Status: Success
Time: 2021-08-17 08:45:59,718 UTC
Data:
  Message = Successfully started job to remove instance ocid1.instance.....<region1>...instance3_uniqueID from DR config
63b36a80-7047-42bd-8b97-8235269e240d
  Job Id = 303b42ff-077c-4504-ac73-25930652f73a
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=303b42ff-077c-4504-ac73-25930652f73a
Command: drGetJob jobId=303b42ff-077c-4504-ac73-25930652f73a
Status: Success
Time: 2021-08-17 08:56:27,719 UTC
Data:
  Type = remove_computeinstance
  Job Id = 303b42ff-077c-4504-ac73-25930652f73a
  Status = finished
  Start Time = 2021-08-17 08:46:00.641212
  End Time = 2021-08-17 07:47:19.142262
  Result = success
  Message = job successfully retrieved
  Response = Successfully removed instance [ocid1.instance.....<region1>...instance3_uniqueID] from DR config
[63b36a80-7047-42bd-8b97-8235269e240d]
```

4. Repeat the drRemoveComputeInstance command with the OCIDs of all the compute instances that you want to remove from the DR configuration.

## Refreshing a DR Configuration

To ensure that the replication information stored in a DR configuration is updated with all the latest changes in your environment, you can refresh the DR configuration.

### Using the Service CLI

1. Look up the ID of the DR configuration you want to refresh (drGetConfigs).
2. Refresh the data stored in the selected DR configuration with the drRefreshConfig command.

Syntax:

```
drRefreshConfig drConfigId=<DR_configuration_id>
```

Example:

```
PCA-ADMIN> drRefreshConfig drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Command: drRefreshConfig drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
```

```
Status: Success
Time: 2021-08-17 10:43:33,241 UTC
Data:
Message = Successfully started job to refresh DR config 63b36a80-7047-42bd-8b97-8235269e240d
Job Id = 205eb34e-f416-41d3-95a5-506a1d891fdb
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=205eb34e-f416-41d3-95a5-506a1d891fdb
Command: drGetJob jobId=205eb34e-f416-41d3-95a5-506a1d891fdb
Status: Success
Time: 2021-08-17 10:51:27,719 UTC
Data:
Type = refresh_config
Job Id = 205eb34e-f416-41d3-95a5-506a1d891fdb
Status = finished
Start Time = 2021-08-17 10:43:34.264828
End Time = 2021-08-17 10:45:12.718561
Result = success
Message = job successfully retrieved
Response = Successfully refreshed DR config [63b36a80-7047-42bd-8b97-8235269e240d]
```

## Deleting a DR Configuration

When you no longer need a DR configuration, you can remove it with a single command. It also removes all site mappings and cleans up the associated storage projects on the ZFS Storage Appliances of the primary and replica system. However, you must stop all compute instances that are part of the DR configuration before you can delete it.

### Using the Service CLI

1. Stop all the compute instances that are part of the DR configuration you want to delete.
2. Look up the ID of the DR configuration you want to delete (`drGetConfigs`).
3. Delete the selected DR configuration with the `drDeleteConfig` command.

Syntax:

```
drDeleteConfig drConfigId=<DR_configuration_id>
```

Example:

```
PCA-ADMIN> drDeleteConfig drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Command: drDeleteConfig drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
Status: Success
Time: 2021-08-17 14:45:19,634 UTC
Data:
Message = Successfully started job to delete DR config 63b36a80-7047-42bd-8b97-8235269e240d
Job Id = d2c1198d-f521-4b8d-a9f1-c36c7965d567
```

4. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=d2c1198d-f521-4b8d-a9f1-c36c7965d567
Command: drGetJob jobId=d2c1198d-f521-4b8d-a9f1-c36c7965d567
Status: Success
Time: 2021-08-17 16:18:33,462 UTC
Data:
Type = delete_config
Job Id = d2c1198d-f521-4b8d-a9f1-c36c7965d567
Status = finished
Start Time = 2021-08-17 14:45:20.105569
End Time = 2021-08-17 14:53:32.405569
```

Result = success  
Message = job successfully retrieved  
Response = Successfully deleted DR config [63b36a80-7047-42bd-8b97-8235269e240d]

# 9

## Native Disaster Recovery

This chapter explains how to configure disaster recovery so that two peered Oracle Private Cloud Appliance systems in different physical locations operate as the fallback for the other system in case an outage occurs at one site.

It is important to understand what is covered under disaster recovery and what is not.

Disaster recovery supports:

- Compute instances
- The block volumes associated with these compute instances

The following limitations apply to the disaster recovery feature:

- File systems are not supported
- Object storage is not supported
- OKE clusters are not supported
- Application and network load balancers are not supported
- SR-IOV instances are not supported

Implementation details and technical background information for this feature can be found in the [Oracle Private Cloud Appliance Concepts Guide](#). Refer to the section "Disaster Recovery" in the chapter [Appliance Administration Overview](#).

## Migrating to the Native Disaster Recovery Service

If you have Private Cloud Appliance environments set up for disaster recovery using Oracle Site Guard, you have the option to migrate to the native Disaster Recovery Service as soon as the systems have been upgraded or patched to appliance software version 3.0.2-b1261765 or later.

First, ensure that both environments are running the minimum required appliance software version. A first-generation disaster recovery setup cannot operate as the standby of a system running the newer native DR service, or the other way around. As soon as a [peer connection](#) has been established, first-generation DR configurations can no longer be used, but you can migrate them to the native DR service.

### Collecting Oracle Site Guard DR Plans

To extract the DR plan data from Oracle Site Guard, you need the `get_em_plan.sh` script, and a valid user name and password for the Oracle Enterprise Manager database schema. During appliance upgrade or patching, the script is saved to the directory `/var/lib/pca-foundation/scripts`.

1. Log in to the system that runs the Oracle Enterprise Manager database, where the data from the Oracle Site Guard plugin is stored.
2. Copy the `get_em_plan.sh` script to this system. Run the script from the command line and specify an output file. The data is stored in CSV format.

**Note**

To retrieve the plan data for a single DR configuration, add its name to the end of the command, separated by a space.

```
$/get_em_plan.sh em_plan_data.csv
```

- When prompted by the script, enter the database user name and password.  
The DR plan data is retrieved from the database and saved to the \*.csv output file you specified.
- Copy the \*.csv file before logging out.

**Switching to the Native DR Service**

Before converting the first-generation DR plans, you must set up peering between both systems, and activate the native DR service.

- Ensure that an active peer connection exists between the primary and standby appliance. See [Establishing a Peer Connection](#).
- Activate the native DR service by updating the existing service configuration. The serial number of the peer system is a required parameter.

```
PCA-ADMIN> drNativeUpdateService peerSerialNumber=<peer_serial>
```

**Converting to Native DR Plans**

You can import the DR plan data from Oracle Site Guard into the native DR service using the Service CLI.

- Ensure that the \*.csv output file with the DR plan data from Oracle Site Guard is in a location that the DR service pod can access.
- Create DR plans from the \*.csv file with the `create DrPlan` command.

For each DR configuration that matches a record in the \*.csv file, the DR plans are converted and saved in the native DR format.

Syntax (entered on a single line):

```
create DrPlan
drPlanDataFile=<csv_file_name>
configsToMigrate=["all"|"config_id_1,config_id_n"]
```

- To convert DR plan data for a subset of DR configurations from the CSV file, specify the IDs of the required configurations:

```
PCA-ADMIN> create DrPlan drPlanDataFile=/home/admin/em_plan_data.csv \
configsToMigrate="config1_id, config2_id, config3_id"
JobId: 146c6422-cd82-4ecf-9e5f-a99acfc9baab
Data: DrPlan id: null::null. Successfully started job for plan creation using /home/admin/em_plan_data.csv
```

- To convert all DR plans for all matching DR configurations from the CSV file:

```
PCA-ADMIN> create DrPlan drPlanDataFile=/home/admin/em_plan_data.csv configsToMigrate="all"
JobId: 862b045a-c4f6-4f8c-8130-9d9b487bd363
Data: DrPlan id: null::null. Successfully started job for plan creation using /home/admin/em_plan_data.csv
```

- Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=862b045a-c4f6-4f8c-8130-9d9b487bd363
Data:
Id = 862b045a-c4f6-4f8c-8130-9d9b487bd363
Type = Job
Associated Work Request Id = 8671211c-68b3-426e-aa7c-d93d9ffb32e4
Done = true
Name = CREATE_TYPE
Progress Message = Plan migration succeeded.
Run State = Succeeded
Transcript = Created job CREATE_TYPE
Username = admin
WorkItemIds 1 = id:08cb4f9a-575b-4536-b832-b85948c90212 type:WorkItem name:
```

Use the work request ID to display more details.

```
PCA-ADMIN> drGetJob jobId=8671211c-68b3-426e-aa7c-d93d9ffb32e4
Data:
Type = drplan_migration
Job Id = 8671211c-68b3-426e-aa7c-d93d9ffb32e4
Status = finished
Result = succeeded
Message = job successfully retrieved
Response =
  config_id: b51a51ac-043e-4b6d-ab29-20f86d905e81: plan_name: fo1 status: success
  config_id: b51a51ac-043e-4b6d-ab29-20f86d905e81: plan_name: sw1 status: success
  config_id: b51a51ac-043e-4b6d-ab29-20f86d905e81: plan_name: pfo1 status: success
```

4. To confirm that the DR plans have been added, use the `drListPlan` command.

```
PCA-ADMIN> drListPlan drConfigId=b51a51ac-043e-4b6d-ab29-20f86d905e81
Data:
id                Plan Name  Operation
--                -
b51a51ac-043e-4b6d-ab29-20f86d905e81::fo1  fo1      failover
b51a51ac-043e-4b6d-ab29-20f86d905e81::sw1  sw1      switchover
b51a51ac-043e-4b6d-ab29-20f86d905e81::pfo1 pfo1     postfailover
```

5. Modify the DR plans as required. For details, see [Working With Disaster Recovery Plans](#).

## Establishing a Peer Connection

The disaster recovery service requires a mutual, symmetrical peer connection between two Oracle Private Cloud Appliance racks. Dedicated cabling must be installed at each site, and the disaster recovery service on each rack must be configured to accept the other rack as its standby system.

### Adding Cable Connections

A peer connection between Oracle Private Cloud Appliance racks requires additional physical connections. Dedicated cabling must be installed between the spine switches and the data center network.

Racks with factory-installed software version 3.0.2-b1261765 or later already have all the necessary internal network interfaces and connections. Only external cabling for the peer connection is required. The ZFS Storage Appliances use the same physical connections for their data replication.

In existing installations where first-generation disaster recovery is configured, an active replication network between the ZFS Storage Appliances exists. When upgrading or patching

to the latest appliance software, the existing replication network remains active. The new physical connections from the spine switches are used for peering traffic only.

### Data Center Cabling for Rack Peering

Direct peering between racks requires dedicated cabling for each participating system. The additional connections between the spine switches and the data center network are the physical basis on which the network tunnels of the peer connection are configured.

For the purpose of peering, port 6 on each spine switch must be connected to the data center network. To provide the required connection speed of 10 or 25 Gbps, a 4-way breakout cable is attached to spine port 6. From the breakout cable, 1 transceiver is connected to the data center network. Cabling must be identical for both spine switches.

### Appliance Internal Cabling for Rack Peering

Appliance rack configurations shipped from the factory before the release of the native DR service do not have the required internal cabling to enable replication through the peer connection tunnels. They lack these crucial components:

- PCIe 25GbE network interface card in some models of the ZFS Storage Appliance Controllers
- Ethernet cabling between ZFS Storage Appliance Controllers and spine switches (port 27)

These components can be added to existing installations, so their hardware configuration is equivalent to racks with factory-installed software version 3.0.2-b1261765 or later. Contact Oracle for assistance and additional information.

### Backward Compatibility

The native DR service supports both cabling layouts:

- **Peering Topology:** combined peer connection and storage replication network through the spine switches
- **Compatibility Topology:** peer connection and physically separated direct replication link between the ZFS Storage Appliances

The compatibility topology provides different options for existing installations after upgrading or patching to software version 3.0.2-b1261765 or later. If you have a first-generation DR setup, you can choose to continue with this configuration, on condition that you do **not** establish a peer connection at the appliance level. However, Oracle recommends migrating the existing configuration to the native DR service, in accordance with your infrastructure design and maintenance schedule. Data center cabling for the peer connection must be added, but you can continue to use the existing storage replication connection. For more information, see [Migrating to the Native Disaster Recovery Service](#).

## Creating a Local Endpoint

Traffic between peered Private Cloud Appliance systems flows through tunnels between endpoints. A rack must have a local endpoint configured before it can participate in a peer connection.

Assuming the network configuration remains the same, a local endpoint is set up once. It remains configured if a peer connection is deleted, so it can be reused for any new connection. These parameters are required to create the local endpoint:

- an IP address for each spine switch
- the IP addresses of the data center gateways to which the spines are connected

- IP address of the capacity ZFS pool and, if present, the high-performance ZFS pool
- the ASN ID of your network environment, if applicable

### Network Configuration Guidelines

A local endpoint requires a /29 address block, which has 6 usable IP addresses. Within that /29 range, the spine switch pair is assigned 3 IPs (one is shared). Each ZFS pool is assigned 1 IP *outside* the spine switch subnet. To allow for additional future peer connections, it is recommended to reserve a data center IP range of at least /25 in size, which corresponds with 16 address blocks of /29 size.

When setting up the local endpoint, you must provide the netmask with the peering network IPs, but not with IPs that have already been configured, such as the data center addresses. The gateway IPs must be provided by the network administrator, and assigned to the data center switches to which the spine switches are connected. Note that spine 1 corresponds with gateway 1, and spine 2 with gateway 2.

### Caution

If your systems are set up with a first-generation disaster recovery configuration, and you are migrating to the native disaster recovery service, perform these steps:

1. Gather the existing disaster recovery configuration details on both appliances, using the `drShowService` command.

```
PCA-ADMIN> drShowService
Data:
Local Ip = 10.100.3.83/28
Local Ip Perf = 10.100.3.84/28
Remote Host = sn01-dr1.exmaple.com
Remote Host Perf = sn02-dr1.exmaple.com
Replication = ENABLED
Replication High = ENABLED
Message = Successfully retrieved site configuration
maxConfig = 12
gateway IP = 10.100.3.81
gateway IP Perf = 10.100.3.81
Job Retention Hours = 48
```

2. Remove the ZFS pool replication IPs from the existing network configuration.

```
PCA-ADMIN> edit networkConfig ZFSCapacityPoolReplicationEndpoint=""
PCA-ADMIN> edit networkConfig ZFSPerfPoolReplicationEndpoint=""
```

3. Use the storage IP addresses (and subnet mask) already assigned to the ZFS Storage Appliance Controllers for replication between the storage pools.

When you have obtained all required IP addresses, create the appliance local endpoint using either the Service CLI or Service Web UI.

### Using the Service CLI

Enter the following command on a single line, replacing the sample IPs with the ones you obtained:

```
PCA-ADMIN> create LocalEndpoint \  
spine1Ip=<10.212.128.3/29> datacenterGateway1Ip=<10.212.128.1> \  
spine2Ip=<10.212.128.4/29> datacenterGateway2Ip=<10.212.128.2> \  
zfsCapacityPoolEndpointIp=<10.212.128.129/29> zfsPerformancePoolEndpointIp=<10.212.128.130/29> \  
localAsn=<136025>
```

Check the local endpoint configuration with the command: `getLocalEndpoint`.

### Using the Service Web UI

Under Disaster Recovery Service, open the Local Endpoint page. In the top-right corner, click Create.

In the pop-up window, enter the IP addresses in the respective fields. Click Create Local Endpoint to apply the settings.

In the Local Endpoint page, the Information tab indicates the endpoint is configured. Click the Configuration tab to display the details.

### Deleting the Local Endpoint

The local endpoint cannot be deleted if it is part of an existing peer configuration. Delete the peer connection first.

### Using the Service CLI

Enter the command: `deleteLocalEndpoint`.

### Using the Service Web UI

Under Disaster Recovery Service, open the Local Endpoint page. In the top-right corner, click Delete.

## Creating the Peer Connection

When two Private Cloud Appliance systems have been cabled correctly, and their local endpoints have been configured, the peer connection can be created.

The peer connection is a symmetrical configuration, meaning the setup must be performed on each connected appliance. The administrators exchange the relevant configuration details of their system, so they can each include the peer details required for creating the connection. A trust relationship between the appliances is established through a CA chain stored in the Secret Service (Vault).

When the first appliance completes its side of the connection setup, it goes into a waiting state. By design, the appliance with the IP address ending with the lowest value initiates the connection. As soon as the entry for the peer appliance is detected, the CA certificates are verified and the mutual trust relationship is confirmed. After successful peering, a pair of secure tunnels is established between the spine switches. These allow the administration services on the appliances to exchange information with each other.

These parameters are required to create a peer connection:

- the IP addresses (4 in total) of the local and the remote endpoint for each tunnel
- the IP addresses of the remote spine switches in the peer appliance
- properties of the peer appliance: domain name, system name, serial number, ASN ID if applicable
- properties of the peer Admin Service: host name, admin user name, admin password, CA chain

The network configuration must allow peer-to-peer connectivity between the replication endpoints, or use routable IPs when both systems are in separate address spaces. Ensure that the new network setup does not overlap with existing connections between the appliance and the data center.

A peer connection requires a /30 subnet, with 2 IPs assigned to each local endpoint. When setting up the connection, you include the netmask for the local endpoint IPs, but not for the remote endpoint IPs and remote spine switch IPs.

When you have obtained all required parameters, create the peer connection using either the Service CLI or Service Web UI.

### Using the Service CLI

Enter the following command on a single line, replacing the sample parameters with the ones you obtained:

```
PCA-ADMIN> create PeerConnection name=<peerconnection1> description=<"my peer connection"> \
peerSerialNumber=<1654BF2465> peerSystemName=<mypca1> peerDomainName=<mydomain.com> \
localEndpoint1Ip=<172.16.21.1/30> remoteEndpoint1Ip=<172.16.21.2> \
localEndpoint2Ip=<172.16.21.5/30> remoteEndpoint2Ip=<172.16.21.6> \
remoteSpine1Ip=<10.212.128.10> remoteSpine2Ip=<10.212.128.11> \
peerAdminHostname=<mypca1.mydomain.com> peerAdminUserName=<admin> peerAdminPassword=<password> \
peerAdminCaChain=<ca_string>
remoteAsn=<136025>
```

Check the peer connection configuration using the following commands:

```
PCA-ADMIN> list PeerConnection
```

Data:

id	Name	Peer Admin Hostname	Peer Rack Serial Number	Lifecycle State
ocid1.drpeerconnection...unique_ID	peerconnection1	mypca1.mydomain.com	1654BF2465	ACTIVE

```
PCA-ADMIN> show peerConnection id=ocid1.drpeerconnection...unique_ID
```

Data:

```
Id = ocid1.drpeerconnection...unique_ID
Type = PeerConnection
Lifecycle Sub State = ACTIVE
Lifecycle State = ACTIVE
Peer Rack Serial Number = 1654BF2465
Local Endpoint 1 Ip = 172.16.21.1/30
Local Endpoint 2 Ip = 172.16.21.5/30
Remote Endpoint 1 Ip = 172.16.21.2
Remote Endpoint 2 Ip = 172.16.21.6
Remote Spine 1 Ip = 10.212.128.10
Remote Spine 2 Ip = 10.212.128.11
Peer Admin CaChain = -----BEGIN CERTIFICATE-----\nMIIIFbjCCA1agAwIBAgIQfMPkn17+ZTNI/jZjYzbpn[...]
Peer Admin Hostname = mypca1.mydomain.com
Peer Rack Domain Name = mydomain.com
Peer Rack System Name = mypca1
Peer Rack Admin User Name = admin
Peer Rack Admin User Password = *****
Remote Asn = 136025
ProgressRecordIds 1 = id:d39144d6-feef-4988-ba71-fac4b046fff8 type:ProgressRecord name:
ProgressRecordIds 2 = id:940b397f-993c-4ab9-9708-909dabb65c47 type:ProgressRecord name:
ProgressRecordIds 3 = id:64b31360-3d0d-4dc2-a925-35164143eb25 type:ProgressRecord name:
ProgressRecordIds 4 = id:7e8d9e2e-74b1-4d31-9098-7a09d719ec6a type:ProgressRecord name:
ProgressRecordIds 5 = id:2309bcde-1689-410b-a93e-528444ada2a5 type:ProgressRecord name:
ProgressRecordIds 6 = id:8a4d5747-d8fe-48e2-96e2-f4c797964cbe type:ProgressRecord name:
Name = peerconnection1
Work State = Normal
```

### Using the Service Web UI

Under Disaster Recovery Service, open the Peer Connections page. In the top-right corner, click Create Peer Connection.

In the pop-up window, enter all parameters in the respective fields. Click Create Peer Connection to apply the settings.

In the Peer Connections page, the table displays a new entry for the connection you created. Click the name in the table to display the detail page of the peer connection, and review its configuration parameters.

### Updating the Peer Connection

There is no CLI command or UI function to modify the peer connection once it's configured. Changing the peer connection requires that you delete it and create a new connection with the updated parameters.

### Deleting the Peer Connection

If a peer connection is no longer used, you can delete it. Ensure that the peer configuration is removed from each connected appliance.

### Using the Service CLI

Look up the ID of the peer connection you want to delete, then enter the delete command as shown.

```
PCA-ADMIN> list PeerConnection
Data:
id          Name          Peer Admin Hostname  Peer Rack Serial Number  Lifecycle State
--          ---          -
ocid1.drpeerconnection....unique_ID peerconnection1 mypca1.mydomain.com 1654BF2465        ACTIVE

PCA-ADMIN> delete peerConnection id=ocid1.drpeerconnection....unique_ID
```

### Using the Service Web UI

Under Disaster Recovery Service, open the Peer Connections page. In the table, click the name of the connection you want to delete. The peer connection detail page is displayed. In the top-right corner, click Delete.

## Setting Up the Disaster Recovery Service

The Private Cloud Appliance native Disaster Recovery Service is configured on top of an active peer connection between two systems located at different sites. When the appliances have been peered, the administrator can set up the DR service.

The DR service is set up with a single command and the peer serial number as a parameter. The system verifies that replication targets are configured correctly, collects the required settings from the peer connection, and stores the information in a metadata file used for enabling the DR service. The setup must be performed on both peered appliances for the DR service to become active.

### Using the Service CLI

Obtain the serial number of the peered appliance (list PeerConnection). To check if a DR service setup already exists, enter the drNativeListService command.

To enable the DR service over an existing peer connection, enter the setup command as follows:

```
PCA-ADMIN> drNativeSetupService peerSerialNumber=<peer_serial>
```

To verify the DR service details, enter the list and show commands as follows:

```
PCA-ADMIN> drNativeListService
Data:
id          Peer Serial Number  Remote Host
```

```
-- -----  
1654BF2465 1654BF2465 mydrhost1.mypca2.mydomain.com  
  
PCA-ADMIN> drNativeShowService peerSerialNumber=1654BF2465  
Data:  
Type = DrSiteConfig  
Local Ip = 192.168.8.17/28  
Remote Host = mydrhost1.mypca2.mydomain.com  
Replication = ENABLED  
Replication High = DISABLED  
Message = Successfully retrieved site configuration  
maxConfig = 20  
Job Retention Hours = 24  
Peer Serial Number = 1654BF2465
```

### Using the Service Web UI

Under Disaster Recovery Service, open the Native Services page. In the top-right corner, click Setup Native Service.

In the pop-up window, enter the serial number of the peered appliance in the designated field. (Max Config and Job Retention Hours are optional parameters.) Click Setup Native Service to apply the settings.

### Updating the Disaster Recovery Service

You can change certain properties of the DR service: the retention period for jobs, and the maximum number of DR configurations (up to 20). Ensure that the service configuration is identical on both peered appliances.

If you are migrating from first-generation DR to the native DR service, the same command is used to update the original metadata file so it can be used by the native DR service.

### Using the Service CLI

Obtain the serial number of the peered appliance (`drNativeListService`). Enter the update command and new parameters on a single line as follows:

```
PCA-ADMIN> drNativeUpdateService peerSerialNumber=<peer_serial> \  
maxConfig=<12> jobRetentionHours=<48>
```

### Using the Service Web UI

Under Disaster Recovery Service, open the Native Services page. Click the DR service. In the top-right corner of the DR service detail page, click Edit. In the popup window, make the necessary changes and save them.

### Deleting the Disaster Recovery Service

To deactivate disaster recovery, delete the DR service setup from both peered appliances.

### Using the Service CLI

Obtain the serial number of the peered appliance (`drNativeListService`). Enter the delete command as follows:

```
PCA-ADMIN> drNativeDeleteService peerSerialNumber=<peer_serial>
```

### Using the Service Web UI

Under Disaster Recovery Service, open the Native Services page. Click the DR service. In the top-right corner of the DR service detail page, click Delete.

# Managing Disaster Recovery Configurations

DR configurations specify the resources that play a critical role in protecting workloads against site-level incidents. These resources include compute instances with their associated block volumes, as well as the compartments they belong to and the network resources that provide their connectivity. Relevant network resources and compartment hierarchies must be set up on both the primary and standby rack, and associated with each other through site mappings.

## Rules and Conditions

When populating DR configurations, respect the following rules regarding compute and storage resources.

- A compute instance must be stopped before it can be added to a DR configuration. There is one exception: when all volumes attached to the instance are also attached to one or more instances already included in the same DR configuration.
- A compute instance must be stopped before it can be removed from a DR configuration. There is one exception: when all volumes attached to the instance are also attached to one or more instances still included in the same DR configuration.
- All compute instances in a DR configuration must be stopped before the DR configuration can be deleted.
- A volume attached to a compute instance might be created from another source volume or volume backup. Such an instance (instance T) can be added to a DR configuration on condition that the source volume is not attached to any instance in any DR configuration. Note that *source volume* also refers to the volume used for the volume backup, and its direct or indirect source.

Alternatively, the instance with the source volume attached can be added to a DR configuration on condition that *instance T* is not added to any DR configuration. Due to the volume source/target relationship, only one of the instances involved can be part of a DR configuration, not both.

- A DR configuration must be refreshed when configuration for any of its instances has changed, including changes for the instance and the attached storage or network resources. This is to ensure the instances on the standby rack after switchover or failover are started with needed attributes preserved.
- Refreshing a DR configuration results in a failure in case a volume **and** the source from which it was created, are both attached to one or more compute instances in any DR configuration.
- Instances in a DR configuration preserve the primary and secondary private IPs upon switchover or failover on condition that the subnet for the IP address is tagged appropriately. The subnet must be assigned a *freeform tag*, with the key `preserve_private_ips` and a value set to anything.

When this freeform tag is applied, any instance with a VNIC attached to this subnet preserves the associated primary and secondary IP addresses after switchover or failover. The precheck on the standby rack verifies whether the IP address is already in use, and fails if this is true.

## Creating a DR Configuration

A DR configuration is the parent object to which you add cloud resources that you want to protect against system outages. You start by creating an empty DR configuration with an associated ZFS project on the ZFS Storage Appliance.

If your system contains optional high-performance disk shelves, two ZFS projects are created for each DR configuration: one in the standard pool and one in the high-performance pool. When you add instances to the DR configuration that have disks running on standard as well as high-performance storage, those storage resources are automatically added to the ZFS project in the appropriate pool.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - a unique name for the DR configuration
  - a unique name for the associated ZFS storage project
2. Create an empty DR configuration with the `create DrConfig` command.

Syntax (entered on a single line):

```
create DrConfig
configName=<DR_configuration_name>
zfsProjectName=<ZFS_storage_project_name>
```

Example:

```
PCA-ADMIN> create DrConfig configName=mydrconf1 zfsProjectName=mydrconf1-project
JobId: 16047c38-1a7e-48ac-9588-297b915d49fc
Data: DrConfig id: d3cd87de-afd4-4718-a3e6-1105b56b42d8. Successfully started job to create config mydrconf1
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=16047c38-1a7e-48ac-9588-297b915d49fc
Data:
Id = 16047c38-1a7e-48ac-9588-297b915d49fc
Type = Job
Associated Work Request Id = 20d20925-39d1-4f9a-beef-5e9d3ca465d5
Done = true
Name = CREATE_TYPE
Progress Message = DrConfig id: d3cd87de-afd4-4718-a3e6-1105b56b42d8. Successfully created DR config mydrconf1:
d3cd87de-afd4-4718-a3e6-1105b56b42d8
Run State = Succeeded
Transcript = Created job CREATE_TYPE
Username = admin
WorkItemIds 1 = id:e82ecec2-7881-45c5-ae84-6a3894ef66c6 type:WorkItem name:
```

4. When the DR configuration is created, the storage project for data replication is set up on the ZFS Storage Appliances. Note the DR configuration ID. You need it for all subsequent commands to modify the configuration.
5. To display a list of existing DR configurations, use the `drGetConfigs` command.

```
PCA-ADMIN> drGetConfigs
Data:
id          Config Name
-----
d3cd87de-afd4-4718-a3e6-1105b56b42d8  mydrconf1
e8291afa-a413-4932-880a-abb8ac22c85d  mydrconf2
7ad05d9f-731c-41b8-b477-35da4b999071  mydrconf3
```

6. To display the status and details of a DR configuration, use the `show DrConfig` command.

```
PCA-ADMIN> show DrConfig id=d3cd87de-afd4-4718-a3e6-1105b56b42d8
Data:
Id = d3cd87de-afd4-4718-a3e6-1105b56b42d8
Type = DrConfig
Config State = Enabled
Config Name = mydrconf1
```

```
Config Id = d3cd87de-afd4-4718-a3e6-1105b56b42d8
Zfs Project Name = mydrconf1-project
Message = Successfully retrieved config data
Message = Successfully retrieved config data
Replica State = OK
Replica Lag In Seconds = 3
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the top-right corner, click Create Configuration.
2. In the Create Configuration window, enter the following information:
  - a unique name for the DR configuration
  - a unique name for the associated ZFS storage project
3. Click Submit. A new, empty DR configuration appears in the table.

Next, you add site mappings and instances to the DR configuration.

### About the DR Configuration State

The normal working state of a DR configuration is "Enabled". However, as part of the process of DR plan execution, the state of the DR configurations changes. In general, an administrator is not expected to manually change the state, but it might be useful, and it is possible to temporarily disable a DR configuration.

The following configuration states occur:

- **Enabled:** The DR configuration is in active working state. The mappings and instances can be modified, and DR plans can be modified and executed.
- **Disabled:** The DR configuration is locked. The mappings and instances cannot be modified, and DR plans cannot be created, edited, or executed.
- **Frozen:** The DR configuration is locked, like in the disabled state. It is no longer usable because it was migrated to the standby appliance.

## Maintaining Site Mappings

Site mappings determine how and where on the standby system the instances should be brought back up in case the primary system experiences an outage. Each site mapping consists of a source object on the primary system and a corresponding target object on the standby system. Ensure that these resources exist on both systems before you add the site mappings to the DR configuration.

### Adding Site Mappings

These are the site mapping types you can include in a DR configuration:

- **Compartment:** specifies that, if a failover or switchover occurs, instances from the source compartment must be brought up in the target compartment on the standby system
- **Subnet:** specifies that, if a failover or switchover occurs, instances connected to the source subnet must be connected to the target subnet on the standby system

### Using the Service CLI

1. Gather the information that you need to run the command:

- DR configuration ID (drGetConfigs)
  - Mapping source and target object OCIDs
- Use the Compute Enclave UI or CLI on the primary and standby system respectively.  
CLI commands:
- oci iam compartment list
  - oci network subnet list --compartment-id "ocid1.compartment.....uniqueID"

2. Add a site mapping to the DR configuration with the drAddSiteMapping command.

Syntax (entered on a single line):

```
drAddSiteMapping
drConfigId=<DR_configuration_id>
objType=[compartment | subnet]
sourceId=<source_object_OCID>
targetId=<target_object_OCID>
```

Examples:

```
PCA-ADMIN> drAddSiteMapping drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8 objType=compartment \
sourceId=ocid1.compartment.....<region1>.....uniqueID targetId=ocid1.compartment.....<region1>.....uniqueID
Data:
Message = Successfully added site mapping to DR config [mydrconf1]
Returned Object Id = 77233b2c-aa4a-4314-8196-4d8dd1d20721
```

```
PCA-ADMIN> drAddSiteMapping drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8 objType=subnet \
sourceId=ocid1.subnet.....<region1>.....uniqueID targetId=ocid1.subnet.....<region2>.....uniqueID
Data:
Message = Successfully added site mapping to DR config [mydrconf1]
Returned Object Id = 4467bb39-7256-4715-9a01-bedd25f73a82
```

3. Repeat the command with the OCIDs of all the source and target objects that you want to include in the site mappings of the DR configuration.
4. To display the list of site mappings included in the DR configuration, use the drGetSiteMappings command. The DR configuration ID is a required parameter.

Syntax:

```
PCA-ADMIN> drGetSiteMappings drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
Data:
id                Object Type
--                -
77233b2c-aa4a-4314-8196-4d8dd1d20721  compartment
6ef9e3b6-6704-4260-b345-f55e364ba9e9  compartment
4467bb39-7256-4715-9a01-bedd25f73a82  subnet
096a5499-9720-4a4f-a8be-ce055f0efd5d  subnet
```

5. To display the status and details of a site mapping included in the DR configuration, use the drGetSiteMapping command. The DR configuration ID and site mapping ID are required parameters.

```
PCA-ADMIN> drGetSiteMapping drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
mappingId=4467bb39-7256-4715-9a01-bedd25f73a82
Data:
Type = DrSiteMapping
Mapping Id = 4467bb39-7256-4715-9a01-bedd25f73a82
Object Type = subnet
Source Id = ocid1.subnet1.....<region1>.....uniqueID
```

Target Id = ocid1.subnet1....<region2>....uniqueID  
 Message = Successfully retrieved mapping data

### Using the Service Web UI

1. Collect the OCIDs of the source and target objects you are going to map. Use the Compute Enclave UI or CLI on the primary and standby system respectively. CLI commands:
  - `oci iam compartment list`
  - `oci network subnet list --compartment-id "ocid1.compartment.....uniqueID"`
2. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration to which you want to add site mappings. The DR Configuration detail page appears.
3. In the Resources section, click Site Mappings.  
 On the right hand side of the Site Mappings box, click Add Site Mapping.
4. In the Add Site Mapping window, enter the following information in the respective fields:
  - **Object Type:** Specify which resource type you are mapping. Enter either compartment or subnet.
  - **Target ID:** Enter the OCID of the target object on the remote system.
  - **Source ID:** Enter the OCID of the source object on the local system.
 Click Add Site Mapping. The new site mapping appears in the resources table.  
 Optionally, click the site mapping to display its detail page.
5. Repeat these steps to add all the required site mappings.  
 As cloud resources across the environment are added and removed over time, verify on a regular basis that the site mappings in the DR configuration are up-to-date.

### Removing Site Mappings

You can remove a site mapping from the DR configuration if it is no longer required.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)
  - Site mapping ID (`drGetSiteMappings`)
2. Remove the selected site mapping from the DR configuration with the `drRemoveSiteMapping` command.

Syntax (entered on a single line):

```
drRemoveSiteMapping
drConfigId=<DR_configuration_id>
mappingId=<site_mapping_id>
```

Example:

```
PCA-ADMIN> drRemoveSiteMapping drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
mappingId=1ce736ef-13f9-40a9-aa6b-3e8960261224
Data:
Message = Successfully removed mapping [1ce736ef-13f9-40a9-aa6b-3e8960261224] from DR config [mydrconf1]
```

- Repeat the command with the IDs of all the site mappings that you want to remove from the DR configuration.

### Using the Service Web UI

- Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to edit site mappings. The DR Configuration detail page appears.
  - In the Resources section, click Site Mappings. All existing site mappings are displayed in the table.
  - In the Actions column, open the quick menu (3 dots) and click Remove Site Mapping. Alternatively, click the site mapping to display its detail page. In the top-right corner, click Remove Site Mapping.
  - Repeat until all obsolete site mappings have been removed.
- As cloud resources across the environment are added and removed over time, verify on a regular basis that the site mappings in the DR configuration are up-to-date.

## Adding and Removing Compute Instances

When the relevant site mappings have been set up in the DR configuration, you add the required compute instances. Their data and disks are stored in the ZFS storage project associated with the DR configuration, and replicated over the peer connection between the Private Cloud Appliance systems.

### Adding Instances

Compute instances must be stopped before you add them to a DR configuration. Instances can be added one at a time, per compartment, or all at once across all compartments.

#### Note

In the Service Web UI, compute instances must be added one by one.

### Using the Service CLI

- Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)
  - Instance OCIDs from the Compute Enclave UI or CLI (`oci compute instance list --compartment-id <compartment_OCID>`)
- Add compute instances to the DR configuration with the `drAddComputeInstance` command.

Syntax (entered on a single line):

```
drAddComputeInstance
drConfigId=<DR_configuration_id>
instanceId="<instance-1_OCID>","<instance-n_OCID>"
compartmentId="<compartment-1_OCID>","<compartment-1_OCID>"
all=True
```

**Note**

When selecting "all", the instance and compartment parameters cannot be included in the command.

**Examples:**

- individual instances: single instance or comma-separated list

```
PCA-ADMIN> drAddComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
instanceId="ocid1.instance1....uniqueID","ocid1.instance2....uniqueID"
Data:
Message = Successfully started job to add instance to config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = dc0787ef-c254-4d91-b7dc-ca4d2412f461
```

- all instances in a compartment: single compartment or comma-separated list

```
PCA-ADMIN> drAddComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
compartmentId="ocid1.compartment1....uniqueID","ocid1.compartment2....uniqueID"
Data:
Message = Successfully started job to add instance to config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = 3801d0be-f2ab-48d5-ad12-cd8afd444bc2
```

- all instances across all mapped compartments

```
PCA-ADMIN> drAddComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8 all=True
Data:
Message = Successfully started job to add instance to config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = d0a11091-f202-4350-b59a-f0cddc9fb035
```

- Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=3801d0be-f2ab-48d5-ad12-cd8afd444bc2
Data:
Type = add_computeinstance
Job Id = 3801d0be-f2ab-48d5-ad12-cd8afd444bc2
Status = finished
Result = success
Message = job successfully retrieved
Response = Successfully added the following instances to mydrconf1
['ocid1.instance1....uniqueID','ocid1.instance2....uniqueID']
```

- Repeat the `drAddComputeInstance` command to add all the required compute instances to the DR configuration.
- To display the list of instances included in the DR configuration, use the `drGetComputeInstances` command. The DR configuration ID is a required parameter.

```
PCA-ADMIN> drGetComputeInstances drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
Data:
id
--
ocid1.instance1....uniqueID
ocid1.instance2....uniqueID
ocid1.instance3....uniqueID
ocid1.instance4....uniqueID
```

- To display the status and details of an instance included in the DR configuration, use the `drGetComputeInstance` command. The DR configuration ID and instance ID are required parameters.

```
PCA-ADMIN> drGetComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
instanceId=ocid1.instance1....uniqueID
Data:
  Compartment Id = ocid1.compartment....uniqueID
  Boot Volume Id = ocid1.bootvolume....uniqueID
  Compute Instance Shape = VM.PCAStandard1.Flex
  Message = Successfully retrieved instance data
```

### Using the Service Web UI

1. Collect the OCIDs of the compute instances that need to be added to the DR configuration. Use the Compute Enclave UI, or these CLI commands:
  - `oci iam compartment list`
  - `oci compute instance list --compartment-id <compartment_OCID>`
2. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration to which you want to add compute instances. The DR Configuration detail page appears.
3. In the Resources section, click Compute Instances.

On the right hand side of the Compute Instances box, click Add Instance.

4. In the Add Compute Instances window, select compute instances using one of these methods:
  - By instance OCID  
Add a single instance by entering its OCID. Add multiple instances by entering their OCIDs as a comma-separated list.
  - By compartment OCID  
Add all instances in a compartment by entering the compartment OCID. Add instances from multiple compartments by entering the compartment OCIDs as a comma-separated list.
  - All  
Add all instances from all mapped compartments.

Click Submit. A DR job is started. When it completes successfully, the instances appear in the resources table.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

5. Check regularly for instances requiring DR protection, and repeat these steps whenever necessary.

As cloud resources across the environment are added and removed over time, verify on a regular basis that the list of compute instances in the DR configuration is up-to-date.

### Removing Instances

Instances can only be part of a single DR configuration. You can remove a compute instance from the DR configuration to which it was added.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)

- Instance OCID (drGetComputeInstances)
2. Remove the selected compute instances from the DR configuration with the drRemoveComputeInstance command.

Syntax (entered on a single line):

```
drRemoveComputeInstance
drConfigId=<DR_configuration_id>
instanceId=<instance_OCID>
compartmentId=<compartment_OCID>
all=True
```

### Note

When selecting "all", the instance and compartment parameters cannot be included in the command.

Examples:

- single instance

```
PCA-ADMIN> drRemoveComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
instanceId=ocid1.instance....uniqueID
Data:
Message = Successfully started job to remove instance(s) from DR config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = 20f6735e-7702-4d27-ba19-41cf08a86d90
```

- all instances in a compartment

```
PCA-ADMIN> drRemoveComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8
compartmentId=ocid1.compartment....uniqueID
Data:
Message = Successfully started job to remove instance(s) from DR config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = c4b0ff83-45e3-4498-98bb-e8a2f6b904e8
```

- all instances in the DR configuration

```
PCA-ADMIN> drRemoveComputeInstance drConfigId=d3cd87de-afd4-4718-a3e6-1105b56b42d8 all=True
Data:
Message = Successfully started job to remove instance(s) from DR config d3cd87de-afd4-4718-a3e6-1105b56b42d8
Job Id = 0a2cdc62-1d13-4774-b3c6-d2a30b75db5e
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drgetjob jobId=20f6735e-7702-4d27-ba19-41cf08a86d90
Data:
Type = remove_computeinstance
Job Id = 20f6735e-7702-4d27-ba19-41cf08a86d90
Status = finished
Result = success
Message = job successfully retrieved
Response = Successfully removed the following instances from mydrconf1 ['ocid1.instance....uniqueID']
```

4. Repeat the drRemoveComputeInstance command to remove all the required compute instances from the DR configuration.

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration from which you want to remove compute instances. The DR Configuration detail page appears.
2. In the Resources section, click Compute Instances. All instances included in the DR configuration are displayed in the table.
3. To quickly remove a single instance, in the Actions column, open the quick menu (3 dots) and click Delete. To remove multiple instances, proceed to the next step.
4. Open the Controls menu in the top-right corner of the DR Configuration detail page and click Remove Instances.
5. In the Remove Compute Instances window, select compute instances for removal using one of these methods:
  - By instance OCID  
Remove a single instance by entering its OCID. Remove multiple instances by entering their OCIDs as a comma-separated list.
  - By compartment OCID  
Remove all instances in a compartment by entering the compartment OCID. Remove instances included from multiple compartments by entering the compartment OCIDs as a comma-separated list.
  - All  
Remove all instances included from all mapped compartments.

Click Submit. A DR job is started. When it completes successfully, the instances are removed from the resources table.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

6. Check regularly for instances that no longer require DR protection, and repeat these steps whenever necessary.

As cloud resources across the environment are added and removed over time, verify on a regular basis that the list of compute instances in the DR configuration is up-to-date.

## Refreshing a DR Configuration

To ensure that the replication information stored in a DR configuration is updated with all the latest changes in your environment, you must refresh the DR configuration after making changes to the protected resources. Refreshing the DR configuration is important for the continuation of service after a DR plan is executed, because it keeps the relevant metadata up-to-date with the current status of the resources under DR protection.

### Using the Service CLI

1. Look up the ID of the DR configuration you want to refresh (`drGetConfigs`).
2. Refresh the data stored in the selected DR configuration with the `drRefreshConfig` command.

```
PCA-ADMIN> drRefreshConfig drConfigId=b51a51ac-043e-4b6d-ab29-20f86d905e81
Data:
Message = Successfully started job to refresh DR config b51a51ac-043e-4b6d-ab29-20f86d905e81
Job Id = ed3a509e-4ae9-4314-9ec5-fb12ff78615a
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> drGetJob jobId=ed3a509e-4ae9-4314-9ec5-fb12ff78615a
Data:
Type = refresh_config
Job Id = ed3a509e-4ae9-4314-9ec5-fb12ff78615a
Status = finished
Result = success
Message = job successfully retrieved
Response = Successfully refreshed DR config [mydrconf1]
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page.
2. In the table, click the DR configuration to open its detail page.
3. In the top-right corner, select Controls, then click Refresh Configuration.
4. A DR job is started. When it completes successfully, the DR configuration has been refreshed.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

## Deleting a DR Configuration

When you no longer need a DR configuration, you can remove it with a single command. However, you must stop all compute instances that are part of the DR configuration before you can delete it. Deleting a DR configuration also removes all site mappings and cleans up the associated storage projects on the local ZFS Storage Appliance.

If required, the storage projects replicated on the standby ZFS Storage Appliance can be cleaned up manually. If the native DR service is unconfigured on both systems, all related storage project on both systems are cleaned up automatically as part of the process.

### Using the Service CLI

1. Stop all the compute instances that are part of the DR configuration you want to delete.
2. Look up the ID of the DR configuration you want to delete (`drGetConfigs`).
3. Delete the selected DR configuration with the `drDeleteConfig` command.

```
PCA-ADMIN> delete DrConfig id=d3cd87de-afd4-4718-a3e6-1105b56b42d8
JobId: a687aafa-7c5c-4a68-99da-62e9384140be
Data: Successfully started job to delete config d3cd87de-afd4-4718-a3e6-1105b56b42d8
```

4. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=a687aafa-7c5c-4a68-99da-62e9384140be
Data:
Id = a687aafa-7c5c-4a68-99da-62e9384140be
Type = Job
Associated Work Request Id = e5f77289-b263-4a5c-9163-7ce4629cc550
Done = true
Name = DELETE_TYPE
Progress Message = Successfully deleted DR config [mydrconf1]
Run State = Succeeded
Transcript = Created job DELETE_TYPE
Username = admin
WorkItemIds 1 = id:622a4f37-9222-4ea4-af2c-351cf28e739b type:WorkItem name:
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. All configurations are displayed in the table.
2. In the Actions column, open the quick menu (3 dots) and click Delete.  
Alternatively, click the DR configuration to display its detail page. In the top-right corner, select Controls, then click Delete.
3. When prompted, click Confirm. A DR job is started. When it completes successfully, the DR configuration is removed from the table.  
To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

## Working With Disaster Recovery Plans

A DR plan describes the operations that must be performed on the resources that are under the protection of the disaster recovery service. A DR plan is associated with a DR configuration, and is executed by an administrator either when a site-level incident is detected (failover), or when one of the sites must be taken offline (switchover). After a failover, when the affected system is back online, postfailover operations are performed to ensure that both systems are ready to run new DR operations.

## About DR Operations and Default Plans

The native DR service provides plans with default steps for each type of operation. DR plan steps can be customized. The built-in plans are configured as follows:

### Switchover Plan

When a switchover is performed, there is no outage, so both peered systems are online. The goal is to move all resources covered in the DR configuration from the primary system (A) to the standby system (B). When completed, system B becomes the primary and system A the standby for the resources in question.

The plan starts with prechecks to ensure that both systems meet the requirements to allow compute instances to be stopped on the primary system and started again on the standby system. The prechecks include site mappings as well as other critical elements, such as tags, security lists, or network security groups. The role reversal precheck specifically ensures that the ZFS Storage Appliance in each rack is in the correct state.

When the prechecks are completed without errors, the DR configuration on the primary system (A) is frozen and its compute instances are stopped, so the role reversal can begin. Based on resource metadata exchanged between the peered systems, and replicated data on the standby ZFS Storage Appliance, the target system (B) is prepared to assume the primary role for the instances in the DR configuration. The replication process is reversed and ready to use the source system (A) as its standby as soon as the switchover is complete.

Using the replicated volumes, the compute instances in the DR configuration are launched on the standby system (B). An identical DR configuration is created on the standby system, with all source and target resources in the site mappings inverted. The metadata of the newly launched instances is stored in the DR configuration. On the primary system (A) a cleanup is performed: the DR configuration is disabled and its compute instances are terminated.

To complete the switchover, data replication from the new primary system (B) to the standby system (A) is started, the DR plans are moved to the new standby system (A), and the storage project and metadata associated with the original DR configuration are deleted from system A.

### Failover Plan

A failover is performed on the standby system, when one of the peered systems goes down. The goal is to recover all resources covered in the DR configuration on the standby system (B), allowing continuation of service. The failover steps are similar to the switchover plan, but none of the operations on the primary system (A) can be performed. The primary system cannot be cleaned up until it comes back online.

The plan starts with prechecks to ensure that the standby system and its ZFS Storage Appliance are in the correct state to bring up the resources covered in the DR configuration. When the prechecks are completed without errors, the role reversal begins.

Using the replicated metadata and resources, the compute instances in the DR configuration are launched on the standby system (B), which assumes the primary role. An identical DR configuration is created on system B, which has become the primary, with inverted site mappings and metadata collected from the newly launched instances. In preparation of the original primary system (A) coming back online, the replication process is reversed and ready to use system A as the standby.

When the original primary system (A) eventually comes online, the remaining steps to return the DR configuration to a correct working state are performed by executing the postfailover plan.

### Postfailover Plan

A postfailover plan is performed after a failover, when the system that experienced an outage comes back online, and the peer connection is restored. The goal is to clean up the DR configuration on the primary system that went down (A), and set it up as the standby for the new primary system (B).

There are no prechecks in a postfailover plan. System A is back online after an outage and needs to be cleaned up: the DR configuration is disabled and its compute instances are terminated. Data replication from the new primary system (B) to the standby system (A) is started, the DR plans are moved to the new standby system (A), and the storage project and metadata associated with the original DR configuration are deleted from system A.

To move resources that were originally hosted on system A back from system B, the administrator must perform a switchover from B to A for the relevant DR configuration(s).

## Creating and Maintaining DR Plans

Each DR plan defines the steps to perform when a DR operation is executed: failover, switchover, or postfailover. Each DR configuration has an associated DR plan file. This file can contain multiple DR plans, each identified by a unique name assigned by the administrator during creation. The DR plan file is stored in JSON format in the directory `/mnt/dr_metadata` on the standby rack.

### Creating a DR Plan

The DR service provides default plans for each standard operation. You create them with a simple command to get started. If necessary, you can customize the plan steps afterward.

### Using the Service CLI

1. Look up the ID of the DR configuration for which you want to create a DR plan (`drGetConfigs`).
2. Create a default DR plan with the `create DrPlan` command.

Typically, a DR configuration has associated DR plans for each operation type. In this example, plans are created for failover, switchover, and postfailover.

```
PCA-ADMIN> create DrPlan drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041 operation=SWITCHOVER
planName=sw1
JobId: eae66f69-7b99-420e-b324-7d8964b2202b
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1. Successfully started job for DR Plan Create for config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041
```

```
PCA-ADMIN> create DrPlan drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041 operation=FAILOVER planName=fo1
```

```
PCA-ADMIN> create DrPlan drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041 operation=POSTFAILOVER
planName=pfo1
```

### 3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=eae66f69-7b99-420e-b324-7d8964b2202b
Data:
Id = eae66f69-7b99-420e-b324-7d8964b2202b
Type = Job
Associated Work Request Id = ec0f39df-6256-4c4c-a839-0d00a8f326dc
Done = true
Name = CREATE_TYPE
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1. Successfully created [sw1] plan on
STANDBY for DR operation [switchover] for config [6e797d8b-7245-4d49-8e68-bf67f2d53041]
Run State = Succeeded
Transcript = Created job CREATE_TYPE
Username = admin
WorkItemIds 1 = id:5ca6d187-e01a-40e2-bc97-3193a9a88742 type:WorkItem name:
```

### 4. To display a list of existing DR plans for a DR configuration, use the drListPlan command.

```
PCA-ADMIN> drListPlan drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041
Data:
id                Plan Name  Operation
--                -
6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1  sw1       switchover
6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1  fo1       failover
6e797d8b-7245-4d49-8e68-bf67f2d53041::pfo1 pfo1      postfailover
```

### 5. To display the status and details of a DR plan, use the show DrPlan command.

```
PCA-ADMIN> show DrPlan id=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Data:
Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Type = DrPlan
Plan Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Plan Name = sw1
Config Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041
Operation = switchover
Steps 1 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck
Steps 1 - Step Name = PRECHECK
Steps 1 - Enabled = true
Steps 1 - Last Status = norun
Steps 1 - Command = None
Steps 1 - Check Only = true
Steps 2 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role_reversal_precheck
Steps 2 - Step Name = ROLE_REVERSAL_PRECHECK
Steps 2 - Enabled = true
Steps 2 - Last Status = norun
Steps 2 - Command = None
Steps 2 - Check Only = true
Steps 3 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::stop_primary
Steps 3 - Step Name = STOP_PRIMARY
```

```

Steps 3 - Enabled = true
Steps 3 - Last Status = norun
Steps 3 - Command = None
Steps 3 - Check Only = false
Steps 4 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role_reversal
Steps 4 - Step Name = ROLE_REVERSAL
Steps 4 - Enabled = true
Steps 4 - Last Status = norun
Steps 4 - Command = None
Steps 4 - Check Only = false
Steps 5 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::start_standby
Steps 5 - Step Name = START_STANDBY
Steps 5 - Enabled = true
Steps 5 - Last Status = norun
Steps 5 - Command = None
Steps 5 - Check Only = false
Steps 6 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::cleanup_primary
Steps 6 - Step Name = CLEANUP_PRIMARY
Steps 6 - Enabled = true
Steps 6 - Last Status = norun
Steps 6 - Command = None
Steps 6 - Check Only = false
Steps 7 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::post_config
Steps 7 - Step Name = POST_CONFIG
Steps 7 - Enabled = true
Steps 7 - Last Status = norun
Steps 7 - Command = None
Steps 7 - Check Only = false

```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration to which you want to add a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans.  
On the right hand side of the Plans box, click Add Plan.
3. In the Create DR Plan window, enter the following information in the respective fields:
  - **Plan Name:** Enter a name for this DR plan.
  - **Operation:** Select a standard operation: switchover, failover, postfailover.
  - **Steps:** Leave blank if you want the default steps for the selected operation. Otherwise, specify the steps to add to the DR plan. The options are: 'precheck', 'role\_reversal\_precheck', 'stop\_primary', 'role\_reversal', 'start\_standby', 'cleanup\_primary', 'post\_config'.
  - **Dr Plan Data File:** Used for migrating first-generation DR plans. Leave blank.
  - **Configurations To Migrate:** Used for migrating first-generation DR plans. Leave blank.
4. Click Create DR Plan. A DR job is started. When it completes successfully, the DR plan appears in the resources table.  
To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.
5. Repeat these steps to add all the required DR plans.
6. Optionally, click a DR plan name to display its detail page.

## Changing a DR Plan

A DR plan has a complex structure with many parameters, which makes it difficult to update from the command line in particular. Instead, you can change the individual steps that make up the DR plan. See [Customizing the Steps in a DR Plan](#).

## Deleting a DR Plan

You can remove a DR plan if it is no longer required. To delete all plans associated with a DR configuration at once, use the command `drDeleteAllPlans`.

## Using the Service CLI

1. Look up the ID of the DR plan you want to delete (`drListPlan`).
2. Remove the selected DR plan with the `delete DrPlan` command.

```
PCA-ADMIN> delete DrPlan id=6e797d8b-7245-4d49-8e68-bf67f2d53041::fo2
JobId: 603d480f-1e0f-4229-b596-aaaf8588e682
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo2. Successfully started job for DR Plan delete for config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=603d480f-1e0f-4229-b596-aaaf8588e682
Data:
Id = 603d480f-1e0f-4229-b596-aaaf8588e682
Type = Job
Associated Work Request Id = 391a0799-235a-4b26-aa99-4b5dd14ba19a
Done = true
Name = DELETE_TYPE
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo2. Successfully deleted [fo2] plan on
STANDBY for config [6e797d8b-7245-4d49-8e68-bf67f2d53041]
Run State = Succeeded
Transcript = Created job DELETE_TYPE
Username = admin
WorkItemIds 1 = id:86d0bd13-5f9c-4513-9404-60d8980b2243 type:WorkItem name:
```

## Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to delete a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans. All existing DR plans are displayed in the table.
3. In the Actions column, open the quick menu (3 dots) and click Delete.  
Alternatively, click the DR plan name to display its detail page. In the top-right corner, click Delete.
4. When prompted, click Confirm. A DR job is started. When it completes successfully, the DR plan is removed from the table.  
To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.
5. Repeat until all obsolete DR plans have been removed.

## Deleting All DR Plans

Instead of deleting DR plans one by one, you can delete all plans associated with a DR configuration at once.

### Using the Service CLI

1. Look up the ID of the DR configuration for which you want to delete all plans (drGetConfigs).
2. Remove the selected DR plan with the drdeleteallPlans command.

```
PCA-ADMIN> drdeleteallPlans drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041
JobId: b595dd62-8046-4ca7-90a0-dcbbf084e663
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::all_plans. Successfully started job for DR Plan delete for
config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=b595dd62-8046-4ca7-90a0-dcbbf084e663
Data:
Id = b595dd62-8046-4ca7-90a0-dcbbf084e663
Type = Job
Associated Work Request Id = d7cfb184-4e38-469d-b189-bb809386f5d4
Done = true
Name = DELETE_TYPE
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::all_plans. Successfully deleted [all_plans] plan
on STANDBY for config [6e797d8b-7245-4d49-8e68-bf67f2d53041]
Run State = Succeeded
Transcript = Created job DELETE_TYPE
Username = admin
WorkItemIds 1 = id:6cfec1c9-4a72-492d-a414-38bf4df6cf3a type:WorkItem name:
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to delete a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans. All existing DR plans are displayed in the table. Verify that all plans in the table should be deleted.
3. In the top-right corner, select Controls, then click Delete All Plans.
4. When prompted, click Confirm. A DR job is started. When it completes successfully, all DR plan are removed from the table.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

## Customizing the Steps in a DR Plan

Customizing DR plan steps is the most convenient way to change a DR plan with default configuration. The administrator can enable or disable steps, configure steps to run in *check-only* mode, insert steps, and remove steps.

#### Note

Custom steps and custom commands are not available in appliance software version 3.0.2-b1261765.

### Changing the Properties of a DR Plan Step

You can enable or disable an existing DR plan step, and decide whether it should be performed in check-only mode.

### Using the Service CLI

1. Look up the ID of the DR plan step you want to change (show DrPlan). If you don't have the DR plan ID, look it up using the commands `drGetConfigs` and `drListPlan`.
2. Change the DR plan step with the `edit DrPlanStep` command. These are the properties you can change in order to affect how a DR plan is executed:
  - `enabled=[True|False]` – determines whether this step is performed or not
  - `checkOnly=[True|False]` – determines whether this step is performed in check-only mode
  - `command=[string]` – (NOT available) specifies a custom command to be run as part of this step

For standard DR plan operations, the command parameter must be set to `None`.

```
PCA-ADMIN> edit DrPlanStep id=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck checkOnly=False
JobId: d281141c-c388-490e-b038-239598488bc6
Data: DrPlanStep id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck. Successfully started job for DR Plan Step
update for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=d281141c-c388-490e-b038-239598488bc6
Data:
Id = d281141c-c388-490e-b038-239598488bc6
Type = Job
Associated Work Request Id = 8f40d0a3-b2ac-4742-bed8-70beb26d669d
Done = true
Name = MODIFY_TYPE
Progress Message = DrPlanStep id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck. Successfully updated plan
step [precheck] in DR Plan [sw1] for config [6e797d8b-7245-4d49-8e68-bf67f2d53041]
Transcript = Created job MODIFY_TYPE
Username = admin
WorkItemIds 1 = id:92a94a5e-4773-4ba8-ac4b-990496c5c2f9 type:WorkItem name:
```

4. To display the status and updated details of a DR plan step, use the `show DrPlanStep` command.

```
PCA-ADMIN> show DrPlanStep id=6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1::precheck
Data:
Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1::precheck
Type = DrPlanStep
Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1::precheck
Step Name = PRECHECK
Config Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041
Plan Name = fo1
Enabled = true
Last Status = norun
Command = None
Check Only = false
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to modify a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans. Click the name of the plan for which you want to edit the steps. The DR Plan detail page appears. The Resources section displays all steps in the plan.

3. In the Actions column, open the quick menu (3 dots) for the step you want to change, and click Edit.
4. In the Update DR Plan Step window, edit the properties as needed:
  - **Step Name:** Do not change the selected step.
  - **Enabled:** Select Yes or No to determine whether this step is performed or not during plan execution.
  - **Insert Type:** Do not change. This field is used only for adding a new step.
  - **Insert Location:** Do not change. This field is used only for adding a new step.
  - **Command:** (NOT available.) Specify a custom command to be run as part of this step.  
For standard DR plan operations, this field must remain empty. Custom commands only apply to custom DR plan steps.
  - **Check Only:** Select Yes or No to determine whether this step is performed or not in check-only mode.
5. Click Update DR Plan Step to apply your changes.

### Inserting a DR Plan Step

Extra steps can be added to a DR plan.

### Using the Service CLI

1. Gather the information that you need to run the command:
  - DR configuration ID (`drGetConfigs`)
  - DR plan name (`drListPlan`)
  - name of the step before or after which the new step must be inserted (`show DrPlan`)
2. Add a step to the DR plan with the `create DrPlanStep` command.

Note the mandatory and optional parameters. If the `insertType` and `insertLocation` parameters are not provided, the new step is added as the final step of the DR plan.

Syntax (entered on a single line):

```
create DrPlanStep
drConfigId=<DR_configuration_id>
planName=<DR_plan_name>
stepName=<DR_plan_new_step_name>
[optional:]
insertType=[BEFORE/AFTER]
insertLocation=<DR_plan_existing_step_name>
enabled=[True/False]
checkOnly=[True/False]
command=<execution_path>
```

**Note**

Custom steps and custom commands are not available in appliance software version 3.0.2-b1261765.

- Do not use the `command` parameter.
- Select a step name from this list: `precheck`, `role_reversal_precheck`, `stop_primary`, `role_reversal`, `start_standby`, `cleanup_primary`, `post_config`.

**Example:**

```
PCA-ADMIN> create DrPlanStep drConfigId=6e797d8b-7245-4d49-8e68-bf67f2d53041 planName=sw1
stepName=ROLE_REVERSAL insertType=BEFORE insertLocation=START_STANDBY
JobId: 7a162b6c-1ddc-410c-b27a-0996fb2d26df
Data: DrPlanStep id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role_reversal. Successfully started job for DR Plan Step
Create for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041
```

**3. Use the job ID to check the status of the operation you started.**

```
PCA-ADMIN> show Job id=7a162b6c-1ddc-410c-b27a-0996fb2d26df
Data:
Id = 7a162b6c-1ddc-410c-b27a-0996fb2d26df
Type = Job
Associated Work Request Id = a4c76ef9-f91d-402e-986f-9795738fb429
Done = true
Name = CREATE_TYPE
Progress Message = Successfully added step [role_reversal] in plan [sw1] for config [6e797d8b-7245-4d49-8e68-
bf67f2d53041]
Run State = Succeeded
Transcript = Created job CREATE_TYPE
Username = admin
WorkItemIds 1 = id:5c376fae-f281-4f7a-984f-b996b19d367d type:WorkItem name:
```

**4. Repeat the create DrPlanStep command to add all the required steps to the DR plan.****5. To display the status and updated details of a DR plan, use the show DrPlan command.**

```
PCA-ADMIN> show DrPlan id=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Data:
Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Type = DrPlan
Plan Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Plan Name = sw1
Config Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041
Operation = switchover
Steps 1 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck
Steps 1 - Step Name = PRECHECK
Steps 1 - Enabled = true
Steps 1 - Last Status = norun
Steps 1 - Command = None
Steps 1 - Check Only = true
Steps 2 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role_reversal_precheck
Steps 2 - Step Name = ROLE_REVERSAL_PRECHECK
Steps 2 - Enabled = true
Steps 2 - Last Status = norun
Steps 2 - Command = None
Steps 2 - Check Only = true
Steps 3 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::stop_primary
```

Steps 3 - Step Name = STOP\_PRIMARY  
 Steps 3 - Enabled = true  
 Steps 3 - Last Status = norun  
 Steps 3 - Command = None  
 Steps 3 - Check Only = false  
 Steps 4 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role\_reversal  
 Steps 4 - Step Name = **ROLE\_REVERSAL**  
 Steps 4 - Enabled = true  
 Steps 4 - Last Status = norun  
 Steps 4 - Command = None  
 Steps 4 - Check Only = false  
 Steps 5 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::start\_standby  
 Steps 5 - Step Name = START\_STANDBY  
 Steps 5 - Enabled = true  
 Steps 5 - Last Status = norun  
 Steps 5 - Command = None  
 Steps 5 - Check Only = false  
 Steps 6 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::cleanup\_primary  
 Steps 6 - Step Name = CLEANUP\_PRIMARY  
 Steps 6 - Enabled = true  
 Steps 6 - Last Status = norun  
 Steps 6 - Command = None  
 Steps 6 - Check Only = false  
 Steps 7 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::post\_config  
 Steps 7 - Step Name = POST\_CONFIG  
 Steps 7 - Enabled = true  
 Steps 7 - Last Status = norun  
 Steps 7 - Command = None  
 Steps 7 - Check Only = false

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to modify a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans. Click the name of the plan to which you want to add one or more steps. The DR Plan detail page appears. The Resources section displays all steps in the plan.
3. In the Resources section of the DR Plan detail page, on the right hand side of the Steps box, click Add Step.
4. In the Add DR Plan Step window, enter the following information in the respective fields:
  - **Step Name:** Select a standard step from the list.
  - **Enabled:** Select Yes or No to determine whether this step is performed or not during plan execution.
  - **Insert Type:** Select Before or After the insert location.
  - **Insert Location:** Select the existing step before or after which you want the new step to be inserted.
  - **Command:** (NOT available.) Specify a custom command to be run as part of this step.  
 For standard DR plan operations, this field must remain empty.
  - **Check Only:** Select Yes or No to determine whether this step is performed or not in check-only mode.

5. Click Add DR Plan Step to insert this new step in the selected location.
6. Repeat until all the required DR plan steps have been added.

### Deleting a DR Plan Step

Steps can be removed from a DR plan.

### Using the Service CLI

1. Look up the ID of the DR plan step you want to delete (show DrPlan). If you don't have the DR plan ID, look it up using the commands `drGetConfigs` and `drListPlan`.
2. Delete the DR plan step with the `delete DrPlanStep` command.

```
PCA-ADMIN> delete DrPlanStep id=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck
JobId: c75a4c71-0525-40b7-9618-c2a4e8fcb051
Data: DrPlanStep id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck. Successfully started job for DR Plan Step
delete for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=c75a4c71-0525-40b7-9618-c2a4e8fcb051
Data:
Id = c75a4c71-0525-40b7-9618-c2a4e8fcb051
Type = Job
Associated Work Request Id = d7bd3873-e5a0-4437-b6c2-8ca0cce83953
Done = true
Name = DELETE_TYPE
Progress Message = DrPlanStep id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::precheck. Successfully deleted step
[precheck] in plan [sw1] for config [6e797d8b-7245-4d49-8e68-bf67f2d53041]
Run State = Succeeded
Transcript = Created job DELETE_TYPE
Username = admin
WorkItemIds 1 = id:b462883b-6e6e-49b6-b455-1ceb54a5e2e3 type:WorkItem name:
```

4. Repeat the `delete DrPlanStep` command to remove all the required steps from the DR plan.
5. To display the status and updated details of a DR plan, use the `show DrPlan` command.

```
PCA-ADMIN> show DrPlan id=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Data:
Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Type = DrPlan
Plan Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
Plan Name = sw1
Config Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041
Operation = switchover
Steps 1 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::stop_primary
Steps 1 - Step Name = STOP_PRIMARY
Steps 1 - Enabled = true
Steps 1 - Last Status = norun
Steps 1 - Command = None
Steps 1 - Check Only = false
Steps 2 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::role_reversal
Steps 2 - Step Name = ROLE_REVERSAL
Steps 2 - Enabled = true
Steps 2 - Last Status = norun
Steps 2 - Command = None
Steps 2 - Check Only = false
Steps 3 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::start_standby
Steps 3 - Step Name = START_STANDBY
```

```
Steps 3 - Enabled = true
Steps 3 - Last Status = norun
Steps 3 - Command = None
Steps 3 - Check Only = false
Steps 4 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::cleanup_primary
Steps 4 - Step Name = CLEANUP_PRIMARY
Steps 4 - Enabled = true
Steps 4 - Last Status = norun
Steps 4 - Command = None
Steps 4 - Check Only = false
Steps 5 - Step Id = 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1::post_config
Steps 5 - Step Name = POST_CONFIG
Steps 5 - Enabled = true
Steps 5 - Last Status = norun
Steps 5 - Command = None
Steps 5 - Check Only = false
```

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to modify a DR plan. The DR Configuration detail page appears.
2. In the Resources section, click Plans. Click the name of the plan for which you want to edit the steps. The DR Plan detail page appears. The Resources section displays all steps in the plan.
3. In the Actions column, open the quick menu (3 dots) for the step you want to remove, and click Delete.

When prompted, click Confirm.

4. Repeat until all obsolete DR plan steps have been removed.

## Executing a DR Plan

A DR plan can be executed from either the standby or the primary Private Cloud Appliance. However, in the case of a failover, the primary rack is down, so a failover plan is always executed from the standby system.

A switchover can be performed for the purpose of testing the disaster recovery setup, or when extensive maintenance is required on the primary system. To return both appliances to their normal working state after a failover, a postfailover plan is executed on each system when the primary is back online. The switchover plan has postfailover steps built in, so it does not require an additional run of the postfailover plan.

As a result of executing a DR plan, resources are moved between peered systems and the primary system changes. Those resources are not automatically moved back to their original host system. To move resources back to their original environment, you must perform another switchover for the relevant DR configuration(s).

### Performing a Switchover

A switchover allows the administrator to move resources away from a system so it can be taken offline, for example in case of planned maintenance. A (second) switchover is also performed to move resources back to their original host system, after they were impacted by a failover or switchover.

## Using the Service CLI

1. Look up the ID of the switchover DR plan you want to execute. Use `drGetConfigs` to find the DR configuration, and display its associated DR plans using `drListPlan`.
2. From the primary or standby appliance, execute the switchover DR plan with the `drExecutePlan` command.

### Note

To run the command in check-only mode, add the parameter `checkOnly=True`. Only the DR plan steps enabled for check-only mode will be performed.

```
PCA-ADMIN> drExecutePlan planId=6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1
JobId: 92b4acc2-2dff-492c-9ba2-0a2ac058baa5
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1. Successfully started job for DR Plan Execute for config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name sw1
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=92b4acc2-2dff-492c-9ba2-0a2ac058baa5
Data:
Id = 92b4acc2-2dff-492c-9ba2-0a2ac058baa5
Type = Job
Associated Work Request Id = c6cca56c-a1cc-421c-9ded-acf0e7cd9da2
Done = false
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1. Successfully started job for DR Plan
Execute for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name sw1
Run State = Active
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:e06881fc-ea57-4835-bb86-e1244d3787c3 type:WorkItem name:
```

4. Ensure that the job completes successfully.

```
PCA-ADMIN> show Job id=92b4acc2-2dff-492c-9ba2-0a2ac058baa5
Data:
Id = 92b4acc2-2dff-492c-9ba2-0a2ac058baa5
Type = Job
Associated Work Request Id = c6cca56c-a1cc-421c-9ded-acf0e7cd9da2
Done = true
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::sw1. DrPlan id: 6e797d8b-7245-4d49-8e68-
bf67f2d53041::sw1. drexecuteplan succeeded for config [6e797d8b-7245-4d49-8e68-bf67f2d53041] Operation: [switchover]
plan_name: [sw1]. Response: [Successfully completed checks for switchover for DR config_id 6e797d8b-7245-4d49-8e68-
bf67f2d53041. Plan Execution Status: [precheck : pass , role_reversal_precheck : pass , stop_primary : norun , role_reversal :
norun , start_standby : norun , cleanup_primary : norun , post_config : norun , ]]
Run State = Succeeded
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:e06881fc-ea57-4835-bb86-e1244d3787c3 type:WorkItem name:
```

After successful completion, all instances included in the DR configuration have been recovered and are running on the standby appliance.

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to perform a switchover. The DR Configuration detail page appears.
2. In the Resources section, click Plans.
3. In the Actions column, open the quick menu (3 dots) for the switchover plan of your choice, and click Execute Plan.

Alternatively, click the DR plan name to display its detail page. In the top-right corner, click Execute Plan.

4. When prompted, choose whether to execute the full plan or a subset of the steps in *check-only* mode.

Click Confirm. A DR job is started. When it completes successfully, all steps in the switchover DR plan have been performed as expected.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

After successful completion, all instances included in the DR configuration have been recovered and are running on the standby appliance.

### Performing a Failover

The native DR service does not provide automated failover. An administrator must confirm that the primary appliance is down, and execute the failover plan from the standby appliance. A failover is meant to allow continuation of service when the primary system experiences an outage.

#### Note

When one appliance is down, the peer rack reports a fault with a name containing "peerconnect" and the rack serial number. Use the Service CLI to check the fault list (list fault <parameters>) and display the details of the peer connection problem. For example:

```
PCA-ADMIN> show fault id=57701191-5764-480b-826c-38c4b1970dde
Data:
Cause = 1742XC3024 : network is not in a CONNECTED state: CONNECTING
Action = Please contact customer support for solution
Health Exporter = peerconnect-checker
Diagnosing Source = peer connect health checker
Faulted Component Type = SOFTWARE
Description = 1749XC302P-- 1742XC3024 : network is not in a CONNECTED state: CONNECTING
Name = 1749XC302P--PCA-8000-UY--peerconnect
```

### Using the Service CLI

1. Look up the ID of the failover DR plan you need to execute. Use `drGetConfigs` to find the DR configuration, and display its associated DR plans using `drListPlan`.
2. From the standby appliance, execute the failover DR plan with the `drExecutePlan` command.

**Note**

To run the command in check-only mode, add the parameter `checkOnly=True`. Only the DR plan steps enabled for check-only mode will be performed.

```
PCA-ADMIN> drExecutePlan planId=6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1
JobId: 49521287-c148-4791-9626-13190fce3d1d
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1. Successfully started job for DR Plan Execute for config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name fo1
```

3. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=49521287-c148-4791-9626-13190fce3d1d
Data:
Id = 49521287-c148-4791-9626-13190fce3d1d
Type = Job
Associated Work Request Id = c8e3b554-a3ef-4e9b-a52c-c9a518f70974
Done = false
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1. Successfully started job for DR Plan
Execute for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name fo1
Run State = Active
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:d7a09483-ef2e-4e03-81bb-fed5ee661428 type:WorkItem name:
```

4. Ensure that the job completes successfully.

```
PCA-ADMIN> show Job id=49521287-c148-4791-9626-13190fce3d1d
Data:
Id = 49521287-c148-4791-9626-13190fce3d1d
Type = Job
Associated Work Request Id = c8e3b554-a3ef-4e9b-a52c-c9a518f70974
Done = true
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::fo1. DrPlan id: 6e797d8b-7245-4d49-8e68-
bf67f2d53041::fo1. drexecuteplan succeeded for config [6e797d8b-7245-4d49-8e68-bf67f2d53041] Operation: [failover]
plan_name: [fo1]. Response: [Successfully completed checks for failover for DR config_id 6e797d8b-7245-4d49-8e68-
bf67f2d53041. Plan Execution Status: [precheck : pass , role_reversal_precheck : pass , role_reversal : pass , start_standby :
pass , ]]
Run State = Succeeded
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:d7a09483-ef2e-4e03-81bb-fed5ee661428 type:WorkItem name:
```

After successful completion, all instances included in the DR configuration have been recovered and are running on the standby appliance.

### Using the Service Web UI

1. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to perform a switchover. The DR Configuration detail page appears.
2. In the Resources section, click Plans.
3. In the Actions column, open the quick menu (3 dots) for the failover plan of your choice, and click Execute Plan.

Alternatively, click the DR plan name to display its detail page. In the top-right corner, click Execute Plan.

4. When prompted, choose whether to execute the full plan or a subset of the steps in *check-only* mode.

Click Confirm. A DR job is started. When it completes successfully, all steps in the switchover DR plan have been performed as expected.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

After successful completion, all instances included in the DR configuration have been recovered and are running on the standby appliance.

### Performing Postfailover Operations

A postfailover is performed after a failover, when the system that experienced an outage comes back online. The plan can be executed from either of the peered systems. During postfailover, the DR configuration is cleaned up on the primary system that went down. The original standby system becomes the primary for the resources covered by the DR configuration, using the original primary as the new target for DR data replication.

### Using the Service CLI

1. After a failover, confirm that the primary appliance is back online and in healthy condition. Ensure that the peering status is *active* and replication is *enabled*. Neither rack should report an active fault with a name containing "peerconnect". (Check with Service CLI command `list fault`.)
2. Look up the ID of the postfailover DR plan you want to execute. Use `drGetConfigs` to find the DR configuration, and display its associated DR plans using `drListPlan`.
3. From the primary or standby appliance, execute the postfailover DR plan with the `drExecutePlan` command.

#### Note

For postfailover operations, the check-only mode does not apply.

```
PCA-ADMIN> drExecutePlan planId=6e797d8b-7245-4d49-8e68-bf67f2d53041::pfo1
JobId: 56d040ba-30a6-4bea-b924-78ebabed2626
Data: DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::pfo1. Successfully started job for DR Plan Execute for config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name pfo1
```

4. Use the job ID to check the status of the operation you started.

```
PCA-ADMIN> show Job id=56d040ba-30a6-4bea-b924-78ebabed2626
Data:
Id = 56d040ba-30a6-4bea-b924-78ebabed2626
Type = Job
Associated Work Request Id = b4ad564b-e385-4688-94ff-11bf5267d72e
Done = false
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::pfo1. Successfully started job for DR Plan
Execute for config_id 6e797d8b-7245-4d49-8e68-bf67f2d53041, plan_name pfo1
Run State = Active
```

```
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:2e4db010-239e-41a1-aa0d-cb97167c64fc type:WorkItem name:
```

#### 5. Ensure that the job completes successfully.

```
PCA-ADMIN> show Job id=56d040ba-30a6-4bea-b924-78ebabed2626
Data:
Id = 56d040ba-30a6-4bea-b924-78ebabed2626
Type = Job
Associated Work Request Id = b4ad564b-e385-4688-94ff-11bf5267d72e
Done = true
Name = OPERATION-EXECUTE_DR_PLAN
Progress Message = DrPlan id: 6e797d8b-7245-4d49-8e68-bf67f2d53041::pfo1. DrPlan id: 6e797d8b-7245-4d49-8e68-
bf67f2d53041::pfo1. drexecuteplan succeeded for config [6e797d8b-7245-4d49-8e68-bf67f2d53041] Operation:
[postfailover] plan_name: [pfo1]. Response: [Successfully completed checks for postfailover for DR config_id
6e797d8b-7245-4d49-8e68-bf67f2d53041. Plan Execution Status: [stop_primary : pass , cleanup_primary : pass ,
post_config : pass , ]]
Run State = Succeeded
Transcript = Created job OPERATION
Username = admin
WorkItemIds 1 = id:2e4db010-239e-41a1-aa0d-cb97167c64fc type:WorkItem name:
```

After successful completion, all instances impacted by the switchover or failover have been restored and are running on the appliance where they were hosted before.

### Using the Service Web UI

1. After a failover, confirm that the primary appliance is back online and in healthy condition.  
Ensure that the peering status is *active* and replication is *enabled*. Neither rack should report an active fault with a name containing "peerconnect". (Display active faults in the Service Web UI.)
2. Under Disaster Recovery Service, open the DR Configurations page. In the table, click the configuration for which you want to perform postfailover operations. The DR Configuration detail page appears.
3. In the Resources section, click Plans.
4. In the Actions column, open the quick menu (3 dots) for the postfailover plan of your choice, and click Execute Plan.  
Alternatively, click the DR plan name to display its detail page. In the top-right corner, click Execute Plan.
5. When prompted, click Confirm.

#### Note

For postfailover operations, the check-only mode does not apply.

A DR job is started. When it completes successfully, all steps in the postfailover DR plan have been performed as expected.

To track progress, under Disaster Recovery Service, select Jobs. The Jobs table reports the status of each job. Click a record in the table to display the job details.

When the job has completed successfully, all instances impacted by the switchover or failover have been restored and are running on the appliance where they were hosted before.

## Tracking Disaster Recovery Metrics

The Native Disaster Recovery (DR) service provides metric data in a different way compared to other microservices. No DR metrics are exposed through Grafana by default, but they appear after the service setup is complete. A prebuilt dashboard is not provided. If you prefer to display DR metrics in a dashboard you can build one that suits your requirements. For more information, see [Using Grafana](#).

The DR metrics are available from the Grafana Explore module. Using the Prometheus data source, you open the Metrics drop-down and select "dr" to display the list.

DR Metric	Available after...	Description
dr_replication	Native DR service setup	indicates that the DR service is enabled (value=1) and displays target details
dr_job_retention_hours	Native DR service setup	indicates how long DR jobs are stored
dr_max_config	Native DR service setup	indicates the maximum number of DR configurations
dr_peering_status	Native DR service setup	indicates that a peer connection with another appliance is enabled (value=1)
dr_replication_sync_status	Native DR service setup	indicates whether data sync between peered systems is successful (value=0) or not (value=1)
dr_configcount	creating DR configurations	indicates the total number of DR configurations present
dr_instances	creating DR configurations	indicates how many instances are included in each DR configuration
dr_operation_success_count	running DR plan operations	indicates the total number of successful DR operations per DR plan
dr_operation_fail_count	running DR plan operations	indicates the total number of failed DR operations per DR plan

### Note

For existing DR configurations, metrics are not generated immediately after the appliance software upgrade that enables the functionality. They appear when a DR operation has been run, or when a new DR configuration is added. This also applies to a DR setup that was migrated to the Native DR service.