

Oracle Private Cloud Appliance Administrator Guide



F74802-08
July 2024



Oracle Private Cloud Appliance Administrator Guide,
F74802-08
Copyright © 2022, 2024, Oracle and/or its affiliates.

Contents

Preface

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

1 Working in the Service Enclave

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

2 Hardware Administration

Displaying Rack Component Details	2-1
Viewing Appliance Details	2-1
Using the Rack Units List	2-2
Changing Passwords for Hardware Components	2-4
Checking Component Health	2-5
Performing Compute Node Operations	2-6
Provisioning a Compute Node	2-7
Providing Platform Images	2-8
Disabling Compute Node Provisioning	2-9
Locking a Compute Node for Maintenance	2-11
Migrating Instances from a Compute Node	2-14
Configuring the Compute Service for High Availability	2-16
Using Instance and Compute Service High Availability Configuration	2-16

Viewing and Setting Compute Service Configuration	2-18
Compute Service Configuration Commands	2-19
Configuring the Recovery State for a Stopped Instance	2-21
Enabling Strict Fault Domain Enforcement	2-21
Starting, Resetting or Stopping a Compute Node	2-22
Deprovisioning a Compute Node	2-23
Configuring the Active Directory Domain for File Storage	2-25
Reconfiguring the Network Environment	2-27
Editing Routing Information	2-28
Editing Management Node Information	2-28
Editing Data Center Uplink Information	2-30
Updating NTP Server Information	2-31
Editing Administration Network Information	2-32
Updating DNS Information	2-35
Updating Public IP Information	2-36
Creating and Managing Exadata Networks	2-37
Creating an Exadata Network	2-38
Enabling Oracle Exadata Access	2-39
List Exadata Networks	2-40
Get Exadata Network Details	2-41
Disabling Oracle Exadata Access	2-41
Deleting an Exadata Network	2-42
Accessing External Interfaces with Your CA Trust Chain	2-42
Create Certificate Signing Requests	2-43
Uploading Your CA Certificates	2-45

3 Administrator Account Management

Creating a New Administrator Account	3-1
Changing Administrator Credentials	3-2
Managing Administrator Privileges	3-3
Working with Authorization Groups	3-5
Working with Authorization Families	3-8
Changing Administrator Account Preferences	3-10
Deleting an Administrator Account	3-11
Federating with Microsoft Active Directory	3-12
Gathering Required Information from ADFS	3-13
Verifying Identity Provider Self-Signed Certificates	3-13
Managing Identity Providers	3-14
Adding Active Directory as an Identity Provider	3-15
Updating an Identity Provider	3-16
Viewing Identity Provider Details	3-16

Listing Identity Providers	3-17
Deleting an Identity Provider	3-17
Working with Group Mappings for an Identity Provider	3-17
Creating Group Mappings	3-17
Updating a Group Mapping	3-18
Viewing Group Mappings	3-18
Deleting a Group Mapping	3-18
Adding Private Cloud Appliance as a Trusted Relying Party in ADFS	3-19
Providing Federated Users Sign In Information	3-21

4 Tenancy Management

Creating a New Tenancy	4-1
Modifying the Configuration of a Tenancy	4-3
Deleting a Tenancy	4-3

5 Status and Health Monitoring

Using Grafana	5-1
Adding Grafana Users	5-2
Using Grafana Dashboards	5-4
Using Grafana Alerts	5-4
Checking the Health and Status of Hardware and Platform Components	5-6
Viewing and Interpreting Monitoring Data	5-7
Monitoring System Capacity	5-8
Viewing CPU and Memory Usage By Fault Domain	5-9
Viewing Disk Space Usage on the ZFS Storage Appliance	5-9
Accessing System Logs	5-10
Viewing Loki Logs	5-10
Audit Logs	5-12
LBaaS Logs	5-13
Using Oracle Auto Service Request	5-13
Understanding Oracle Auto Service Request	5-14
Oracle Auto Service Request Prerequisites	5-14
Registering Private Cloud Appliance for Oracle Auto Service Request	5-15
Testing Oracle Auto Service Request Configuration	5-16
Unregistering Private Cloud Appliance for Oracle Auto Service Request	5-17
Disabling Oracle Auto Service Request	5-17
Enabling Oracle Auto Service Request	5-17
Viewing Admin Service Health Data	5-18
Compute Node CPU and Memory Utilization Faults	5-21
Storage Utilization Faults	5-23

Hardware Run State Faults	5-24
Health Checker Notification Faults	5-24
Manually Clearing Faults	5-25
Using Support Bundles	5-26
Using the asrInitiateBundle Command	5-26
Using the support-bundles Command	5-26
Uploading Support Bundles to Oracle Support	5-33

6 Backup and Restore

Activating Standard Daily Backup	6-1
Executing a Backup Operation	6-3
Restoring the System from a Backup	6-4

7 Disaster Recovery

Enabling Disaster Recovery on the Appliances	7-1
Collecting System Parameters for Disaster Recovery	7-2
Connecting the Components in the Disaster Recovery Setup	7-2
Setting Up Peering Between the ZFS Storage Appliances	7-3
Setting Up Peering Between the ZFS Storage Appliances Before 302-b892153	7-4
Setting Up Peering Between the ZFS Storage Appliances	7-8
Managing Disaster Recovery Configurations	7-14
Creating a DR Configuration	7-15
Adding Site Mappings to a DR Configuration	7-16
Removing Site Mappings from a DR Configuration	7-18
Adding Instances to a DR Configuration	7-19
Removing Instances from a DR Configuration	7-21
Refreshing a DR Configuration	7-22
Deleting a DR Configuration	7-22

Preface

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

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

Audience

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

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

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

Feedback

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

Conventions

The following text conventions are used in this document:

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

Documentation Accessibility

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

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 Displaying Rack Component Details .

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">• Routing uplink gateway, VLAN, and HSRP group, and spine virtual IP• Management nodes IPs and hostnames• Uplink port speed, count, port FEC, VLAN MTU, and netmask and spine IPs• NTP servers IP addresses• Admin network status• DNS servers IP addresses• Public IP ranges and object storage IP <p>For more information, see Reconfiguring the Network Environment.</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">• Name• Rack unit type• State• Rack elevation <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 Displaying Rack Component Details.</p> <p>For each compute node in the list, you can see additional information:</p> <ul style="list-style-type: none">• Provisioning state• Maintenance lock• Provisioned lock <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 Performing Compute Node Operations.</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"> • Name • Description • Action menu <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 Tenancy Management.</p>
Identity Providers	<p>Read-only list of identity providers and the following information for each:</p> <ul style="list-style-type: none"> • Name • Force Authentication • Encrypt Assertion • Action menu <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 Federating with Microsoft Active Directory.</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"> • Name • IDP Group Name • Admin Group Name • Description • Action menu <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 Federating with Microsoft Active Directory.</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"> • Name • Authorization Group • Default User • Action menu <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 Administrator Account Management.</p>
Jobs	<p>Read-only list of jobs that ran and the following information for each:</p> <ul style="list-style-type: none"> • Object type • Start and end times • Run status - Active, Succeeded, Failed, or Aborted • Action menu <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"> • Job name • Request and job IDs • Start and end times • Command name • Result - Passed, Failed, Not Run, Canceled, or None <p>A Create Upgrade or Patch button, where you can select:</p> <ul style="list-style-type: none"> • Upgrade Request - includes several types of upgrades, such as compute node, host, ILOM, Kubernetes, and platform. • Patch Request - includes several types of patches, such as compute node, host, ILOM, Kubernetes, OCI Images, and platform. <p>For more information, refer to the Oracle Private Cloud Appliance Upgrade Guide and Oracle Private Cloud Appliance Patching Guide.</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 Using Oracle Auto Service Request.</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
changeIlomPassword
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 `changeilompasword ?` command.

```
PCA-ADMIN> changepassword ?
ComputeNode
LeafSwitch
ManagementNode
ManagementSwitch
SpineSwitch
ZFSAppliance

PCA-ADMIN> changeilomPassword ?
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 `changeilompassword` command.

Syntax (entered on a single line):

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

Example:

```
PCA-ADMIN> changeilomPassword
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 `getcomputeIloMHealth` and `getmgmtIloMHealth` 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> getcomputeIloMHealth

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> getmgmtIloMHealth

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

```

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

```

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.

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.
 - 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
```

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
  --                               ----           -
  -----
```


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	
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

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>

```

Creating and Managing Exadata Networks

Oracle Private Cloud Appliance supports direct connectivity to Oracle Exadata clusters.

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

In order to *use* an Exadata 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 Exadata CIDR as destination and the DRG as target.

Note:

If an Oracle Exadata 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 Oracle Exadata network is in the **available** or **terminated** state.

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

Creating an Exadata Network

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

 **Note:**

When connecting more than one Oracle Exadata system to the Private Cloud Appliance rack, you must be careful that the *configured IP address ranges do not overlap*. This cannot be detected or prevented any other way.

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	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 Exadata network to the customer's data center servers. advertiseNetwork=true is only available for dynamic routing configurations.

 **Note:**

When the Exadata network with 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 Oracle Exadata 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
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

Using the Service Web UI

1. Determine the Exadata network parameters listed in the table above.
2. In the Dashboard, click the Rack Units quick action tile.
3. In the PCA Config navigation menu on the Rack Units page, click Exadata Networks.
4. In the top-right corner above the table, click Create Exadata Network.
5. Fill out the Exadata 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 Exadata Networks table and its Lifecycle State changes to Available when the configuration has been applied successfully.
7. Next, add a subnet to the Exadata network. See [Enabling Oracle Exadata Access](#).

Using the Service CLI

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

```
PCA-ADMIN> exaDataCreateNetwork 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"
Command: exaDataCreateNetwork 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" gateway="10.nn.nn.27"
speed="10"
Status: Success
Time: 2021-11-22 06:10:05,260 UTC
Data: ocid1.exadata.unique_id
```

3. Next, add a subnet to the Exadata network. See [Enabling Oracle Exadata Access](#).

Enabling Oracle Exadata Access

Enabling access from a subnet to the Exadata network must be done through the Service CLI.

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

Using the Service CLI

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

```
PCA-ADMIN> exaDataEnableAccess exadataNetworkId=ocidl.exadata.unique_id \
subnetId=ocidl.subnet.unique_id
Command: exaDataEnableAccess exadataNetworkId=ocidl.exadata.unique_id \
subnetId=ocidl.subnet.unique_id
Status: Success
Time: 2021-11-17 18:56:45,251 UTC
Data:
  id
  --
  ocidl.vcn.unique_id
```

3. If you are using a secondary NIC to access the Exadata network, you must add a route to the Exadata CIDR address range for interface `eth1` (the secondary NIC).

```
PCA-ADMIN> route add Exadata-CIDR-address-range gw Addr-default-router eth1
```

For example, if the Exadata address range is `192.168.0.0/24` and the gateway is `192.168.0.1`:

```
PCA-ADMIN> route add 192.168.0.0/24 gw 192.168.0.1 eth1
```

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		eth1

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

List Exadata 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 Exadata Networks. The table contains all configured Exadata networks.

Using the Service CLI

1. Use the `exaDataListNetwork` command to display configured Exadata networks, including their OCIDs.

```
PCA-ADMIN> exaDataListNetwork
Command: exaDataListNetwork
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 Exadata Network Details

Using the Service Web UI

1. Navigate to the Exadata Network page.
2. In the overview table, click the name (OCID) of the network for which you want to display details.

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

```

PCA-ADMIN> exaDataGetNetwork exadataNetworkId=ocid1.exadata.unique_id
Command: exaDataGetNetwork exadataNetworkId=ocid1.exadata.unique_id
Status: Success
Time: 2021-11-22 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 2 = 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

```

Disabling Oracle Exadata Access

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

Subnets that have been granted access, appear in the Exadata 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 Exadata network you want to disable, using the `exaDataListNetwork` command.
2. Get the OCID of the subnet ID for the Exadata network using the `exaDataGetNetwork` command.
3. Disable access to a configured Exadata network.

```
PCA-ADMIN> exaDataDisableAccess exadataNetworkId=ocidl.exadata.unique_id \  
subnetId=ocidl.subnet.unique_id  
Command: exaDataDisableAccess exadataNetworkId=ocidl.exadata.unique_id \  
subnetId=ocidl.subnet.unique_id  
Status: Success  
Time: 2021-11-02 11:29:49,873 UTC  
PCA-ADMIN> exaDatadisableAccess exadataNetworkId=ocidl.exadata.unique_id \  
subnetId=ocidl.subnet.unique_id \  
Command: exaDatadisableAccess exadataNetworkId=ocidl.exadata.unique_id \  
subnetId=ocidl.subnet.unique_id \  
Status: Success  
Time: 2021-12-15 11:26:40,344 UTC  
Data:  
  id  
  --  
  ocidl.vcn.unique_id \  
PCA-ADMIN>
```

Deleting an Exadata Network

Using the Service Web UI

1. Make sure that, for the Exadata network you intend to delete, access has been disabled first.
2. Navigate to the Exadata Network page.
3. Choose one of these options to delete the Exadata 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 Exadata network detail page, then click the Delete button in the top-right corner.

Using the Service CLI

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

```
PCA-ADMIN> exaDatadeleteNetwork exadataNetworkId=ocidl.exadata.unique_id  
Command: exaDatadeleteNetwork exadataNetworkId=ocidl.exadata.unique_id  
Status: Success  
Time: 2021-11-16 05:59:54,177 UTC
```

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>
- Regular uplink
 - auth.<domain_name>
 - autoscaling.<domain_name>
 - console.<domain_name>
 - containerengine.<domain_name>
 - dns.<domain_name>
 - filestorage.<domain_name>
 - iaas.<domain_name>
 - identity.<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.

 **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
  --                               ----
  587fc90d-3312-41d9-8be3-1ce21b8d9b41  MonitorGroup
  c18cc6af-4ef8-4b1c-b85d-ee3b065f503e  DrAdminGroup
  8f03faf2-c321-4455-af21-75cbffc269ef  AdminGroup
  5ac65f5d-1f8c-42ea-alde-95a1941f009f  Day0ConfigGroup
  365ece7b-0a09-4a04-853c-7a0f6c4789f0  InternalGroup
  7da8be67-758c-4cd6-8255-e9d2900c788e  SuperAdminGroup
```

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-alde-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

```



```

changeIllumPassword: 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
alphabetizeMode	YES, NO	Enable this setting to display any managed object's attributes in alphabetical order. The default setting is "No".
attributeDisplay	DISPLAYNAME, ATTRIBUTENAME	Use this setting to control whether the name of each object's attribute is displayed. The default setting is "displayName".
endLineCharsDisplayValue	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".
outputMode	VERBOSE, SPARSE, XML	Specify the CLI output format. The default setting is "Sparse".
wsCallMode	SYNCHRONOUS, ASYNCHRONOUS	Use this setting to determine whether the CLI output from a command is invoked synchronously or asynchronously. The default setting is "Asynchronous".

Setting	Options	Description
wsTimeoutInSec nds	<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-fbald85e049b 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-fbald85e049b 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://<yourservename>/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://<yourservename>/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 `pcam01`.
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/python3.6/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://<yourservername>/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.

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

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`, 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 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 [Accessing System Logs](#). 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 Container Engine for Kubernetes (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-a676-01b3d73cffee
```

Accessing System Logs

Logs are collected from all over the system and aggregated in Loki. The log data can be queried, filtered, and displayed using Grafana.

Viewing 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 [Viewing 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 [Viewing 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 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.

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 username=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 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 register 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 5-1](#) are reported through ASR.

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 5-1 Admin Service Fault Detection Configuration Summary

Fault Type	Fault Detection Frequency (seconds)	Fault Detection Delay (seconds)	Data Source	Method of Detection
Compute Node CPU and Memory Utilization Faults	60	< 20	Admin calls ComputeNode service	Faults are raised by fault task based on the compute node object attributes stored in the database.
Storage Utilization Faults	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.
Hardware Run State Faults	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.

Table 5-1 (Cont.) Admin Service Fault Detection Configuration Summary

Fault Type	Fault Detection Frequency (seconds)	Fault Detection Delay (seconds)	Data Source	Method of Detection
Health Checker Notification Faults	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 PCA X9-2 Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.1) for a list of Private Cloud Appliance X9-2 events that are actionable by ASR.
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 PCA X9-2 Appliance: Automatic Service Request (ASR) Event Coverage (Doc ID 2833567.1) for a list of Private Cloud Appliance X9-2 events that are actionable by ASR.

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
RackUnitIloRunStateFaultStatusFault (pcacn005) Cleared Critical
f7d22180-aeae-4159-b5c8-5e55a7906a78
RackUnitIloRunStateFaultStatusFault (pcacn004) Cleared Critical
a4fef907-8e54-4750-9fac-6829fbade90d
ComputeNodeCpuFaultStatusFault (pcacn006) Cleared Minor
f8d93384-da30-43cd-9396-6e6671d240e2
RackUnitIloRunStateFaultStatusFault (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
eecd5ef2-4a71-4137-be96-54c028212d2f
ComputeNodeMemoryFaultStatusFault (pcacn004) Cleared Minor
cf68d2ee-e483-e573-b46e-c31bc8e968 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
RackUnitIloRunStateFaultStatusFault (pcacn007) Cleared Critical
b4c66a7c-def3-42c2-8842-d4763afc5184
RackUnitIloRunStateFaultStatusFault (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
RackUnitIloRunStateFaultStatusFault (pcacn001) Cleared Critical
2b5ece1c-50fc-436a-81b3-da0c5b418fe3
RackUnitIloRunStateFaultStatusFault (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
PcaSystemReSyncHwStatusFault (null) Cleared Warning
28c5ba93-6b4e-42f3-ad61-90734b46bf30 SPENV-8000-RU--ORACLE SERVER
E5-2L Cleared Critical
3d932188-0120-489f-a512-1a244ec01e49
RackUnitIloRunStateFaultStatusFault (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
RackUnitIloRunStateFaultStatusFault (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
zfssa_pool_total zfssa_pool_free	Pool usage percentage is calculated using the following formula for each pool: $\frac{(zfssa_pool_total - zfssa_pool_free)}{zfssa_pool_total}$	If the pool usage percentage is above a pre-configured value, a major fault is raised. The default value is 80 percent.
zfssa_pool_status	The zfssa_pool_status metric can have the following values: <ul style="list-style-type: none"> • 0 - exported • 1 - degraded • 2 - online • -1 - offline • -2 - faulted • -3 - unavailable • -4 - removed 	A major fault is raised for any pool/ zfssa_node 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
SERVICE REQUIRED	Major	Active
<i>other</i>	Not applicable	Cleared

Health Checker Notification Faults

Health Checker faults are raised from notifications from the ZFSSA and Network Health Checker components. For every notification it receives, the Admin service will raise a fault.

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 ZFSSA 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

1. Using SSH, log into the management node VIP as `admin`.

```
# ssh -l admin 100.96.2.32 -p 30006
```

2. 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>
```

3. 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 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](#).

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. This topic describes the available modes. The following is the recommended way to use this command:

1. Start data collection by specifying `triage` mode to understand the preliminary status of the Private Cloud Appliance.
2. If `NOT_HEALTHY` appears in the `triage` mode results, then do one of the following:
 - Use `time_slice` mode to collect data by time slots. These results can be further narrowed by specifying pod name, job, and `k8s_app` label.
 - Use `smart` mode to query data from specific health-checkers.

The `support-bundles` command requires a mode (`-m`) option. Some modes have additional options.

The following table lists the options that are common to all modes of the `support-bundles` command.

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>		

For most modes, the `support-bundles` command produces a single archive file. The output archive file is named `[SR_number]pca-support-bundle.current-time.tgz`. The `SR_number` is used if you provided the `-sr` option. If you are creating the support bundle for a service request, you should specify the `SR_number`.

For `native` mode, the `support-bundles` command produces a directory of archive files.

The archive files are stored in `/nfs/shared_storage/support_bundles/` on the management node.

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
llapi (systemd:llapi): 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
<code>header.json</code>	Time stamp and command line to generate this bundle.
<code>compute_node_info.json</code>	Pods running in the compute node.
<code>management_node_info.json</code>	Pods running in the management node.
<code>rack_info.json</code>	Rack installation time and build version.
<code>loki_search_results.log.n</code>	Chunk files in json.

Time Slice Mode

In time slice mode, data is collected by specifying start and end timestamps.

If you do not specify either the `-j` or `--all` option, then data is collected from all health checker jobs.

You can narrow the data collection by specifying any of the following:

- Loki job label
- Loki `k8s_app` label
- Pod name

```
[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.

The time slice mode of the `support-bundles` command has the following options in addition to the mode and service request number options listed at the beginning of this topic.

- 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.

Option	Description	Required
-j <i>job_name</i>	Loki job name. Default value: .+checker	no
--job_name <i>job_name</i>	See Label List Query below.	
--all	Queries all job names except for jobs known for too much logging, such as audit, kubernetes-audit, and vault-audit and k8s_app label pcacoredns.	no
--k8s_app <i>label</i>	The k8s_app label value to query within the k8s-stdout-logs job. See Label List Query below.	no
-l <i>level</i>	Message level	no
--levelname <i>level</i>		
-s <i>timestamp</i>	Start date in format yyyy-mmm-ddTHH:mm:ss	yes
--start_date <i>timestamp</i>	The minimum argument is yyyy-mmm-dd	
-e <i>timestamp</i>	End date in format yyyy-mmm-ddTHH:mm:ss	yes
--end_date <i>timestamp</i>	The minimum argument is yyyy-mmm-dd	
--pod_name <i>pod_name</i>	The pod name (such as kube or network-checker) to filter output based on the pod. Only the starting letters are necessary.	no

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', 'kubelet', '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-llapi', '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-brn', '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.oracledx.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.n	Chunk files in json.

Smart Mode

In smart mode, health checkers are queried for recent NOT_HEALTHY status. By default, two days of logs are collected. If you need more than two days of logs, specify the `--force` option. Use the `-hc` option to specify a health checker.

```
[root@pcamn01 ~]# support-bundles -m smart
```

See more examples below.

The smart mode of the `support-bundles` command has the following options in addition to the mode and service request number options listed at the beginning of this topic.

If only the start date or only the end date is given, the time is calculated and queried two days prior to the given end date or two days after the given start date. If only the start date is given and under the two day time range, the default most recent unhealthy time is used.

Option	Description	Required
-hc <i>health_checker_name</i>	Loki health checker name.	no
--health_checker <i>health_checker_name</i>	See the health checker log files table below.	
--errors_only	Level name filtering takes place only on Error, Critical, and Severe.	no
--force	Force the start date to override the two-day time range limit.	no
-s <i>timestamp</i>	Start date in format <i>yyyy-mm-ddTHH:mm:ss</i>	no
--start_date <i>timestamp</i>	The minimum argument is <i>yyyy-mm-dd</i> Default value: End date minus 2 days	
-e <i>timestamp</i>	End date in format <i>yyyy-mm-ddTHH:mm:ss</i>	no
--end_date <i>timestamp</i>	The minimum argument is <i>yyyy-mm-dd</i> Default value: Most recent unhealthy time	

The following table lists the log files for each health checker.

Health Checker	Supporting Log Files
L0_hw_health-checker	<ul style="list-style-type: none"> pca.log, pca.health.log, pca.l1api.log, pacemaker.log pca-platform-l1api pca-healthchecker pacemaker pca-platform-l0
cert-checker	No logs - only certificate and expiry date (from the checker)
etcd-checker	<ul style="list-style-type: none"> etcd-container.log
flannel-checker	k8s-stdout-logs: filter by pod (flannel), node, and container
kubernetes-checker	k8s-stdout-logs: filter by pod (kube-apiserver), node, and container
l0-cluster-services-checker	<ul style="list-style-type: none"> pacemaker.log, corosync.log corosync pcsd
mysql-cluster-checker	<ul style="list-style-type: none"> mysqld
network-checker	<ul style="list-style-type: none"> HMS
registry-checker	messages (registry itself does not produce logs)
vault-checker	<ul style="list-style-type: none"> hc-vault-audit.log hc-vault-audit.log
zfssa-checker	<ul style="list-style-type: none"> zfssa-checker zfssa-log-exporter (log = alert audit pcalog)

Examples:

No -hc. Query unhealthy data from all health checkers.

```
[root@pcamn01 ~]# support-bundles -m smart -sr 3-xxxxxxxxxxx
```

Use -hc to specify one health checker.

```
[root@pcamn01 ~]# support-bundles -m smart -sr 3-xxxxxxxxxxx -hc network-checker
```


Timestamps with `--force`.

```
[root@pcamn01 ~]# support-bundles -m smart -sr 3-xxxxxxxxxxx -s "2022-01-11/00:00:00" -e
"2022-01-15/23:59:59" --force
```

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.n	Chunk files in json.

Native Mode

Unlike other support bundle modes, the native bundle command returns immediately and the bundle collection runs in the background. Native bundles might take hours to collect. Collection progress information is provided in the `native_collection.log` in the bundle directory.

Also unlike other support bundle modes, the output of native bundles is not a single archive file. Instead, a bundle directory is created in the `/nfs/shared_storage/support_bundles/` area on the management node. The directory contains the `native_collection.log` file and a number of `tar.gz` files.

```
[root@pcamn01 ~]# support-bundles -m native -t bundle_type [-c component_name] [-sr
SR_number]
```

The native mode of the `support-bundles` command has the following options in addition to the mode and service request number options listed at the beginning of this topic.

Option	Description	Required
<code>-t bundle_type</code> <code>--type bundle_type</code>	Bundle type: <code>sosreport</code> or <code>zfs-bundle</code>	yes
<code>-c component_name</code> <code>--component component_name</code>	Component name This option only applies to type <code>sosreport</code> .	no

ZFS Bundle

When type is `zfs-bundle`, a ZFS support bundle collection starts on both ZFS nodes and downloads the new ZFS support bundles into the bundle directory.

```
[root@pcamn01 ~]# support-bundles -m native -t zfs-bundle
2021-11-16T22:49:30.982 - support_bundles - INFO - Starting Support Bundles
2021-11-16T22:49:31.037 - support_bundles - INFO - Locating filter-logs Pod
2021-11-16T22:49:31.064 - support_bundles - INFO - Executing command - ['python3',
'/usr/lib/python3.6/site-packages/filter_logs/native.py', '-t', 'zfs-bundle']
2021-11-16T22:49:31.287 - support_bundles - INFO - LAUNCHING COMMAND: ['python3',
'/usr/lib/python3.6/site-packages/filter_logs/native_app.py', '-t', 'zfs-bundle', '--
target_directory', '/support_bundles/zfs-bundle_20211116T224931267']
ZFS native bundle collection running to /nfs/shared_storage/support_bundles/zfs-
bundle_20211116T224931267
Monitor /nfs/shared_storage/support_bundles/zfs-bundle_20211116T224931267/
native_collection.log for progress.
```

```
2021-11-16T22:49:31.287 - support_bundles - INFO - Finished running Support Bundles
```

SOS Report Bundle

When `type` is `sosreport`, the `component_name` is a management node or compute node. If `component_name` is not specified, the report is collected from all management and compute nodes.

```
[root@pcamn01 ~]# support-bundles -m native -t sosreport -c pcacn003 -sr SR_number
```

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` 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 split 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.

 **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.

6

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.

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
```

- 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
```

- 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.

- 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      0 0 * * *      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.wtczj8r3z17wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nrmrcuyo4vathc8rj
brs-job-1698401412-backup   layer0,zfs,vault,mysql,loki,sauron
ocid1.brs-
job.PCA3X62D9C1.mypca.089a7b8cuqz0cam2r7xexo4i4p7j7ia7sql9f8w89dyp9q3y10dnbaac6mu
brs-job-1698401607-backup   layer0

PCA-ADMIN> getBackupJob backupJobId=ocid1.brs-
job.PCA3X62D9C1.mypca.wtczj8r3z17wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nrmrcuyo4vathc8rj
Data:
  Type = BackupJob
  Job Id = ocid1.brs-
job.PCA3X62D9C1.mypca.wtczj8r3z17wjftxbkskaj6j60x6xqai3lpjoxp7ywi7nrmrcuyo4vathc8rj
```



```
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/ocidl.brs-job.unique_ID/backup_timestamp/
```

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.

7

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:

- Computer 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

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 in 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-xxxxxf69410a
Status: Success
Time: 2023-08-16 12:31:55,828 UTC
Data:
  Type = update_service
  Job Id = 1d6826ac-04db-49f9-aa27-xxxxxf69410a
  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-drl.rack1.example.com remoteHostPerf=sn02-drl.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 private IPs upon switchover or failover if the subnet for the IP address is tagged 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 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-affa-
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

```

6. 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

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 replica system respectively.
CLI commands:

```
- oci iam compartment list
- oci network subnet list --compartment-id
  "ocidl.compartment.....uniqueID"
- oci network nsg list --compartment-id "ocidl.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 | networksecuritygroup]
sourceId=<source_object_OCID>
targetId=<target_object_OCID>
```

Examples:

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=compartment \
sourceId="ocidl.compartment.....<region1>...uniqueID" \
targetId="ocidl.compartment.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
objType=compartment sourceId="ocidl.compartment.....<region1>...uniqueID"
targetId="ocidl.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="ocidl.subnet.....<region1>...uniqueID" \
targetId="ocidl.subnet.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
objType=subnet sourceId="ocidl.subnet.....<region1>...uniqueID"
targetId="ocidl.subnet.....<region2>...uniqueID"
Status: Success
Time: 2021-08-17 09:07:24,957 UTC
Data:
  dlbf2cf2-d8c7-4271-b8b6-cdf757648175
```

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=networksecuritygroup \
sourceId="ocidl.nsg.....<region1>...uniqueID" \
targetId="ocidl.nsg.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
objType=networksecuritygroup sourceId="ocidl.nsg.....<region1>...uniqueID"
targetId="ocidl.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
```

- 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]
```

4. Repeat the `drAddComputeInstance` command with the OCIDs of all the compute instances that you want to add to the DR configuration.
5. 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
```

6. 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=ocidl.instance.....<region1>...instance1_uniqueID
Command: drGetComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
instanceId=ocidl.instance.....<region1>...instance1_uniqueID
Status: Success
Time: 2021-08-17 08:34:42,413 UTC
Data:
  Type = ComputeInstance
  Compartment Id = ocidl.compartment.....uniqueID
  Boot Volume Id = ocidl.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

1. 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=ocidl.instance.....<region1>...instance3_uniqueID
Command: drRemoveComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
instanceId=ocidl.instance.....<region1>...instance3_uniqueID
Status: Success
Time: 2021-08-17 08:45:59,718 UTC
Data:
  Message = Successfully started job to remove instance
ocidl.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]
```