Oracle® Private Cloud Appliance

Administrator Guide for Release 3.0.1

ORACLE

Oracle Legal Notices

Table of Contents

Preface

This publication is part of the customer documentation set for Oracle Private Cloud Appliance Release 3.0.1. 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 Windows operating systems.

Feedback

Provide feedback about this documentation at:

http://www.oracle.com/goto/docfeedback

Conventions

The following text conventions are used in this document:

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

Document Revision

Document generated on: 2022-01-31 (revision: 1388)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

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

Access to Oracle Support for Accessibility

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

Diversity and Inclusion

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

Chapter 1 Working in the Service Enclave

Table of Contents

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.

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.

1.1 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.

1.1.1 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 Oracle 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 Oracle Private Cloud Appliance dashboard displays with quick action tiles.

   > **Note**
   >
   > If this is the first log in after an Oracle Private Cloud Appliance installation, the dashboard displays the ASR Phone Home page so you can register your system with Oracle Support.
   >
   > For more information, see Section 5.5.3, "Registering Oracle Private Cloud Appliance for ASR".

1.1.2 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 Section 5.1, "Using Grafana".

In the top bar of the dashboard you can locate the realm and the system and domain names for your Oracle 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 Oracle 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 Section 2.1, "Displaying Rack Component Details".
Network Environment	Read-only network configuration information and an Edit button that opens a Configure Network Params wizard where you can modify: • 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

Component or Resource	Information Provided
	For more information, see Section 2.4, "Reconfiguring the Network Environment".
Rack Units	Read-only list of all hardware components installed in the rack and detected by the appliance software and the following information for each: • Name • Rack unit type • State • Rack elevation 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. For more information, see Section 2.1, "Displaying Rack Component Details". For each compute node in the list, you can see additional information: • Provisioning state • Maintenance lock • Provisioned lock 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 Section 2.2, "Performing Compute Node Operations".
Tenancies	Read-only list of all tenancies in the system and the following information for each: • Name • Description

Component or Resource	Information Provided
	• Action menu Contains options to view a tenancy's details page, edit a tenancy's description, or delete a tenancy. You can also edit or delete a tenancy from its details page. A Create Tenancy button. For more information, see Chapter 4, *Tenancy Management*.
Identity Providers	Read-only list of identity providers and the following information for each: • Name • Force Authentication • Encrypt Assertion • Action menu Contains options to view an identity provider's details page and edit or delete the identity provider. You can also edit or delete an identity provider from its details page. A Create Identity Provider button. For more information, see Section 3.4, "Federating with Microsoft Active Directory".
IDP Group Mappings	Read-only list of IDP group mappings in the system and the following information for each: • Name • IDP Group Name • Admin Group Name • Description • Action menu Contains options to view read-only information on an IDP group mapping details page. MORE...

Component or Resource	Information Provided
	A Create Group Mapping button. For more information, see Section 3.4, "Federating with Microsoft Active Directory".

Users	Read only list of users in the system and the following information for each:
	
	• Name
	
	• Authorization Group
	
	• Default User
	
	• Action menu
	
	  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.
	
	A Create User button.
	
	For more information, see Chapter 3, *Administrator Account Management*.

Jobs	Read-only list of jobs that ran and the following information for each: • Object type • Start and end times • Run status - Active, Succeeded, Failed, or Aborted • Action menu Contains an option to view read-only information on a job's details page, which includes the user account that the job ran from.
Upgrade & Patching	Read-only list of upgrade and patching jobs that ran and the following information for each: • Job name • Request and job IDs • Start and end times

Component or Resource	Information Provided
	• Command name • Result - Passed, Failed, Not Run, Cancelled, or None A Create Upgrade or Patch button, where you can select: • 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. For more information, see Chapter 7, *System Upgrade*.
ASR Phone Home	Read-only auto service request information and a Register button where you can register your Oracle Private Cloud Appliance. For more information, see Section 5.5, "Using Auto Service Requests".

1.2 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 you establish an SSH connection to TCP port 30006 on one of the management nodes. Log in with an authorized administrator account. After successful authentication, the `PCA-ADMIN>` prompt is displayed.

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

You are now in an interactive, closed shell environment where you perform administrative operations by entering commands at this prompt. The command syntax and completion functions are described below. To terminate your CLI session, enter the `exit` command.

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)>
    ----------------------------------
    abort:  Job
    asrClientDisable:  ASRPhonehome
    asrClientEnable:  ASRPhonehome
    asrClientRegister:  ASRPhonehome
    asrClientUnregister:  ASRPhonehome
    changePassword:  User
    createAdminAccount:
    createUserInGroup:  User
    [...]
    upgradePlatform:  UpgradeRequest
    upgradeSwitch:  UpgradeRequest
    upgradeVault:  UpgradeRequest
    upgradeZfssa:  UpgradeRequest

PCA-ADMIN> showcustomcmds ?
                            ASRPhonehome
                            ComputeNode
                            DrConfig
                            Event
                            ExadataNetwork
                            FaultManager
                            Job
                            NetworkConfig
                            PcaSystem
                            PurgeManager
                            UpgradeJob
                            UpgradeJobList
                            UpgradeRequest
                            User
PCA-ADMIN> showcustomcmds ComputeNode
    provisioningLock
    provisioningUnlock
    maintenanceLock
    maintenanceUnlock
    provision
    deprovision
    migrateVm
    reset
    start
    stop
```

Chapter 2 Hardware Administration

Table of Contents

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.

2.1 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.

2.1.1 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 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 = ad1
  Realm = 1742XC3024
  Region = oraclepca
  ASR Reminder = true
  Name = pca
  Work State = Normal
  FaultIds 1 = id:55f8de1e-ab25-4fc6-b6f4-a9ddd283605b  type:Fault  name:PcaSystemInitialHwDiscoveryStatusS
  FaultIds 2 = id:5c532489-6dad-45e1-a065-6c7649514ce1  type:Fault  name:PcaSystemReSyncHwStatusStatusFault
```

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=my
Command: edit PcaSystem name=myPca description="My Private Cloud" domainName=my.example.com systemName=mycl
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
```

2.1.2 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 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 Section 2.2, "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
```

2.2 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.

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.

2.2.1 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
```

2.2.2 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
```

2.2.3 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.

Using the Service Web UI

1. Make sure that provisioning has been disabled on the compute node.

   See Section 2.2.2, "Disabling Compute Node Provisioning".

2. In the navigation menu, click Rack Units.

3. In the Rack Units table, click the host name of the compute node that requires maintenance.

   The compute node detail page appears.

4. 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.

5. 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. Make sure that provisioning has been disabled on the compute node.

    See Section 2.2.2, "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
```

2.2.4 Migrating Instances from a Compute Node

Several compute node operations can only be performed if there are no running compute instances. As an administrator you can migrate all running instances away from a compute node, for example to perform maintenance.

Using the Service Web UI

1. Make sure that provisioning has been disabled on the compute node.

   See Section 2.2.2, "Disabling Compute Node Provisioning".

2. In the navigation menu, click Rack Units.

3. 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.

4. In the top-right corner of the page, click Controls and select the Migrate All Vms command.

   The system migrates running instances to other compute nodes.

Using the Service CLI

1. Display the list of compute nodes.

   Copy the ID of the compute node you want to migrate running compute instances away from.

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

2. Make sure that provisioning has been disabled on the compute node.

   See Section 2.2.2, "Disabling Compute Node Provisioning".

3. Migrate all running compute instances off the compute node with this command:

```
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
[...]
  Done = true
  Name = MODIFY_TYPE
  Run State = Succeeded
```

After successful completion, all running compute instances have been migrated away from the current compute node.

2.2.5 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 Section 2.2.3, "Locking a Compute Node for Maintenance".

2. In the navigation menu, click Rack Units.

3. In the Rack Units table, locate the compute node you want to start, reset or stop.

4. Click the Action menu (three vertical dots) and select the appropriate action: Start, Reset, or Stop.

5. When the confirmation window appears, click the appropriate action button to proceed.

    A pop-up window appears for a few seconds to confirm that the compute node is starting, stopping, or restarting.

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

Using the Service CLI

1. Display the list of compute nodes.

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

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

2. Make sure that the compute node is locked for maintenance.

    See Section 2.2.3, "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.

2.2.6 Deprovisioning a Compute Node

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

Using the Service Web UI

1. In the navigation menu, click Rack Units.

2. In the Rack Units table, click the host name of the compute node you want to deprovision.

   The compute node detail page appears.

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

   When the confirmation window appears, click Lock to proceed.

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

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

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

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

   When the confirmation window appears, click Deprovision to proceed.

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

Using the Service CLI

1. Display the list of compute nodes.

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

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

2. Set a provisioning lock on the compute node.

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

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

```
PCA-ADMIN> show Job id=ed4a4646-6d73-41f9-9cb0-73ea35e0d766
[...]
  Done = true
  Name = MODIFY_TYPE
  Run State = Succeeded
```

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

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

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

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

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

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

5. Deprovision the compute node with this command:

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

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

```
PCA-ADMIN> show Job id=9868fdac-ddb6-4260-9ce1-c018cf2ddc8d
[...]
```

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

6. Confirm that the compute node has been deprovisioned.

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

2.3 Configuring the Active Directory Domain for File Storage

The file storage service in Oracle Private Cloud Appliance enables users of 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 Windows instances and Oracle 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.

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
```

2.4 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 Oracle Private Cloud Appliance and can potentially disrupt system operations.

2.4.1 Editing Routing Information

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

2.4.2 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>
```

2.4.3 Editing Data Center Uplink Information

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

> **Caution**
>
> Reconfiguring the Oracle 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>
```

2.4.4 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:

   - 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>
```

2.4.5 Editing Administration Network Information

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

Using the Service Web UI

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.

   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 administration network parameters:

   - Admin Network

   - Admin Port Speed

   - Admin Port Count

   - Admin Port MTU

   - Admin Port FEC

   - Admin VLAN

   - Admin Gateway IP

   - Admin Netmask

   - Admin CIDR

   - Admin Spine1 IP

   - Admin Spine2 IP

   - Admin Spine VIP

   ```
   PCA-ADMIN> edit NetworkConfig adminPortSpeed=25
   Command: edit NetworkConfig adminPortSpeed=25
   Time: 2021-11-24 11:01:00,605 UTC
   JobId: 62f8137f-772a-4a52-98f4-1cfv9c1f24te
   PCA-ADMIN>
   ```

2.4.6 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>
```

2.4.7 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>
```

2.5 Creating and Managing Exadata Networks

Oracle Private Cloud Appliance supports direct connectivity to Exadata clusters. 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, and a route rule in the enabled subnet that has the Exadata

CIDR as a destination and the DRG as target. For more information about Exadata Integration, see the *Network Infrastructure* section of *Hardware Overview*.

This section describes creating and managing Exadata networks from the Service CLI. Before you can create an Exadata network, you must physically connect your Oracle Private Cloud Appliance to an Oracle Exadata rack. For instructions, see the *Optional Connection to Exadata* section in the *Configuring the Oracle Private Cloud Appliance*.

in order to USE an Exadata network you need a DRG and a route rule in the enabled subnet that has the exadata CIDR as destination and the DRG as Target.

2.5.1 Creating an Exadata Network

Using the Service CLI

1. Determine the parameters for the Exadata network.

Table 2.1 Parameters Required for Connecting to Exadata

Parameter	Example Value	Description
vlan	3062	Choose a VLAN from 2 to 3899 that is not in use by the uplink VLAN or other Exadata VLANs. (VLAN 3900 to 3967, and VLAN 3968 to 4095 are reserved).
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.
cidr	10.*nn*.*nn*.0/24	Choose a valid CIDR range that is within the CIDR range of the Exadata.
ports	7/1,7/2	Valid ports are '7/1','7/2','7/3','7/4','8/1','8/2','8/3','8/4', '9/1','9/2','9/3','9/4','10/1','10/2','10/3','10/4'.
advertiseNetwork	True	True or False - enables or disables the visibility of the Exadata network to the customer's datacenter servers. `advertiseNetwork=true` is only available for dynamic routing configurations.

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"
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 Section 2.5.2, "Enabling Exadata Access".

2.5.2 Enabling Exadata Access

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=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Command: exaDataEnableAccess exadataNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Status: Success
Time: 2021-11-17 18:56:45,251 UTC
Data: id -- ocid1.vcn.unique_id
```

2.5.3 List Exadata Networks

Use the exaDataListNetwork command to view Exadata networks, including their OCIDs.

Using the Service CLI

1. Display configured Exadata networks.

```
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
```

2.5.4 Get Exadata Network Details

Use the exaDataGetNetwork command to find details about a specific Exadata network, including the state of the network, subnet and VCN IDs.

Using the Service CLI

1. Get the OCID of the Exadata network for which you want details, using the exaDataListNetwork command.

2. Display Exadata network details.

```
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
  Lifecycle State = AVAILABLE
```

2.5.5 Disabling Exadata Access

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=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id
Command: exaDataDisableAccess exadataNetworkId=ocid1.exadata.unique_id  \
subnetId=ocid1.subnet.unique_id
Status: Success
Time: 2021-11-02 11:29:49,873 UTC
PCA-ADMIN> exaDatadisableAccess exadataNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id \
Command: exaDatadisableAccess exadataNetworkId=ocid1.exadata.unique_id \
subnetId=ocid1.subnet.unique_id \
Status: Success
Time: 2021-12-15 11:26:40,344 UTC
Data:
  id
  --
  ocid1.vcn.unique_id \
PCA-ADMIN>
```

2.5.6 Deleting an Exadata Network

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=ocid1.exadata.unique_id
Command: exaDatadeleteNetwork exadataNetworkId=ocid1.exadata.unique_id
Status: Success
Time: 2021-11-16 05:59:54,177 UTC
```

Chapter 3 Administrator Account Management

Table of Contents

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*.

3.1 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-a1de-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=36
Command: createUserInGroup name=testadmin password=***** confirmPassword=***** authGroup=365ece7b-0a09-4a04
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:
```

3.2 Modifying an Administrator Account

This section describes which administrator account settings can be changed, and how you can change them.

3.2.1 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
```

3.2.2 Managing Administrator Privileges

The privileges associated with an administrator account depend on its membership of an authorization group. More information can be found in the *Oracle Private Cloud Appliance Concepts Guide*: refer to the section *Administrator Access* in the chapter *Appliance Administration Overview*. An authorization group is specified during account creation, but the Service CLI allows you to change which authorization groups an administrator belongs to.

Using the Service CLI

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

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

PCA-ADMIN> list AuthorizationGroup
Command: list AuthorizationGroup
Status: Success
Time: 2021-08-25 08:38:58,632 UTC
Data:
  id                                    name
  --                                    ----
  587fc90d-3312-41d9-8be3-1ce21b8d9b41  MonitorGroup
  c18cc6af-4ef8-4b1c-b85d-ee3b065f503e  DrAdminGroup
```

```
    8f03faf2-c321-4455-af21-75cbffc269ef    AdminGroup
    5ac65f5d-1f8c-42ea-a1de-95a1941f009f    Day0ConfigGroup
    365ece7b-0a09-4a04-853c-7a0f6c4789f0    InternalGroup
    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-3312-41d9-8be
Command: add User id=682ebc19-8493-4e9a-817c-148acea4b1d4 to AuthorizationGroup id=587fc90d-3312-41d9-8be3-
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
```

3.2.3 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".
wsTimeoutInSeconds	<value>	When the CLI is set to "Synchronous" call mode, use this setting to determine how many seconds the CLI waits for a job returned by an operation to complete.

Using the Service CLI

1. Display your current account preferences.

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

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

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

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

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

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

3.3 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
```

3.4 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 Oracle 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 *Oracle Private Cloud Appliance Concepts Guide*.

You can federate multiple Active Directory (AD) accounts with Oracle 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 Oracle 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 Oracle Private Cloud Appliance.

- Created users in Active Directory who will sign into the Oracle Private Cloud Appliance Service Web UI.

> **Note**
>
> Consider naming Active Directory groups that you intend to map to Oracle 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.

3.4.1 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 Oracle Private Cloud Appliance groups.

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

   ```
   https://<yourservername>/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 Oracle Private Cloud Appliance groups.

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

3.4.2 Verifying Identity Provider Self-Signed Certificates

> **Important**
>
> You can skip this section if your ADFS certificate is signed by a known certificate authority because they should already exist in the Oracle 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 Oracle 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://<yourservername>/FederationMetadata/2007-06/FederationMetadata.xml
   ```

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

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

3. Log on to management node 1 whose default name is `pcamn01`.

4. Navigate to `/etc/pca3.0/vault` and create a new directory named `customer_ca`.

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

5. In the `customer_ca` directory, create a new file in PEM format.

6. Copy the certificate from the FederationMetadata.xml file, which is located inbetween the
    `<X509Certificate>` and `</X509Certificate>` tag set 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
    ```

3.4.3 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.

3.4.3.1 Adding Active Directory as an Identity Provider

To federate with Active Directory 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 Oracle 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, Oracle 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 Oracle 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 metadat 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 Oracle Private Cloud Appliance and Active Directory.

To set up group mappings, see Section 3.4.4.1, "Creating Group Mappings".

3.4.3.2 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 .

4. Click Update Identity Provider.

3.4.3.3 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.

3.4.3.4 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.

3.4.3.5 Deleting an Identity Provider

If you want to remove the option for federated users to log into Oracle 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.

3.4.4 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.

- Oracle 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.

3.4.4.1 Creating Group Mappings

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

3.4.4.2 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.

3.4.4.3 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.

3.4.4.4 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.

3.4.5 Adding Oracle Private Cloud Appliance as a Trusted Relying Party in ADFS

Important

The Oracle Private Cloud Appliance certificate bundle must be added to Active Directory, so that ADFS can trust the Oracle Private Cloud Appliance certificate. If you do not do this, user logins will fail. For more information about the Oracle Private Cloud Appliance certificate bundle, see *Obtaining the Certificate Authority Bundle* in the *Oracle Private Cloud Appliance User Guide*.

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

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 Oracle Private Cloud Appliance as a trusted relying party.

   a. Under AD FS, right-click Relying Party Trusts and the 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 Oracle 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 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 Oracle 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 
        ```

    c. Click Finish.

The Issuance Transform Rules dialog displays the new rule.

3.4.6 Providing Federated Users Sign In Information

Before federated users can log into Oracle 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 Oracle Private Cloud Appliance.

Chapter 4 Tenancy Management

Table of Contents

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*.

4.1 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 purpos
adminUserName=testadmin adminPassword=************ confirmPassword=************
Command: create Tenant name=myTestTenancy serviceNamespace=test description="A tenancy for testing purposes
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.AK00661530.scasg01.jrgyo2w39riz38jhzredwz7s4zglm4slu6m6u37ok4odx5vfszak0
  AssociatedObj Type = Tenant
  AssociatedObj Id = ocid1.tenancy.AK00661530.scasg01.jrgyo2w39riz38jhzredwz7s4zglm4slu6m6u37ok4odx5vfszak0
  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.AK00661530.scasg01.r9l0nzgsm3vvtd6ugyrbx8em0pqogxp0x524yi7z3h1dztk6fuak00090146   myTenancy
  ocid1.tenancy.AK00661530.scasg01.iyalhgadxg2d71ej6qx8fs8n9v0dey8wqd7firgs6djbontjvcak00090146   myTenancy
  ocid1.tenancy.AK00661530.scasg01.9ax6fcf0bhe7an2b0m90e2t5uojkmfd1e47mkvye59e1u46ly6ak00090146   myTenancy
  ocid1.tenancy.AK00661530.scasg01.g7or03paq3k6j9hixsahhp6fh4ta4ntjz8x5yispcix5xeviu9ak00090146   myTestTen

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.AK00661530.scasg01.jrgyo2w39riz38jhzredwz7s4zglm4slu6m6u37ok4odx5vfszak00090146
  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.

4.2 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>`

4.3 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.AK00661530.scasg01.r9l0nzgsm3vvtd6ugyrbx8em0pqogxp0x524yi7z3h1dztk6fuak00090146   myTen
  ocid1.tenancy.AK00661530.scasg01.iyalhgadxg2d71ej6qx8fs8n9v0dey8wqd7firgs6djbontjvcak00090146   myTen
  ocid1.tenancy.AK00661530.scasg01.9ax6fcf0bhe7an2b0m90e2t5uojkmfd1e47mkvye59e1u46ly6ak00090146   myTen
  ocid1.tenancy.AK00661530.scasg01.g7or03paq3k6j9hixsahhp6fh4ta4ntjz8x5yispcix5xeviu9ak00090146   myTes
```

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.AK00661530.scasg01.r9l0nzgsm3vvtd6ugyrbx8em0pqogxp0x524yi7z3h1dztk6fuak00090146   myTenancy
  ocid1.tenancy.AK00661530.scasg01.iyalhgadxg2d71ej6qx8fs8n9v0dey8wqd7firgs6djbontjvcak00090146   myTenancy
  ocid1.tenancy.AK00661530.scasg01.9ax6fcf0bhe7an2b0m90e2t5uojkmfd1e47mkvye59e1u46ly6ak00090146   myTenancy
```

Chapter 5 Status and Health Monitoring

Table of Contents

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. Customers might prefer to expand and customize this setup, but this is beyond the scope of the Oracle Private Cloud Appliance documentation.

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*.

5.1 Using Grafana

With Grafana, Oracle Private Cloud Appliance offers administrators a single, visually oriented 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.

To access the Grafana home page

1. Open the Service Web UI and log in.

2. On the right-hand side of the dashboard, click the Monitoring tile.

   The Grafana home page opens in a new browser tab. Enter your user name and password when prompted.

When logs and metrics are stored in Prometheus they are given a time stamp based on the time and time zone settings of the appliance. However, Grafana displays the time based on user preferences, which may result in an offset because you are in a different time zone. It might be preferable to synchronize the time line in the Grafana visualizations with the time zone of the appliance.

To change the Grafana time line display

1. Open the Grafana home page.

2. In the menu bar on the left hand side, click your user account icon (near the bottom) to display your account preferences.

3. In the Preferences section, change the Time Zone setting to the same time zone as the appliance.

4. Click the Save button below to apply the change.

The pre-defined dashboards for Oracle Private Cloud Appliance are not directly accessible from the Grafana home page, although you can star your most used dashboards to appear on your home page later. Dashboards are organized in folders, which you access through the Dashboards section of the main menu.

To browse the Grafana dashboards

1. In the menu bar on the left hand side, point to Dashboards and select Manage.

   The list of folders, or dashboard sets, is displayed.

2. Click a folder to display the list of dashboards it contains. Click a dashboard to display its contents.

3. To navigate back to the list of folders and dashboards, use the menu bar as you did in step 1.

With the exception of the *My Sauron (Read Only)* dashboard set, all pre-defined dashboards and panels are editable by design. You can modify them or create your own using the specific metrics you want to monitor. The same applies to the alerts.

Alerts are managed in a separate area. Oracle has pre-defined a series of alerts for your convenience.

To access the alerting rules and notifications

1. In the menu bar on the left hand side, click Alerting (the bell icon).

   A list of all defined alert rules is displayed, including their current status.

2. Click an alert rule to display a detail panel and see how its status has evolved over time and relative to the alert threshold.

3. To navigate back to the list of alert rules, use the menu bar as you did in step 1.

4. To configure alert notifications, go to the Notification Channels tab of the Alerting page.

> **Note**
>
> If you wish to configure custom alerts using your own external notification channel, you must first configure the proxy for Grafana using the Sauron API endpoint. To do so, 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.<mypca>.example.com/v1/grafana/proxy/config?http-proxy=<proxy_fqdn>:<proxy
> Enter host password for user '<admin_user_name>':
> Grafana proxy config successfully updated!
> ```

Finally, Grafana also provides access to the appliance logs, which are aggregated through Loki. For more information, see Section 5.4, "Accessing System Logs".

5.2 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 pre-defined Grafana dashboards allow you to check the status of those essential low-level components, and drill down into 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 troubleshooting in case 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 countless ways.

Grafana Folder	Dashboard	Description
Service Monitoring	Server Stats	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. Some panels in this dashboard display a large number of *time series* in a single graph, so note that you can click to display a single one, or hover over the graph to view detailed data at a specific point on the time axis.
PCA 3.0 Service Advisor	Platform Health Check	This dashboard integrates the appliance health check mechanisms into the centralized approach that Grafana provides for logging and monitoring. By default, the Platform Health Check dashboard displays the failures for all health check services. You can change the panel display by selecting a health checker from the list of platform services, and you can choose to display healthy, unhealthy or all results. 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, simply click the time stamp.
My Sauron (Read Only)	Node Exporter Full	This dashboard displays a large number of detailed metric panels for a single compute or management node. Select a host from the list to display its data. 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. All dashboards in the My Sauron (Read Only) 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.

5.3 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 countless ways.

Grafana Folder	Dashboard	Description
Service Monitoring	ClusterLabs HA Cluster Details	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 pre-existing distributed data provided by the cluster components' tools. The monitoring data includes Pacemaker cluster summary, nodes and resource stats, and Corosync ring errors and quorum votes.
Service Monitoring	MySQL Cluster Exporter	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.
Service Monitoring	Service Level	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.
Service Monitoring	VM Stats	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. The panels in this dashboard display a large number of *time series* in a single graph, so note that you can click to display a single one, or hover over the graph to view detailed data at a specific point on the time axis.
PCA 3.0 Service Advisor	Kube Endpoint	This dashboard focuses specifically on the Kubernetes endpoints and provides endpoint alerts. These alerts can be sent to a notification channel of your choice.
PCA 3.0 Service Advisor	Kube Ingress	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.
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.

Grafana Folder	Dashboard	Description
Kubernetes Monitoring	(all)	These folders contains a large and diverse collection of dashboards with a wide range of monitoring data. covering practically all aspects of your Kubernetes cluster.
Kubernetes Monitoring Containers		The data covers Kubernetes at the cluster, node, pod and container levels. Metrics provide insights into deployment, ingress, usage of CPU, disk, memory and network, and much more.
Kubernetes Monitoring Node		

5.4 Accessing System Logs

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

To view the Loki logs

1. Open the Grafana home page.

2. In the menu bar on the left hand side, click Explore (the compass icon).

   By default, the Explore page's data source is set to "Prometheus".

3. At the top of the page near the left hand side, select "Loki" from the data source list.

4. Use the Log Labels list to query and filter the logs.

The logs are categorized with labels, which you can query in order to display log entries of a particular type or category. The principal log label categories used within Oracle Private Cloud Appliance are the following:

- `job`

  The log labels in this category are divided into three categories:

  - Platform: logs from services and components running in the foundation layers of the appliance architecture.

    Log labels in this category include: `"him"`/`"has"`/`"hms"` (hardware management), `"api-server"`, `"vault"`/`"etcd"` (secret service), `"corosync"`/`"pacemaker"`/`"pcsd"` (management cluster), "messages" (RabbitMQ)`"pca-platform-l0"`, `"pca-platform-l1api"`, and so on.

  - Infrastructure services: logs from the user-level cloud services and administrative services deployed on top of the platform. These services are easier to identify by their name.

    Log labels in this category include: `"brs"` (backup/restore), `"ceui"` (Compute Web UI), `"seui"` (Service Web UI), `"compute"`, `"dr-admin"` (disaster recovery), `"filesystem"`, `"iam"` (identity and access management), `"pca-upgrader"`, and so on.

  - Standard output: logs that the containerized infrastructure services send to the `stdout` stream. This output is visible to users when they execute a UI operation or CLI command.

Use the log label `job="k8s-stdout-logs"` to filter for the standard output logs. The log data comes from the microservices' Kubernetes containers, and can be filtered further by specifying a pod and/or container name.

- **k8s_app**

  Log labels in this category allow you to narrow down the standard output logs (`job="k8s-stdout-logs"`). That log data comes from the microservices' Kubernetes containers, and can be filtered further by selecting the label that corresponds with the name of the specific service you are interested in.

You navigate through the logs by selecting one of the `job` or `k8s_app` log labels. You pick the label that corresponds with the service or application you are interested in, and the list of logs is displayed in reverse chronological order. You can narrow your search by zooming in on a portion of the time line shown above the log entries. Color coding helps to identify the items that require your attention; for example: warnings are marked in yellow and errors are marked in red.

Audit Logs

The audit logs can be consulted as separate categories. From the Log Labels list, you can select these audit labels:

- `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 field next to the Log Labels list:

```
{job=~"(admin|api-server)"} | json tag="tag" | tag=~"(api-audit.log|audit.log)"
```

To execute, either click the Run Query button in the top-right corner or press `Shift` + `Enter`.

5.5 Using Auto Service Requests

Oracle Private Cloud Appliance is qualified for Oracle Auto Service Request (ASR). ASR is a software feature for support purposes. It is integrated with My Oracle Support and helps resolve problems faster by

automatically opening service requests when specific hardware failures occur. Using ASR is optional: the service must be registered and enabled for your appliance.

5.5.1 Understanding Oracle Auto Service Request (ASR)

ASR is designed to automatically open service requests when specific Oracle Private Cloud Appliance hardware faults occur. To enable this feature, the Oracle 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. It 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 Oracle Private Cloud Appliance to notify them of the creation of the service request. A service request may not be filed automatically on some occasions. This can happen because of the unreliable nature of the SNMP protocol or a loss of connectivity to ASR. Oracle recommends that customers continue to 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: http://www.oracle.com/technetwork/systems/asr/overview/index.html.

- Oracle Auto Service Request user documentation: http://docs.oracle.com/cd/E37710_01/index.htm.

5.5.2 ASR Prerequisites

Before you register for the ASR service, make sure that the prerequisites in this section are met.

Verifying ASR Prerequisites

1. Make sure that you have a valid My Oracle Support account.

   If necessary, create an account at https://support.oracle.com.

2. Ensure that the following are set up correctly in My Oracle Support:

   - technical contact person at the customer site who is responsible for Oracle Private Cloud Appliance

   - valid shipping address at the customer site where the Oracle 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.

5.5.3 Registering Oracle Private Cloud Appliance for ASR

To register the ASR client, the Oracle Private Cloud Appliance must be configured to send hardware fault telemetry to Oracle in one of three ways; 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 you register your Oracle Private Cloud Appliance for ASR, the 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 username and password, then complete the fields for the Phone Home configuration that you choose.

   - **Username*:** Enter your Oracle Single Sign On (SSO) credentials, which can be obtained from My Oracle Support.

   - **Password*:** Enter the password for your SSO account.

   - **Proxy Username:** To use a proxy host, enter a username 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:** Optionally, if you use an aggregation point, or other endpoint for ASR data consolidation, enter that endpoint in this format `http://<host>[:<port>]/asr`

     *Required fields

Using the Service CLI

Configure ASR Directly to https://transport.oracle.com

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

   ```
   # 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.us.oracle.com/ \
   Command: asrClientRegister username=asr-pca3_ca@example.com \
   password=*****  confirmPassword=***** \
   endpoint=https://transport.us.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.us.oracle.com/
PCA-ADMIN>
```

Configure ASR to a Proxy Host

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

```
# 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 `root`.

```
# 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.us.oracle.com/ \
Command: asrClientRegister username=oracle_email@example.com \
password=*****  confirmPassword=***** \
endpoint=https://transport.us.oracle.com/
Status: Success
Time: 2021-07-12 18:47:14,630 UTC
```

5.5.4 Testing ASR Configuration

Once configured, you can 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 will confirm if the test is successful or not.

4. If the test is not succesful, confirm your ASR configuration information and repeat the test.

Using the Service CLI

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

```
# 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>
```

5.5.5 Unregistering Oracle Private Cloud Appliance for ASR

When you unregister your Oracle Private Cloud Appliance for ASR, the service is automatically disabled, so you do not need to perform that 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 root.

```
# 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>
```

5.5.6 Disabling ASR

During system maintenance, or other circumstances, you may want to temporarily disable ASR on your appliance to halt the flow of fault messages to your configured endpoint, without unregistering the system.

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 root.

```
# 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>
```

5.5.7 Enabling ASR

If you have disabled ASR on your appliance, use one of these methods to restart the ASR service.

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 `root`.

```
# 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>
```

Chapter 6 Backup and Restore

Table of Contents

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.

Implementation details and technical background information for this feature can be found in the *Oracle Private Cloud Appliance Concepts Guide*. Refer to the section *Backup and Restore* in the chapter *Appliance Administration Overview*.

6.1 Activating Standard Daily Backup

System backups are not available by default. To activate it, the administrator must set up a Kubernetes CronJob by running the applicable script from the management node that owns the virtual IP of the cluster.

> **Caution**
>
> Make sure that daily backups are activated after system initialization. If this procedure is omitted, there will be no backup data to restore a component or service from a last known good state.

Execute these steps when the system initialization process has been completed.

1. Log on to one of the management nodes.

```
# ssh root@pcamn01
```

2. Retrieve the name of the Kubernetes pod that runs the backup and restore service. Use the following command:

```
# kubectl get pods -A | grep brs
default    brs-5bdc556546-gxtx9          3/3     Running   0    17d
```

3. Execute the `default-backup` script as shown below to set up the Kubernetes CronJob to make a daily backup.

```
kubectl exec brs-5bdc556546-gxtx9 -c brs -- /usr/sbin/default-backup
```

4. Verify that the CronJob has been added in the default namespace.

```
# kubectl get cronjobs -A
NAMESPACE       NAME                          SCHEDULE         SUSPEND   ACTIVE   LAST SCHEDULE   AGE
default         brs-cronjob-1629969790-backup 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    l0-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
```

Backups are created on the ZFS storage appliance at this 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....`*uniqueID*`/backup_<timestamp>/`

6.2 Executing a Backup Operation

It is critical that the standard daily backups are activated on your appliance. In addition, it is possible to initiate a system backup manually, if necessary.

Execute these steps to manually initiate a system backup.

1. Log on to one of the management nodes.

   ```
   # ssh root@pcamn01
   ```

2. Retrieve the name of the Kubernetes pod that runs the backup and restore service. Use the following command:

   ```
   # kubectl get pods -A | grep brs
   default    brs-5bdc556546-gxtx9          3/3    Running    0    17d
   ```

3. Execute the `default-backup` script with the "backup-now" option, as shown below.

   ```
   kubectl exec brs-5bdc556546-gxtx9 -c brs -- /usr/sbin/default-backup backup-now
   ```

4. Verify that the backup job is executed, and that it is completed successfully.

   ```
   # kubectl get pods -A | grep brs
   default    brs-5bdc556546-gxtx9                3/3    Running      0    17d
   default    brs-job-1641877703-backup-jkwx7     0/2    Running      0    8m40s

   # kubectl get pods -A | grep brs
   default    brs-5bdc556546-gxtx9                3/3    Running      0    17d
   default    brs-job-1641877703-backup-jkwx7     0/2    Completed    0    8m40s
   ```

Backups are created on the ZFS storage appliance at this 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....`*uniqueID*`/backup_<timestamp>/`

6.3 Restoring the System from a Backup

Restoring system data from a backup is a procedure that must be performed by Oracle-qualified support personnel. Please contact your Oracle representative for assistance.

Chapter 7 System Upgrade

Table of Contents

This chapter explains how an administrator upgrades the Oracle Private Cloud Appliance or one of its components.

Do not install or upgrade individual packages on the appliance components. Only upgrades as described in this chapter are supported. Security and other updates are provided through patches. Patching is separate from the upgrade functionality and uses a ULN mirror to download supported packages to the management nodes' shared storage.

Implementation details and technical background information for the upgrade and patching functionality can be found in the *Oracle Private Cloud Appliance Concepts Guide*. Refer to the sections *Upgrade* and *Patching* in the chapter *Appliance Administration Overview*.

Patching instructions are provided in a separate document. Refer to the *Oracle Private Cloud Appliance Patching Guide*.

7.1 Upgrade Requirements

Before you start an upgrade procedure, make sure that you have the required permissions and have downloaded the ISO image to a suitable location.

7.1.1 Verifying Permissions

To be able to execute an upgrade, you must have an administrator account to log in to the Service Enclave. You must be a member of one of these authorization groups: SuperAdmin, Admin, or DR Admin. For more information, see Chapter 3, *Administrator Account Management*.

When you log in to the Service CLI, you can verify that the upgrade commands are available to you by displaying all custom commands. The list of commands is filtered based on your access profile. If the upgrade commands are listed, it means you have permission to execute them.

```
PCA-ADMIN> showallcustomcmds
    Operation Name: <Related Object(s)>
    ----------------------------------
    [...]
    getUpgradeJob:  UpgradeJob
```

```
getUpgradeJobs:  UpgradeJobList
getUpgradeRequests:  UpgradeRequest
killUpgradeJob:  UpgradeJob
[...]
upgradeCN:  UpgradeRequest
upgradeEtcd:  UpgradeRequest
upgradeFullMN:  UpgradeRequest
upgradeHost:  UpgradeRequest
upgradeIlom:  UpgradeRequest
upgradeKubernetes:  UpgradeRequest
upgradeMySQL:  UpgradeRequest
upgradePlatform:  UpgradeRequest
upgradeSwitch:  UpgradeRequest
upgradeVault:  UpgradeRequest
upgradeZfssa:  UpgradeRequest
```

7.1.2 Preparing the ISO Image

Software versions and upgrades for Oracle Private Cloud Appliance are made available for download through My Oracle Support. The ISO file contains all the files and packages required to upgrade the appliance hardware and software components to a given release. All the items within the ISO file have been tested to work with each other and qualified for installation on your rack system.

To be able to use an ISO file to upgrade your appliance, you only need to download the file to a location from where a web server can make it available to the Oracle Private Cloud Appliance management nodes. If you have set up a bastion host connected to the internal administration network of the appliance, it is convenient to store the ISO file on that machine and run a web server to make the ISO file accessible over http.

When you execute an upgrade command on the appliance, you provide the path to the ISO file as a parameter. At that point, the ISO file is copied to the shared storage mounted on all three management nodes, and unpacked into a well-defined directory structure. You do not need to perform these steps manually in advance.

7.1.3 Ensuring the System Is In Ready State

Upgrades can be performed with limited impact on the system. No downtime is required, and user workloads continue to run while the underlying infrastructure is being upgraded in stages. However, it is considered good practice to ensure that backups are created of the system and the resources in your environment.

Every upgrade operation is preceded by a set of pre-checks. The upgrade will only begin if all pre-checks are passed. You are not required to execute the pre-checks manually; they are built into the upgrade code and will report an error if the system is not in the required state for the upgrade.

It is important to note that concurrent upgrade operations are not supported. An upgrade job must be completed before a new one can be started.

7.2 Performing a Full Management Node Upgrade

A full management node upgrade is a convenient way to upgrade all the required components on all three management nodes using just a single command. As part of this process, the following components are upgraded, in this specific order:

1. the host operating system

2. the MySQL cluster database

3. the secret service (including Etcd and Vault)

4. the Kubernetes container orchestration packages

5. the containerized microservices

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type.

   For a full management node upgrade, select Upgrade MN.

4. Fill out the upgrade request parameters:

   - **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   - **ISO Checksum:** Enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

   - **Image Location:** Enter the path to the location where the ISO image is stored.

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. Gather the information that you need to run the command:

   - the location of the ISO image to upgrade from

   - the checksum used to verify that the ISO image is valid

2. Enter the upgrade command.

   Syntax (entered on a single line):

   ```
   upgradeFullMN
   imageLocation=<path-to-iso>
   isoChecksum=<iso-file-checksum>
   ```

   Example:

   ```
   PCA-ADMIN> upgradeFullMN \
   imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
   isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
   Command: upgradeFullMN imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksum=24
   Status: Success
   Time: 2021-09-24 06:56:31,871 UTC
   Data:
     Service request has been submitted. Upgrade Request Id = UWS-39329657-1051-4267-8c5a-9314f8e63a64
   ```

3. Use the request ID to check the status of the upgrade process.

As the full management node upgrade is a multi-component upgrade process, there are multiple upgrade jobs associated with the upgrade request. You can filter for those jobs based on the request ID. Using the job ID, you can drill down into the details of each upgrade job.

```
PCA-ADMIN> getUpgradeJobs requestId=UWS-39329657-1051-4267-8c5a-9314f8e63a64
Command: getUpgradeJobs requestId=UWS-39329657-1051-4267-8c5a-9314f8e63a64
Status: Success
Time: 2021-09-24 17:32:31,595 UTC
Data:
  id                                 upgradeRequestId                                commandName   result
  --                                 ----------------                                -----------   ------
  1634578760906-platform-66082       UWS-39329657-1051-4267-8c5a-9314f8e63a64        platform      Passed
  1634578263434-kubernetes-63574     UWS-39329657-1051-4267-8c5a-9314f8e63a64        kubernetes    Passed
  1634578012353-vault-51696          UWS-39329657-1051-4267-8c5a-9314f8e63a64        vault         Passed
  1634577380954-etcd-46337           UWS-39329657-1051-4267-8c5a-9314f8e63a64        etcd          Passed
  1634577341291-mysql-40127          UWS-39329657-1051-4267-8c5a-9314f8e63a64        mysql         Passed
  1634576985926-host-36556           UWS-39329657-1051-4267-8c5a-9314f8e63a64        host          Passed
  1634576652071-host-27088           UWS-39329657-1051-4267-8c5a-9314f8e63a64        host          Passed
  1634576191050-host-24909           UWS-39329657-1051-4267-8c5a-9314f8e63a64        host          Passed

PCA-ADMIN> getUpgradeJob upgradeJobId=1634576652071-host-27088
Command: getUpgradeJob upgradeJobId=1634576652071-host-27088
Status: Success
Time: 2021-09-24 17:35:59,946 UTC
Data:
  Upgrade Request Id = UWS-39329657-1051-4267-8c5a-9314f8e63a64
  Composition Id = 1
  Name = host
  Start Time = 2021-09-24T07:04:12
  End Time = 2021-09-24T07:05:22
  Pid = 27088
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_host_os_2021_09_24-07.04.12.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":"100.96.2.35","result_over
  Status = Passed
  Execution Time(sec) = 139
  Tasks 1 - Name = Validate Image Location
  Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
  Tasks 1 - Time = 2021-10-18T17:04:16
[...]
```

The output of the `getUpgradeJob` command provides detailed information about the tasks performed during the upgrade procedure. It displays descriptions, time stamps, duration, and success or failure. Whenever an upgrade operation fails, the command output indicates which task has failed. For in-depth troubleshooting you can search the log file at the location provided near the start of the command output.

7.3 Upgrading Individual Components

The granular upgrade mechanism allows you to perform upgrade procedures for individual hardware and software components. Besides the components included in the management node upgrade, you can also upgrade different categories of firmware, and the operating system and appliance-specific software on the compute nodes.

7.3.1 Upgrading the Management Node Operating System

The Oracle Linux host operating system of the management nodes must be upgraded one node at a time; a rolling upgrade of all management nodes is not possible. This upgrade process, which involves updating the kernel and system packages, must always be initiated from the management node that holds the cluster virtual IP. Thus, in a three-management-node cluster, when you have upgraded two management

nodes, you must reassign the cluster virtual IP to one of the upgraded management nodes and execute the final upgrade command from that node.

You must upgrade management nodes one at a time, using each one's internal IP address as a command parameter. To obtain the host IP addresses, use the Service CLI command `show ManagementNode name=<node_name>` and look for the `Ip Address` in the output.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Host.

4. Fill out the upgrade request parameters:

   - **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   - **Host IP:** Enter the management node's assigned IP address in the internal administration network. This is an IP address in the internal 100.96.2.0/23 range.

   - **ISO Checksum:** Enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

   - **Image Location:** Enter the path to the location where the ISO image is stored.

   - **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. Gather the information that you need to run the command:

   - the location of the ISO image to upgrade from

   - the checksum used to verify that the ISO image is valid

   - the IP address of the management node for which you intend to upgrade the host operating system

2. Run the Service CLI from the management node that holds the management cluster virtual IP.

   a. Log on to one of the management nodes and check the status of the cluster.

   ```
   # ssh root@pcamn01
   # pcs status
   Cluster name: mncluster
   Stack: corosync
   Current DC: pcamn02 (version 1.1.23-1.0.1.el7-9acf116022) - partition with quorum

   Online: [ pcamn01 pcamn02 pcamn03 ]

   Full list of resources:
   ```

```
  scsi_fencing         (stonith:fence_scsi):         Stopped (disabled)
  Resource Group: mgmt-rg
      vip-mgmt-int     (ocf::heartbeat:IPaddr2):      Started    pcamn02
      vip-mgmt-host    (ocf::heartbeat:IPaddr2):      Started    pcamn02
      vip-mgmt-ilom    (ocf::heartbeat:IPaddr2):      Started    pcamn02
      vip-mgmt-lb      (ocf::heartbeat:IPaddr2):      Started    pcamn02
      vip-mgmt-ext     (ocf::heartbeat:IPaddr2):      Started    pcamn02
      l1api            (systemd:l1api):               Started    pcamn02
      haproxy          (ocf::heartbeat:haproxy):      Started    pcamn02
      pca-node-state   (systemd:pca_node_state):      Started    pcamn02
      dhcp             (ocf::heartbeat:dhcpd):        Started    pcamn02
      hw-monitor       (systemd:hw_monitor):          Started    pcamn02

Daemon Status:
  corosync: active/enabled
  pacemaker: active/enabled
  pcsd: active/enabled
```

In this example, the command output indicates that the node with host name `pcamn02` currently holds the cluster virtual IP.

b. Log in to the management node with the virtual IP and launch the Service CLI.

```
# ssh pcamn02
# ssh admin@localhost -p 30006
PCA-ADMIN>
```

3. Enter the upgrade command.

Syntax (entered on a single line):

```
upgradeHost
imageLocation=<path-to-iso>
isoChecksum=<iso-file-checksum>
hostIp=<management-node-ip>
```

Example:

```
PCA-ADMIN> upgradeHost hostIp=100.96.2.35 \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc4d79
Command: upgradeHost hostIp=100.96.2.35 imageLocation="http://host.example.com/pca-<version>-<build>.iso" i
Status: Success
Time: 2021-09-25 05:47:02,735 UTC
Data:
  Service request has been submitted. Upgrade Job Id = 1632990827394-host-56156 Upgrade Request Id = UWS-1a
```

4. Use the request ID and the job ID to check the status of the upgrade process.

```
PCA-ADMIN> getUpgradeJobs
  id                             upgradeRequestId                         commandName   result
  --                             ----------------                         -----------   ------
  1632990827394-host-56156       UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755 host          Passed

PCA-ADMIN> getUpgradeJob upgradeJobId=1632990827394-host-56156
Command: getUpgradeJob upgradeJobId=1632990827394-host-56156
Status: Success
Time: 2021-09-25 05:54:28,054 UTC
Data:
  Upgrade Request Id = UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755
  Composition Id = 1
  Name = host
  Start Time = 2021-09-25T05:47:02
  End Time = 2021-09-25T05:48:38
```

```
  Pid = 56156
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_host_os_2021_09_25-05.47.02.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":"100.96.2.35","result_
  Status = Passed
  Execution Time(sec) = 96
  Tasks 1 - Name = Validate Image Location
  Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
  Tasks 1 - Time = 2021-09-25T05:47:02
  Tasks 2 - Name = Validate Image Location
[...]
```

5. When the first management node host operating system upgrade has completed successfully, execute the same command for the next management node.

```
PCA-ADMIN> upgradeHost hostIp=100.96.2.33 \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
```

6. When the second management node host operating system upgrade has completed successfully, exit the Service CLI and move the cluster virtual IP to one of the upgraded nodes.

```
PCA-ADMIN> exit
Connection to localhost closed.
# pcs resource move mgmt-rg pcamn01
# pcs status
Cluster name: mncluster
Stack: corosync
[...]
 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
[...]
```

Moving the cluster virtual IP to another management node should only take a number of seconds.

7. Log in to the management node with the virtual IP and launch the Service CLI to execute the host operating system upgrade for the final management node.

```
# ssh pcamn01
# ssh admin@localhost -p 30006
PCA-ADMIN> upgradeHost hostIp=100.96.2.34 \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
```

When this upgrade has completed successfully, the operating system on all management nodes is up-to-date.

7.3.2 Upgrading the MySQL Cluster Database

The MySQL Cluster database is upgraded independently of the management node host operating system; the MySQL packages are deliberately kept separate from the Oracle Linux upgrade.

It is assumed that the database upgrade is performed after the management node host operating system upgrade. As the ISO image has already been unpacked on shared storage during the operating system upgrade, the ISO path and checksum are not considered mandatory parameters for the database upgrade command. However, if the database upgrade is not preceded by a host operating system upgrade, you must include the ISO path and checksum parameters to be able to launch the database upgrade.

The MySQL Cluster database upgrade is a rolling upgrade: with one command the upgrade is executed on each of the three management nodes.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

    The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade MySQL.

4. If required, fill out the upgrade request parameters:

    *   **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

    *   **ISO Checksum:** If you have not executed a previous upgrade command specifying this parameter, enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

    *   **Image Location:** If you have not executed a previous upgrade command specifying this parameter, enter the path to the location where the ISO image is stored.

    *   **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

    The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. If you have not previously upgraded the management node host operating system, gather the information that you need to run the command:

    *   the location of the ISO image to upgrade from

    *   the checksum used to verify that the ISO image is valid

2. Enter the upgrade command.

    Syntax (entered on a single line):

    ```
    upgradeMySQL
    imageLocation=<path-to-iso>
    isoChecksum=<iso-file-checksum>
    ```

    Example:

    ```
    PCA-ADMIN> upgradeMySQL \
    imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
    isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc4d79
    Command: upgradeMySQL imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksum=240420c
    Status: Success
    Time: 2021-09-25 09:21:16,264 UTC
    Data:
      Service request has been submitted. Upgrade Job Id = 1632995409822-mysql-83013 Upgrade Request Id = UWS-7
    ```

    If the upgrade ISO image has already been unpacked on shared storage, simply enter the command without parameters: `upgradeMySQL`.

3. Use the request ID and the job ID to check the status of the upgrade process.

```
PCA-ADMIN> getUpgradeJobs
  id                             upgradeRequestId                          commandName  result
  --                             ----------------                          -----------  ------
  1632995409822-mysql-83013      UWS-77bc0c30-7ff5-4c50-ad09-6f96907e22e1  mysql        Passed
  1632926926773-host-32993       UWS-fef3b663-45b7-4177-a041-26f73e68848d  host         Passed
  1632990827394-host-56156       UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755  host         Passed
  1632990493570-host-6646        UWS-4c78f3ef-ac42-4f32-9483-bb43a309faa3  host         Passed

PCA-ADMIN> getUpgradeJob upgradeJobId=1632995409822-mysql-83013
Command: getUpgradeJob upgradeJobId=1632995409822-mysql-83013
Status: Success
Time: 2021-09-25 09:24:27,874 UTC
Data:
  Upgrade Request Id = UWS-77bc0c30-7ff5-4c50-ad09-6f96907e22e1
  Name = mysql
  Start Time = 2021-09-25T09:21:16
  End Time = 2021-09-25T09:22:04
  Pid = 83013
  Host = pcamn01
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_mysql_cluster_2021_09_25-09.21.16.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override"
  Status = Passed
  Execution Time(sec) = 48
  Tasks 1 - Name = Validate Image Location
  Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
  Tasks 1 - Time = 2021-09-25T09:21:16
[...]
```

7.3.3 Upgrading the Secret Service

The secret service contains two components that need to be upgraded separately: Etcd and Vault. The order in which you upgrade them is not relevant.

It is assumed that the secret service upgrade is performed after the management node host operating system upgrade. As the ISO image has already been unpacked on shared storage during the operating system upgrade, the ISO path and checksum are not considered mandatory parameters for the secret service upgrade commands. However, if the host operating system upgrade has not yet been executed, you must include the ISO path and checksum parameters.

The Etcd and Vault upgrades are rolling upgrades: each upgrade is executed on all three management nodes with one command.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Etcd.

4. If required, fill out the upgrade request parameters:

   • **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   • **ISO Checksum:** If you have not executed a previous upgrade command specifying this parameter, enter the checksum that allows the system to verify that the ISO image is valid for this upgrade.

The checksum is provided alongside the ISO image; its file name is the ISO image name with
`.sha512sum` appended.

- **Image Location:** If you have not executed a previous upgrade command specifying this parameter,
enter the path to the location where the ISO image is stored.

- **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is
"Information". For maximum detail, select "Debug".

5. Click Create Request.

The new upgrade request appears in the Upgrade Jobs table.

6. When the Etcd upgrade has completed successfully, repeat this procedure to create an upgrade
request for Vault.

Using the Service CLI

1. If you have not previously upgraded the management node host operating system, gather the
information that you need to run the command:

- the location of the ISO image to upgrade from

- the checksum used to verify that the ISO image is valid

2. Enter the two upgrade commands. Wait until one upgrade is finished before entering the second
command.

Syntax (entered on a single line):

```
upgradeEtcd
imageLocation=<path-to-iso>
isoChecksum=<iso-file-checksum>
```

```
upgradeVault
imageLocation=<path-to-iso>
isoChecksum=<iso-file-checksum>
```

Example:

```
PCA-ADMIN> upgradeEtcd \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc4d79
Command: upgradeEtcd imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksum=240420cf
Status: Success
Time: 2021-09-25 10:24:52,177 UTC
Data:
  Service request has been submitted. Upgrade Job Id = 1632826770954-etcd-26973 Upgrade Request Id = UWS-fe

PCA-ADMIN> upgradeVault \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc4d79
Command: upgradeVault imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksum=240420c
Status: Success
Time: 2021-09-25 10:38:25,417 UTC
Data:
  Service request has been submitted. Upgrade Job Id = 1632850933353-vault-16966 Upgrade Request Id = UWS-3
```

If the upgrade ISO image has already been unpacked on shared storage, simply enter the commands
without parameters: `upgradeEtcd` and `upgradeVault`.

3. Use the request ID and the job ID to check the status of the upgrade process.

```
PCA-ADMIN> getUpgradeJobs
  id                              upgradeRequestId                        commandName   result
  --                              ----------------                        -----------   ------
  1632995409822-mysql-83013       UWS-77bc0c30-7ff5-4c50-ad09-6f96907e22e1  mysql         Passed
  1632850933353-vault-16966       UWS-352df3d1-c21f-441b-8f6e-9381ac075906  vault         Passed
  1632826770954-etcd-26973        UWS-fec15d32-fc2b-48bd-9ae0-62f49587a284  etcd          Passed
  1632926926773-host-32993        UWS-fef3b663-45b7-4177-a041-26f73e68848d  host          Passed
  1632990827394-host-56156        UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755  host          Passed
  1632990493570-host-6646         UWS-4c78f3ef-ac42-4f32-9483-bb43a309faa3  host          Passed

PCA-ADMIN> getUpgradeJob upgradeJobId=1632850933353-vault-16966
Command: getUpgradeJob upgradeJobId=1632850933353-vault-16966
Status: Success
Time: 2021-09-25 10:39:31,308 UTC
Data:
  Upgrade Request Id = UWS-352df3d1-c21f-441b-8f6e-9381ac075906
  Name = vault
  Start Time = 2021-09-25T10:38:25
  End Time = 2021-09-25T10:39:07
  Pid = 16966
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_vault_2021_09_25-10.38.25.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override"
  Status = Passed
  Execution Time(sec) = 42
  Tasks 1 - Name = Check Vault Running Status
  Tasks 1 - Description = Check vault service running status is healthy
  Tasks 1 - Time = 2021-09-25T10:38:25
[...]
```

7.3.4 Upgrading Firmware

Firmware is included in the ISO image for all component ILOMs, for the ZFS Storage Appliance, and for the switches. Select the instructions below for the component type you want to upgrade.

Upgrading ILOMs

ILOM upgrades can be applied to management nodes and compute nodes. Firmware packages may be different per component type, so make sure you select the correct one from the firmware directory. You must upgrade ILOMs one at a time, using each one's internal IP address as a command parameter.

To obtain the ILOM IP addresses, use the Service CLI command `show ComputeNode name=<node_name>` or `show ManagementNode name=<node_name>` and look for the `ILOM Ip Address` in the output.

> **Caution**
>
> You must NOT upgrade the ILOM of the management node that holds the management virtual IP address, and thus the primary role in the cluster. To upgrade its ILOM, first reboot the management node in question so that another node in the cluster takes over the primary role. Once the node has rebooted completely, you can proceed with the ILOM upgrade.
>
> To determine which management node has the primary role in the cluster, log in to any management node and run the command `pcs status`.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Ilom.

4. Fill out the upgrade request parameters:

   • **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   • **Host IP:** Enter the component's assigned IP address in the ILOM network. This is an IP address in the internal 100.96.0.0/23 range.

   • **Image Location:** Enter the path to the location where the firmware package is stored. ILOM firmware is stored as a `*.pkg` file in the `/pca_firmware/<component>/` subdirectory of the unpacked ISO image.

   • **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. Gather the information that you need to run the command:

   • the IP address of the ILOM for which you intend to upgrade the firmware

   • the path to the firmware package file in the unpacked ISO image

2. Enter the upgrade command.

   Syntax (entered on a single line):

   ```
   upgradeIlom
   imageLocation=<path-to-firmware>
   hostIp=<ilom-ip>
   ```

   Example:

   ```
   PCA-ADMIN> upgradeIlom \
   imageLocation="http://host.example.com/iso/pca_firmware/X9-2/.../ILOM-<version>-ORACLE_SERVER_X9-2-rom.pkg"
   hostIp=100.96.0.66
   Command: upgradeIlom imageLocation="http://host.example.com/iso/pca_firmware/X9-2/.../ILOM-<version>-ORACLE
   Status: Success
   Time: 2021-09-24 11:18:31,044 UTC
   Data:
     Service request has been submitted. Upgrade Job Id = 1620921089806-ilom-21480 Upgrade Request Id = UWS-73
   ```

3. Use the request ID and the job ID to check the status of the upgrade process.

   ```
   PCA-ADMIN> getUpgradeJobs
     id                             upgradeRequestId                      commandName   result
     --                             ----------------                      -----------   ------
     1620921089806-ilom-21480       UWS-732d6fce-9f06-4329-b972-d093bee40010   ilom          Passed
     1632926926773-host-32993       UWS-fef3b663-45b7-4177-a041-26f73e68848d   host          Passed
     1632990827394-host-56156       UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755   host          Passed
     1632990493570-host-6646        UWS-4c78f3ef-ac42-4f32-9483-bb43a309faa3   host          Passed

   PCA-ADMIN> getUpgradeJob upgradeJobId=1620921089806-ilom-21480
   ```

```
Command: getUpgradeJob 1620921089806-ilom-21480
Status: Success
Time: 2021-09-24 11:24:49,243 UTC
Data:
  Upgrade Request Id = UWS-732d6fce-9f06-4329-b972-d093bee40010
  Name = ilom
  Start Time = 2021-09-24 11:18:32
  End Time = 2021-09-24 11:21:18
  Pid = 21480
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_ilom_firmware_2021_09_24-11.18.31.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":"100.96.0.66","result_
  Status = Passed
  Execution Time(sec) = 166
  Tasks 1 - Name = Validate Image Location
  Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
  Tasks 1 - Time = 2021-09-24T11:18:32
[...]
```

At the end of the upgrade, the ILOM itself is rebooted automatically. However, the server component also needs to be rebooted for all changes to take effect. It is the administrator's responsibility to manually reboot the management node or compute node after a successful ILOM upgrade.

Upgrading the ZFS Storage Appliance Operating Software

To upgrade the operating software of the appliance's ZFS Storage Appliance, you only need to provide the path to the firmware package in the unpacked ISO image. The IP addresses of the storage controllers are known, and a single upgrade command initiates a rolling upgrade of both controllers. If a new ILOM firmware version is included for the two controllers, it will be installed as part of the storage appliance upgrade process.

> **Caution**
>
> Do not make storage configuration changes while an upgrade is in progress. While controllers are running different software versions, configuration changes made to one controller are not propagated to its peer controller.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Zfssa.

4. Fill out the upgrade request parameters:

   • **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   • **Image Location:** Enter the path to the location where the firmware package is stored. ZFS Storage Appliance operating software is stored as a `*.pkg` file in the `/pca_firmware/zfs/` subdirectory of the unpacked ISO image.

   • **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. Gather the information that you need to run the command: the path to the AK-NAS firmware package in the unpacked ISO image.

2. Enter the upgrade command.

   Syntax:

   ```
   upgradeZfssa imageLocation=<path-to-firmware>
   ```

   Example:

   ```
   PCA-ADMIN> upgradeZfssa imageLocation="file:///nfs/shared_storage/pca_firmware/zfs/ak-nas-<version>.pkg"
   Command: upgradeZfssa imageLocation="file:///nfs/shared_storage/pca_firmware/zfs/ak-nas-<version>.pkg"
   Status: Success
   Time: 2021-09-27 11:15:07,453 UTC
   Data:
     Service request has been submitted. Upgrade Job Id = 1632914107346-zfssa-83002 Upgrade Request Id = UWS-8
   ```

3. Use the request ID and the job ID to check the status of the upgrade process.

   ```
   PCA-ADMIN> getUpgradeJobs
     id                              upgradeRequestId                          commandName   result
     --                              ----------------                          -----------   ------
     1632914107346-zfssa-83002       UWS-881af57f-5dfb-4c75-8026-9f00cf3eb7c9   zfssa         Passed
     1632926926773-host-32993        UWS-fef3b663-45b7-4177-a041-26f73e68848d   host          Passed
     1632990827394-host-56156        UWS-1a97a8d9-54ef-478d-a0c0-348a17ba6755   host          Passed
     1632990493570-host-6646         UWS-4c78f3ef-ac42-4f32-9483-bb43a309faa3   host          Passed

   PCA-ADMIN> getUpgradeJob upgradeJobId=1632914107346-zfssa-83002
   Command: getUpgradeJob upgradeJobId=1632914107346-zfssa-83002
   Status: Success
   Time: 2021-09-27 11:42:10,729 UTC
   Data:
     Upgrade Request Id = UWS-881af57f-5dfb-4c75-8026-9f00cf3eb7c9
     Name = zfssa
     Start Time = 2021-09-29T11:15:07
     End Time = 2021-09-29T11:26:42
     Pid = 83002
     Host = pcamn02
     Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_zfssa_ak_2021_09_29-11.15.07.log
     Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override":nul
     Status = Passed
     Execution Time(sec) = 697
     Tasks 1 - Name = Validate Image Location
     Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
     Tasks 1 - Time = 2021-09-29T11:15:08
   [...]
   ```

Upgrading the Switch Software

The appliance rack contains three categories of Cisco Nexus switches: a management switch, two leaf switches, and two spine switches. They all run the same Cisco NX-OS network operating software. The preferred upgrade order is as follows: leaf switches first, then spine switches, and finally the management switch.

When upgrading their firmware, use the same binary file with each upgrade command. Only one command per switch category is required, meaning that the leaf switches and the spine switches are upgraded in pairs.

Some versions of the network operating software consist of two files: a binary file and an additional EPLD (electronic programmable logic device) image. If the new firmware includes an EPLD file, make sure it is also added to the upgrade request.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Switch.

4. Fill out the upgrade request parameters:

   - **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   - **EPLD:** If required for this firmware version, enter the path to the location where the EPLD image file is stored. If present, an EPLD file is an `*.img` file stored alongside the NX-OS operating software in the `/pca_firmware/network/cisco/` subdirectory of the unpacked ISO image.

   - **Image Location:** Enter the path to the location where the firmware package is stored. Cisco NX-OS network operating software is stored as a `*.bin` file in the `/pca_firmware/network/cisco/` subdirectory of the unpacked ISO image.

   - **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

   - **Switch Type:** Select the switch type you intend to upgrade. The preferred upgrade order is as follows: leaf switches first, then spine switches, and finally the management switch.

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

6. When the upgrade has completed successfully, but other switches in the system still need to be upgraded, repeat this procedure for any other type of switch that requires upgrading.

Using the Service CLI

1. Gather the information that you need to run the command:

   - the type of witch to upgrade (spine, leaf, management)

   - the path to the firmware binary file in the unpacked ISO image

   - if present with the new firmware version, the path to the EPLD upgrade file in the unpacked ISO image

2. Enter the upgrade command.

   Syntax (entered on a single line):

```
upgradeSwitch
switchType=[MGMT, SPINE, LEAF]
imageLocation=<path-to-firmware>
(epld=<path-to-epld-file>)
```

Example:

```
PCA-ADMIN> upgradeSwitch switchType=LEAF \
imageLocation="http://host.example.com/iso/pca_firmware/network/cisco/nxos.<version>.bin" \
epld="http://host.example.com/iso/pca_firmware/network/cisco/n9000-epld.<version>.img"
Command: upgradeSwitch switchType=LEAF imageLocation="file:///nfs/shared_storage/pca_firmware/network/cisco
Status: Success
Time: 2021-09-24 14:16:54,704 UTC
Data:
  Service request has been submitted. Upgrade Job Id = 1630511206512-cisco-20299 Upgrade Request Id = UWS-4
```

3. Use the request ID and the job ID to check the status of the upgrade process.

```
PCA-ADMIN> getUpgradeJobs
  id                          upgradeRequestId                              commandName  result
  --                          ----------------                              -----------  ------
  1632914107346-zfssa-83002   UWS-881af57f-5dfb-4c75-8026-9f00cf3eb7c9      zfssa        Passed
  1630511206512-cisco-20299   UWS-44688fe5-b4f8-407f-a1b5-8cd1b685c2c3      cisco        Passed
  1620921089806-ilom-21480    UWS-732d6fce-9f06-4329-b972-d093bee40010      ilom         Passed

PCA-ADMIN> getupgradeJob upgradeJobId=1630511206512-cisco-20299
Command: getupgradeJob upgradeJobId=1630511206512-cisco-20299
Status: Success
Time: 2021-09-24 15:48:08,455 UTC
Data:
  Upgrade Request Id = UWS-44688fe5-b4f8-407f-a1b5-8cd1b685c2c3
  Name = cisco
  Start Time = 2021-09-24T14:46:46
  End Time = 2021-09-24T14:59:44
  Pid = 20299
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_cisco_firmware_2021_09_24-14.46.46.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override":nul
  Status = Passed
  Execution Time(sec) = 777
  Tasks 1 - Name = Validate Image Location
  Tasks 1 - Description = Verify that the image exists at the specified location and is correctly named
  Tasks 1 - Time = 2021-09-24T14:46:47
[...]
```

7.3.5 Upgrading the Kubernetes Cluster

The Kubernetes container orchestration environment upgrade is also kept separate from the operating system. With a single command, all Kubernetes packages, such as kubeadm, kubectl and kubelet, are upgraded on the three management nodes and all the compute nodes. Note that this upgrade does not include the microservices running in Kubernetes containers.

For dependency reasons, Kubernetes must be upgraded after the management node host operating system. Thus, the upgrade ISO image has already been unpacked on shared storage; the Kubernetes upgrade command has no mandatory parameters.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Kubernetes.

4. If required, fill out the upgrade request parameters:

- **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

- **ISO Checksum:** If you have not executed a previous upgrade command specifying this parameter, enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

- **Image Location:** If you have not executed a previous upgrade command specifying this parameter, enter the path to the location where the ISO image is stored.

- **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. If you have not previously upgraded the management node host operating system, gather the information that you need to run the command:

   - the location of the ISO image to upgrade from

   - the checksum used to verify that the ISO image is valid

2. Enter the upgrade command.

   Syntax (entered on a single line):

   ```
   upgradeKubernetes
   imageLocation=<path-to-iso>
   isoChecksum=<iso-file-checksum>
   ```

   Example:

   ```
   PCA-ADMIN> upgradeKubernetes \
   imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
   isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
   Command: upgradeKubernetes imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksu
   Status: Success
   Time: 2021-09-26 17:20:09,423 UTC
   Data:
     Service request has been submitted. Upgrade Job Id = 1632849609034-kubernetes-35545 Upgrade Request I
   ```

   If the upgrade ISO image has already been unpacked on shared storage, simply enter the command without parameters: `upgradeKubernetes`.

3. Use the request ID and the job ID to check the status of the upgrade process.

   ```
   PCA-ADMIN> getUpgradeJobs
     id                            upgradeRequestId                            commandName   result
     --                            ----------------                            -----------   ------
     1632849609034-kubernetes-35545   UWS-edfa3b32-c32a-4b67-8df5-2357096052bf   kubernetes    Passed
     1632826770954-etcd-26973         UWS-fec15d32-fc2b-48bd-9ae0-62f49587a284   etcd          Passed
     1632850933353-vault-16966        UWS-352df3d1-c21f-441b-8f6e-9381ac075906   vault         Passed

   PCA-ADMIN> getUpgradeJob upgradeJobId=1632849609034-kubernetes-35545
   Command: getUpgradeJob upgradeJobId=1632849609034-kubernetes-35545
   ```

```
Status: Success
Time: 2021-09-26 17:43:38,443 UTC
Data:
  Upgrade Request Id = UWS-edfa3b32-c32a-4b67-8df5-2357096052bf
  Name = kubernetes
  Start Time = 2021-09-26T17:20:09
  End Time = 2021-09-26T17:21:52
  Pid = 35545
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_kubernetes_cluster_2021_09_26-17.20.09.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override":nul
  Status = Passed
  Execution Time(sec) = 249
  Tasks 1 - Name = Retrieving Cluster Status
  Tasks 1 - Description = Retrieving cluster status and upgrade data from the kubernetes nodes
  Tasks 1 - Time = 2021-09-26T17:20:10
[...]
```

7.3.6 Upgrading the Microservices

The microservices upgrade covers both the internal services of the platform layer, and the administrative and user-level services exposed through the infrastructure services layer.

The containerized microservices have their own separate upgrade mechanism. A service is upgraded if a new Helm deployment chart and container image are found in the ISO image. When a new deployment chart is detected during the upgrade process, the pods running the services are restarted with the new container image.

It is assumed that the microservices upgrade is performed after the management node host operating system upgrade. As the ISO image has already been unpacked on shared storage during the operating system upgrade, the ISO path and checksum are not considered mandatory parameters for the microservices upgrade command. However, if the host operating system upgrade has not yet been executed, you must include the ISO path and checksum parameters.

Using the Service Web UI

1. In the navigation menu, click Upgrade & Patching.

2. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

3. Select the appropriate upgrade request type: Upgrade Platform.

4. If required, fill out the upgrade request parameters:

   • **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

   • **ISO Checksum:** If you have not executed a previous upgrade command specifying this parameter, enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

   • **Image Location:** If you have not executed a previous upgrade command specifying this parameter, enter the path to the location where the ISO image is stored.

   • **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

5. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

Using the Service CLI

1. If you have not previously upgraded the management node host operating system, gather the
   information that you need to run the command:

   • the location of the ISO image to upgrade from

   • the checksum used to verify that the ISO image is valid

2. Enter the upgrade command.

   Syntax (entered on a single line):

   ```
   upgradePlatform
   imageLocation=<path-to-iso>
   isoChecksum=<iso-file-checksum>
   ```

   Example:

   ```
   PCA-ADMIN> upgradePlatform \
   imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
   isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
   Command: upgradePlatform imageLocation="http://host.example.com/pca-<version>-<build>.iso" isoChecksum=
   Status: Success
   Time: 2021-09-26 20:48:41,452 UTC
   Data:
     Service request has been submitted. Upgrade Job Id = 1632850650836-platform-68465 Upgrade Request Id
   ```

   If the upgrade ISO image has already been unpacked on shared storage, simply enter the command
   without parameters: `upgradePlatform`.

3. Use the request ID and the job ID to check the status of the upgrade process.

   ```
   PCA-ADMIN> getUpgradeJobs
     id                                upgradeRequestId                              commandName   result
     --                                ----------------                              -----------   ------
     1632850650836-platform-68465      UWS-26dba234-9b52-426d-836c-ac11f37e717f      platform      Passed
     1632849609034-kubernetes-35545    UWS-edfa3b32-c32a-4b67-8df5-2357096052bf      kubernetes    Passed
     1632826770954-etcd-26973          UWS-fec15d32-fc2b-48bd-9ae0-62f49587a284      etcd          Passed
     1632850933353-vault-16966         UWS-352df3d1-c21f-441b-8f6e-9381ac075906      vault         Passed

   PCA-ADMIN> getUpgradeJob upgradeJobId=1632850650836-platform-68465
   Command: getUpgradeJob 1632850650836-platform-68465
   Status: Success
   Time: 2021-09-26 21:03:16,264 UTC
   Data:
     Upgrade Request Id = UWS-26dba234-9b52-426d-836c-ac11f37e717f
     Name = kubernetes
     Start Time = 2021-09-26T20:48:41
     End Time = 2021-09-26T20:59:34
     Pid = 68465
     Host = pcamn02
     Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_platform_services_2021_09_26-20.48.41.lo
     Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":null,"result_override"
     Status = Passed
     Execution Time(sec) = 653
     Tasks 1 - Name = Check All Ingress Endpoints
     Tasks 1 - Description = Check whether all ingress endpoints are up and running
     Tasks 1 - Time = 2021-09-26T20:48:42
   [...]
   ```

7.3.7 Upgrading a Compute Node

The compute node upgrade is similar to the management node host operating system upgrade: it ensures that the latest Oracle Linux kernel and user space packages are installed, as well as the `ovm-agent` package with appliance-specific optimizations. Compute nodes must be locked and upgraded one at a time; concurrent upgrades are not supported. After successful upgrade, when a compute node has rebooted, the administrator must manually remove the locks to allow the node to return to normal operation.

To obtain the host IP addresses of a compute node, use the Service CLI command `show ComputeNode name=<node_name>` and look for the `Ip Address` in the output.

Using the Service Web UI

1. Set the provisioning and maintenance locks for the compute node you are about to upgrade.

   For more information, refer to Section 2.2, "Performing Compute Node Operations".

   a. In the navigation menu, click Rack Units. In the Rack Units table, click the name of the compute node you want to upgrade to display its detail page.

   b. In the top-right corner of the compute node detail page, click Controls and select the Provisioning Lock command.

   c. When the provisioning lock is set, click Controls again and select the Maintenance Lock command.

2. In the navigation menu, click Upgrade & Patching.

3. In the top-right corner of the Upgrade Jobs page, click Create Upgrade or Patch.

   The Create Request window appears. Choose *Upgrade* as the Request Type.

4. Select the appropriate upgrade request type: Upgrade CN.

5. Fill out the upgrade request parameters:

   • **Host IP:** Enter the compute node's assigned IP address in the internal administration network. This is an IP address in the internal 100.96.2.0/23 range.

   • **ISO Checksum:** If you have not executed a previous upgrade command specifying this parameter, enter the checksum that allows the system to verify that the ISO image is valid for this upgrade. The checksum is provided alongside the ISO image; its file name is the ISO image name with `.sha512sum` appended.

   • **Image Location:** If you have not executed a previous upgrade command specifying this parameter, enter the path to the location where the ISO image is stored.

   • **Log Level:** Optionally, select a specific log level for the upgrade log file. The default log level is "Information". For maximum detail, select "Debug".

   • **Advanced Options JSON:** Optionally, add a JSON string to provide additional command parameters.

6. Click Create Request.

   The new upgrade request appears in the Upgrade Jobs table.

7. When the compute node has been upgraded successfully, release the provisioning and maintenance locks.

For more information, refer to Section 2.2, "Performing Compute Node Operations".

   a. Open the compute node detail page.

   b. In the top-right corner of the compute node detail page, click Controls and select the Maintenance Unlock command.

   c. When the maintenance lock has been released, click Controls again and select the Provisioning Unlock command.

Using the Service CLI

1. Gather the information that you need to run the command:

   • the location of the ISO image to upgrade from

   • the checksum used to verify that the ISO image is valid

   • the IP address of the compute node you intend to upgrade

2. Set the provisioning and maintenance locks for the compute node you are about to upgrade.

   For more information, refer to Section 2.2, "Performing Compute Node Operations".

```
PCA-ADMIN> list ComputeNode
Data:
  id                                      name       provisioningState   provisioningType
  --                                      ----       ----------------    ----------------
  363a26f4-fa34-4e4c-8e17-a1671a0b77d1    pcacn001   Provisioned         KVM
  9e8745c7-52e3-4aae-984c-e198869ee2cc    pcacn002   Provisioned         KVM
  56a9ecda-2402-427f-92d1-7f9be57dba36    pcacn003   Provisioned         KVM

PCA-ADMIN> provisioningLock id=363a26f4-fa34-4e4c-8e17-a1671a0b77d1

PCA-ADMIN> maintenanceLock id=363a26f4-fa34-4e4c-8e17-a1671a0b77d1
```

3. Enter the upgrade command.

   Syntax (entered on a single line):

```
upgradeCN
imageLocation=<path-to-iso>
isoChecksum=<iso-file-checksum>
hostIp=<compute-node-ip>
```

   Example:

```
PCA-ADMIN> upgradeCN hostIp=100.96.2.64 \
imageLocation="http://host.example.com/pca-<version>-<build>.iso" \
isoChecksum=240420cfb9478f6fd026f0a5fa0e998e086275fc45e207fb5631e2e99732e192e8e9d1b4c7f29026f0a5f58dadc
Command: upgradeCN hostIp=100.96.2.64 imageLocation="http://host.example.com/pca-<version>-<build>.iso"
Status: Success
Time: 2021-09-26 06:35:38,884 UTC
Data:
  Service request has been submitted. Upgrade Job Id = 1630938939109-compute-7545 Upgrade Request Id =
```

4. Use the request ID and the job ID to check the status of the upgrade process.

```
PCA-ADMIN> getUpgradeJobs
  id                              upgradeRequestId                          commandName   result
  --                              ----------------                          -----------   ------
  1630938939109-compute-7545      UWS-61736806-7e5a-4648-9259-07c54c39cacb   compute       Passed
```

```
  1632850650836-platform-68465      UWS-26dba234-9b52-426d-836c-ac11f37e717f    platform      Passed
  1632849609034-kubernetes-35545    UWS-edfa3b32-c32a-4b67-8df5-2357096052bf    kubernetes    Passed

PCA-ADMIN> getupgradejob upgradeJobId=1630938939109-compute-7545
Command: getupgradejob upgradeJobId=1630938939109-compute-7545
Status: Success
Time: 2021-09-26 08:15:03,208 UTC
Data:
  Upgrade Request Id = UWS-61736806-7e5a-4648-9259-07c54c39cacb
  Name = compute
  Start Time = 2021-09-26T06:35:39
  End Time = 2021-09-26T06:45:55
  Pid = 7545
  Host = pcamn02
  Log File = /nfs/shared_storage/pca_upgrader/log/pca-upgrader_compute_2021_09_26-06.35.39.log
  Arguments = {"verify_only":false,"upgrade":false,"diagnostics":false,"host_ip":"100.96.2.64","result_over
  Status = Passed
  Execution Time(sec) = 616
  Tasks 1 - Name = Copy Scripts
  Tasks 1 - Description = Copy scripts to shared storage
  Tasks 1 - Time = 2021-09-26T06:35:39
[...]
```

5. When the compute node upgrade has completed successfully and the node has rebooted, release the locks.

    For more information, refer to Section 2.2, "Performing Compute Node Operations".

```
PCA-ADMIN> maintenanceUnlock id=363a26f4-fa34-4e4c-8e17-a1671a0b77d1
PCA-ADMIN> provisioningUnlock id=363a26f4-fa34-4e4c-8e17-a1671a0b77d1
```

6. Proceed to the next compute node and repeat this procedure.

Chapter 8 Disaster Recovery

Table of Contents

This chapter explains how an administrator configures disaster recovery so that two Oracle Private Cloud Appliance systems in different physical locations operate as each other's fallback in case an outage occurs at one site.

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*.

8.1 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.

8.1.1 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:

• Region ID

  Each Oracle Private Cloud Appliance system is defined as an independent region. You need the region IDs of the two systems that will become each other's fallback. To retrieve the region ID, use the `show pcaSystem` command in the Service CLI.

• 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.

To complete the Oracle Site Guard configuration, you need the following details:

• The endpoints of both Oracle Private Cloud Appliance systems, where API calls are received. These are URIs are formatted as follows:

```
http(s)://<myRegion>.<myDomain>:<port>
```

The required ports are 30004 (https) and 30005 (http).

For example:

```
https://myprivatecloud.example.com:30004
http://myprivatecloud.example.com:30005
```

- An administrative user name and password for authentication with the Oracle Private Cloud Appliance services and authorization of the disaster recovery API calls. These credentials are securely stored within Oracle Enterprise Manager.

8.1.2 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.

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 Oracle 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 Oracle 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.

8.1.3 Setting Up Peering Between the ZFS Storage Appliances

Once 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 they operate as each other's replica.

The replication parameters for standard storage are mandatory with the setup command. If your Oracle 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, re-run the setup command to add the high-performance storage pool at a later time, after it has been configured on the ZFS storage appliances.

> **Caution**
>
> Verify all parameters carefully before executing the command. They cannot be changed or corrected once the configuration scripts are launched.

Syntax (entered on a single line):

```
drSetupService
region=<primary_system_region_id>
localIp=<primary_system_standard_replication_ip>
remoteIp=<replica_system_standard_replication_ip>
localIpPerf=<primary_system_performance_replication_ip>
remoteIpPerf=<replica_system_performance_replication_ip>
```

Examples:

- With only standard storage configured:

  system 1

  ```
  PCA-ADMIN> drSetupService region=pca1region \
  localIp=10.50.7.31 remoteIp=10.50.7.33
  ```

  system 2

  ```
  PCA-ADMIN> drSetupService region=pca2region \
  localIp=10.50.7.33 remoteIp=10.50.7.31
  ```

- With both standard and high-performance storage configured:

  system 1

  ```
  PCA-ADMIN> drSetupService region=pca1region \
  localIp=10.50.7.31 remoteIp=10.50.7.33 \
  localIpPerf=10.50.7.32 remoteIpPerf=10.50.7.34
  ```

  system 2

  ```
  PCA-ADMIN> drSetupService region=pca2region \
  localIp=10.50.7.33 remoteIp=10.50.7.31 \
  localIpPerf=10.50.7.34 remoteIpPerf=10.50.7.32
  ```

The script configures both ZFS storage appliances. After successful configuration of the replication interfaces, two more actions must be performed:

- generating CA certificates and uploading them to the ZFS storage appliance

- enabling replication over the interfaces you just configured

Configuring ZFS Storage Appliance Certificates

The configuration of the replication interfaces for disaster recovery introduces new IP addresses to the core configuration of the system. To allow these new hosts to be authenticated and authorized correctly, an updated CA certificate must be uploaded to the ZFS storage appliances.

After configuring the replication interfaces, but *before* enabling replication between the two storage appliances, generate a new CA certificate and upload it to the storage appliances.

1. Log in to one of the management nodes.

2. Execute the following commands:

```
# python3 -c "from pca_foundation.manager.mgmt_node import upload_system_cert; upload_system_cert()"
# /var/lib/pca-foundation/scripts/pca_factory_init.py --restart-http-zfs
```

Enabling Replication for Disaster Recovery

> **Caution**
>
> Before enabling replication, please refer to the *Oracle Private Cloud Appliance Release Notes*. Make sure you have read and understand this section: Enabling DR Replication Fails on Second System.

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 region=pca1region \
    localIp=10.50.7.31 remoteIp=10.50.7.33 \
    enableReplication=True remotePassword=********
    ```

    system 2

    ```
    PCA-ADMIN> drSetupService region=pca2region \
    localIp=10.50.7.33 remoteIp=10.50.7.31 \
    enableReplication=True remotePassword=********
    ```

*   With both standard and high-performance storage configured:

    system 1

    ```
    PCA-ADMIN> drSetupService region=pca1region \
    localIp=10.50.7.31 remoteIp=10.50.7.33 \
    localIpPerf=10.50.7.32 remoteIpPerf=10.50.7.34 \
    enableReplication=True remotePassword=********
    ```

    system 2

    ```
    PCA-ADMIN> drSetupService region=pca2region \
    localIp=10.50.7.33 remoteIp=10.50.7.31 \
    localIpPerf=10.50.7.34 remoteIpPerf=10.50.7.32 \
    enableReplication=True remotePassword=********
    ```

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 Section 8.2, "Managing Disaster Recovery Configurations".

8.2 Managing Disaster Recovery Configurations

This section explains how to configure disaster recovery settings on the two Oracle Private Cloud Appliance systems you intend to set up as each other's fallback.

8.2.1 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
   ```

8.2.2 Adding Instances to a DR Configuration

Once a DR configuration has been created, 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 Oracle 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
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 conf
  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 [63b36a
```

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=ocid1.instance.....<region1>...instance1_uniqueID
Command: drGetComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d instanceId=ocid1.instance
Status: Success
Time: 2021-08-17 08:34:42,413 UTC
Data:
  Type = ComputeInstance
```

```
Compartment Id = ocid1.compartment........uniqueID
Boot Volume Id = ocid1.bootvolume........uniqueID
Compute Instance Shape = VM.PCAStandard1.8
Work State = Normal
```

8.2.3 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=ocid1.instance.....<region1>...instance3_uniqueID
    Command: drRemoveComputeInstance drConfigId=63b36a80-7047-42bd-8b97-8235269e240d instanceId=ocid1.instance.
    Status: Success
    Time: 2021-08-17 08:45:59,718 UTC
    Data:
      Message = Successfully started job to remove instance ocid1.instance.....<region1>...instance3_uniqueID f
      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 conf
    ```

4. Repeat the drRemoveComputeInstance command with the OCIDs of all the compute instances that
    you want to remove from the DR configuration.

8.2.4 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 "ocid1.compartment.....uniqueID"`

     - `oci network nsg list --compartment-id "ocid1.compartment.....uniqueID"`

2. Add a site mapping to the DR configuration with the drAddSiteMapping command.

   Syntax (entered on a single line):

   ```
   drAddSiteMapping
   drConfigId=<DR_configuration_id>
   objType=compartment|subnet|network_security_group
   sourceId=<source_object_OCID>
   targetId=<target_object_OCID>
   ```

   Examples:

   ```
   PCA-ADMIN> drAddSiteMapping \
   drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
   objType=compartment \
   sourceId="ocid1.compartment.....<region1>...uniqueID" \
   targetId="ocid1.compartment.....<region2>...uniqueID"
   Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=compartment sourceId=
   Status: Success
   Time: 2021-08-17 09:07:24,957 UTC
   Data:
     9244634e-431f-43a1-89ab-5d25905d43f9

   PCA-ADMIN> drAddSiteMapping \
   drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
   objType=subnet \
   sourceId="ocid1.subnet.....<region1>...uniqueID" \
   targetId="ocid1.subnet.....<region2>...uniqueID"
   Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=compartment sourceId=
   Status: Success
   Time: 2021-08-17 09:07:24,957 UTC
   Data:
     d1bf2cf2-d8c7-4271-b8b6-cdf757648175
   ```

```
PCA-ADMIN> drAddSiteMapping \
drConfigId=63b36a80-7047-42bd-8b97-8235269e240d \
objType=network_security_group \
sourceId="ocid1.nsg.....<region1>...uniqueID" \
targetId="ocid1.nsg.....<region2>...uniqueID"
Command: drAddSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d objType=compartment sourceId="oci
Status: Success
Time: 2021-08-17 09:07:24,957 UTC
Data:
  422f8892-ba0a-4a89-bc37-61b5c0fbbbaa
```

3. Repeat the command with the OCIDs of all the source and target objects that you want to include in the site mappings of the DR configuration.

4. To display the list of site mappings included in the DR configuration, use the drGetSiteMappings command. The DR configuration ID is a required parameter.

   Syntax:

   ```
   drGetSiteMappings drConfigId=<DR_configuration_id>
   ```

   Example:

   ```
   PCA-ADMIN> drGetSiteMappings drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
   Command: drGetSiteMappings drConfigId=63b36a80-7047-42bd-8b97-8235269e240d
   Status: Success
   Time: 2021-08-17 09:19:22,580 UTC
   Data:
     id                                      name
     --                                      ----
     d1bf2cf2-d8c7-4271-b8b6-cdf757648175    null
     9244634e-431f-43a1-89ab-5d25905d43f9    null
     422f8892-ba0a-4a89-bc37-61b5c0fbbbaa    null
   ```

5. To display the status and details of a site mapping included in the DR configuration, use the drGetSiteMapping command.

   Syntax (entered on a single line):

   ```
   drGetSiteMapping
   drConfigId=<DR_configuration_id>
   mappingId=<site_mapping_id>
   ```

   Example:

   ```
   PCA-ADMIN> drGetSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=d1bf2cf2-d8c7-4271-b8
   Command: drGetSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=d1bf2cf2-d8c7-4271-b8b6
   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
   ```

8.2.5 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-
   Command: drRemoveSiteMapping drConfigId=63b36a80-7047-42bd-8b97-8235269e240d mappingId=422f8892-ba0a-4a
   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.

8.2.6 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]
```

8.2.7 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.

Using the Service CLI

1. Look up the ID of the DR configuration you want to delete (drGetConfigs).

2. 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
   ```

3. 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]
   ```