ORACLE
PeopleSoft

PeopleSoft Cloud Manager for Oracle Cloud Infrastructure

August 2025

ORACLE

Documentation Accessibility

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

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit https://docs.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit https://docs.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Contents

Preface

Understanding the PeopleSoft Online Help and PeopleBooks

The PeopleSoft Online Help is a website that enables you to view all help content for PeopleSoft applications and PeopleTools. The help provides standard navigation and full-text searching, as well as context-sensitive online help for PeopleSoft users.

Hosted PeopleSoft Online Help

You can access the hosted PeopleSoft Online Help on the Oracle Help Center. The hosted PeopleSoft Online Help is updated on a regular schedule, ensuring that you have access to the most current documentation. This reduces the need to view separate documentation posts for application maintenance on My Oracle Support. The hosted PeopleSoft Online Help is available in English only.

To configure the context-sensitive help for your PeopleSoft applications to use the Oracle Help Center, see Configuring Context-Sensitive Help Using the Hosted Online Help Website.

Locally Installed PeopleSoft Online Help

If you're setting up an on-premises PeopleSoft environment, and your organization has firewall restrictions that prevent you from using the hosted PeopleSoft Online Help, you can install the online help locally. Installable PeopleSoft Online Help is made available with selected PeopleSoft Update Images and with PeopleTools releases for on-premises installations, through the Oracle Software Delivery Cloud.

Your installation documentation includes a chapter with instructions for how to install the online help for your business environment, and the documentation zip file may contain a README.txt file with additional installation instructions. See *PeopleSoft 9.2 Application Installation* for your database platform, "Installing PeopleSoft Online Help."

To configure the context-sensitive help for your PeopleSoft applications to use a locally installed online help website, see Configuring Context-Sensitive Help Using a Locally Installed Online Help Website.

Downloadable PeopleBook PDF Files

You can access downloadable PDF versions of the help content in the traditional PeopleBook format on the Oracle Help Center. The content in the PeopleBook PDFs is the same as the content in the PeopleSoft Online Help, but it has a different structure and it does not include the interactive navigation features that are available in the online help.

Common Help Documentation

Common help documentation contains information that applies to multiple applications. The two main types of common help are:

- Application Fundamentals

- Using PeopleSoft Applications

Most product families provide a set of application fundamentals help topics that discuss essential information about the setup and design of your system. This information applies to many or all applications in the PeopleSoft product family. Whether you are implementing a single application, some combination of applications within the product family, or the entire product family, you should be familiar with the contents of the appropriate application fundamentals help. They provide the starting points for fundamental implementation tasks.

In addition, the *PeopleTools: Applications User's Guide* introduces you to the various elements of the PeopleSoft Pure Internet Architecture. It also explains how to use the navigational hierarchy, components, and pages to perform basic functions as you navigate through the system. While your application or implementation may differ, the topics in this user's guide provide general information about using PeopleSoft applications.

Field and Control Definitions

PeopleSoft documentation includes definitions for most fields and controls that appear on application pages. These definitions describe how to use a field or control, where populated values come from, the effects of selecting certain values, and so on. If a field or control is not defined, then it either requires no additional explanation or is documented in a common elements section earlier in the documentation. For example, the Date field rarely requires additional explanation and may not be defined in the documentation for some pages.

Typographical Conventions

The following table describes the typographical conventions that are used in the online help.

Typographical Convention	***Description***
Key+Key	Indicates a key combination action. For example, a plus sign (+) between keys means that you must hold down the first key while you press the second key. For **Alt+W,** hold down the **Alt** key while you press the **W** key.
. . . (ellipses)	Indicate that the preceding item or series can be repeated any number of times in PeopleCode syntax.
{ } (curly braces)	Indicate a choice between two options in PeopleCode syntax. Options are separated by a pipe (|).
[ ] (square brackets)	Indicate optional items in PeopleCode syntax.
& (ampersand)	When placed before a parameter in PeopleCode syntax, an ampersand indicates that the parameter is an already instantiated object. Ampersands also precede all PeopleCode variables.

Typographical Convention	*Description*
⇒	This continuation character has been inserted at the end of a line of code that has been wrapped at the page margin. The code should be viewed or entered as a single, continuous line of code without the continuation character.

ISO Country and Currency Codes

PeopleSoft Online Help topics use International Organization for Standardization (ISO) country and currency codes to identify country-specific information and monetary amounts.

ISO country codes may appear as country identifiers, and ISO currency codes may appear as currency identifiers in your PeopleSoft documentation. Reference to an ISO country code in your documentation does not imply that your application includes every ISO country code. The following example is a country-specific heading: "(FRA) Hiring an Employee."

The PeopleSoft Currency Code table (CURRENCY_CD_TBL) contains sample currency code data. The Currency Code table is based on ISO Standard 4217, "Codes for the representation of currencies," and also relies on ISO country codes in the Country table (COUNTRY_TBL). The navigation to the pages where you maintain currency code and country information depends on which PeopleSoft applications you are using. To access the pages for maintaining the Currency Code and Country tables, consult the online help for your applications for more information.

Region and Industry Identifiers

Information that applies only to a specific region or industry is preceded by a standard identifier in parentheses. This identifier typically appears at the beginning of a section heading, but it may also appear at the beginning of a note or other text.

Example of a region-specific heading: "(Latin America) Setting Up Depreciation"

Region Identifiers

Regions are identified by the region name. The following region identifiers may appear in the PeopleSoft Online Help:

- Asia Pacific
- Europe
- Latin America
- North America

Industry Identifiers

Industries are identified by the industry name or by an abbreviation for that industry. The following industry identifiers may appear in the PeopleSoft Online Help:

- USF (U.S. Federal)

- E&G (Education and Government)

Translations and Embedded Help

PeopleSoft 9.2 software applications include translated embedded help. With the 9.2 release, PeopleSoft aligns with the other Oracle applications by focusing our translation efforts on embedded help. We are not planning to translate our traditional online help and PeopleBooks documentation. Instead we offer very direct translated help at crucial spots within our application through our embedded help widgets. Additionally, we have a one-to-one mapping of application and help translations, meaning that the software and embedded help translation footprint is identical—something we were never able to accomplish in the past.

Using and Managing the PeopleSoft Online Help

Select About This Help in the left navigation panel on any page in the PeopleSoft Online Help to see information on the following topics:

- Using the PeopleSoft Online Help.
- Managing hosted Online Help.
- Managing locally installed PeopleSoft Online Help.

PeopleSoft Cloud Manager Related Links

PeopleSoft Cloud Manager Home Page

PeopleSoft Hosted Online Help

PeopleSoft Information Portal

PeopleSoft Spotlight Series

My Oracle Support

Oracle Help Center

Contact Us

Send your suggestions to psoft-infodev_us@oracle.com.

Please include the applications update image or PeopleTools release that you're using.

Follow Us

Icon	*Link*
	Watch PeopleSoft on YouTube
	Follow @PeopleSoft_Info on X.
	Read PeopleSoft Blogs
	Connect with PeopleSoft on LinkedIn

Chapter 1

Getting Started with PeopleSoft Cloud Manager

Understanding PeopleSoft Cloud Manager on Oracle Cloud Infrastructure

PeopleSoft Cloud Manager is an orchestration framework to provision and manage PeopleSoft environments on Oracle Cloud Infrastructure (OCI). The PeopleSoft Cloud Manager can help creating task specific environments that can last as long as the task is needed. PeopleSoft Cloud Manager will enable you to focus more on business and less on infrastructure management by taking away all the complexities involved in acquiring and managing the infrastructure to run PeopleSoft on OCI.

PeopleSoft Cloud Manager is an application available on the Oracle Cloud Marketplace. Any existing PeopleSoft customer can use it by taking advantage of the Oracle Cloud Service resources.

OCI is a set of complementary cloud services that enable you to build and run a wide range of applications and services in a highly available hosted environment.

Common Abbreviations

Term	***Description***
DPK	PeopleSoft Deployment Packages
PCM	PeopleSoft Cloud Manager
PI	PeopleSoft Update Image
PRP	PeopleSoft Release Patchset
PUM	PeopleSoft Update Manager
OCI	Oracle Cloud Infrastructure
AD	Availability Domain
VCN	Virtual Cloud Network

Term	*Description*
TDE	Transparent Data Encryption
OCID	Oracle Cloud ID

Requirements for PeopleSoft Cloud Manager

To review the requirements for PeopleSoft Cloud Manager:

- See the PeopleTools support information on PeopleSoft Cloud Manager Home Page (My Oracle Support, Doc ID 2231255.2).
- See the requirements for Oracle Cloud Infrastructure accounts, subscriptions, tenancy, and compartments in the Cloud Manager Installation tutorials at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

PeopleSoft Cloud Manager – An Overview

Cloud Manager provides a framework for customers to provision and administer the life cycle of PeopleSoft environments on OCI. Cloud Manager brings in the agility to rapidly bring up PeopleSoft environments on demand, based on your infrastructure requirements.

Features of PeopleSoft Cloud Manager

PeopleSoft Cloud Manager provides the ability to:

- Provision PeopleSoft environments on OCI.
- Automate migration of on-premises environment to OCI.
- Support lifting and shifting of Unicode or non-Unicode database.
- Support lifting and shifting of PeopleSoft application environments that have TDE encrypted databases. The on-premises environments must be TDE encrypted before migrating.
- Support lifting and shifting of Transparent Data Encryption, or TDE, and Real Application Clusters, or RAC, for Database Systems.
- Orchestrate deployment of PeopleSoft 9.2 and IH 9.1 applications on OCI.
- Take advantage of the subscription model to auto download application PeopleSoft Update Images and PRPs.
- Create repeatable deployment templates.
- Perform self-service provisioning of PeopleSoft environments.
- Fully automate deployment that is immune to manual errors and process delays.
- Manage multiple environments from a single page.
- Enable application lifecycle management in Oracle Cloud.

- Clone environments from running instances.
- Access log files through UI for easy troubleshooting.
- View the status of environment provisioning.
- Automate PRP updates for a Cloud Manager instance.
- Define and configure web, app domains in topology/template definitions.
- Refresh DB System environments managed by Cloud Manager including PS_APP_HOME, PS_CUST_HOME, and database, or just the database, from a backup.
- Automate PeopleTools upgrade and updates.
- Perform self update of Cloud Manager to the most current update image from the prior update image.
- Enable creating or reusing file system service (FSS) and mount target from PeopleSoft Cloud Manager UI.
- Import environments from PeopleSoft environments that are running on Oracle Cloud.
- Use Zero Downtime Migration (ZDM) to migrate database and middle tier into Cloud Manager.
- Use multiple PeopleSoft Client environments that are supported.
- Use Windows middle tier with Process Scheduler to support nVision.
- Take advantage of Database systems that support multiple Oracle Database versions.
- Take advantage of Database systems that support Exadata shapes.
- Stop and start environment in OCI from Cloud Manager. Customer is not billed for an environment while it is stopped.
- Back up an environment and restore an environment from the backup from Cloud Manager UI.
- Share PS_HOME, PS_APP_HOME and PS_CUST_HOME in multiple middle tiers using File System Service.
- Resume provisioning when a recoverable failure occurs.
- Take advantage of Auto scale support using Oracle Data Science.
- Define and manage policies for all managed environments.
- Apply critical patch updates to the infrastructure components such as Java, Tuxedo, WebLogic, and Oracle Database client.

Chapter 2

Configuring Cloud Manager

Configuring Cloud Manager

Installation documentation for OCI is posted on the PeopleSoft Cloud Manager Home Page (My Oracle Support DOC ID: 2231255.2), Installation and Implementation tab. PeopleSoft Cloud Manager Home Page

Pages Used to Configure Cloud Manager for OCI

Page Name	*Definition Name*	*Usage*
Cloud Manager Settings tile	ECL_CMCONFG_FL_GBL (Content reference for the tile.)	To access the Cloud Manager Settings page.
Cloud Manager Settings page	ECL_CMCFG_OCI_FL	To change the system settings as per requirements in OCI.
Infrastructure Settings page	ECL_OCICFG_OCI_FL	To configure OCI-related settings for environment provisioning and management.
File Server Configuration page	ECL_CMFILESERV_FL	To configure file server as repository for Cloud Manager in OCI.
Manage PUM Connections	ECL_CMUPDATE_FL	To configure a PUM source for updating the Cloud Manager application.
Manage Updates	ECL_CMSELFUPD_FL	To manage application updates delivered through PeopleSoft IH Updates and PRPs in OCI.
Logs	ECL_CM_FSLOGS_FL	To view Cloud Manager logs.
Data Science Settings	ECL_CM_DS_SETNG_FL	To configure the Auto Scaling feature.
AutoScale Settings	ECL_MLNOTF_SET_FL	To control notifications for monitoring and scaling.
Advisory Settings	ECL_CMRECOM_SET_FL	To enable advisories.

Page Name	*Definition Name*	*Usage*
Role Based Security	ECL_ROL_BAS_FL	To delegate access based on tags.

Cloud Manager Settings Tile

Use the Cloud Manager Settings tile (ECL_CMCONFG_FL_GBL) to access the Cloud Manager Settings page.

Note: Only a Cloud Manager Administrator can view this tile on the Cloud Manager home page.

Navigation:

The Cloud Manager Settings tile is delivered as part of the Cloud Manager home page.

This example illustrates the Cloud Manager Settings tile.

Configuring Cloud Manager Settings for OCI

The steps involved in Cloud Manager configuration for OCI are:

- Configuring Cloud Manager settings (Required).
- Configuring Infrastructure settings (Required).
- Configuring File Server (Required).
- Managing Updates.
- Data Science settings.
- Auto Scale Settings.
- Advisory settings.
- Role Based Security settings.

Cloud Manager Settings Page

Use the Cloud Manager Settings page (ECL_CMCFG_OCI_FL) to change the system settings as per requirements.

Navigation:

Click the Cloud Manager Settings tile on the delivered Cloud Manager Fluid home page. Cloud Manager Settings page is displayed. By default, the details that were provided during Cloud Manager bootstrap process are displayed.

This example illustrates the fields and controls on the Cloud Manager Settings page. You can find definitions for the fields and controls later on this page.

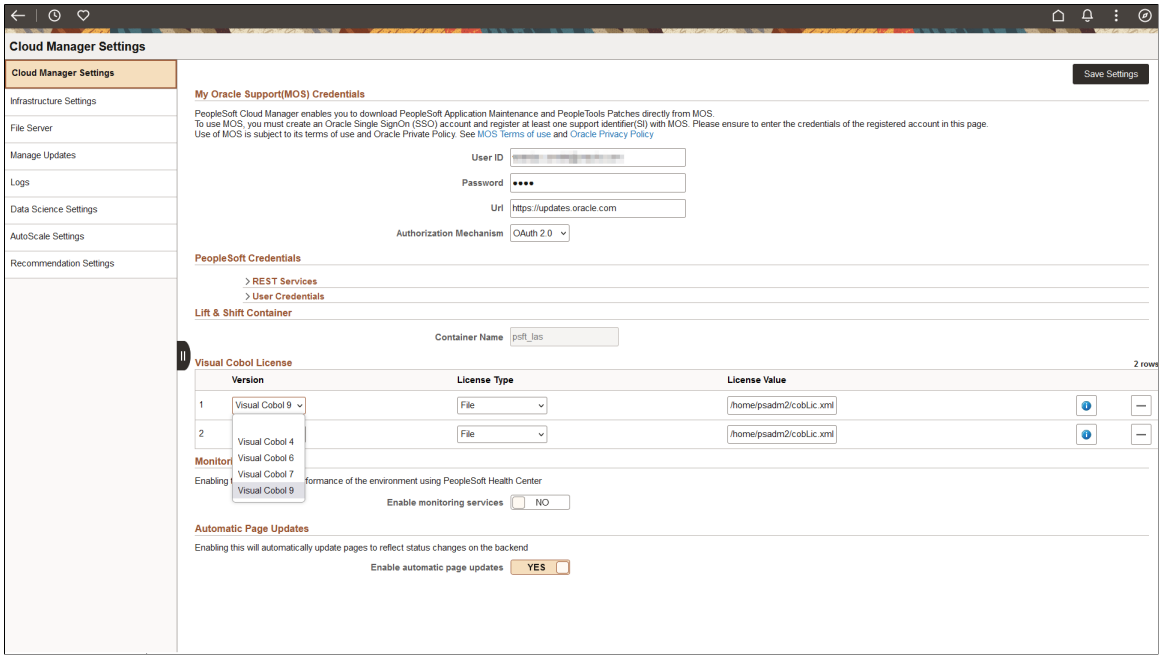

My Oracle Support (MOS) Credentials

This refers to My Oracle Support (MOS) user name and password inputs. Using this credential, Cloud Manager downloads the required updates, PIs and PRPs from MOS. The MOS credentials were provided through the Resource Manager Stack.

Field or Control	***Description***
User ID	Enter the user ID for your My Oracle Support account.
URL	Enter the URL: https://updates.oracle.com.
Password	Enter the password for your My Oracle Support account.

Field or Control	*Description*
Authorization Mechanism	The available authorization mechanisms are OAuth 2.0 and BasicAuth. OAuth 2.0 is the preferred authorization mechanism.

Note: Read the MOS License information. Click the links to understand My Oracle Support terms of use and privacy policy.

PeopleSoft Credentials for REST Services

REST services are standard IB REST services available in the Cloud Manager instance. These REST services are used internally by Cloud Manager modules to send/receive the results of long-running, asynchronous activities.

Important! User credentials must be manually updated on the Cloud Manager instance before updating here. Updating credentials here does not update the Cloud Manager instance.

Field or Control	*Description*
User Name	Enter the delivered Cloud Manager Administrator user name.
Password	Enter the Cloud Manager Administrator password.

Expand the User Credentials section and enter all the necessary passwords.

This example illustrates the fields for REST Services - User Credentials.

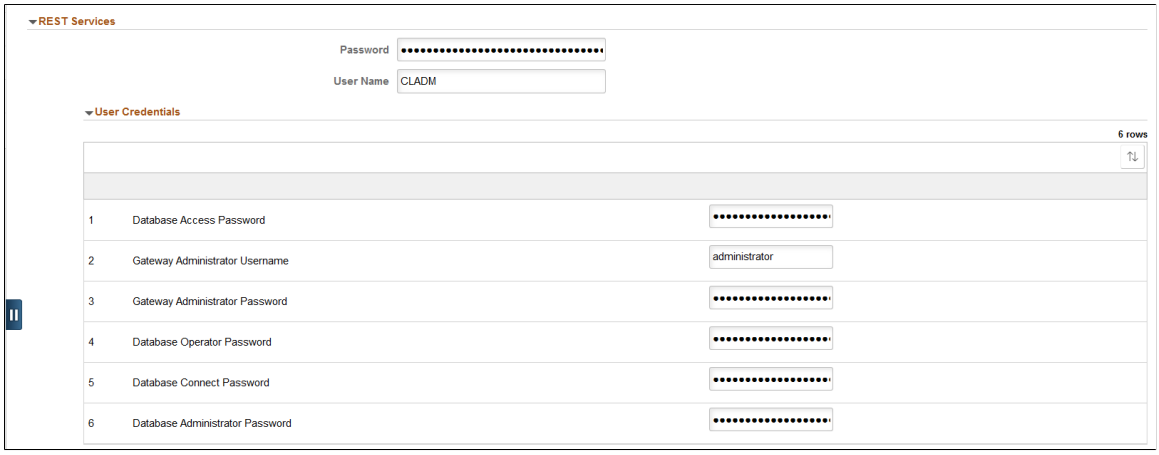

Lift and Shift Container

This section refers to the Oracle Cloud Storage Container name in which the lifted DPKs (Lifted DPK means migrated environment from your on premise environment through Lift process) are stored. It is from this container that the list of lifted environments are displayed on the Lift and Shift page.

Field or Control	*Description*
Container Name	Displays the container name. In the current version of Cloud Manager this name cannot be changed.

Visual COBOL License

Starting with PeopleTools 8.58 Visual COBOL is the only supported COBOL compiler for Cloud Manager environments.

Use this section to provide COBOL license details. Cloud Manager supports up to two COBOL versions. COBOL installation is enabled on the topology by selecting COBOL field value as *Yes* in the Features section of Edit Node modal window. For details on topology, see Adding Nodes with COBOL Enabled. To enable COBOL in the template, the topology for the template must have COBOL enabled. See Configuring Custom Attributes.

Note: Oracle is the exclusive reseller of the Rocket Software Visual COBOL compiler for use with PeopleSoft applications. Contact your Oracle sales representative for a license.

Visual COBOL

Field or Control	*Description*
Version	Select the COBOL version. Visual COBOL 4, Visual COBOL 6, Visual COBOL 7, and Visual COBOL 9 are supported. **Note:** In order to be applied, the selected Visual COBOL version must be supported by the PeopleTools version. Visual COBOL versions 6, 7, and 9 are supported for PeopleTools 8.61.
License Type	License Type can be: • Authorization Code • File • Server

Field or Control	*Description*
License Value	The license depends on the license type: • Authorization Code Enter the authorization code. • File The license file must be copied to the Cloud Manager VM and placed in a location that is accessible to psadm2 user. Enter the path to the license file. • Server Enter the hostname or IP address of the license server. The server must be accessible from the current machine.

Monitoring Services

Select Yes to enable performance monitoring of the environment using PeopleSoft Health Center. The health of your PeopleSoft application is determined by providing historical and real-time analysis of performance and load.

Note: Monitoring must be enabled to use auto scaling with Data Science.

Automatic Page Updates

Select Yes to enable automatic page updates to reflect status changes on the back-end.

Infrastructure Settings Page

Use Cloud Manager Settings – Infrastructure Settings page (ECL_OCICFG_OCI_FL) to configure OCI related settings for instance provisioning and management.

Navigation:

Click the Cloud Manager Settings tile on the delivered Cloud Manager Fluid home page. Cloud Manager Settings page is displayed. On the Cloud Manager Settings page, click the Infrastructure Settings link displayed on the left panel.

This example illustrates the fields and controls on the Infrastructure Settings page. You can find definitions for the fields and controls later on this page.

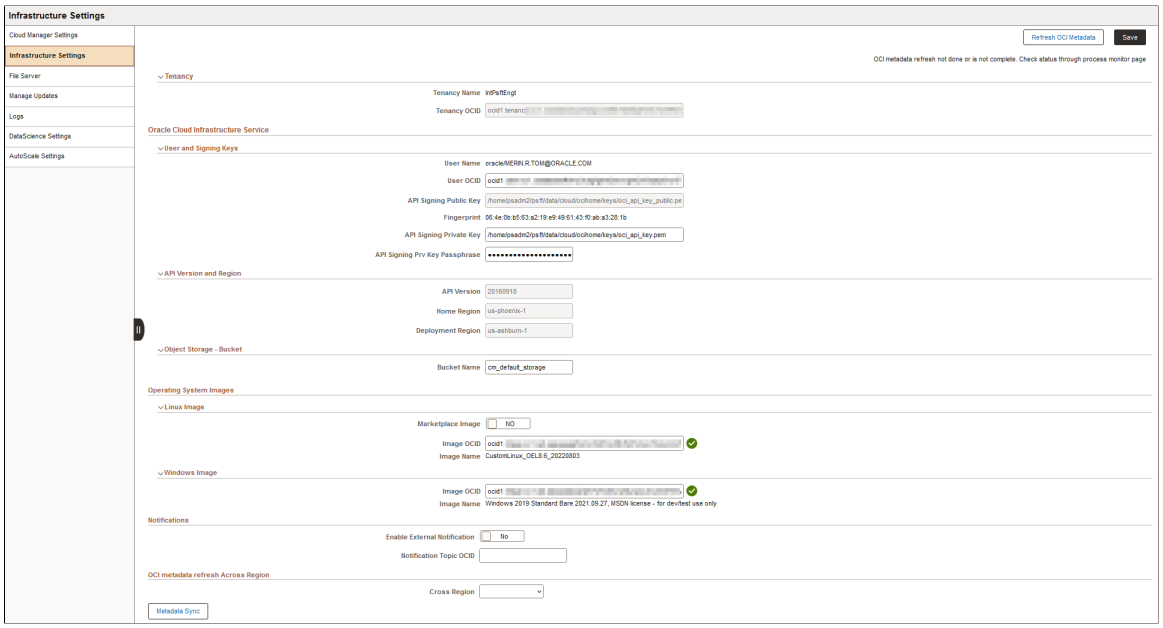

Field or Control	***Description***
Tenancy OCID	Unique Oracle Cloud Identifier (OCID) for the tenancy. Tenancy is the root compartment that contains all your organization's compartment and other OCI Cloud resources. If you use the Oracle Cloud Infrastructure API, you will need your tenancy's OCID in order to sign the API requests. You will also use the tenancy ID in some of the IAM API operations. You can find your tenancy's OCID displayed at the bottom of the Oracle Cloud Infrastructure Console pages. See Locating OCI Credentials.
User OCID	Unique OCID for the user. You can find the user's OCID in the Oracle Cloud Infrastructure Console page showing the user's details. See Locating OCI Credentials.

Field or Control	*Description*
API Signing Public Key and **API Signing Private Key**	RSA key pair in PEM format. Your API requests will be signed with your private key, and Oracle Cloud Infrastructure will use the public key to verify the authenticity of the request. **Note:** For details on the creation and usage of the API signing keys, refer the *PeopleSoft Cloud Manager Installation tutorials*. **Important!** It is not recommended to modify these values without completely understanding the impact. If in case the public keys are required to be changed, then manually update the public keys for the user using the OCI Console.
API Signing Prv Key Passphrase	Displays the API signing private key encrypted with a passphrase.
API Version	API version is the Rest API version for OCI. The base path of the endpoint includes the desired API version (for example, 20160918).
Home Region	When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region. Your home region is where your IAM resources are defined. When you subscribe to a new region, your IAM resources are replicated in the new region, however, the definitions reside in your home region and can only be changed there.
Deployment Region	The region where the PeopleSoft environments will be provisioned by Cloud Manager. Cloud Manager and the file server instance also reside on this same region.
Object Storage Bucket	Accept the default name or enter a new name. The bucket is required for compare reports generated when upgrading the PeopleTools version of a provisioned environment.
Save	Click the **Save** button to save your settings.
Refresh OCI Metadata	Once all the Infrastructure settings are entered and saved, click the **Refresh OCI Metadata** button. When this button is clicked, the Cloud Manager will run a process scheduler job (Process Name: ECL_OCI_SYNC) which will fetch all the OCI-specific metadata required for the Cloud Manager to function properly.

Operating System Image

This refers to OS images in Oracle Cloud that CM uses to provision VMs during environment creation. Refer to the Cloud Manager Installation tutorials at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

.

For Linux Image, you can either:

- Obtain the Oracle Linux Image customized for PeopleSoft from Marketplace.

  See tutorial *Prepare to Install PeopleSoft Cloud Manager*.

- Create a custom Linux Image for PeopleSoft Cloud Manager.

  See tutorial *Create a Custom Linux Image for PeopleSoft Cloud Manager*.

For Windows Image, you can either:

- Use an Oracle platform image for Microsoft Windows for PeopleSoft Cloud Manager.

  See tutorial *Prepare to Install PeopleSoft Cloud Manager*.

- Create a custom Windows image for PeopleSoft Cloud Manager in Oracle Cloud Infrastructure.

  See tutorial *Create a Custom Windows Image for PeopleSoft Cloud Manager in Oracle Cloud Infrastructure*.

This example illustrates the fields and controls on the Operating System Image. You can find definitions for the fields and controls later on this page.

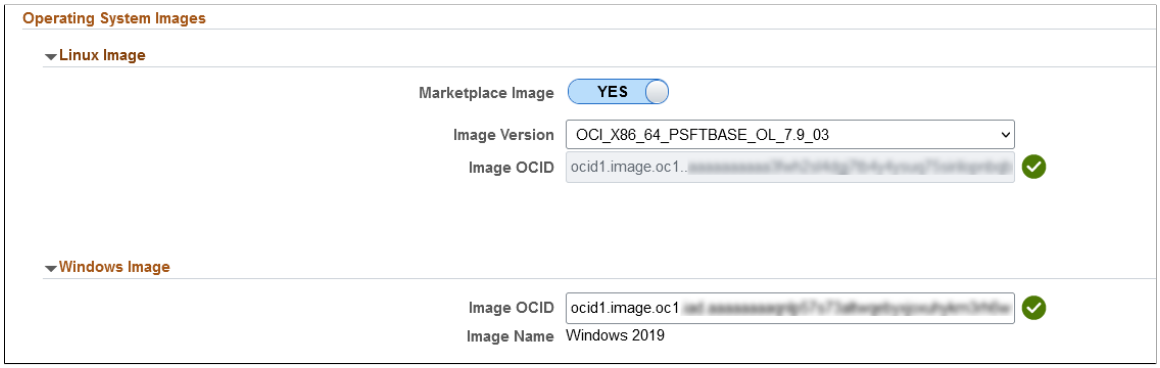

Field or Control	***Description***
Linux Image **Image OCID**	Select whether or not the Linux Image was obtained from Marketplace. The OCID is automatically populated if the image is obtained from Marketplace. If a custom image is used enter the OCID for the Linux Image. See tutorial *Create a Custom Linux Image for PeopleSoft Cloud Manager*.

Field or Control	*Description*
Windows Image **Image OCID**	Enter the OCID for the Windows Image. The image name will be displayed.

Notifications

When using the Upgrade PeopleTools feature in the current Cloud Manager update image, users can view compare reports such as DDDAUDIT, SYSAUDIT and Alter Audit. Users can choose to receive email notification when compare reports are available.

In order to enable notification, the Cloud Manager administrator must complete the prerequisites. Refer to the *Prepare to Upgrade PeopleTools for a Provisioned Environment* section in the tutorial *Prepare to Install PeopleSoft Cloud Manager* at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

This example illustrates the Notifications section on the Infrastructure Settings page. You can find definitions for the fields and controls later on this page.

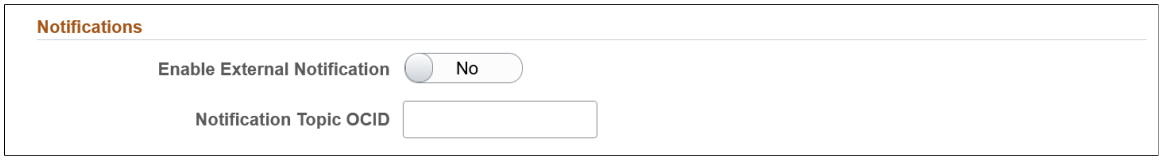

Field or Control	*Description*
Enable External Notification	Select Yes to allow end users performing a PeopleTools upgrade to enable report notifications on the Cloud Manager Upgrade PeopleTools page for an environment. This applies to upgrades to PeopleTools 8.60.x. See Upgrading PeopleTools. This enables Cloud Manager to send you weekly recommendation email. This also enables Cloud Manager to send daily and weekly summary email for scaling and monitoring events, based on the Auto Scale Settings. See Configuring AutoScale Settings

Field or Control	*Description*
Notification Topic OCID	Enter the OCID for the topic created in Oracle Cloud Infrastructure Notification Service. **Note:** Creating the topic is one of the prerequisites listed in the *Prepare to Upgrade PeopleTools for a Provisioned Environment* section in the tutorial *Prepare to Install PeopleSoft Cloud Manager* at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials

OCI Metadata Refresh Across Regions

OCI provides a disaster recovery service to ensure higher availability of an application by switching over to an identical copy of the application instance in another region. The region where an application instance runs is called a Primary region and the region to which a switchover potentially happens is called a Standby region.

See Enabling Disaster Recovery.

This example illustrates the OCI Metadata Refresh Across Regions section on the Infrastructure Settings page. You can find definitions for the fields and controls later on this page.

Field or Control	*Description*
Enable Disaster Recovery	Select Yes to enable the creation of a standby region for a primary region to ensure business continuity in the event of an outage.
Standby Region	Select the standby region to which you need to migrate the components of an application such as database, middle-tier and load balancer.

File Server Page

Use Cloud Manager Settings – File Server page (ECL_CMFILESERV_FL) to configure file server as repository for Cloud Manager.

Use the Cloud Manager File Server page to select or configure a File Storage Service (FSS) file system.

See tutorial *Use File Storage Service for PeopleSoft Cloud Manager Repository*.

The following use cases apply while creating the file server:

- Create a new file server with new Mount Target
- Create file system and export with existing Mount Target
- Use existing file system (Mount Target and Export Path)
- Upgrade existing file server to FSS in upgrade environment

The following ports need to be opened in the FSS mount target's subnet to allow NFS connections:

Field or Control	***Description***
TCP Ports	111 2048 2049 2050
UDP Ports	111 2048

Navigation:

Click the Cloud Manager Settings tile on the Cloud Manager home page. On the Cloud Manager Settings page, click the File Server link displayed on the left panel.

This example illustrates the fields and controls on the File Server page. You can find definitions for the fields and controls later on this page.

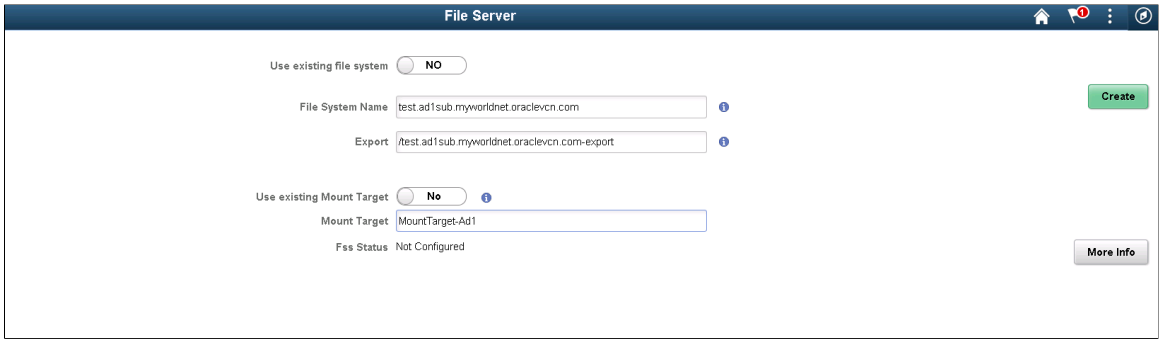

The File Server settings provides the options to set up a new file system.

By default, options *Use existing file system* and *Use existing Mount Target* are set to No.

Field or Control	*Description*
Use exiting file system	Set to No when creating a new file system.
File System Name	Name for the file system to be created.
Export Path	Path for instances to mount the file system through mount target.
Use existing Mount Target	Set to No when creating a new file system.
Mount Target	IP address or DNS name that is used in the mount command to connect NFS clients to a file system. A single mount target can export many file systems. By default, you can create two mount targets per account per availability domain.
Create button	Once the inputs are provided, click the Create button to create the file system, mount target and export path.

Creating a New File System, Mount Target and Export

To create a new File system, mount target and export:

1. Enter the File System Name.
2. Enter the Export Path.
3. Enter the Mount Target.
4. Select No for Use existing file system.
5. Select No for Use existing Mount Target.
6. Click Create.

When the file system becomes available, you can subscribe to your desired download channels in the repository.

Creating New File System and Export with Existing Mount Target

As a mount target can export many file systems, a new File System can be created using an existing Mount Target.

To create a new file system and export with existing mount:

1. Select Infrastructure Settings from the left-side menu and click the Refresh OCI Metadata button.
2. Select File Server from the left-side menu.
3. Select Yes to use existing Mount Target.
4. Select the exiting Mount Target from the drop-down list.

5. Click Create.

This example illustrates the fields and controls on the File Server page for creating new file system and export with existing Mount Target.

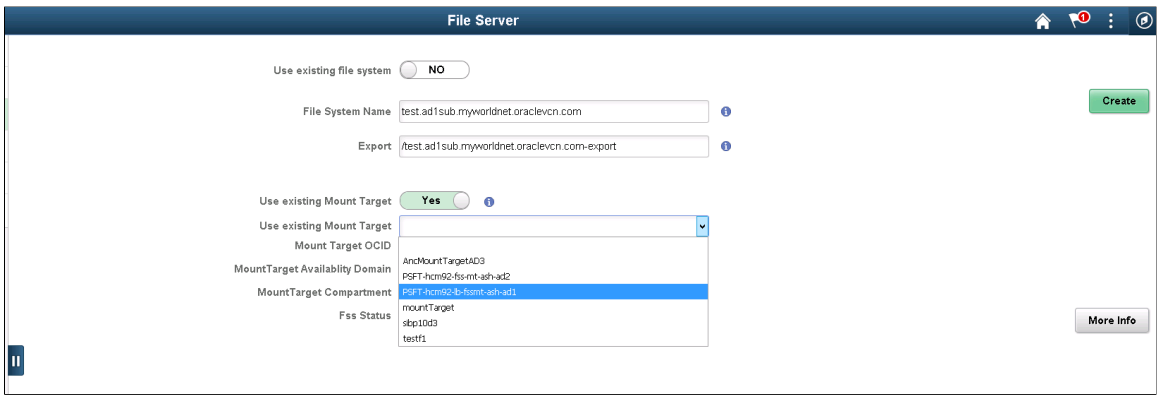

After selecting the existing mount target the availability domain and compartment are displayed.

Using Existing File System

Instead of creating a new file system, an existing file system can be used. Select Use existing file system and enter the export path for the target FSS.

To create a file server from an existing file system:

1. Select Yes to Use existing file system.
2. Enter the Existing Mount path.
3. Click Create.

This example illustrates the fields and controls on the File Server page for using existing file system.

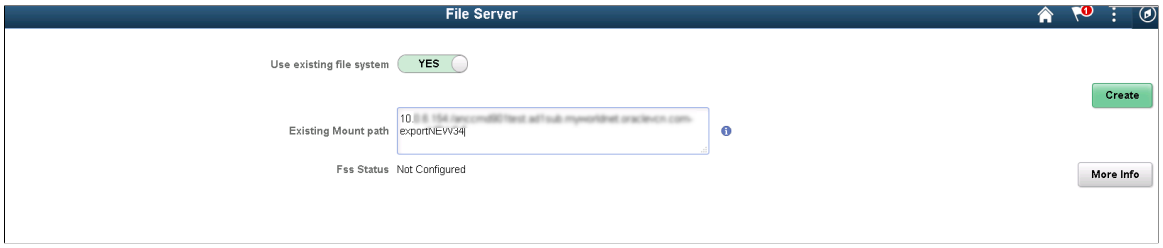

Completed FSS

This example illustrates the fields and controls on the Completed File Server page.

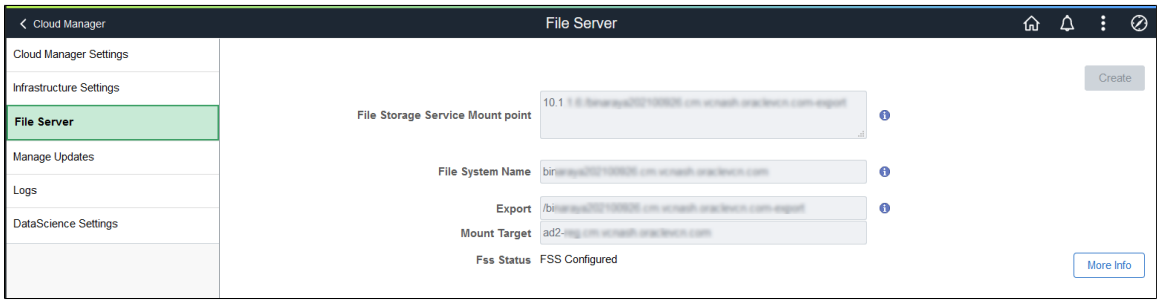

Field or Control	*Description*
File Storage Service Mount Point	IP address or DNS name that is used in the mount command to connect NFS clients to a file system.
File System Name	Name of the File System.
Export	Export path.
Mount Target	IP address or DNS name that is used in the mount command to connect NFS clients to a file system. A single mount target can export many file systems. By default, you can create two mount targets per account per availability domain.
FSS Status	File Server status. Different statuses are: Not Configured, FSS Configured, and Failed. • **Not Configured**: The FSS is not created or configured to store files. • **FSS Configured**: FSS is created and is ready to store files. • **Failed**: FSS creation failed. Check the More infos page to correct the errors and continue. Reset option can also be used to re-create the FSS.
More Info	Select to view the progress of the migration. More Info contains the list of tasks that can be continued or retried in case of a failure or error.

Manage Updates Page

Use the Manage Updates page (ECL_CMSELFUPD_FL) to apply Cloud Manager updates delivered through PeopleSoft IH Updates and PRPs.

Note: This feature is meant for the Cloud Manager application update only.

Navigation:

Click the Cloud Manager Settings tile on the delivered Cloud Manager Fluid home page. The Cloud Manager Settings page is displayed. On the Cloud Manager Settings page, click the Manage Updates link displayed on the left panel.

For information on updating Cloud Manager see Updating Cloud Manager Overview.

This example illustrates the fields and controls on the Manage Updates page when no updates are available.

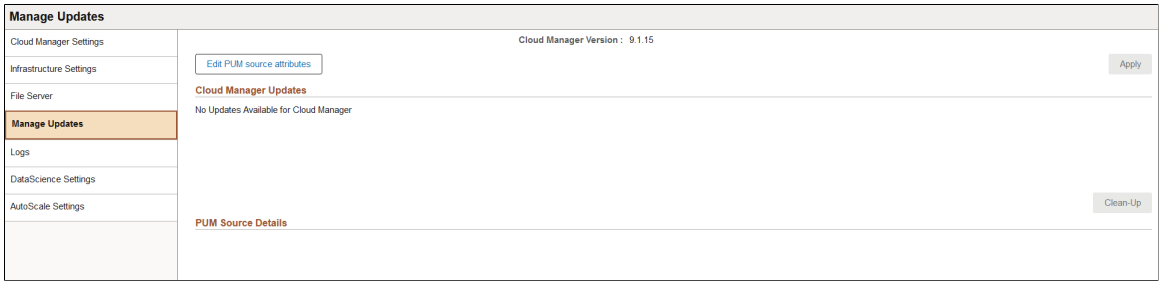

Logs Page

The Logs page (ECL_CM_FSLOGS_FL) enables you to view the logs with respect to FSS creation and its mounting to Cloud Manager.

Navigation:

Click the Cloud Manager Settings tile on the delivered Cloud Manager Fluid home page. Select the Logs tab in the left panel of the Cloud Manager home page.

This example illustrates the fields and controls on the Logs page. You can find definitions for the fields and controls later on this page.

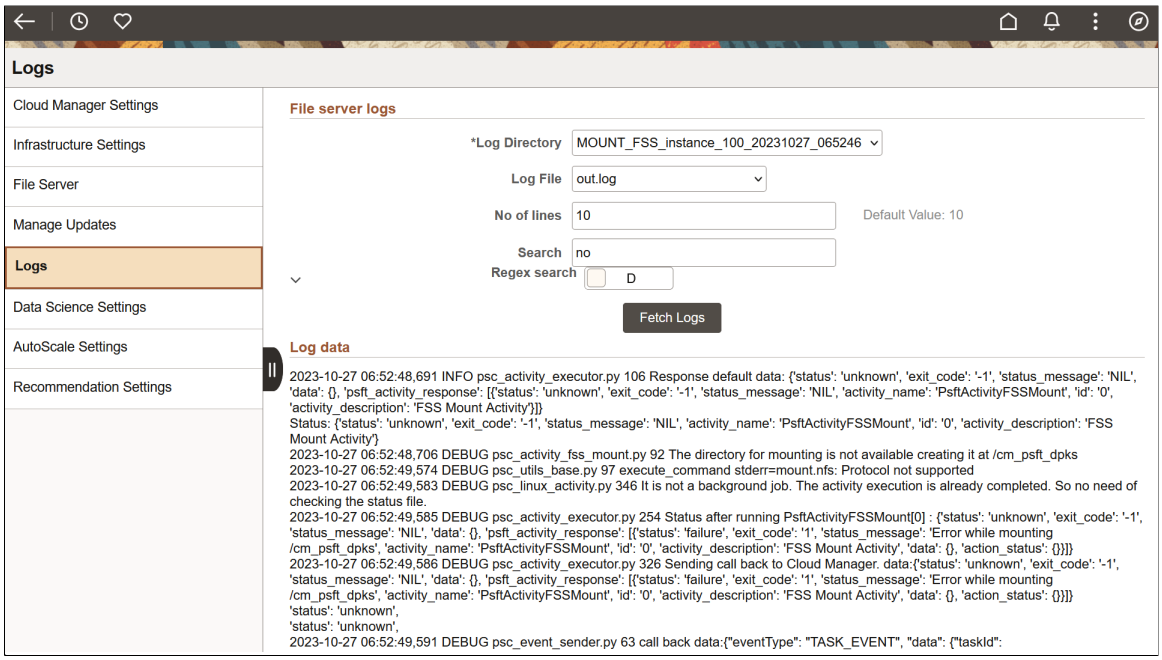

Field or Control	***Description***
Log Directory	Name of the file directory.
Log File	Log files from the selected directory. Select an appropriate log file in this field.
No of lines	Indicates how many lines of the selected log file to be displayed.
Search	Used to search for specific keywords in the log file. When user inputs a keyword, such as "ERROR" as an example, then only those lines are displayed which has an Error string in it. Here, only the specified number of lines are displayed.
Regex Search	Select E to enable advanced searching, where a user can provide UNIX-style regular expressions.
Fetch Logs button	Click this button to fetch log data based on the input provided by the user on the Logs page.
Log Data	Data from the logs.

Data Science Settings Page

The Data Science page (ECL_CM_DS_SETNG_FL) is required in order to use the Auto Scaling feature with Oracle Data Science. Integration to Data Science is optional.

Navigation:

Click the Cloud Manager Settings tile on the delivered Cloud Manager Fluid home page. The Cloud Manager Settings page is displayed. On the Cloud Manager Settings page, click the Data Science Settings link displayed on the left panel.

This example illustrates the fields and controls on the DataScience Settings page.

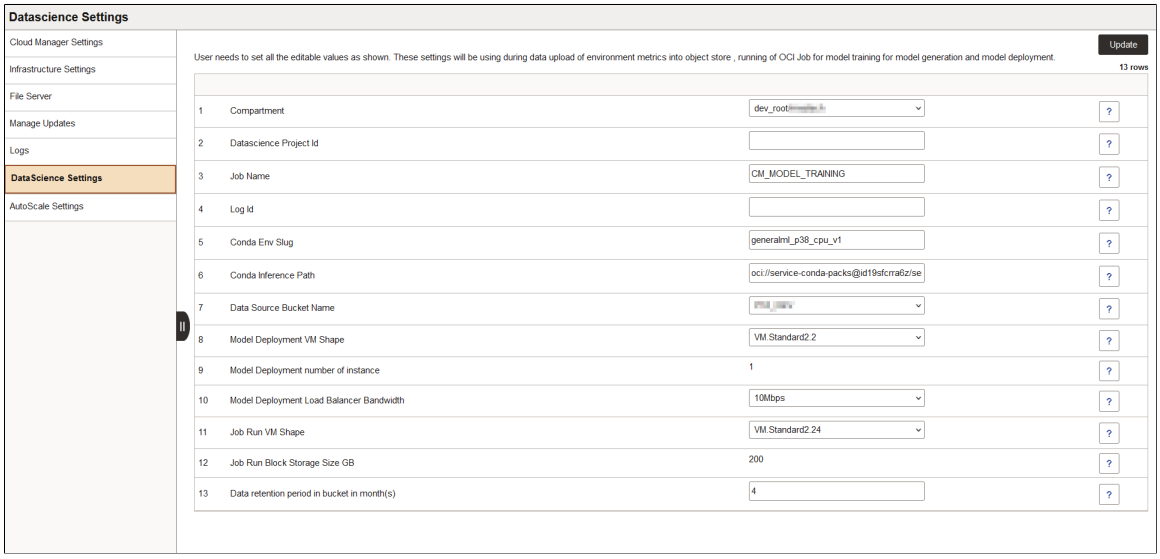

Field or Control	***Description***
Compartment	(Required) Compartment where Data Science is installed.
VCN	(Required) VCN for Data Science resource.
Subnet	(Required) The subnet used for Data Science must be a private subnet.
Data Science Project ID	(Optional) The project ID of the project created under selected compartment and subnet. If the project ID is not supplied the project name will be defaulted to CM_MODEL_TRAINING.
Job Name	(Required) A Name for the data science Job. This has a default value as CM_MODEL_TRAINING.
Log ID	(Optional) Specify the log ID where you want to direct the OCI logs. Keeping this empty will create a group name mltraininglogs and default log name as MODEL_TRAINING _<YYY-MM-DD>.

Field or Control	*Description*
Conda Env Slug	The version of prebuilt Data Science conda environment. This value should not be modified.
Conda Inference Path	The path of the prebuilt Data Science conda environment package. This path is used for setting up a conda environment for model training. This value should not be modified.
Data Source Bucket	(Required) Specify the bucket that will be used for data upload and model training.
Model Deployment VM Shape	(Required) Specify VM shape to be used for model deployment.
Model Deployment number of instance	This is a read only field. Block storage to be used for Data Science job run.
Model Deployment Load Balancer Bandwidth	(Required) Specify load balancer bandwidth to be used for model deployment.
Job Run VM Shape	(Required) Specify the VM shape to be used for Data Science job run.
Job Run Block Storage Size GB	This is a read only field. Block storage to be used for Data Science job run.
Data Retention period in bucket in month(s)	(Required) Specify the data retention period in months. Any data file beyond this period will be purged.

See the tutorial *Create Data Science Resources for Auto Scaling in Cloud Manager (Optional)* at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

AutoScale Settings Page

Use the AutoScale Settings page (ECL_MLNOTF_SET_FL) to control the notifications for the following event types:

- Monitoring: This notification event is triggered when an erroneous condition occurs during prediction flow.
- Scaling: This event is triggered when you need to take an action or an action is already taken by the prediction flow.

This example illustrates the fields and controls on the AutoScale Settings page.

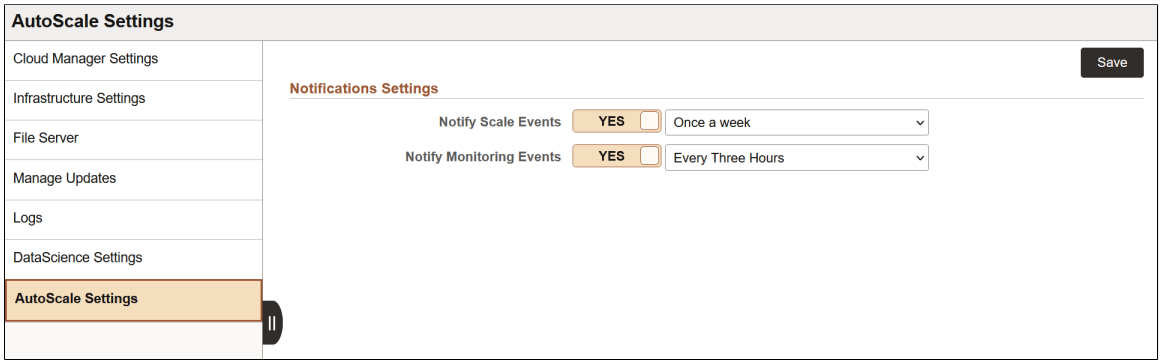

Auto Scale settings under Cloud Manager Settings are applicable to all environments by default. You can override this settings in Environment specific Auto Scale Settings. See Configuring AutoScale Settings.

Cloud Manager sends you an email notification about all the bookkeeping events at the end of the day at around 11:50 p.m.

You can select the following intervals in Notify Scale Events:

- Every Six Hours
- Every Three Hours
- Every Time
- Once Daily
- Once a week

When you select the notification interval as Once a week, you will receive email notification only on Friday. You need to configure the Notification Topic OCID to receive the email notification.

See Infrastructure Settings Page.

Advisory Settings Page

Use the Advisory Settings page (ECL_CMRECOM_SET_FL) to enable notifications for the advisories generated.

Whenever an action is completed or a specified issue is detected, Cloud Manager generates advisories to suggest the best course of action to be taken. The advisory settings on this page apply to all provisioned environments in Cloud Manager. See Configuring and Reviewing Advisories for information on enabling and viewing advisories at the environment level.

This example illustrates the fields and controls on the Advisory Settings page. Definitions for the fields and controls appear following the example.

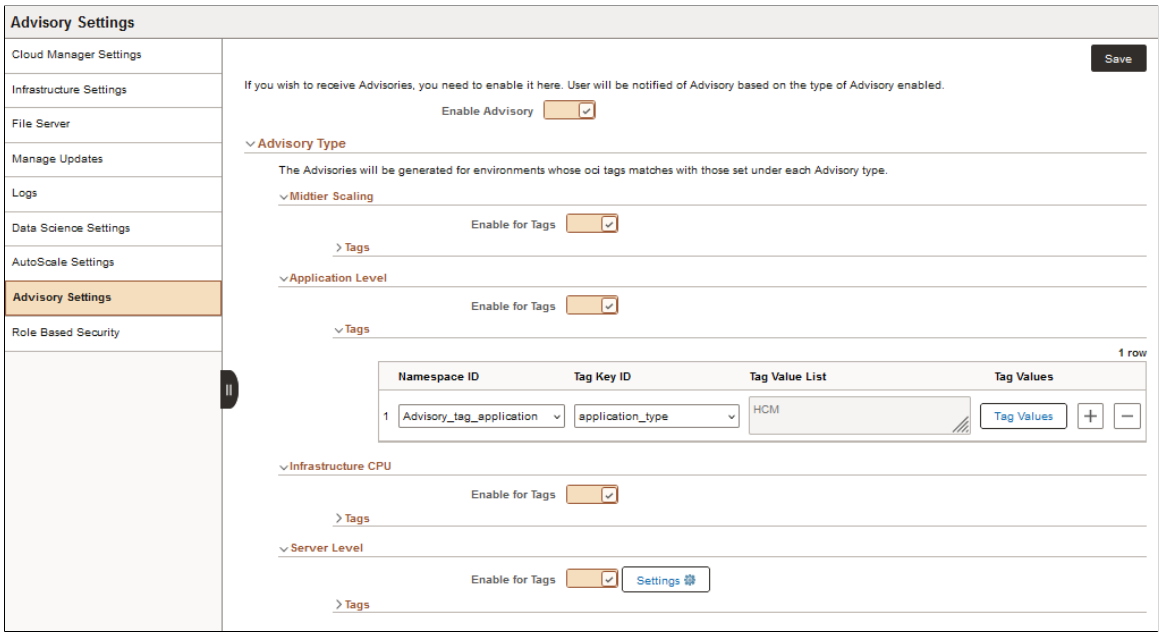

Field or Control	***Description***
Enable Advisory	Select to enable notifications for all advisories. This is the global setting for enabling or disabling advisories. The types of advisories are: • Infrastructure CPU • Midtier Scaling • Application Level • Server Level
Enable for Tags	Select to enable generation of advisories for environments whose tag matches any of the configured tags. Cloud Manager generates advisories only for environments that have tags.
Tags	Choose tags to identify the environments for which you want advisories. The latest values of Tag Namespace, Tag Key, and Tag Value are displayed in the Tags section for each advisory type, if you select **Enable for Tags**. See Configuring Tagging. Tag Value List is a non-editable field. On clicking Tag Values, you can set the tag value.

Field or Control	*Description*
Settings	Click to open the Settings dialog box for Server Level advisories. Enter a value for the boot volume threshold. When the available boot volume space drops below this value, a Server Level advisory is generated.

Role Based Security Page

The Role Based Security page (ECL_ROL_BAS_FL) in Cloud Manager Settings tile provides the ability to delegate access to a group of users on environments that are logically grouped using environment tags.

This feature allows Cloud Manager administrators to map roles with tags and permissions. A user assigned to a role can thus access a set of actions associated with permissions on environments with a set of tags. On the Role Based Security page, Cloud Manager administrator maps each user with a particular role, and then each role with a tag is given corresponding permissions.

This example illustrates the fields and controls on the Role Based Security page. Definitions for the fields and controls appear following the example.

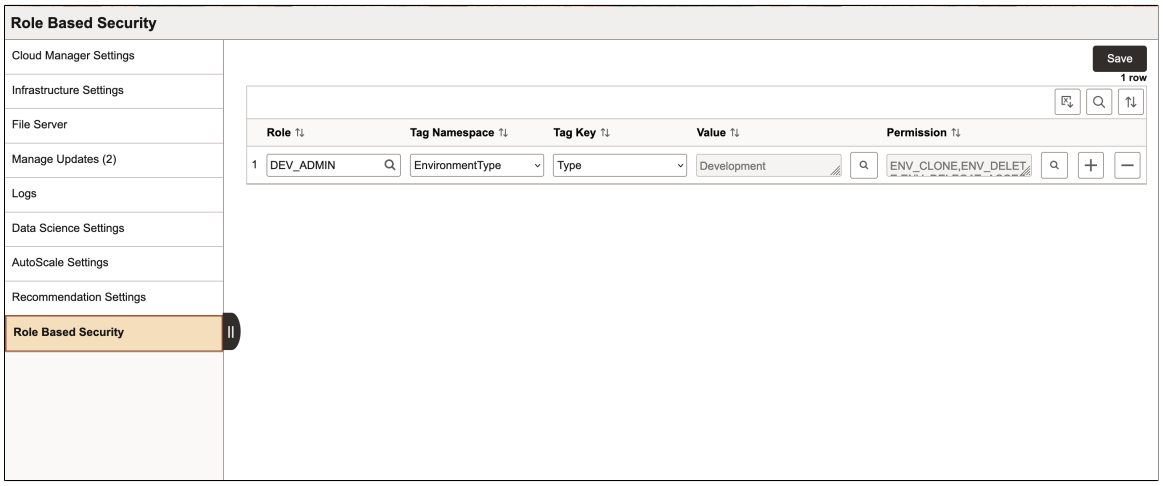

Field or Control	*Description*
Role	Indicates the roles that have access to the environments.
Tag Namespace	Created in OCI as a container for the tag keys. This is a natural grouping that can be used to apply a policy.
Tag Key	Tag keys are created within a Tag Namespace. This is the name that refers to the tag.
Value	Value that the user applying the tag adds to the tag key. The tag values are Development, Test, and Production.

Field or Control	*Description*
Permission	The permission names correspond to permission lists that determine which pages and menu items are visible to the user.

This example illustrates the functionality of Role Based Security.

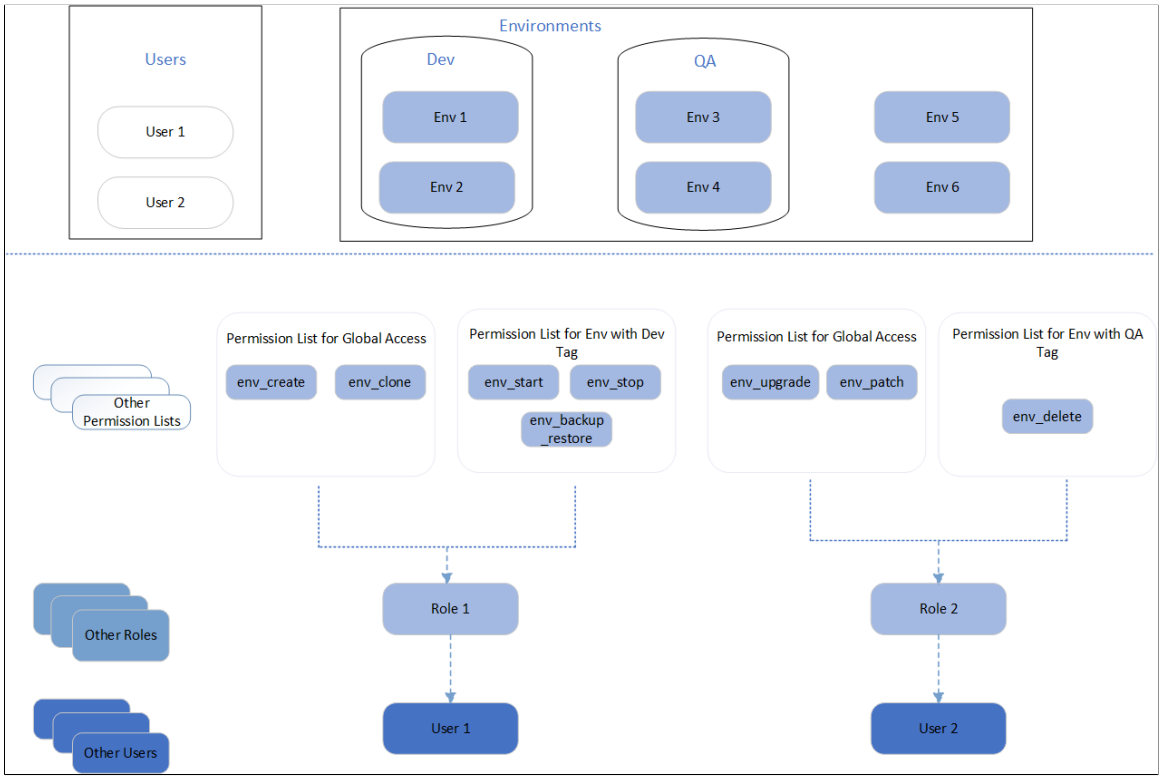

In this example, User 1 is mapped to Role 1 and User 2 is mapped to Role 2. User 1 gets create and clone access on all environments and Start, Stop and Backup-Restore access on environments with Dev tag. Similarly, User 2 gets Upgrade and Patch access on all environments and Delete access on environments with QA tag.

Note: Global permission for roles is given through Menu > PeopleTools > Security > Roles > Definition of a Role. Role Based Security page is used to provide specific permissions to a combination of roles and environments.

The following permission names correspond to permission lists that determine which pages and menu items are visible to the user.

Permission Name	*Permission List*	*Description*
ENV_ATTRIBUTES	PACL_ENV_ATTRIBUTES	Manage environment attributes.
ENV_BKUP_RESTOR	PACL_ENV_BKUP_RESTOR	Backup and restore the environment.

Permission Name	*Permission List*	*Description*
ENV_CLONE	PACL_ENV_CLONE	Clone the environment.
ENV_DELETE	PACL_ENV_DELETE	Delete the environment.
ENV_DETAILS	PACL_ENV_DETAILS	View environment details.
ENV_DR	PACL_ENV_DR	Create a standby environment for disaster recovery.
ENV_IMPORT	PACL_ENV_IMPORT	Import an environment.
ENV_IMPORT_NODE	PACL_ENV_IMPORT_NODE	Import an environment node.
ENV_INFRA_PATCH	PACL_ENV_INFRA_PATCH	Apply an Infrastructure CPU patch to the environment.
ENV_LB_SETTING	PACL_ENV_LB_SETTING	Manage load balancer settings for the environment.
ENV_MANAGE_NODE	PACL_ENV_MANAGE_NODE	Manage environment nodes.
ENV_MANAGE_PUM	PACL_ENV_MANAGE_PUM	Manage PUM environments.
ENV_MANAGE_TAG	PACL_ENV_MANAGE_TAG	Manage Tags for the environment.
ENV_PASSWORD	PACL_ENV_PASSWORD	Manage passwords on the environment.
ENV_PATCH	PACL_ENV_PATCH	Apply a patch to the environment.
ENV_POLICY	PACL_ENV_POLICY	Manage policies for the environment.
ENV_REFRESH	PACL_ENV_REFRESH	Refresh the environment.
ENV_START	PACL_ENV_START	Start the environment
ENV_STOP	PACL_ENV_STOP	Stop the environment
ENV_UPGRADE	PACL_ENV_UPGRADE	Upgrade the environment.

Configuring My Settings

Use the My Settings tile (ECL_INFO_HOME_FL_GBL) to access My SSH Public Key page and Password Groups page.

Navigation:

My Settings tile is delivered as part of the Cloud Manager home page.

This example illustrates the My Settings tile.

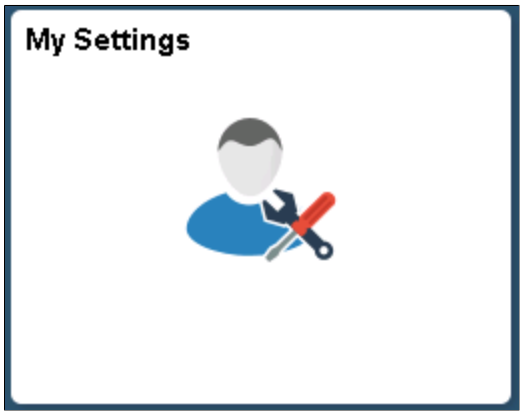

My SSH Public Key

Use the My Settings page (ECL_INFO_HOME_FL) to enter or edit the public SSH key. The SSH key provided here can be used to input SSH keys belonging to you or any administrator to help manage or troubleshoot issues by connecting over SSH.

Note: The SSH key will be automatically added into all environments that will be created by the user after adding this key here.

Navigation:

Click the My Settings tile on the delivered Cloud Manager Fluid home page. The My SSH Public Key page is displayed.

This example illustrates the fields and controls on the My SSH Public Key page.

Field or Control	*Description*
My SSH Public Key	Enter the SSH public key value.

Click **Save** to save the details.

Note: To edit existing key details, click the **Edit** button and replace the text; then click **Save**.

Password Groups

Use the Password Groups (ECL_PWDMANAGER_FL) page to centrally store and reuse the passwords that are already created as secret OCIDs in the secure password storage mechanism called OCI Vault, so that you need not manually enter the passwords while performing environment provisioning, importing environment, or during shift operation.

OCI Vault lets you securely store master encryption keys and secrets that you might otherwise store in configuration files or in code. See Vault Overview to know more about using keys and managing secrets. The Password Groups page on Cloud Manager lets you leverage this facility in OCI Vault, so that you need not manually enter the passwords.

Note: You can also manually enter the passwords while provisioning an environment.

This example illustrates the fields and controls on the Password Groups page. You can find definitions for the fields and controls later on this page.

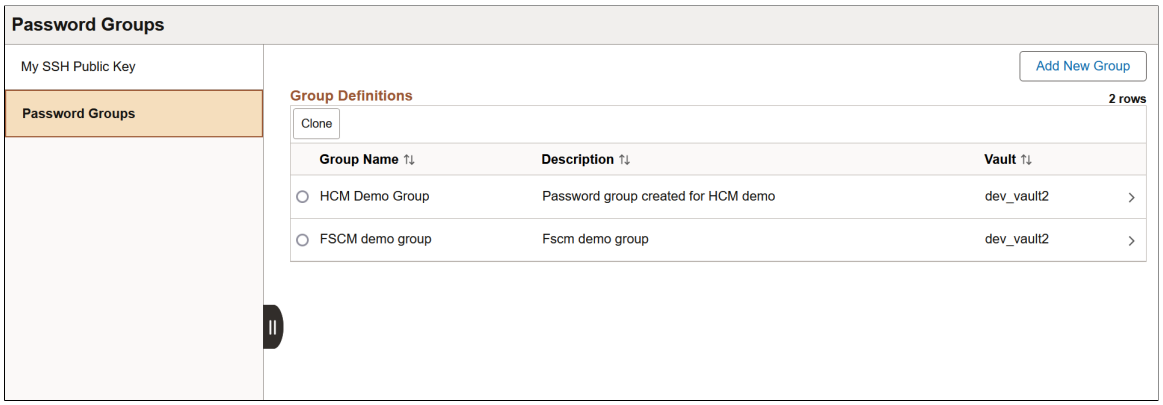

Field or Control	*Description*
Clone	Click the button to clone an existing password group listed on the Password Groups page.
Add New Group	Click the button to add a new password group.

The existing password groups are listed under the Group Definitions section on the Password Groups page. When you click any row listed on the Group Definitions section, the password group information is displayed.

Cloning an Existing Password Group

You can copy all the password details of an existing password group by clicking Clone button on the Password Groups page.

Note: You can edit or delete the newly cloned password group as long as it is not associated with an environment.

This example illustrates the fields and controls on the Password Groups page, when you clone an existing password group. You can find definitions for the fields and controls later on this page.

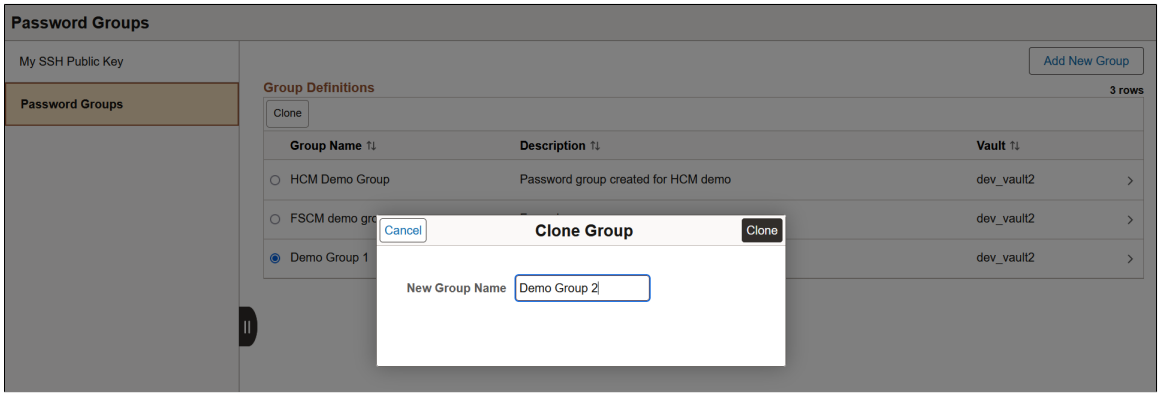

Field or Control	*Description*
New Group Name	Enter a unique name for the cloned password group.

The following passwords can be mapped to an existing secret in the selected OCI Vault:

- ADB Wallet Password
- Database Access Password
- Database Administrator Password
- Database Connect Password
- Database Operator Password
- Gateway Administrator Password
- RMAN Backup Encryption Password
- Search Proxy Password
- Source TDE Keystore (Wallet) Password
- TDE Master Key secret password
- Web Profile Password for user PTWEBSERVER
- Weblogic Administrator Password
- Windows Administrator Password

This example illustrates the fields and controls on the Password Groups page, when you clone password group. You can find definitions for the fields and controls later on this page.

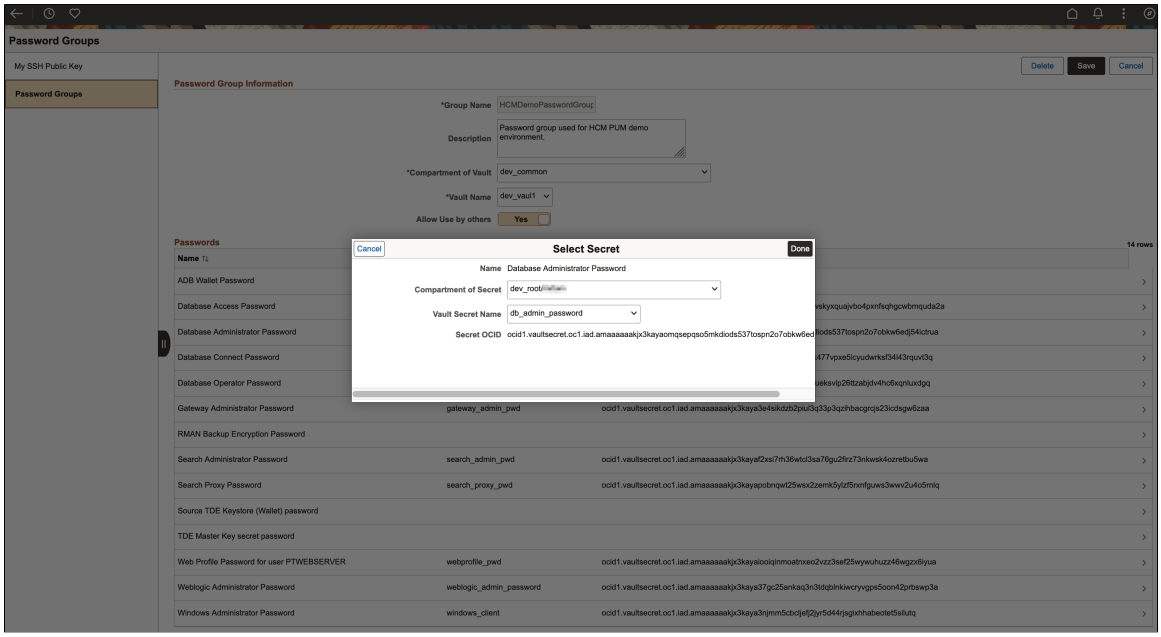

Field or Control	*Description*
Group Name	Enter a unique name for the password group. This is a required field.
Description	Enter a brief description of the new password group.
Compartment of Vault	Select the compartment from OCI Vault, which contains the secrets. This is a required field.
Vault Name	Select the OCI Vault Name from the drop down list. This is a required field.
Allow Use by others	Select Yes to allow the password group to be accessed by other users. The default value is No, which keeps the password group private.
Compartment of Secret	Select the compartment from OCI Vault, which contains the selected password.
Vault Secret Name	Select the secret name contained in OCI Vault from the drop down list.

The selected secret details are displayed in Passwords section. You can also choose to leave secrets blank, if you do not wish to set them.

Adding New Password Groups

You can create a new password group to add the secret OCIDs that are already created in the secure password storage mechanism called OCI Vault.

The following passwords can be mapped to an existing vault secret:

- ADB Wallet Password
- Database Access Password
- Database Administrator Password
- Database Connect Password
- Database Operator Password
- Gateway Administrator Password
- RMAN Backup Encryption Password
- Search Proxy Password

- Source TDE Keystore (Wallet) Password
- TDE Master Key secret password
- Web Profile Password for user PTWEBSERVER
- Weblogic Administrator Password
- Windows Administrator Password

This example illustrates the fields and controls on the Password Groups page, when you create a new password group.

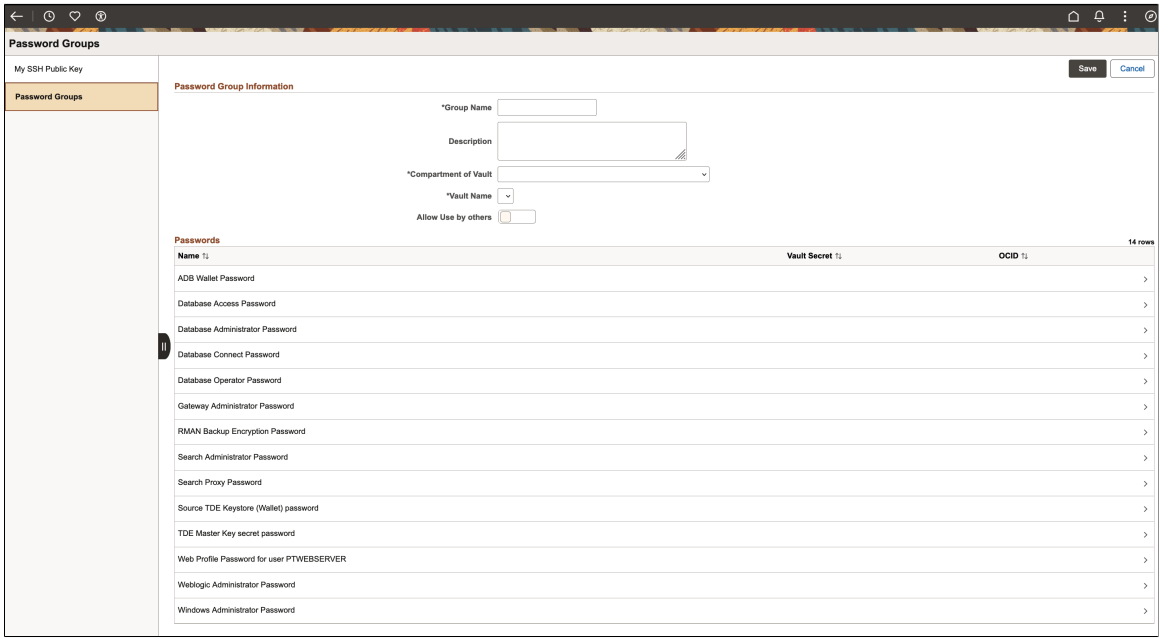

Click Save to save the password group information. You can edit the password details except name, and empty secrets, even after a password group is associated with an environment. However, you cannot delete the password group if it is associated with an environment.

When you create a new password group, the compartment of the secret is selected by default, which is the same as the vault compartment. You can change the selection if required.

This example illustrates the fields and controls on the Password Groups page on selecting secret for a newly created password group.

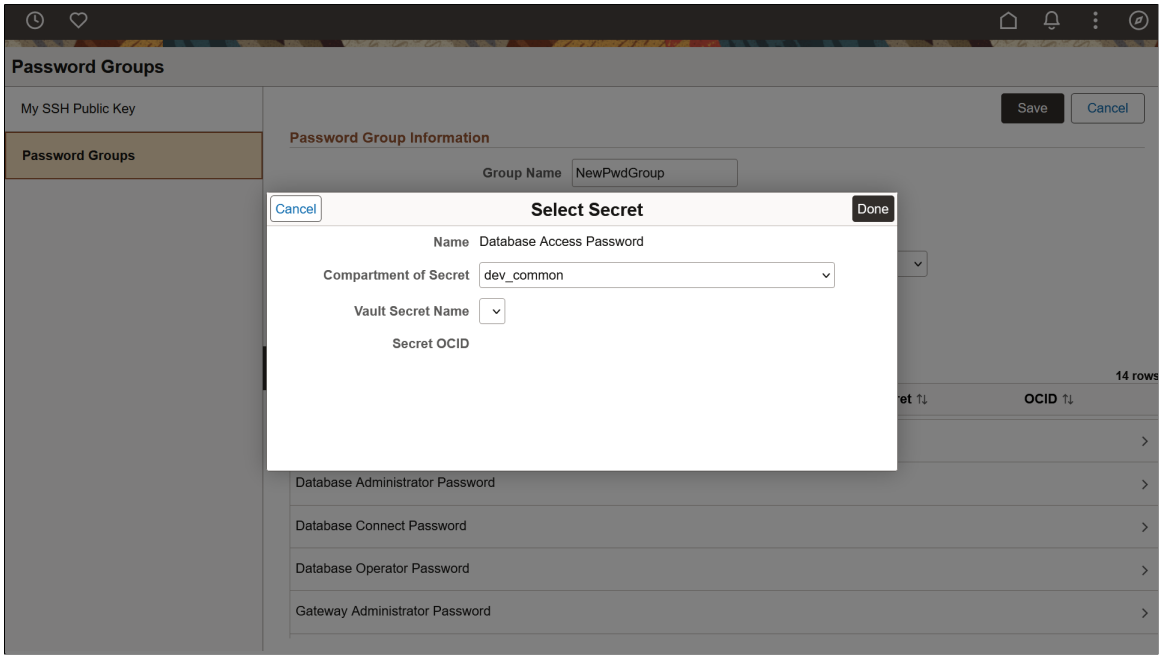

Chapter 3

Managing Repository

Repository Overview

Cloud Manager provides an easy way to automatically download and manage PeopleSoft Update Images (PIs), PeopleSoft Release Patchsets (PRPs), PeopleTools Product Patches and PeopleSoft Custom Update Images. Cloud Manager uses the file server to store downloaded artifacts from MOS. To streamline and automate downloads of various PeopleSoft application update images and PRPs, Cloud Manager has introduced the concept of Subscription Channels. Each PeopleSoft application has an associated Channel, which an administrator can choose to subscribe in order to download the latest PeopleSoft Update Images and PRPs for that particular PeopleSoft application. Cloud Manager is delivered with channels for PeopleSoft applications, which are available after you complete the installation and configuration. An administrator can subscribe to multiple channels and download all necessary PIs and PRPs.

Cloud Manager uses an application called Download Manager to download updates from MOS, which is invoked through process scheduler in asynchronous mode every time a channel is subscribed.

On the Repository tile, Administrators can:

- Subscribe to release channels for latest PeopleSoft application updates.
- Manage downloaded PeopleSoft Update Images and PRPs.
- Upload custom scripts and other objects into the repository.

Pages Used to Manage Cloud Manager Repository as an Administrator

Page Name	***Definition Page***	***Usage***
Repository Tile	ECL_REPOSITORY_FL_GBL (Content reference for the tile.)	Access the various features such as, Channel Subscriptions and Download History, and functions such as, downloading logs and deleting downloads.
My Downloads Page	ECL_REPO_AMYDLS_FL	View the PRPs and PIs downloaded. New entries are added as soon as new artifacts are downloaded.
Download Subscriptions Page	ECL_REPO_BCHNL_FL	Create download channels and subscribe them to initiate downloads. You can also use predefined download channels to initiate downloads.

Page Name	*Definition Page*	*Usage*
Download History Page	ECL_REPO_BDLHIS_FL	View the history of downloads, such as PIs and PRPs downloaded.
Logs Page	ECL_REPO_MLOG_FL	View the download manager logs.
Upload Custom Scripts Page	ECL_UPLD_CUST_SCR	Upload custom scripts.

Repository Tile

Use the Repository tile to access Repository landing page.

Navigation:

The Repository tile (ECL_REPOSITORY_FL_GBL) is delivered as part of the Cloud Manager home page.

This example illustrates the Repository Tile.

Working with the Repository

Use the Repository to:

- View downloaded artifacts
- Subscribe to channels
- Download history
- Download logs
- Filter and delete downloads
- Manage custom scripts

My Downloads Page

Use the My Downloads page (ECL_REPO_AMYDLS_FL) to view the artifacts downloaded. New entries are added as soon as new artifacts are downloaded.

Navigation:

Click the Repository tile on the delivered Cloud Manager Fluid home page. My Downloads page is displayed by default.

This example illustrates the fields and controls on the My Downloads page. You can find definitions for the fields and controls later on this page.

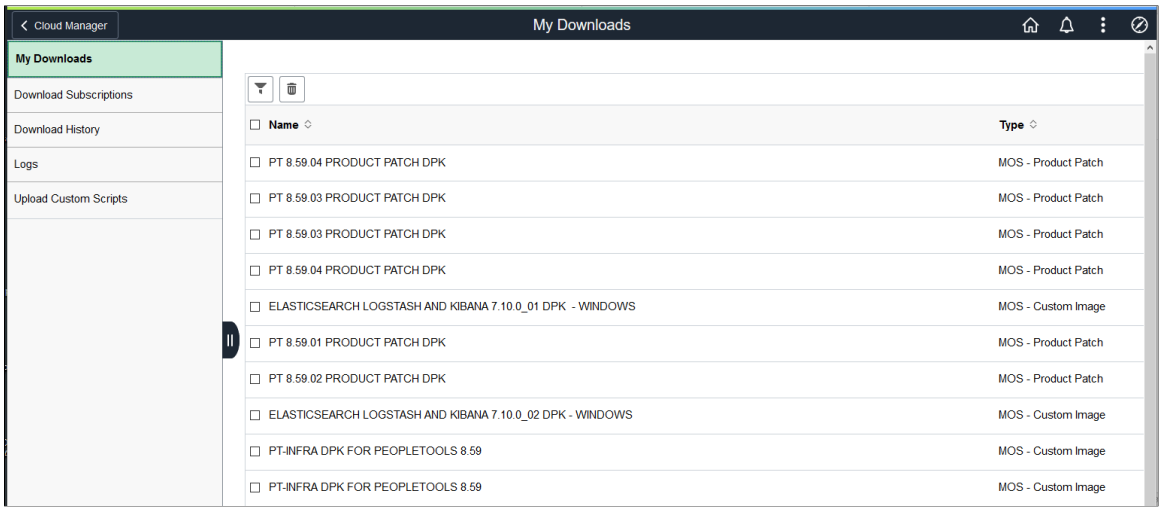

Note: Clicking an item in the My Downloads page displays additional details of the downloaded artifact.

Field or Control	***Description***
	Use the Filter icon to refine the search results based on search criteria.
	Use the Delete icon to delete downloaded PIs and PRPs. Select the check box corresponding to the row you want to delete, and then click **Delete** button.
Name	Name of the downloaded artifact.
Type	Indicates the artifact type such as PeopleSoft Update Image, PRP, Custom Image, and so on.
Product	Indicates the PeopleSoft application product pillar.
Release	Indicates the PeopleSoft application release.
Version	Indicates the application PeopleSoft Update Image version.
Platform	Indicates the Operating System platform, such as Linux, or Windows.

Field or Control	*Description*
Size	Total size of the PeopleSoft Update Image or PRP.

Download Subscriptions Page

Use the Download Subscriptions page (ECL_REPO_BCHNL_FL) to subscribe to download channels and initiate downloads.

Note: Cloud Manager delivers default channels and those channels are available in the unsubscribed list of the Download Subscriptions page.

Navigation:

Click the Repository tile on the delivered Cloud Manager Fluid home page. Select the Download Subscriptions tab in the left panel of the Cloud Manager home page.

This example illustrates the fields and controls on the Download Subscriptions page. You can find definitions for the fields and controls later on this page.

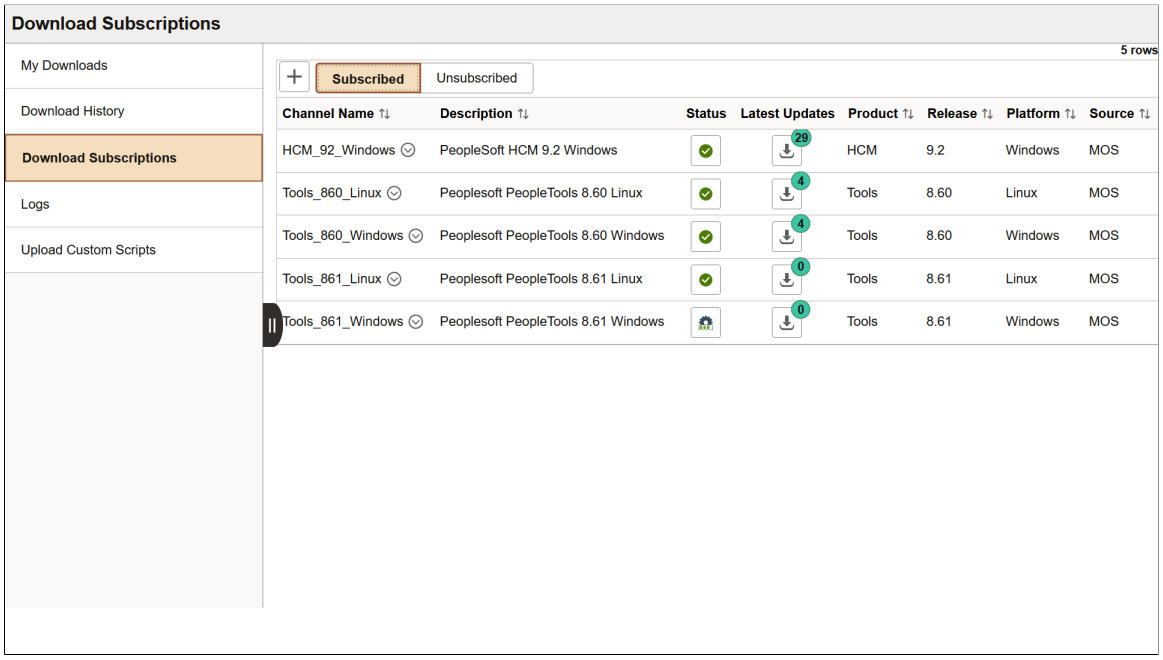

Field or Control	*Description*
	To subscribe or unsubscribe channel, click the Related Actions button corresponding to the channel name. If you select the **Subscribe** option, Cloud Manager starts monitoring for any new PIs or PRPs and downloads them from My Oracle Support. If you select the **Unsubscribe** option, Cloud Manager will no longer monitor or download latest PeopleSoft Update Images or PRPs. When a release channel is subscribed, Cloud Manager invokes the download manager application, which connects to MOS and downloads latest updates for the release channel. Note that artifacts, such as Update Images, are large in size and can take a few hours to download. User can view the status of active downloads from the Download History page.
Subscribed tab	Click this tab to view a list of subscribed channels. When you select the Related Action to subscribe to a channel, that channel will be added to the Subscribed tab. **Note:** This operation will renew the channel subscriptions for all channels present in the Subscribed tab. This means that Cloud Manager will check for updates and download them for all channels present in the Subscribed tab.
Unsubscribed tab	Click this tab to view a list of unsubscribed channels. By default, newly created download channels are listed under the Unsubscribed tab.
Status   	Status will indicate current state. • Success The Success icon indicates the download was successful. No further action is necessary. • In-progress The In-progress icon indicates the update is downloading. Click on the icon to view the status of the download. • Error The Error icon indicates the download failures. Click the icon to open the Download Error page.

This example illustrates the fields and controls on the Download Subscriptions – Unsubscribed page. You can find definitions for the fields and controls later on this page.

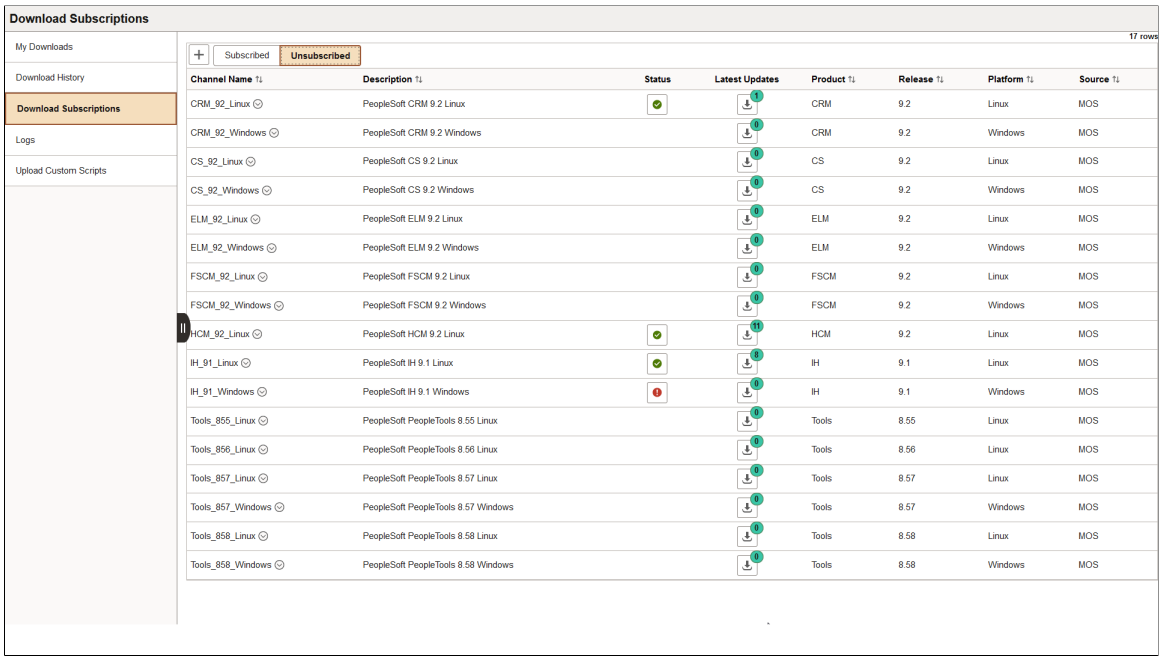

Use the related actions button to subscribe to a channel.

Download Error Page

Click the failed icon on the Download page to view the Download Error pop-up page. There are two types of error pop-up pages, one for standard errors and another for invalid password. The invalid password pop-up page only applies to password protected downloads.

This example illustrates the fields and controls on the Download Error page. You can find definitions for the fields and controls later on this page.

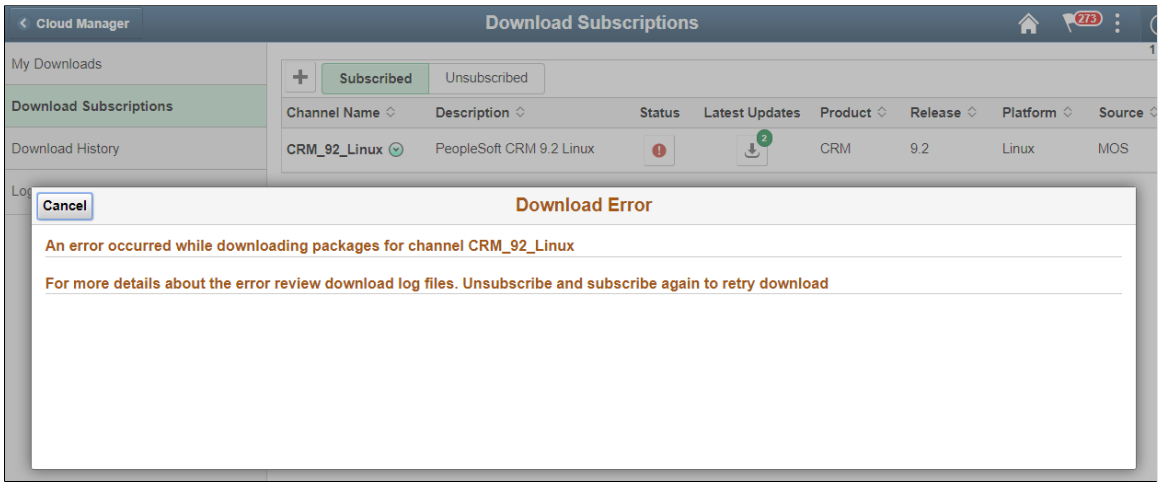

The download error message will indicate the error and direct the user to review the download log files. Error could be due to issues occurring during the download, such as network connection disruption (unable to reach MOS) or if there is no space available on the file server to save downloads.

Password Protected Subscription Password Error

When you subscribe to the a password protected download channel, the DownloadAssist file is downloaded from MOS to Cloud Manager. The DownloadAssist file contains the password required to download the packages.

This example illustrates a Download Error page displaying a password error.

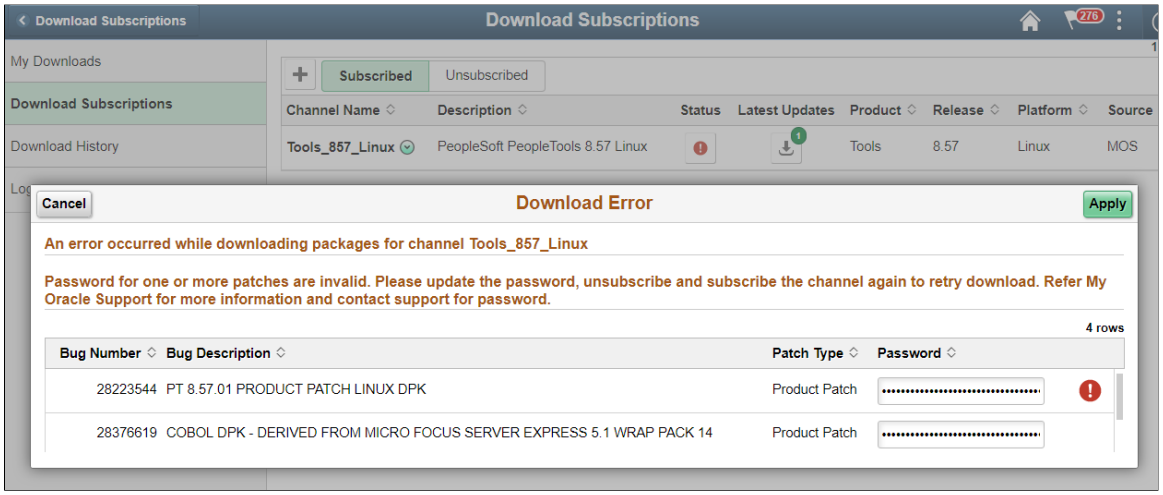

This message will only appear if Cloud Manager is unable to retrieve the password due to an error in MOS. If this does occur, the user will need to request the password from Support and update it manually on this page.

If the issue is due to an expired password, user can do a unsubscribe to the channel and then re-subscribe to the channel. Any password changes will be updated in the DownloadAssist file that is posted in MOS. On re-subscribing to a channel, the new password should get automatically updated in Cloud Manager.

If a re-subscribing doesn't solve the problem, then there could be issues in Cloud Manager that is failing to retrieve and update the password automatically or an unforeseen issue in MOS which is not allowing the download of the DownloadAssist file. In such a scenario, the user is expected to get the password from support and update the same on the Download Error page.

Note: The next scheduled download will renew the subscription and include any password changes in the DownloadAssist file. See Changing Download Interval

---

Downloading PeopleTools Patches

Cloud Manager can download previous PeopleTools patch releases for the PeopleTools channel. But for Application channels, only latest patches get downloaded.

In case of Tools channel subscription, you are presented with a modal window for selecting the patch version you want to download.

Navigation:

Click the Unsubscribed tab. Select any Tools channel. Click the Related Options menu and select Subscribe.

This example illustrates the fields and controls in Download Filter modal window.

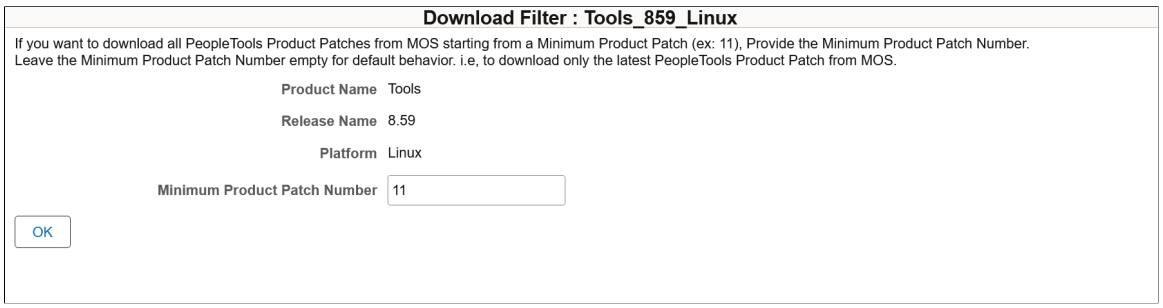

Enter the required product patch version in the Minimum Product Patch Number field. For example, if user enters 11 in this field, then CM will download tools patches 8.59.11, 8.59.12, 8.59.13 up to the latest.

Download History Page

Use the Download History page (ECL_REPO_BDLHIS_FL) to view the history of downloads.

Note: The entries in Download History page are updated based on the download interval. Clicking an entry in the Download History page displays the current state of the download channel (that is, a list of files already downloaded, another list of files in the download queue and those that are currently downloading).

Navigation:

Click the Repository tile on the delivered Cloud Manager Fluid home page. Select the Download History tab in the left panel of the Cloud Manager home page.

This example illustrates the fields and controls on the Download History page. You can find definitions for the fields and controls later on this page.

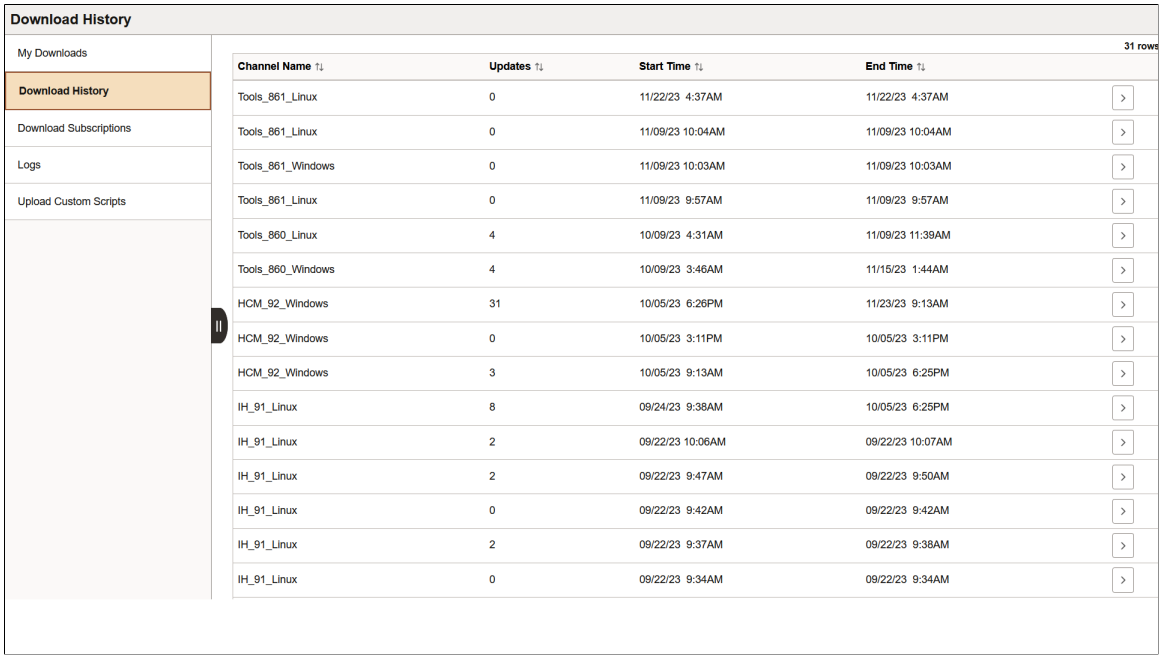

Field or Control	*Description*
Channel Name	Name of the download channel.
Updates	Number of updates downloaded.
Start Time and **End Time**	Indicates the time when downloads are started/finished for the release channel.

Logs Page

Use the Logs page (ECL_REPO_MLOG_FL) to view the download logs corresponding to the subscribed channels. It displays download logs for all the files that get downloaded.

Navigation:

Click the Repository tile on the delivered Cloud Manager Fluid home page. Select the Logs tab in the left panel of the Cloud Manager home page.

This example illustrates the fields and controls on the Logs page. You can find definitions for the fields and controls later on this page.

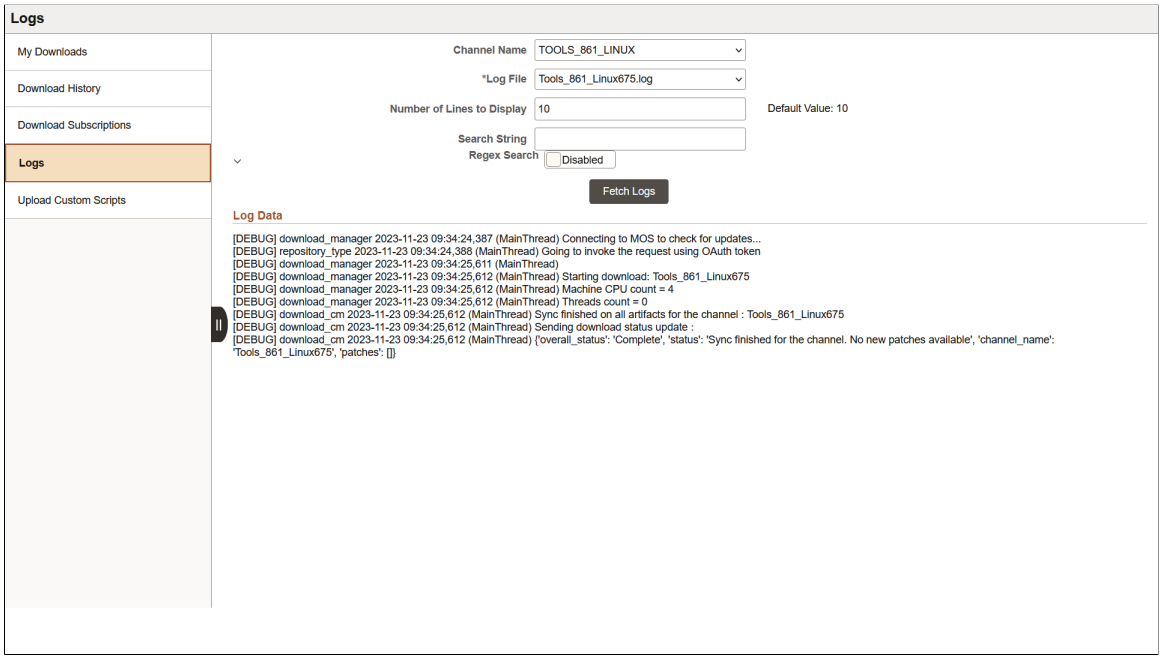

Field or Control	*Description*
Channel Name	Name of the subscribed channel.
Log File	Log files are generated when a channel is subscribed. Select an appropriate log file in this field.

Field or Control	*Description*
Number of Lines to Display	Indicates how many lines of the selected log file to be displayed.
Search String	Used to search for specific keywords in the log file. When user inputs a keyword, such as "ERROR" as an example, then only those lines are displayed which has an Error string in it. Here, only the specified number of lines are displayed.
Regex Search	Enables advanced searching, where a user can provide Unix style regular expressions.
Fetch Logs button	Click this button to fetch log data based on the input provided by the user in the Logs page.
Log Data	Data from the logs.

Re-synchronizing Repository Data with Downloaded List

Sometimes even after subscribing to a channel, you may not able to see some of the downloaded patches in Repository > My Downloads page. The logs may show that the downloads are being skipped. This indicates a situation where the patches are already downloaded but their metadata is not synced with Cloud Manager. In such situations re-sync the downloaded patches metadata with Cloud Manager using the following steps.

1. Go to Repository > Download Subscriptions page and unsubscribe all channels that are currently subscribed.
2. Navigate to **PeopleTools** > **Process Scheduler** > **Schedule Process Requests**.
3. Enter a Run Control.
4. Run the process "ECL_REPODM".
5. After the process finishes, you should be able to see the missing patches in the Repository > My Downloads page.

Subscribing Channels using the Cloud Manager Repository

This section details the process flow for subscribing to channels using the Cloud Manager Repository.

Note: Cloud Manager has a process scheduler recurring job defined, which invokes the download manager for all the subscribed release channels once a week. This will make sure that latest updates for all the subscribed release channels are downloaded every week without any user interaction.

Prerequisites

The administrator needs to define My Oracle Support credentials prior to subscribing channels using the Cloud Manager Repository. For this, perform the following:

1. Select the Cloud Manager Settings tile.
2. Edit the value in the User ID field and My Oracle Support password field in the My Oracle Support (MOS) Credentials section.
3. Click Save Settings to save the details. For details on the Cloud Manager settings, see Cloud Manager Settings Page.

Note: This is a one-time setup.

Perform the following steps to subscribe to channels using the Cloud Manager Repository tile:

1. Click the Repository home page available on the Cloud Manager home page.
2. Select Download Subscriptions on the left panel. The Download Subscriptions page is displayed.
3. Click Unsubscribed.
4. To subscribe to the release channel, perform the following:
    a. Click the Related Actions button corresponding to the channel name.
    b. Select Subscribe action. If there are any new updates, then the system starts downloading the new updates.

Changing Download Interval

By default, Cloud Manager polls My Oracle Support for new updates every week. The recurrence definition for download channel subscriptions is **CloudManager Repository Update**. To modify the download schedule to meet your organizational needs, modify the recurrence pattern to a pattern that meets your needs.

For example, if you want to poll for updates on alternate days, you would perform the following:

1. In Cloud Manager, navigate to PeopleTools, Process Scheduler, Process Scheduler Recurrences.
2. Select the recurrence CloudManager Repository Update.
3. Select Recurrence Pattern to *Weekly*.
4. Select the alternating days as in Monday, Wednesday, and Friday or Sunday, Tuesday, Thursday.

    Note: It is recommended to select alternate days.

5. Save the page.

This example illustrates the fields and controls on the Recurrence Definition – Cloud Manager Repository Update where the occurrence is set to alternate days.

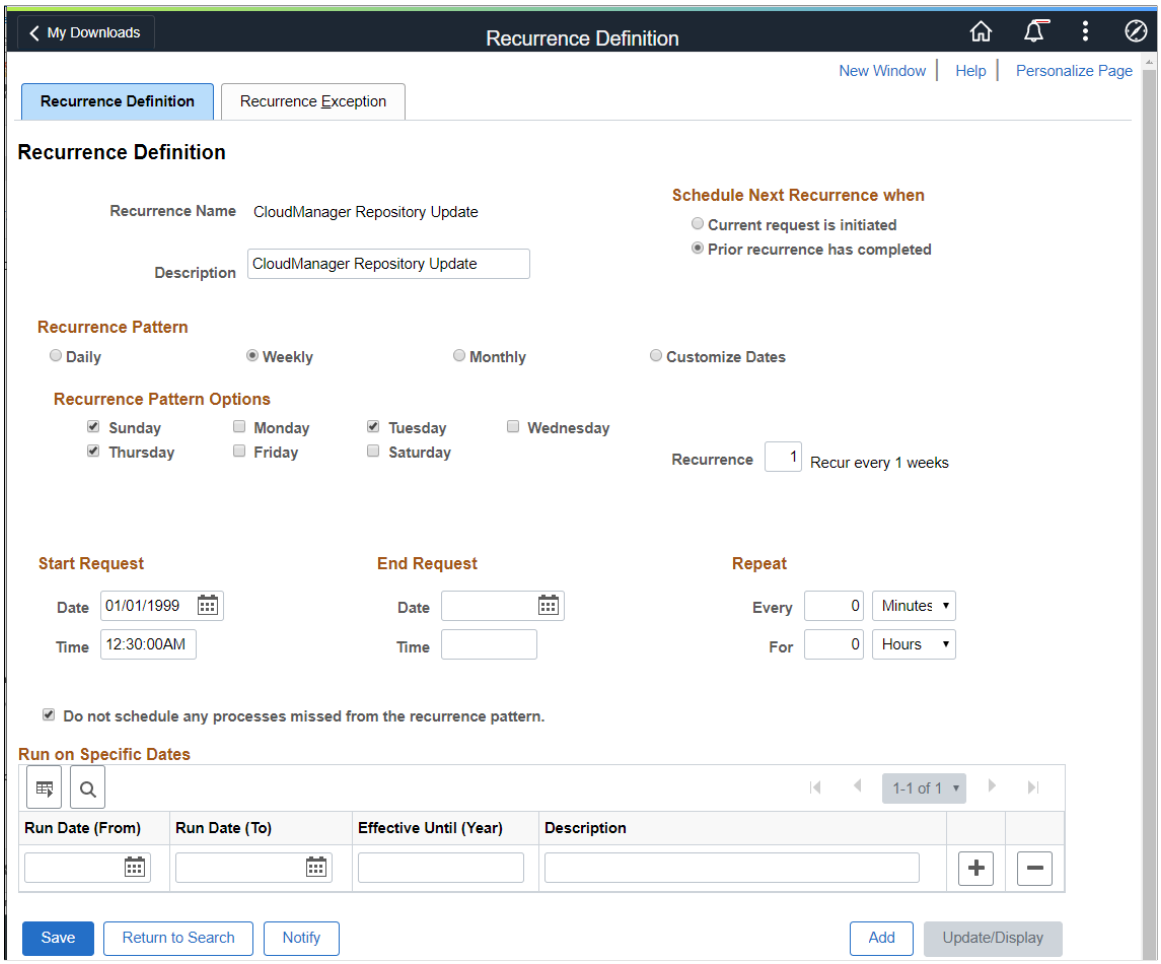

For more information on setting up recurring schedules, see Product documentation for *PeopleTools: Process Scheduler,* “Defining PeopleSoft Process Scheduler Support Information”, Defining Recurrence Definitions.

Upload Custom Scripts Page

Starting with Cloud Manager Update Image 8 custom scripts can be run before or after provisioning a PeopleSoft environment.

Note: After a PeopleTools upgrade, you must edit the webserver property (configuration.properties) file to add or modify the HttpRepositoryPath as shown here: `HttpRepositoryPath=/opt/oracle/psft/customscripts`

This table lists the script types and which tiers they can run on.

Script Type	*Where you can run the script*
Python	Full Tier Middle Tier Database Tier PeopleSoft Client Search Stack Database as a service (DBaaS)
Shell	Full Tier Middle Tier Database Tier Search Stack Database as a service (DBaaS)
Batch	PeopleSoft Client and Windows Middle Tier
PowerShell	Windows Client and Windows Middle Tier

Scripts can contain instructions to call and run one or more other scripts. For example, a shell script can invoke multiple shells, Python, or other scripts. You can also upload any kind of supporting files that are used by the scripts.

To upload a custom file:

1. From Cloud Manager homepage, select the Repository tile.
2. Select Upload Custom Scripts.
3. Click the Add icon.
4. The File Attachment pop up appears where you can choose files from object storage, your local device, or a public Git repository by clicking on the corresponding option.

Note: You must enable firewall connectivity between Cloud Manager and your Git repository. You must also ensure that the customisation artifacts uploaded through Git, object storage or local drive are robust and secure.
Cloud Manager does not capture the data regarding modification of settings using custom scripts and custom DPK, and the responsibility of ensuring that such modifications do not adversely affect other lifecycle activities of Cloud Manager rests with you.

This example illustrates the options for uploading custom files on the Upload Custom Scripts page.

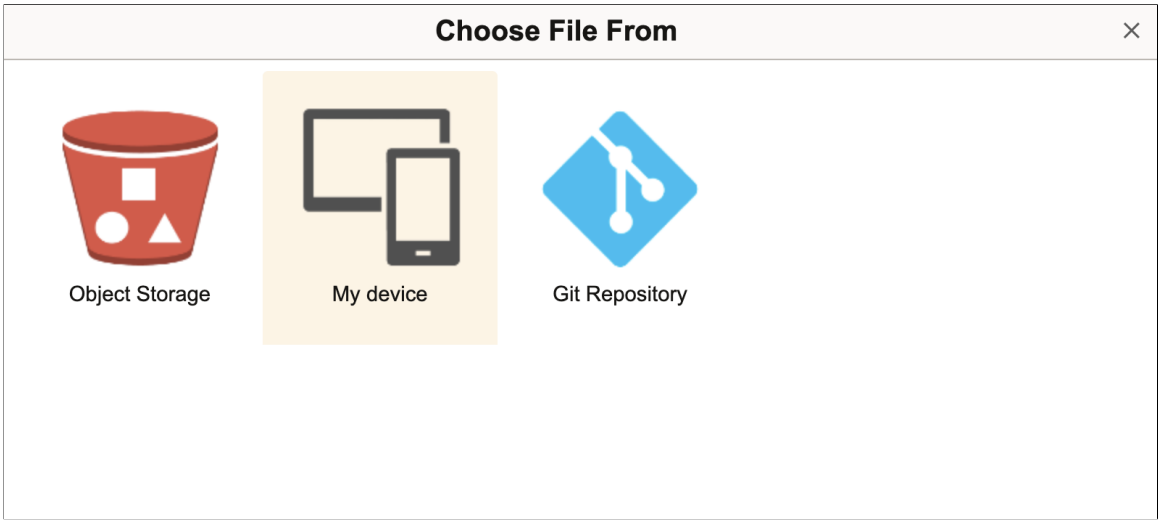

This example illustrates the fields and controls on the File Attachment pop up when you select Object Storage as the option for uploading custom files. You can find definitions for the fields and controls later on this page.

Field or Control	*Description*
Compartment Id	Compartment where custom file is stored.
Bucket Name	Name of the bucket that contains the custom file.
Object name	Name of the object that contains the custom file.

This example illustrates the fields and controls on the File Attachment pop up for uploading custom files directly from your device.

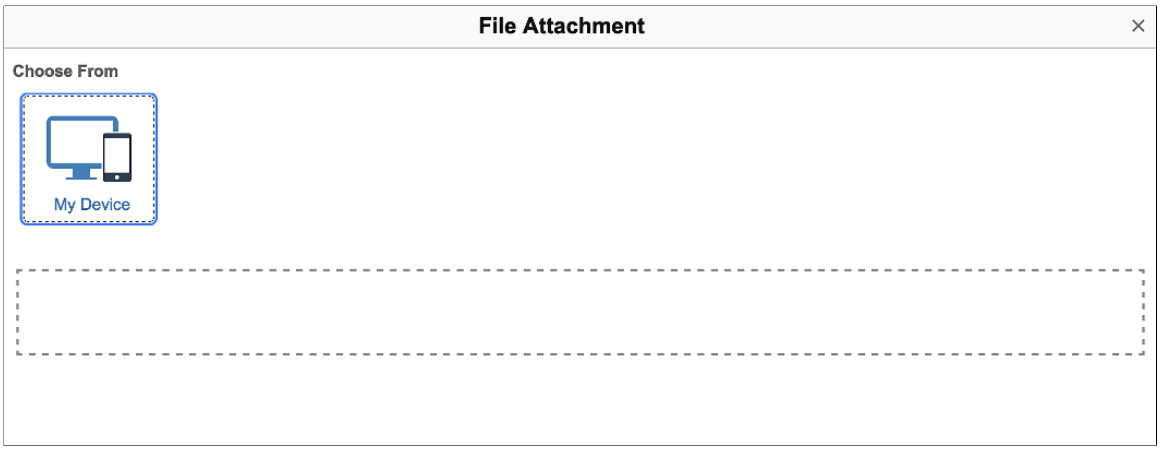

Note: The maximum supported file size for uploading custom files is 2 GB.

This example illustrates the fields and controls on the File Attachment pop up for sourcing input from Git. You can find definitions for the fields and controls later on this page.

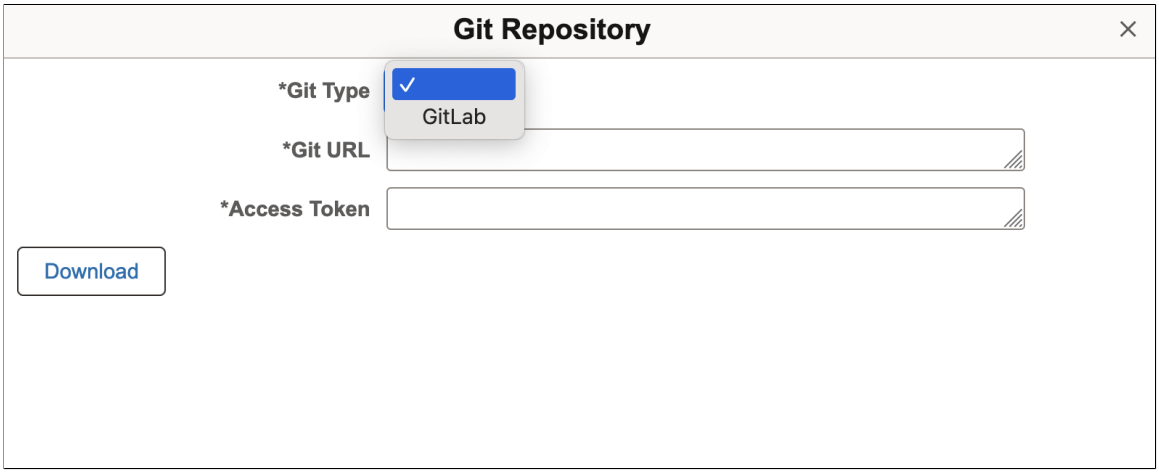

Field or Control	***Description***
Git Type	Git Type used for sourcing input. Supported types are GitLab and GitHub.
Git URL	Raw file download REST APIs URL of your Git repository.
Access Token	Access token required by Cloud Manager for using GIT to download artifacts. Cloud Manager needs access to only the repository's raw file download API. **Note:** Access token is optional for public repository.

5. Once you have selected the file, the Upload option appears.
6. Click the Upload button to upload the file to Cloud Manager.
7. When upload is complete, click Done to complete the process.

This example illustrates the fields and controls on the Upload Custom Scripts page. You can find definitions for the fields and controls later on this page.

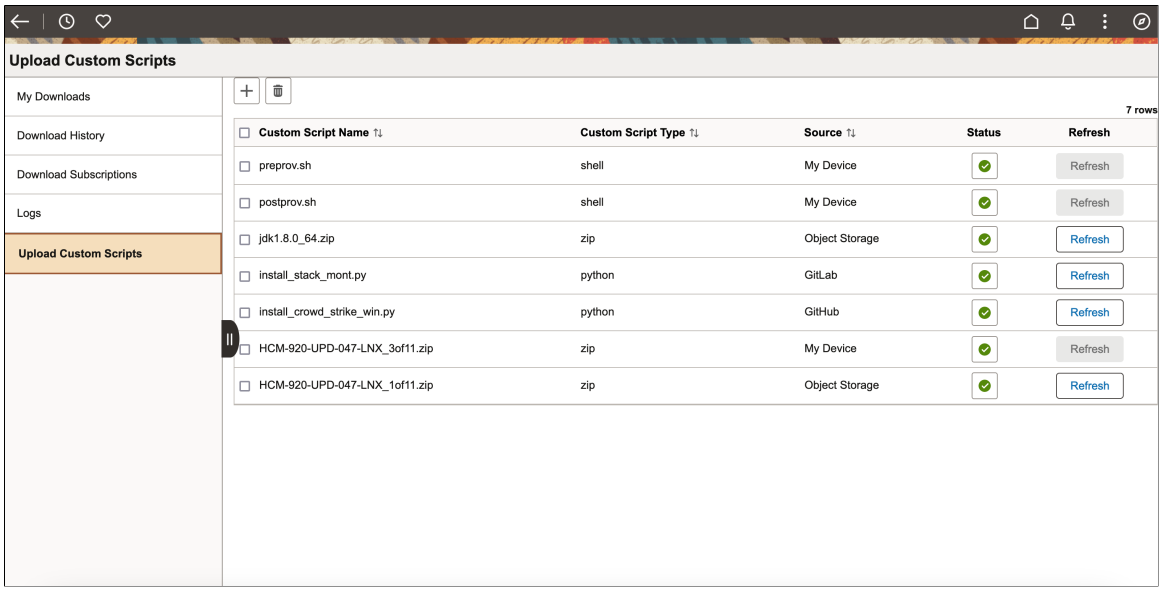

Field or Control	***Description***
Custom Script Name	Name of the custom script.
Custom Script Type	Type of the custom script. Possible values are Shell, Python and Zip.
Source	Source of the uploaded custom file. Possible values are Object Storage, your local device, GitLab, and GitHub.
Status	Status of the custom script upload. File metadata of the successfully uploaded files can be viewed by clicking the Status button corresponding to the file.
Refresh	On clicking the Refresh button, Cloud Manager is updated to the latest version for Git and Object Storage sources. If the script is relocated after you upload it or if the access token is renewed, Refresh can fail by throwing an error message related to HTTP status 401 or 404. When this happens, you must enter the new details again as you would do for a new script or file upload.

Environment Variables Allowed in Custom Scripts

Custom scripts provide an ability to extend Cloud Manager's provisioning flow. During provisioning orchestration, Cloud Manager provides access to following environment variables inside the custom scripts. The environment variables allowed in custom scripts and their sample values are as follows:

PSFT_APP_APPDOM01_JOLT_LISTENER_PORT=9033

PSFT_INST_PS_CUST_HOME=/u01/app/oracle/product/pt/ps_cust_home

PSFT_ES_PORT=9200

PSFT_INFRA_AVAILDOMAINPRIMARY=evQs:US-ASHBURN-AD-1

PSFT_INFRA_COMPARTMENT_ID=ocid1.compartment.oc1..aaaaaaaaoxxxxxxxxxxxxxxxxxxxxxx7era

PSFT_CONNECT_ID=people

PSFT_APP_APPDOM01_PSPUBHND_DFLT_MAX_INSTANCES=1

PSFT_ULIMIT_CUSTOMIZATION=true

PSFT_PRCS_PRCSDOM01_MSTRSRV=Yes

PSFT_ADMIN_PWD=xxx

PSFT_INST_DB_LOCATION=/u01/app/oracle/product/db

PSFT_APP_APPDOM01_PSPUBHND_DFLT_MIN_INSTANCES=1

PSFT_INST_PS_APP_HOME=/u01/app/oracle/product/pt/ps_app_home

PSFT_INST_PS_CFG_HOME=/u01/app/oracle/product/sample09-fulltierlinux-1/ps_cfg_home

PSFT_ADB_CM_WALLET_LOCATION=xxx

PSFT_CUST_SCRIPT=TestPost.sh

PSFT_PRCS_PRCSDOM01_PSAESRV_MIN_INSTANCES=2

PSFT_WEB_WEBSERVER01_WEBSERVER_HTTP_PORT=8000

PSFT_APP_APPDOM01_JOLT_LISTENER_MAX_INSTANCES=3

PSFT_WEB_WEBSERVER01_AUTH_TOKEN_DOMAIN=.ft.vcnnet.oraclevcn.com

PSFT_APP_APPDOM01_JOLT_LISTENER_MIN_INSTANCES=3

PSFT_INFRA_TENANCY_ID=ocid1.tenancy.oc1..aaaaaaaayy35xxxxxxxxxxxxxxxxxxxxxxma5a

PSFT_PI_IMAGE=17

PSFT_COBOL_TYPE=VisualCobol

PSFT_PRCS_DOMAINS=['PRCSDOM01']

PSFT_WEBLOGIC_ADMIN_USER=system

PSFT_INFRA_FAULT_DOMAIN=OCI Default

PSFT_PRCS_PRCSDOM01_APPLICATION_ENGINE=1

PSFT_ADB_WALLET_PASSWORD=xxx

PSFT_INFRA_PRIVATE_IP=10.1.10.124

PSFT_PRCS_PRCSDOM01_MAXAPIAWARE=1

PSFT_PSFT_BASE=/u01/app/oracle/product

PSFT_PRCS_PRCSDOM01_COBOL_SQL=1

PSFT_ACCESS_ID=SYSADM

PSFT_WEB_WEBSERVER01_PEOPLESOFT01_APPSERVER_CONNECTIONS=sample09-fulltierlinux-1.ft.vcnnet.oraclevcn.com:9033

PSFT_ES_INSTALL_DIRECTORY=/u01/app/oracle/product

PSFT_GW_ADMIN_USER=administrator

PSFT_SEARCH_PROXY_USER=people

PSFT_OS_USERS={'psft_search_user_name': 'esuser', 'psft_app_install_user_name': 'psadm3', 'remote_login_user': 'opc', 'psft_install_user_name': 'psadm1', 'psft_runtime_user_name': 'psadm2', 'oracle_user_name': 'oracle2'}

PSFT_INST_EYAML_PRIVATE_KEY=/u01/app/oracle/product/sample09-fulltierlinux-1/dpk/puppet/secure/keys/private_key.pkcs7.pem

PSFT_PRCS_PRCSDOM01_PSAESRV_MAX_INSTANCES=2

PSFT_WEBLOGIC_ADMIN_PWD=xxx

PSFT_KERNEL_CUSTOMIZATION=true

PSFT_PRE_PROVISION_CUST_SCRIPT=xxx

USER=root

PSFT_ACCESS_PWD=xxx

PSFT_OPR_ID=VP1

PSFT_INFRA_NSG_IDS=[]

PSFT_CLOUD_INIT_COMMANDS=echo "# testing01 " >> /etc/test.sh;

echo "# testing02 " >> /etc/test.sh;

PSFT_KERNEL_DATA=xxx

PSFT_IS_ADB_BASED_ENV=False

PSFT_SEARCH_ADMIN_USER=esadmin

PSFT_EM_AGENT=xxx

PSFT_INST_TOOLS_CLIENT_HOME=/u01/app/oracle/product/pt/tools_client

PSFT_ULIMIT_DATA=xxx

PSFT_PRCS_PRCSDOM01_PSDSTSRV_MAX_INSTANCES=2

PSFT_INFRA_VCN_ID=ocid1.vcn.oc1.iad.amaaaaaaxxxxxxxxxxxxxxxxxxxxxxxxxxxxxdaq

PSFT_OS_USER_GROUPS={'oracle_runtime_group_name': 'dba', 'oracle_install_group_name': 'oinstall', 'psft_app_install_group_name': 'appinst', 'psft_runtime_group_name': 'psft'}

PSFT_DB_TYPE=DEMO

PSFT_PRCS_PRCSDOM01_APPENG=Yes

PSFT_ES_CLUSTER_NAME=ESCL

SUDO_USER=opc

PSFT_ES_HOSTNAME=xxx

PSFT_INST_PUPPET_BASE_LOCATION=/u01/app/oracle/product/sample09-fulltierlinux-1/dpk/puppet

PSFT_APP_APPDOM01_PSSUBHND_DFLT_MIN_INSTANCES=1

SHLIB_PATH=:/home/opc/cloud/lnx_python/lib

PSFT_INST_DPK_LOCATION=/u01/app/oracle/product/sample09-fulltierlinux-1/dpk

PSFT_APP_APPDOM01_PSQRYSRV_MAX_INSTANCES=1

PSFT_CUSTOMER_DPK_FILE=xxx

PSFT_APP_APPDOM01_PSAPPSRV_MIN_INSTANCES=2

PSFT_CUSTOMIZATION=xxx

PSFT_PI_NUMBER=17

PSFT_WEBPROFILE_USER_PWD=xxx

PSFT_ES_PROXY_PASSWD=xxx

PSFT_APP_APPDOM01_PSBRKHND_DFLT_MIN_INSTANCES=1

PSFT_APP_APPDOM01_WORKSTATION_LISTENER_PORT=7000

PSFT_INST_PUPPET_PROD_LOCATION=/u01/app/oracle/product/sample09-fulltierlinux-1/dpk/puppet/production

PSFT_WEB_WEBSERVER01_WEBSERVER_HTTPS_PORT=8443

PSFT_PRCS_PRCSDOM01_SQR_REPORT=1

PSFT_SEARCH_PROVIDER=xxx

PSFT_DB_IS_ML=xxx

PSFT_PRCS_PRCSDOM01_OPTIMIZATION_ENGINE=1

PSFT_INFRA_PUBLIC_IP=129.80.206.203

PSFT_ELASTIC_SETUP=false

PSFT_INST_PT_LOCATION=/u01/app/oracle/product/pt

PSFT_ELASTIC_SEARCH=N

PSFT_APP_APPDOM01_PSSAMSRV_MIN_INSTANCES=1

PSFT_APP_TYPE=IH

PSFT_CUSTOM_INPUT=xxx

PSFT_INST_LOG_DIR=/home/opc/log/1715748988842831561/
psft_activity_custom_post_provisioning_lnx_4_20240515_060212

PSFT_INST_EYAML_KEYS_DIR=/u01/app/oracle/product/sample09-fulltierlinux-1/dpk/puppet/secure/
keys

PSFT_PRCS_PRCSDOM01_XML_PUBLISHER=1

PSFT_PIA_DOMAINS=['WEBSERVER01']

PSFT_COBOL_SETUP=false

PSFT_APP_APPDOM01_PSBRKHND_DFLT_MAX_INSTANCES=1

PSFT_INFRA_REGION=us-ashburn-1

PSFT_POST_PROVISION_CUST_SCRIPT=TestPost.sh

PSFT_CONNECT_PWD=xxx

PSFT_HOST_NAME=${envname}-${nodetype}${ostype}-${instno}

SHLVL=4

PYTHONPATH=/home/opc/cloud

PSFT_WEB_WEBSERVER01_SITE_NAMES=peoplesoft01

PSFT_DB_SERVICE_NAME=PSPDB

PSFT_APP_APPDOM01_PSAPPSRV_MAX_INSTANCES=2

PSFT_PRCS_PRCSDOM01_SQR_PROCESS=1

PSFT_APP_APPDOM01_PSSAMSRV_MAX_INSTANCES=1

PSFT_ES_DISCOVERY_HOST_NAME=127.0.0.1

PSFT_INFRA_HOST=sample09-fulltierlinux-1.ft.vcnnet.oraclevcn.com

PSFT_INFRA_HOST_OS=linux

LOGNAME=root

CV_ASSUME_DISTID=OL7

PSFT_APP_APPDOM01_PSSUBHND_DFLT_MAX_INSTANCES=1

PSFT_PRCS_PRCSDOM01_PSDSTSRV_MIN_INSTANCES=2

PATH=/home/opc/cloud/lnx_python/bin:/sbin:/bin:/usr/sbin:/usr/bin

PSFT_DB_NAME=PSPDB

PSFT_INFRA_SUBNET_ID=xxx

PSFT_APP_APPDOM01_PSQRYSRV_MIN_INSTANCES=1

PSFT_GW_ADMIN_USER_PWD=xxx

PSFT_INST_EYAML_PUBLIC_KEY=xxx

PSFT_ENV_TYPE=fulltier

PSFT_COBOL_VERSION=xxx

PSFT_DB_INSTANCE_TYPE=Compute

PSFT_APP_DOMAINS=['APPDOM01']

PSFT_DB_PORT=1522

PSFT_LOGSTASH=N

PSFT_OPR_PWD=xxx

PSFT_KIBANA_PORT=5601

How to Access Environment Variables in or from Custom Scripts

This section shows some examples of using the environment variables.

Example for Python:

```
crowdstrike_installer=os.environ.get('CS_BIN_PATH')+'/'+os.environ.get('CS_BINARY')⇒
```

Example for Shell:

```
crowdstrike_installer="${CS_BIN_PATH}/{CS_BINARY}"
```

Example for Batch:

```
set crowdstrike_installer=%CS_BIN_PATH%/%CS_BINARY%
```

Example for puppet script (Ruby script):

```
crowdstrike_installer=ENV["CS_BIN_PATH"]+"/"+ENV["CS_BINARY"]
```

Example for puppet script (puppet script):

```
$psft_site =  hiera('peoplesoft_site_name')
```


Sample Batch File

This sample batch file will print the PATH and PSFT_DB_NAME and store it in a file.

```
set file=C:\temp\script_batch.log

@echo off

@echo Starting the provisioning script > %file%

@echo The value for variable PATH is:  >> %file%

@echo %PATH% >> %file%

@echo The value for variable PSFT_DB_NAME is:  >> %file%

@echo %PSFT_DB_NAME% >> %file%

@echo Ending the provisioning script >> %file%
```

Sample Python Script

This sample python script will print the PATH and PSFT_DB_NAME and store it in a file.

```
import os, time

millis = int(round(time.time() * 1000))

file_name = '/tmp/post-provision_python_{}.log'.format(millis)

f = open(file_name, 'w+')

f.write('Starting post-provisioning script')

f.write('The variable PATH value is: {}\n'.format(str(os.environ.get('PATH'))))

f.write('The variable PSFT_DB_NAME value is: {}\n'.format(str(os.environ.get('PSFT_⇒

DB_NAME'))))

f.write('Ending post-provisioning script')

f.close()
```

Sample PowerShell Script

This sample PowerShell script reads and writes the environment variables PSFT_ACCESS_ID and PATH to a file.

```
$file="C:\script.log"

Add-Content $file "Starting the provisioning script"

Add-Content $file "The value for variable PATH is:"

Add-Content $file (Get-ChildItem Env:PATH).Value

Add-Content $file "The value for variable PSFT_ACCESS_ID is:"

Add-Content $file (Get-ChildItem Env:PSFT_ACCESS_ID).Value

Add-Content $file "Ending the provisioning script"
```


Sample Shell Script

This sample shell script will print the PATH and PSFT_DB_NAME and store it in a file.

```
now=$(date +%d%m%Y%H%M%S)
file=/tmp/prov_$now.log

echo "Starting the provisioning script" > $file

echo "The value for variable PATH is: " >> $file

echo $PATH >> $file

echo "The value for variable PSFT_DB_NAME is: " >> $file

echo $PSFT_DB_NAME >> $file

echo "Ending the provisioning script" >> $file
```

Creating Custom DPK

You must create the custom DPK file by following directory structure and file structure similar to PeopleSoft DPK format. It must be in zip file format (for example, customerDPK.zip).

This example illustrates the custom DPK structure.

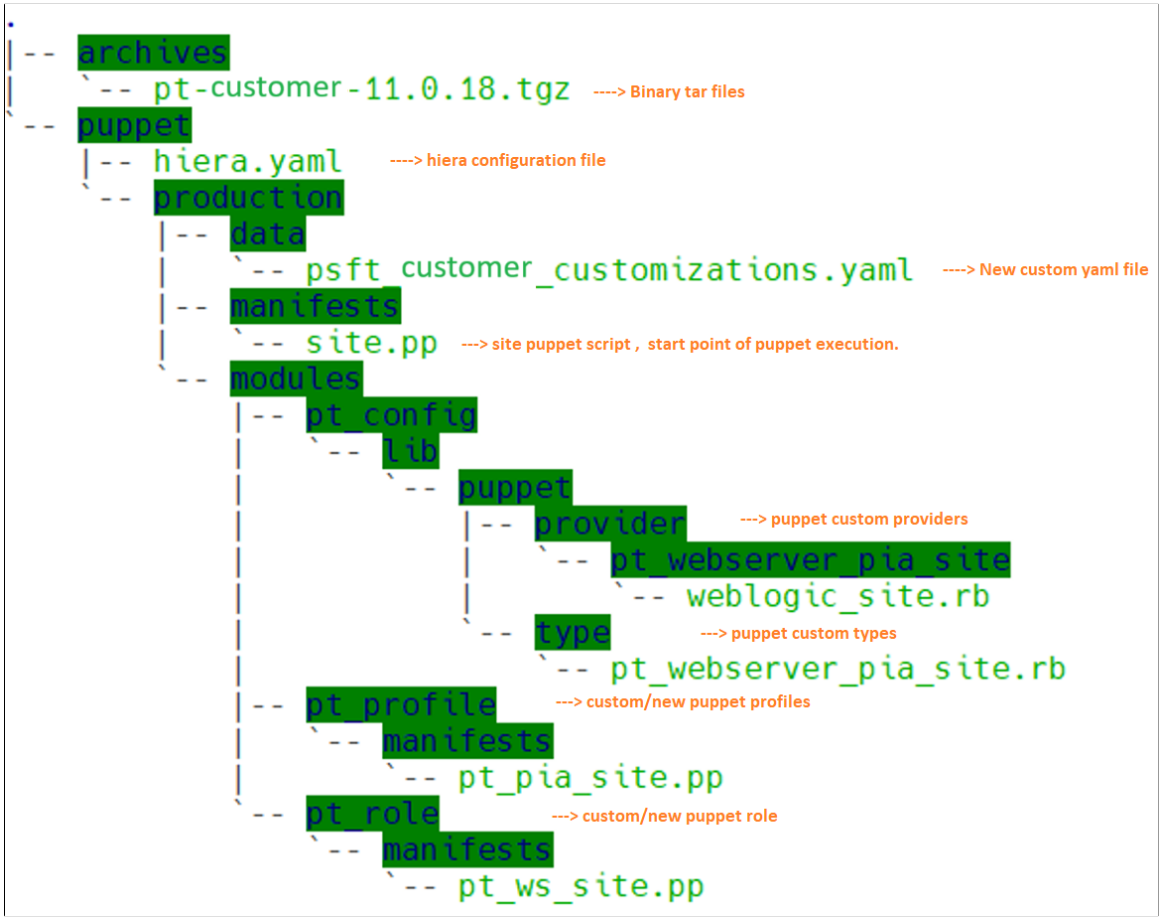

During the deployment of custom DPK, Cloud Manager overwrites the existing instance archives, YAML files, and puppet scripts. Since use cases such as clone, restore, add node, PeopleTools Patch, and

PeopleTools Update use the existing puppet script present in the environment, you must ensure that the modifications using custom scripts do not interfere with the existing logic.

The Puppet process starts from the site.pp file (current customerDPK.zip/puppet/production/manifests/site.pp) in the custom DPK. Therefore, it is important to update the correct puppet role details in the site.pp file. The following sample script invokes the pt_ws_site puppet role, by using the puppet role script (customerDPK.zip/puppet/production/modules/pt_role/manifests/pt_ws_site.pp) file.

```
node default {
  include  ::pt_role::pt_ws_site
}
```

The puppet role script can be used for implementing different puppet profiles in sequence. You can write custom implementations using any combination of the following options:

- Puppet profiles (customeredDPK.zip/puppet/production/modules/pt_profile/manifests/pt_pia_site.pp)
- Puppet custom providers (customerDPK.zip/puppet/production/modules/pt_config/lib/puppet/pt_config/lib/puppet/provider/pt_webserver_pia_site/weblogic_site.rb)
- Custom types (customerDPK.zip/puppet/production/modules/pt_config/lib/puppet/pt_config/lib/puppet/type/pt_webserver_pia_site.rb)

You can pack the binary files and installer.exe in a tgz file in the archive directory (for example: customerDPK.zip/archives/pt-customer-11.0.18.tgz). Puppet custom script can be used for deploying the binaries. Use the Puppet Customizations field to provide new custom values in the psft_customizations.yaml. See Configuring DB Systems Settings.

Cloud Manager usually keeps all the dpk-specific files in a hostname based directory (for example: /u01/app/oracle/product/usernameft07-fulltierlinux-1/dpk). To update any values hiera.yaml file in CustomDPK, use factor global variables (%{*::hostname*}) for resolving the hostname. This example illustrates a sample hiera.yaml file:

```
:backends:
  - eyaml
  - yaml

:hierarchy:
  - defaults
  - psft_customizations
  - psft_unix_system
  - psft_deployment
  - psft_configuration
  - psft_patches

:yaml:
  :datadir: /u01/app/oracle/product/%{::hostname}/dpk/puppet/production/data

:eyaml:
  :datadir:   /u01/app/oracle/product/%{::hostname}/dpk/puppet/production/data
  :extension: 'yaml'

  :pkcs7_private_key: /u01/app/oracle/product/%{::hostname}/dpk/puppet/secure/keys/⇒

private_key.pkcs7.pem
  :pkcs7_public_key:  /u01/app/oracle/product/%{::hostname}/dpk/puppet/secure/keys/⇒

public_key.pkcs7.pem
```


Note: When you use custom DPK for ad hoc changes in the environment, you must ensure that the settings are specified accordingly and that proper cleanup is done. Cloud Manager does not track or consider the changes done using custom DPK. You can use a post provisioning script, which is run after the custom DPK, for correcting the puppet script and for cleanup operation.

Custom DPK is applied to the environment as part of the final tasks in provisioning use case. If the use case fails, you can retry the failed steps from **Environment** > **Details** > **Provision Task Status**. If a custom DPK file (Custom Puppet Script) is sourced from a Git repository or object storage, the latest file is pulled from the repository during provisioning.

Chapter 4

Managing Topology

Topology Overview

Topology defines the infrastructure layout that will be created on Oracle Cloud by Cloud Manager. Essentially, a topology defines a set of nodes, which is an abstraction of a virtual machine. While defining a node, you can set the values for node attributes, such as operating system, VM shape, disk capacity, and PeopleSoft components to be installed.

The PeopleSoft administrators create topologies for PeopleSoft applications as per their organization requirements. By default, the Cloud Manager is delivered with the following topologies:

- Lift and Shift
- Lift and Shift - DBaaS
- PUM Fulltier

Note: Users are not allowed to delete lift and shift topologies that are used for lift and shift process.

Pages Used to Manage Topology as an Administrator

Page Name	*Definition Name*	*Usage*
Topology Tile	ECL_TOPOLOGY_FL (Content reference for the tile.)	To access the Topology landing page.
Topology Definitions Page	ECL_TOPO_COMP_FL	To create new topologies, edit or clone existing topologies.
Topology Information Page	ECL_TOPO_COMP_FL	Create a new topology.

Topology Tile

Use the Topology tile to access the Topology landing page.

Navigation:

The Topology tile is delivered as part of the Cloud Manager home page.

This example illustrates the Topology Tile.

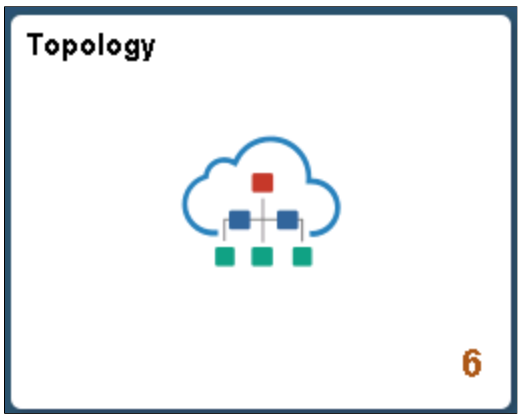

Topology Definitions Page

Use the Topology Definitions page (ECL_TOPOLOGY_FL) to perform the following:

- Create a new topology
- Edit an existing topology
- Clone an existing topology
- Delete an existing topology

Navigation:

Click the Topology tile on the delivered Cloud Manager home page. The Topology Definitions page is displayed.

This example illustrates the fields and controls on the Topology Definitions page. You can find definitions for the fields and controls later on this page.

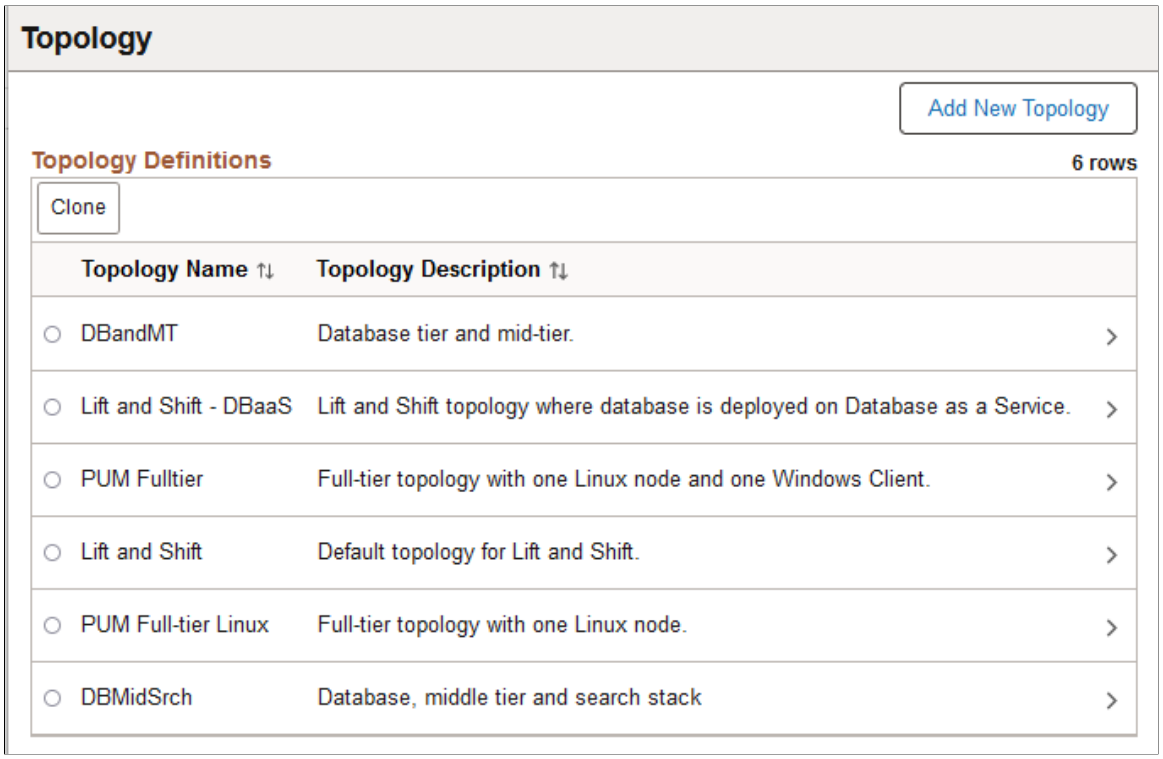

Field or Control	***Description***
Topology Name	Name of the topology.
Topology Description	Description for the topology. Click the arrow on the right (>) to view and edit the definition.
Add New Topology	Click to define a new topology.
Clone	Select an existing topology and click this button to make a copy that you can then edit.

Creating a New Topology

Use the Topology page to create a new topology.

To create a new topology:

1. Click the **Add New Topology** button available on the upper-right corner of the Topology Definitions page.
2. In the Topology Information page, enter the topology name and the corresponding description.

3. Click the **Add Node** button to create a node. This opens the Add Node page.

   Use the Add Node page to set the values for node attributes like Operating System, sizing parameter, disk to be attached, and the PeopleSoft component to be installed.

4. Add additional nodes as needed.

5. Click **Save** to save the details.

Validation Rules for Topology

The following are the set of current validation rules for topology:

If there is a full tier node, then you:

- Cannot have another full tier node.
- Cannot have a middle tier node.
- Cannot have a database node.
- Cannot have a Database as a Service node.

If there is a mid-tier node, then you:

- Cannot have a full tier node.
- Must have either a Database as a Service node or a database (on Compute) node.

If there is a database node, then you:

- Cannot have another database node.
- Cannot have a Database as a Service node.
- Cannot have a full tier node.

If there is a Database as a Service node, then you:

- Cannot have another Database as a Service node.
- Cannot have database node.
- Cannot have a full tier node.

Apart from this, you may have a Windows Client Node in all the above mentioned cases and an optional Search Stack Node.

Topology Information Page

Use the Topology Information page (ECL_TOPO_COMP_FL) to create a new topology.

Navigation:

Click the **Add New Topology** button on the upper-right corner of the Topology Definitions page to access the Topology Information page. Click the **Add Node** button to add one or more nodes. The following sections include descriptions of various node definitions.

This example illustrates the fields and controls on the Topology Information page.

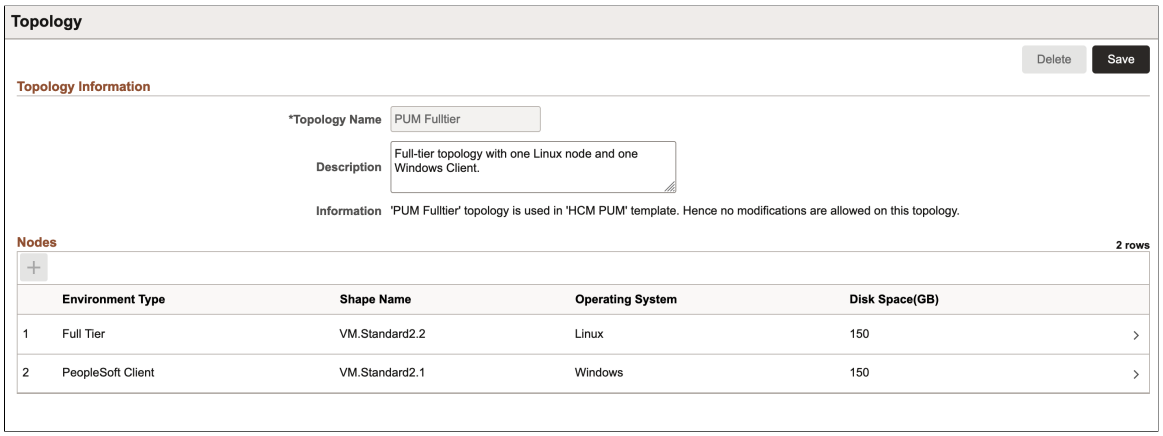

Add Node Page

Use Add Node page to add nodes for creating a topology.

This example illustrates the fields and controls on the Add Node page. You can find definitions for the fields and controls later on this page.

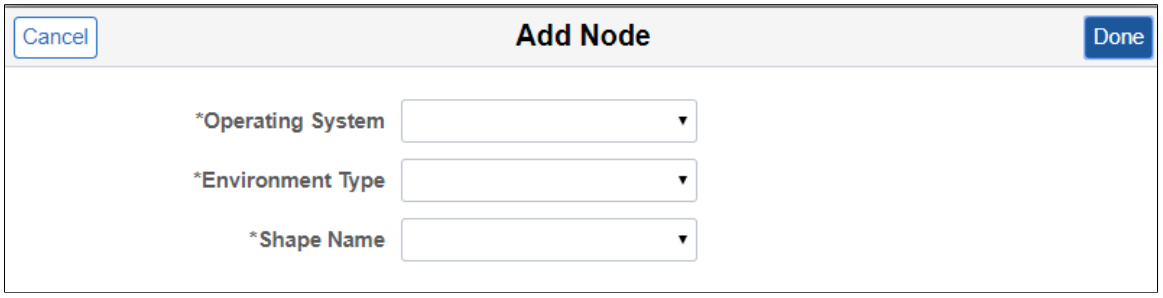

Field or Control	*Description*
Operating System	Select the operating system (Linux or Windows) used to create the topology.

Field or Control	*Description*
Environment Type	Select the PeopleSoft software components to be deployed on the node. The following environment types are available for Linux operating system: • DB Systems: Database systems are dedicated instances running Oracle Linux, optimized for running one or more Oracle databases. A DB System is an Oracle Base Database Service resource. Cloud Manager supports provisioning of databases on OCI DB Systems. Cloud Manager provisions 1 and 2 node Database systems on virtual machines. **Note:** You may also see Database as a Service, or DBaaS, in the documentation and Cloud Manager user interface. The shape for DB systems can be VM or Exadata. See Adding DB Systems Node • Database Tier: Deploys the PeopleSoft database on a compute node. • Search Stack: Sets up OpenSearch, Elasticsearch (ES), and optionally OpenSearch Dashboards or Kibana on the node. See Adding Search Stack Nodes • Full Tier: Deploys Database, Appserver, Webserver and Process Scheduler on the node. Additional features available for Full Tier include COBOL, OpenSearch, and OpenSearch Dashboards (or Elasticsearch and Kibana). • Middle Tier: Deploys Appserver, Webserver and Process Scheduler on the node. If Process Scheduler is selected, you can add the COBOL feature. See Adding Middle Tier Nodes The following environment types are available for Windows operating system: **Note:** For applying PeopleTools patch to an environment, it is mandatory to have a PeopleSoft client or a Windows middle tier node defined. • PeopleSoft Client: Deploys Windows client components on the node. • Middle Tier: Deploys Windows client with process scheduler. Optionally you can add nVision. See Adding Windows Middle Tier Nodes

Field or Control	*Description*
Shape Name	Select the required VM shape. For DB systems, VM shapes and Exadata are supported. **Note:** The Exadata DB System must already exist on OCI. When you select Exadata, the environment (shift) will create the databases within the Exadata DB system. For non-DB system nodes, the list of VM shapes depends on the custom Linux and Windows images that are specified in the "Infrastructure Settings" page. In OCI, whenever a user creates a custom Linux or Windows image, a set of shapes get associated with that image. CM shows that set of shapes, when the end user creates or modifies the nodes in a topology. **Note:** The list of shapes will not appear until you do a Refresh of OCI Metadata after configuring the Operating System images in the Settings page. Some shapes may not be available in new tenancies.
Bootable Volume Size (GB)	Enter the size in GB for the boot volume. The default is the size set by the image. If this option is not enabled, the default is set to 150 GB.
Use Block Volume	Enable this option to use a block volume in addition to the instance boot volume for the node. This option is enabled by default. Disable this option to use only the instance boot volume for the node. **Note:** This option is not available for DB Systems. For a definition of boot and block volume, see Overview of Block Volume in the Oracle Cloud Infrastructure documentation.

Field or Control	*Description*
Disk Space (GB)	Select the amount of disk space attached to the VM instance. **Note:** Assume that if the lifted DPK is K size, then the disk size should be 2.5 times K. **Note:** For DB System, only a limited set of pre-defined disk sizes are supported. The allowed disk sizes are: • 256 GB • 512 GB • 1024 GB • Multiples of 1024 GB **Note:** For BM or Exadata DB system shapes, this field is not visible.
Attachment Type	Select iSCSI or Paravirtualization as the method to attach a block volume to the VM instance. Paravirtualized attachments simplify the process of configuring your block storage by removing the extra commands that are required before connecting to an iSCSI-attached volume. The trade-off is that IOPS performance for iSCSI attachments is greater than that for paravirtualized attachments. For information on these methods, see Volume Attachment Types in the Oracle Cloud Infrastructure documentation.

Adding DB Systems Node

Available shapes for Database Systems (DB Systems) are:

- VM
- Exadata

Note: Before performing an environment shift, you must modify the Lift and Shift - DBaaS topology with the required size and capacity of the database.

Database System Node on Exadata

When you select the shape name as Exadata, you are instructing Cloud Manager to create the environment (shift) with the databases on your existing Exadata DB System on OCI.

When you select Exadata for the DB Systems, disk space is not displayed.

Important! Exadata database is a RAC system which supports multiple nodes (VMs). You need to add the SSH public key for the Cloud Manager user on all nodes. See Adding SSH Keys to a VM Cluster in the Oracle Cloud Infrastructure documentation.

This example illustrates the fields and controls on the Add Node page for DB systems on Exadata.

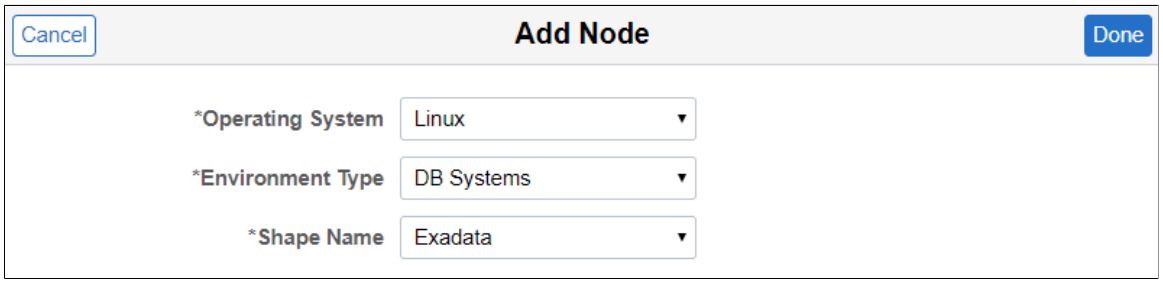

Adding Middle Tier Nodes

When you add a middle tier node, the tiers section is displayed.

This example illustrates the fields and controls on the Add Node page for Middle Tier. You can find definitions for the fields and controls later on this page.

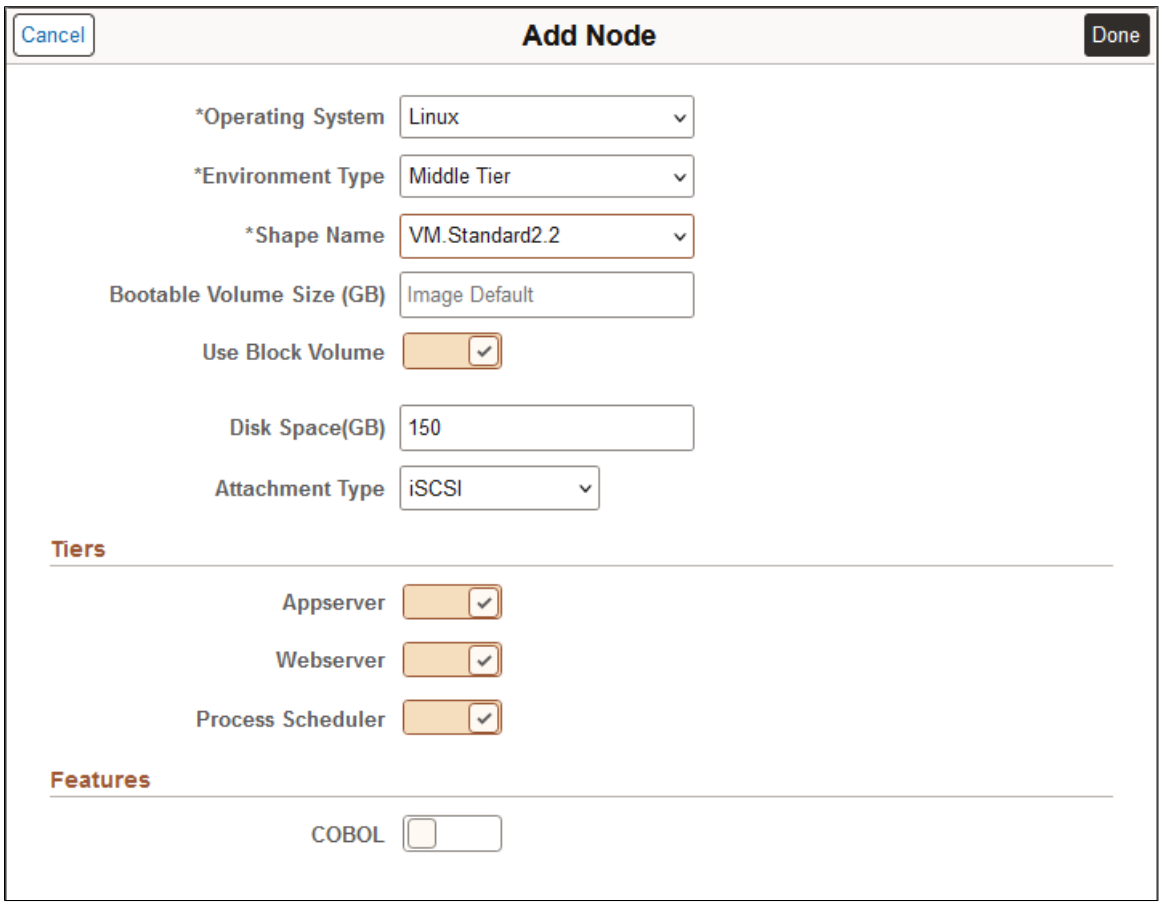

Select the tier or tiers for this node.

This table lists the supported combinations for the tiers:

Application Server	*Web Server (PIA)*	*Process Scheduler*
Yes	Yes	Yes
Yes	No	Yes
Yes	No	No
No	Yes	No
No	No	Yes

You can select one of the available Flex shapes for Linux Middle Tier and Full Tier nodes. A node with a Flex shape requires additional settings.

This example illustrates the fields and controls on the Add Node page for VM.Standard.E3.Flex shape. You can find definitions for the fields and controls later on this page.

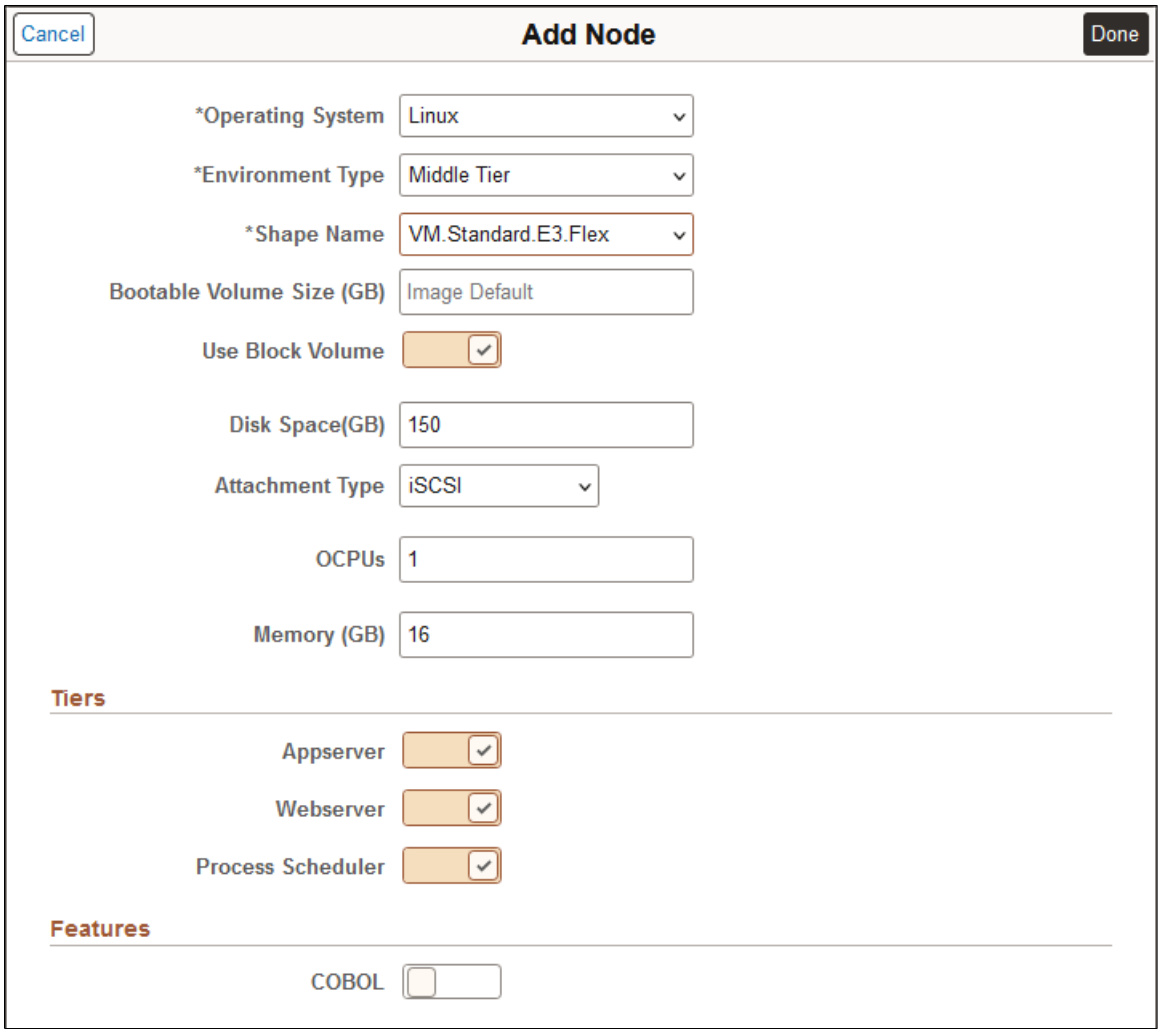

Field or Control	*Description*
OCPUs	Select from 1 to 64 OCPUs.
Memory (GB)	For each OCPU, you can select up to 64 GB of memory, with a maximum of 1024 GB total. The minimum amount of memory allowed is either 8 GB or a value matching the number of OCPUs, whichever is greater.

Adding Nodes with COBOL Enabled

COBOL can be enabled in the node only when environment type is Full Tier or Middle Tier with Process Scheduler enabled.

This example illustrates a middle tier node with Process Scheduler and COBOL enabled.

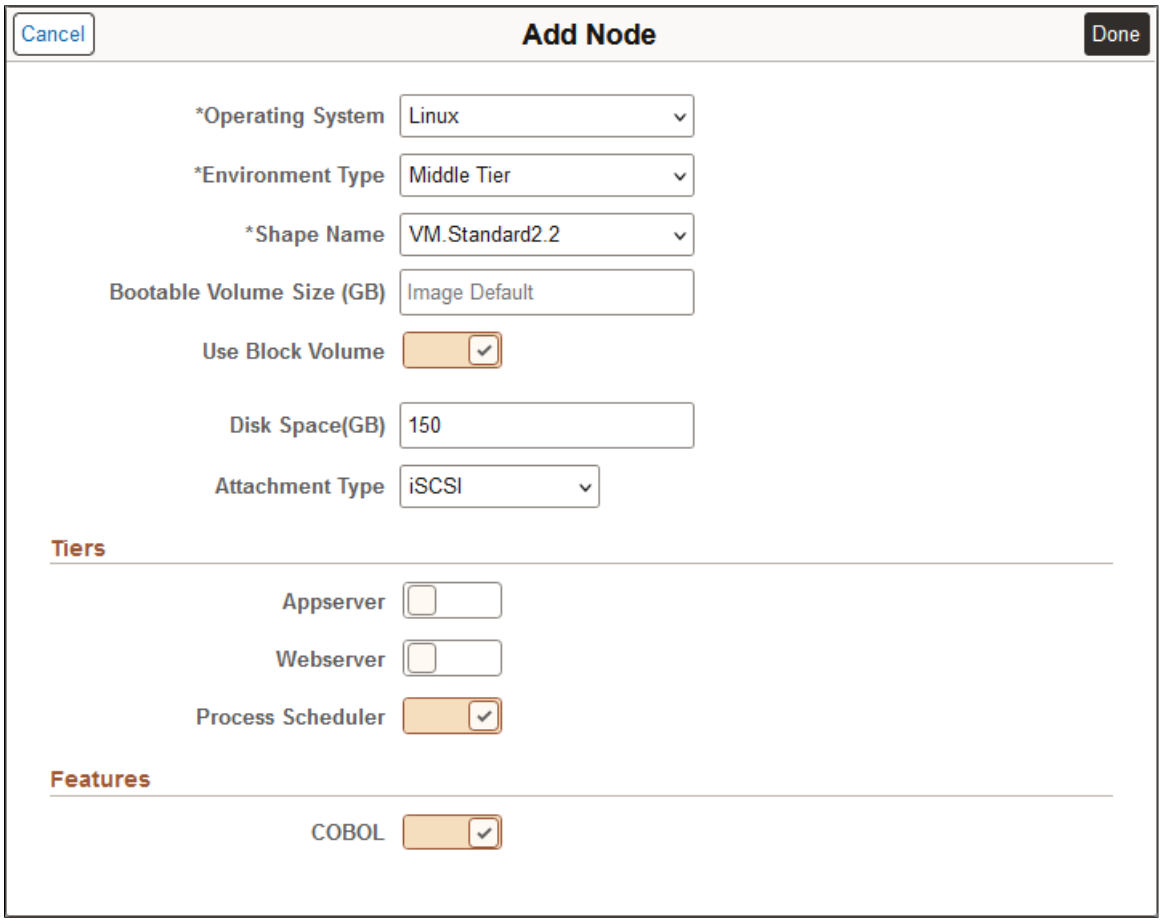

Note: Process Scheduler must be enabled in order to enable COBOL.

Adding Search Stack Nodes

Note the following regarding the components for Search Stack nodes:

- The node configuration of search component is automatically done by Cloud Manager.
- You can select either OpenSearch or Elasticsearch as the search provider, based on the PeopleTools release.

  Select Search to enable either OpenSearch or Elasticsearch, and Dashboards to enable either OpenSearch Dashboards or Kibana. The selection for the search provider is part of Template configuration.

  See Configuring Search Stack General Settings

  For information on how to deploy and configure OpenSearch refer to PeopleSoft Search and Insights Home Page on My Oracle Support (Doc ID 2205540.2) and *PeopleTools: Search Technology*.
- In order to provision Kibana, you must select Elasticsearch.
- In order to provision OpenSearch Dashboards, you must select OpenSearch.
- For environments configured with OpenSearch, OpenSearch Dashboards will be automatically installed as part of the PeopleTools upgrade to 8.60.07.
- OpenSearch and OpenSearch Dashboards are only available with PeopleTools 8.60.07 or later, in addition to Elasticsearch and Kibana. Support is also available for PeopleTools 8.59.21 patch or later. For more information on support, see the PeopleSoft Cloud Manager Home Page (My Oracle Support, Doc ID 2231255.2).
- PeopleTools 8.61 supports only OpenSearch and OpenSearch Dashboards.

This example illustrates the fields and controls on the Add Node page for Search Stack.

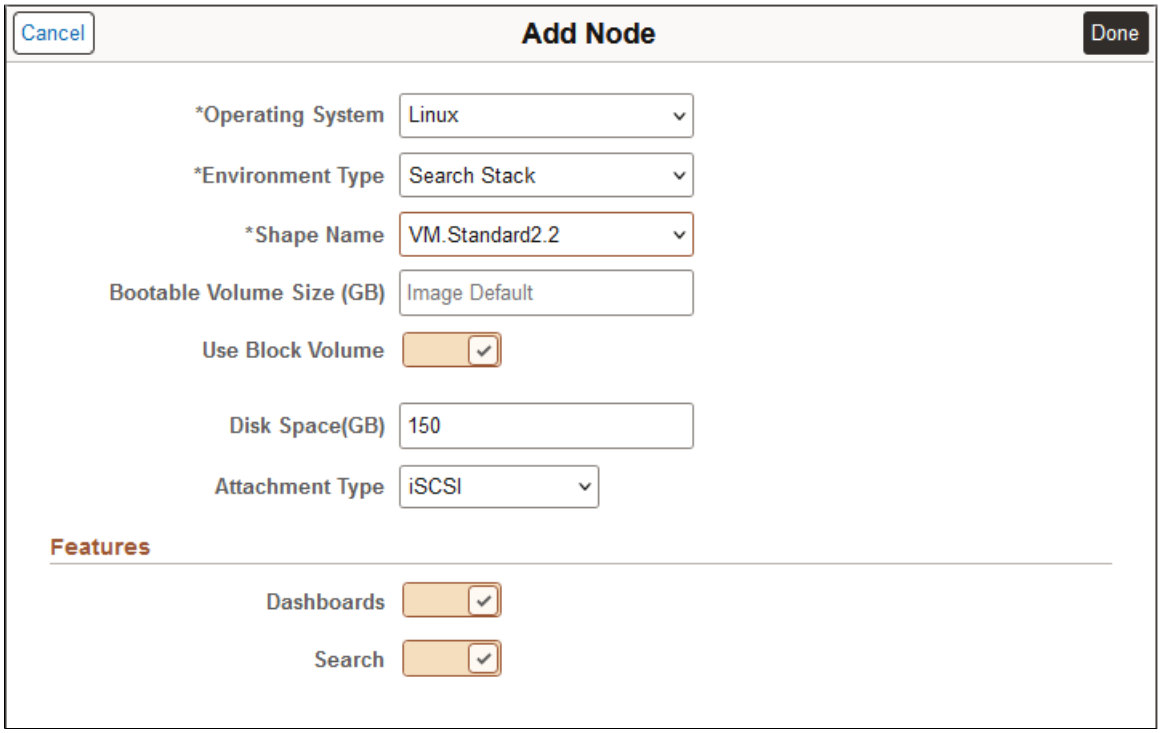

Adding Windows Middle Tier Nodes

Starting with Cloud Manager Update Image 10, you can create multiple Windows middle tier nodes. You can use either a custom Windows image or an OCI platform image. The image is selected on the Infrastructure Settings page.

Prerequisites:

- Configure the Windows Image in Infrastructure Settings.

  See Infrastructure Settings Page

- Subscribe to the Windows channel for your application.

  See Download Subscriptions Page

Use the Add Node page to create a new Windows node.

This example illustrates the fields and controls on the Add Node page with Windows Operating System selected. You can find definitions for the fields and controls later on this page.

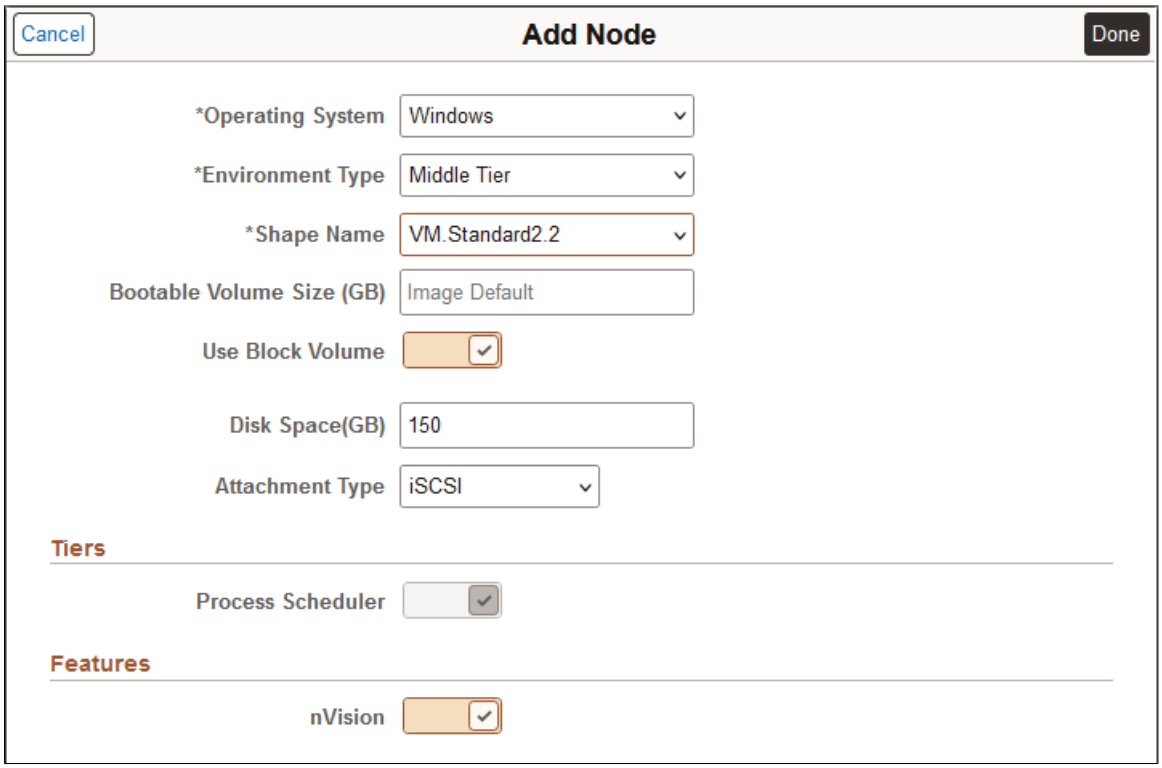

Field or Control	*Description*
Operating System	Select Windows.
Environment Type	Select Middle Tier.
Shape Name	Select the required VM shape.

Field or Control	*Description*
Bootable Volume Size (GB)	Enter the size in GB for the boot volume. The default is the size set by the image. If this option is not enabled, the default is set to 150 GB.
Use Block Volume	Enable this option to use a block volume in addition to the instance boot volume for the node. This option is enabled by default. Disable this option to use only the instance boot volume for the node. For a definition of boot and block volume, see Overview of Block Volume in the Oracle Cloud Infrastructure documentation.
Disk Space (GB)	Select the amount of disk space attached to the VM.
Attachment Type	Select iSCSI or Paravirtualization.
Tiers	Only Process Scheduler domain is supported for Windows middle tier Node. Windows Process Scheduler is required for running nVision reports.
Features	The nVision feature can be set either on or off.

Creating Multi Domain and Multi Middle Tier Node Configurations

Cloud Manager Update Image 8 and above supports both horizontal and vertical elasticity by allowing multi node and multi domain configuration for the PeopleSoft environment.

Cloud Manager supports multiple middle tier and PIA domains on single node as well on multiple nodes. It also provides an option to enable Integration Broker in one application domain in an environment. This Domain Configuration feature extends the existing provisioning environment feature in Cloud Manager.

See Configuring AppServer Tier Domain Settings.

This example illustrates the fields and controls on the Topology page for a topology with multiple middle tiers.

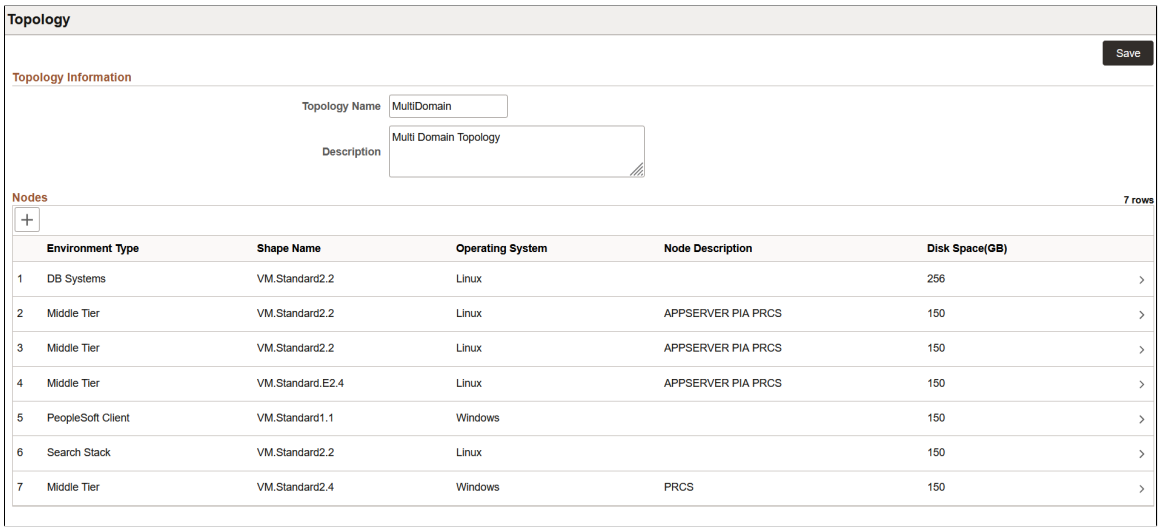

Editing an Existing Topology

To edit an existing topology, perform the following:

1. Click any existing topology in the Topology page. This displays the definition page for the topology that you want to edit.

   Note: No modifications are allowed to Lift and Shift - DBaaS topology.

   This example illustrates the fields and controls on the Topology definition page. You can find definitions for the fields and controls later on this page.

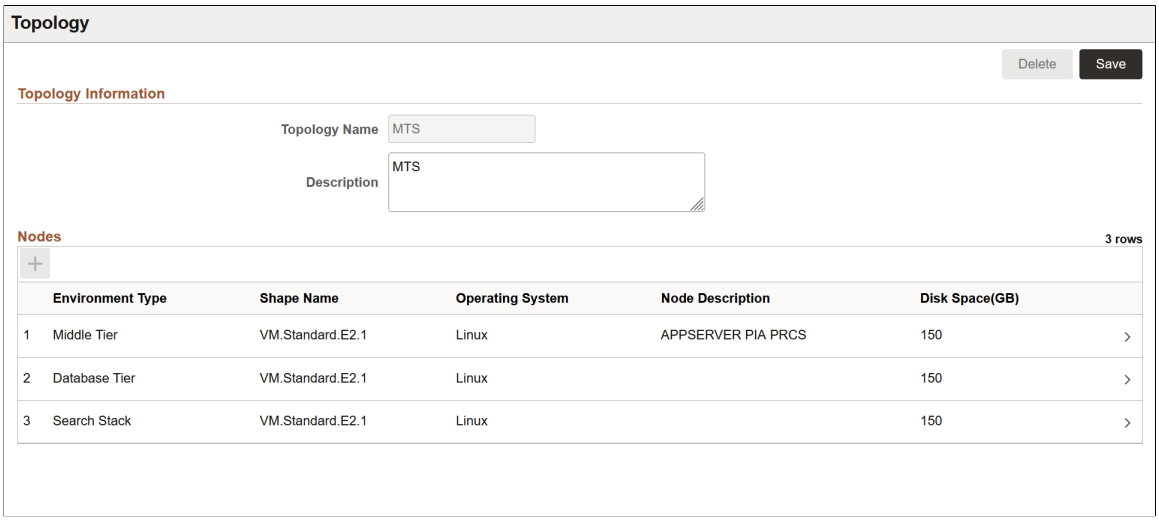

2. You can edit the description, if required.

3. Click + to add new nodes.

See Add Node Page.

4. To edit any node attribute value, click on any node row. This displays the Edit Node window.

   This example illustrates the fields and controls on the Edit Node window for a Middle Tier node.

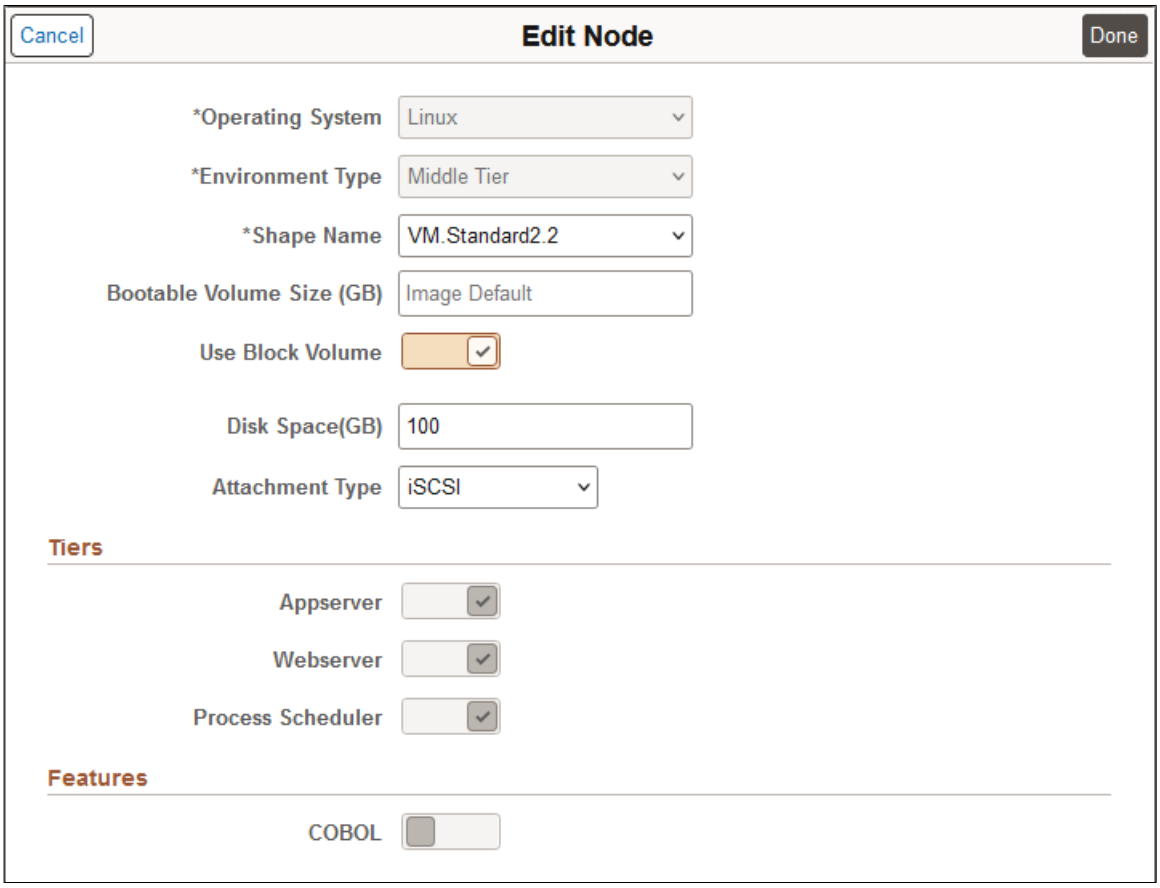

5. Edit the fields to satisfy your requirements.
6. Click Done to save the edited details.

Cloning an Existing Topology

To clone an existing topology, perform the following:

1. Select the radio button corresponding to a topology that you want to clone.
2. Click Clone button in the Topology page. This displays the Clone Topology window.

This example illustrates the fields and controls on the Clone Topology window.

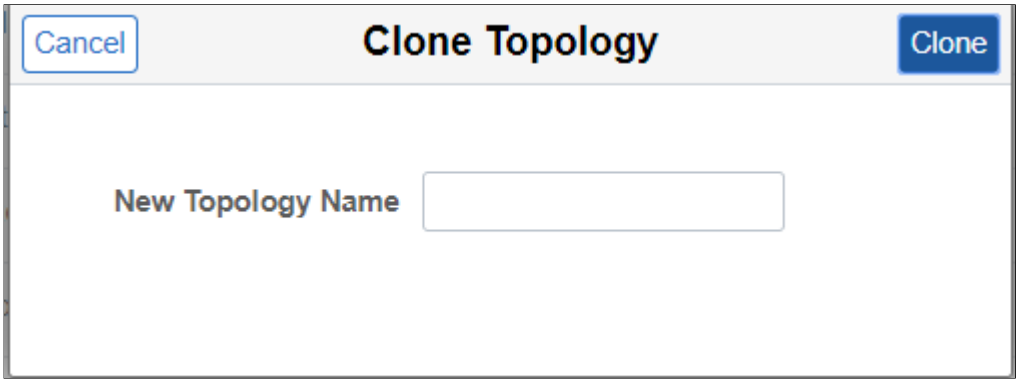

3. Enter a new topology name and click Clone. The new topology is added to the topology list.

Deleting an Existing Topology

To delete an existing topology, perform the following:

1. Click any existing topology in the Topology page. This displays the Topology Definition page of the topology.
2. Click Delete, to delete the topology.

This example illustrates the fields and controls on the Example Delete Topology page.

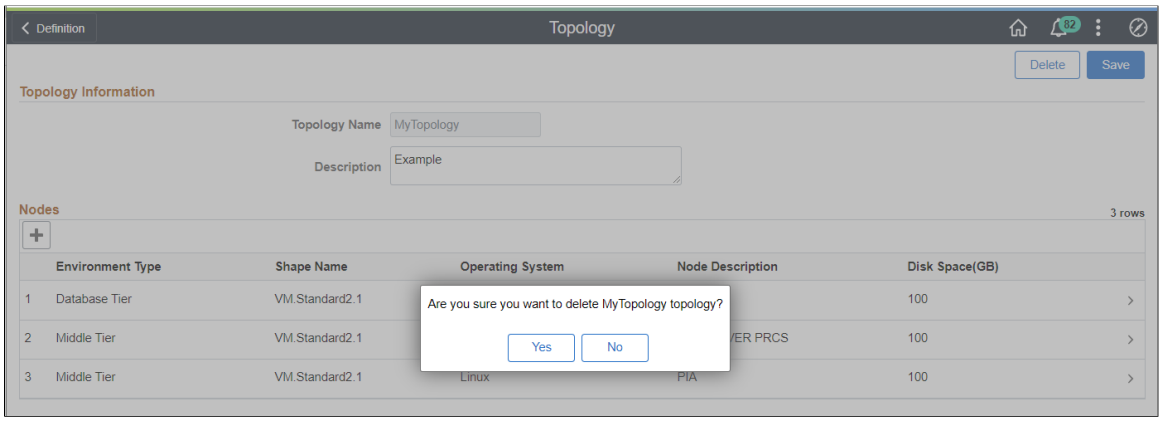

Chapter 5

Managing Templates

Template Overview

An environment template is a repeatable blueprint that is used to deploy PeopleSoft environments using Cloud Manager. A template defines the topology to be used when deploying the PeopleSoft application DPK, which gets downloaded to the Repository. A template also defines environment attributes to enable streamlined deployments. Access to templates can be managed by defining security attributes of the templates.

Page Name	***Definition Name***	***Usage***
Environment Template Tile	ECL_TEMPLATE_LP_FL_GBL (Content reference for the tile.)	Access the Environment Template landing page.
Environment Template Page	ECL_TEMPLATE_FL	Create new templates or edit, delete or clone existing templates.
Environment Template – General Details Page	ECL_TEMPL_GEN_FL	Enter the template name, description, and selecting a database.
Environment Template – Select Topology Page	ECL_TEMPL_TOP_FL	Select the topology that you have already defined.
Environment Template – Security and Policies Page	ECL_TEMPL_SEC_FL	Associate zones in which the environment is created, the roles that have access to the template, and policies that will be auto-enabled for the environment.
Environment Template – Summary Page	ECL_TEMPL_REV_FL	Displays the summary of the environment template that the user is about to create.

Environment Template Tile

Use the Environment Template tile (ECL_TEMPLATE_LP_FL_GBL) to access Environment Template landing page.

Navigation:

The Environment Template tile is delivered as part of the Cloud Manager home page.

This example illustrates the Environment Template Tile.

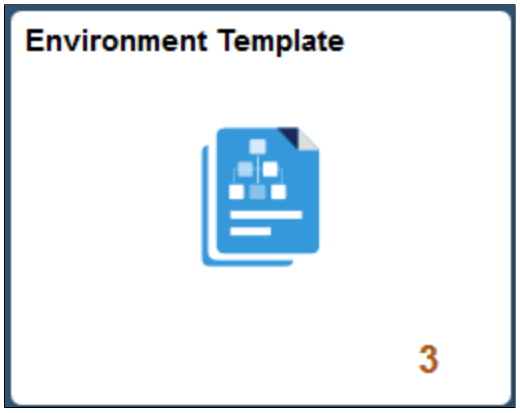

Environment Template Page

Use the Environment Template page (ECL_TEMPLATE_FL) to create a new template, and edit, delete or clone an existing template.

Navigation:

Click the Environment Template tile on the delivered Cloud Manager Fluid home page. The Environment Template page is displayed by default.

This example illustrates the fields and controls on the Environment Template page. You can find definitions for the fields and controls later on this page.

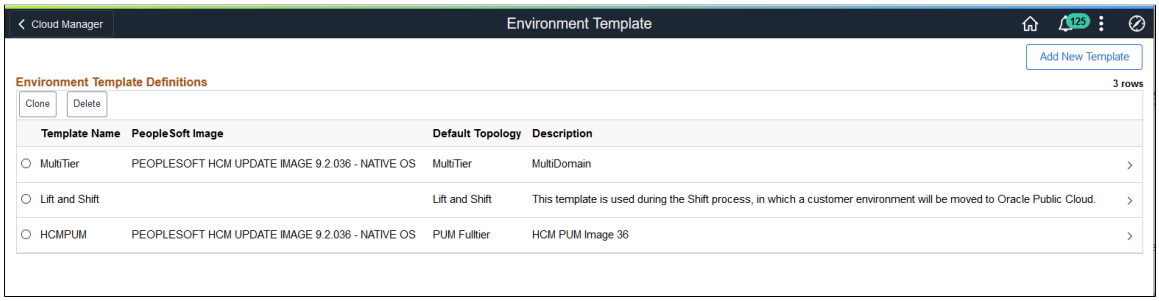

Note: The Lift And Shift template is the default template displayed in the Environment Template page with no database associated with it.

Field or Control	***Description***
Template Name	Name of the template.
Database	Indicates the PeopleSoft application DPK that gets installed when the template is deployed.
Default Topology	Default topology associated with the template.

Field or Control	*Description*
Description	Add a description for the template.

Edit, Delete, or Clone an Existing Template

User can edit, delete or clone the existing templates using the Environment Template landing page.

Note: It is recommended to recreate the existing templates to ensure that the new custom attributes are available in the template.

- To edit an existing template details, click a row and modify the details as per requirement.
- To delete an existing template, select the radio button corresponding to the template which you want to delete and click the Delete button. Users cannot delete a template, if it is already used for defining an environment.
- To clone an existing template, select the radio button corresponding to the template which you want to clone and click the Clone button available on the Environment Template landing page. The Clone Template modal window is displayed, wherein you can enter the new template name and click the Clone button. The new template is added to the template list.

Default Environment Templates

A default template is provided for Lift and Shift, which is used during environment shifting by default. This Lift and Shift template and its associated topology must be modified such that it is suitable for the environment being shifted. The Lift and shift topology is fixed in terms of number of nodes, but the shape and disk space parameters can be modified. For any environment to be provisioned in Cloud Manager, the administrator creates a template and a user uses that template to provision. In case of Lift and Shift, a default template is provided out of the box and there is no need to create any templates. When an administrator creates an environment on the Lift and Shift page, the process automatically chooses the default Lift and Shift template. This Lift and Shift template must be modified to suit the environment being shifted. For more details, see Understanding the Lift and Shift Process

Creating a Template

Use the Environment Template wizard to create a new template using a step by step guided process.

By default, the create template guided process involves the following steps:

1. Entering general details.

   See Environment Template – General Details Page

2. Selecting topologies.

   See Environment Template – Select Topology Page

3. Defining security and policies.

See Environment Template – Security and Policies Page

4. Submitting the details.

   See Environment Template – Summary Page

Environment Template – General Details Page

Use the Environment Template – General Details page to enter the template name, description, and select a database.

Navigation:

Click the Add New Template button on the Environment Template landing page.

This example illustrates the fields and controls on the Environment Template – General Details page. You can find definitions for the fields and controls later on this page.

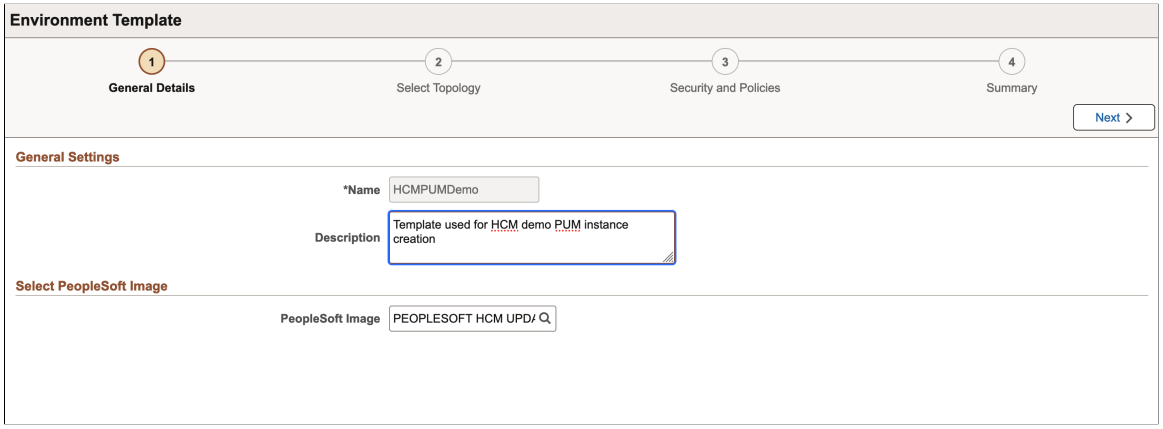

Field or Control	*Description*
Name	Name of the template which you want to create. It is a required field.
Description	Add a description for the template.
PeopleSoft Image	Select a PeopleSoft application DPK from the list of DPKs available in the Repository.

Environment Template – Select Topology Page

Use the Environment Template – Select Topology page to select the topology that you have already defined. You may edit the default attributes associated with the selected topology.

Navigation:

- Click Next in the Environment Template – General Details page.

- Click Step 2, Select Topology, at the top of the page.

This example illustrates the fields and controls on the Select Topology page. You can find definitions for the fields and controls later on this page.

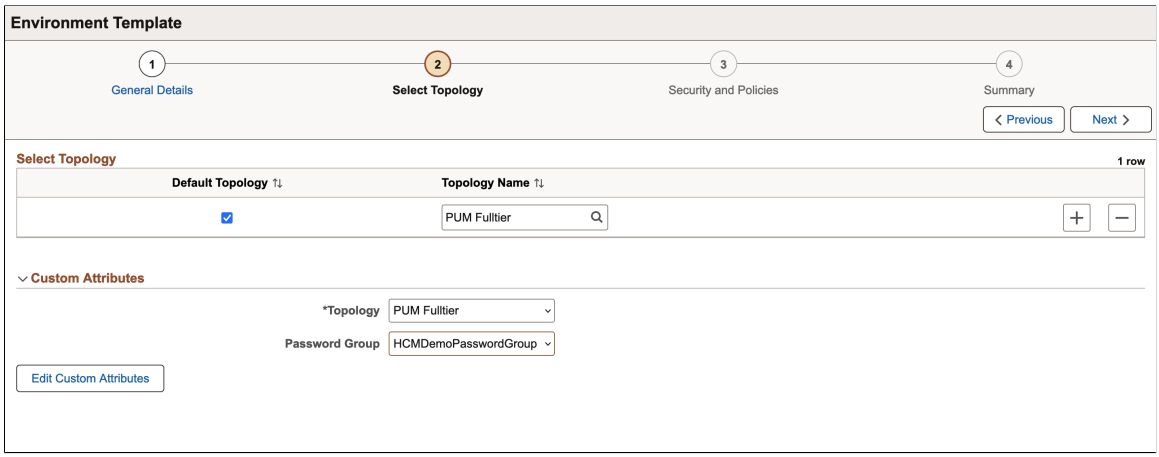

Field or Control	*Description*
Default Topology	Users can mark one of the topology associated with the template as the default topology. During the environment creation process using a template, you can override this default topology and select any other topology associated with that template. If you don't want to override, then the default topology will get used automatically. Click + to add more topologies. A new row of empty fields appears below the existing record. You can configure the fields based on the requirements. **Note:** Be sure to select the topology under the Override Topology section and then continue with the template creation.
Topology Name	Select the required topology that you want to include in the template. **Note:** While selecting a topology, the custom attributes associated with the selected topology is displayed. It is possible to override the default attributes based on the requirements.

If the selected topology that includes the Search Stack is selected, a Monitoring Services section is available to enable monitoring.

Note: Monitoring must be enabled for auto scaling. See Setting Up Auto Scaling.

This example illustrates the fields and controls on the Select Topology page when the topology includes Search Stack.

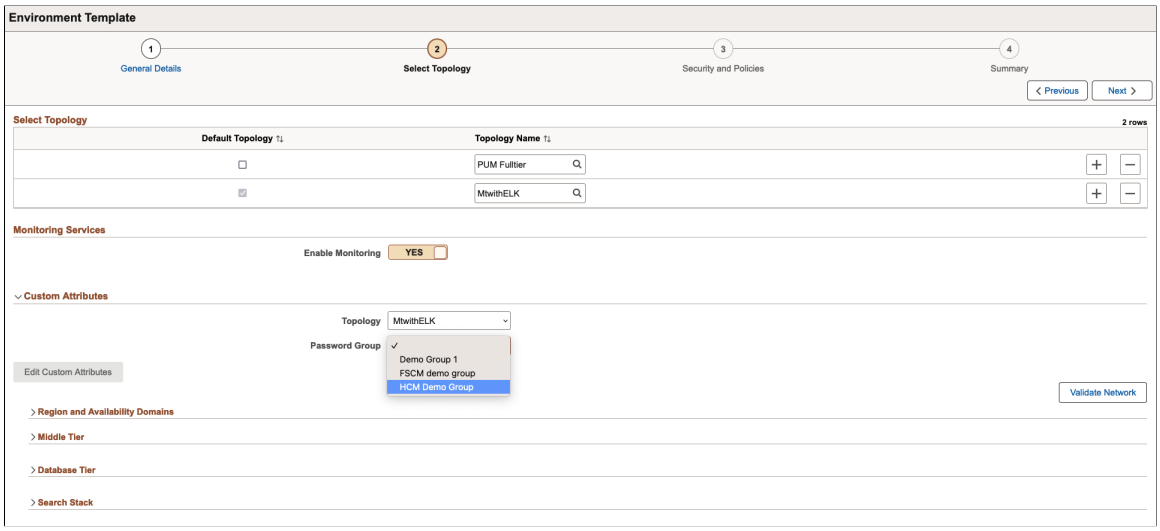

Configuring Custom Attributes

1. Expand the Custom Attributes section.
2. Select the required topology.
3. Select the required password group. See Password Groups.
4. Enter the required attributes and click Next.

   Note: Cloud Manager allows users to add customization during template creation under Edit Custom Attributes section. This customization can be added only to middle tier and database tier. The customization will be available to users when they select this template. This facilitates the user to define custom attribute values for the environment being deployed.

5. After entering the required attributes, click the Validate Network button to ensure your infrastructure settings are correct and the network is valid.

   This validates whether the port is open for an incoming/outgoing connection across different subnets and VCN. The connection can be:

   a. From Cloud Manager to VM

   b. From VM to Cloud Manager

   c. From VM to VM

   Some of the validations are based on user input (Jolt Port, WSL Port, Database Server Port, HTTP PIA Port, HTTPS PIA Port).

   The following implicit validations are performed:

From	*To*
Every subnet	Cloud Manager subnet: NFS ports TCP - 2049, 111, 892, 32803
Cloud Manager subnet	Every subnet (including itself) except Windows VM subnet: ssh port 22
Cloud Manager subnet	Windows VM's subnet: WinRM TCP Ports 5985, 5986
Cloud Manager subnet	Windows VM's subnet: CIFS ports TCP 445, 139, 138, 137

The following sections must be configured in every template.

- Region and Availability. See Configuring Region and Availability Domains
- Network for node. See Configuring Network Settings
- Network Security Groups (optional) for node. See Configuring Network Security Group Settings
- Fault domain for the node. See Configuring the Fault Domain

  Note: DB Systems node does not have fault domain setting.

- Advanced section. See Configuring Advanced Section

  Note: Add Node operation and Search nodes do not include an Advanced section.

The General Setting and Domain Settings depend on the type of node.

- Fulltier. See Configuring Full Tier Template Settings.
- DB Systems. See Configuring DB Systems Settings.
- Web Server. See Configuring Web Server Tier Settings.
- Application Server. See Configuring AppServer Tier Domain Settings.
- Process Scheduler. See Configuring Process Scheduler General Setting and Configuring Process Scheduler Domain Settings.
- Windows. See Configuring Windows Middle Tier General Settings.
- PeopleSoft Client. See Configuring PeopleSoft Client General Settings.
- Database Tier. See Configuring Database Tier.
- Search Stack. See Configuring Search Stack General Settings.

Configuring Region and Availability Domains

This example illustrates the fields and controls on the Region and Availability Domains section. You can find definitions for the fields and controls later on this page.

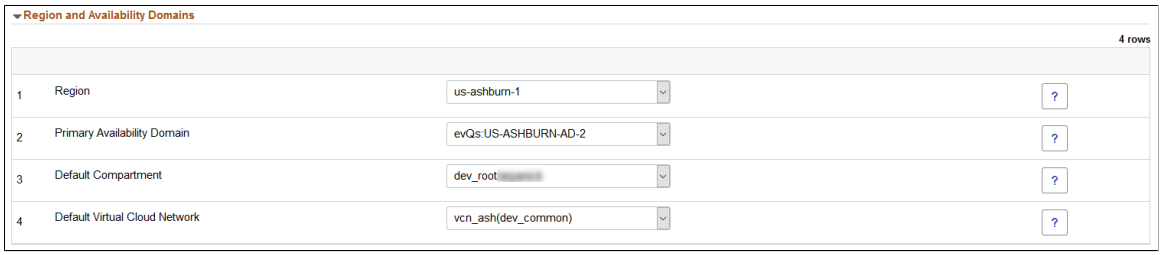

Field or Control	*Description*
Region	A region is a localized geographic area, and an availability domain is one or more data centers located within a region. A region is composed of several availability domains. **Note:** Cloud Manager will provision and manage environments only in the region where it is deployed. See Infrastructure Settings Page.
Primary Availability Domain	Availability domain in OCI.
Default Compartment	Compartments allow you to organize and control access to your cloud resources. A compartment is a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. Select a default compartment, this compartment will be used for all tiers in the template.
Default Virtual Cloud Network	Virtual Cloud Network within OCI. A virtual cloud network is a virtual version of a traditional network—including subnets, route tables, and gateways on which your instances run. Select the default Virtual Cloud Network, this will be the default VCN used for all tiers in the template.

For details on setting up the OCI environment, refer to the tutorial Prepare to Install PeopleSoft Cloud Manager.

Note: In OCI, the templates will not have any default values for Region and Availability Domains section. All templates must be updated with these settings before the template can be used to deploy an environment. See Creating an Environment.

Configuring Tagging

Tags are created in a Compartment in OCI. Tags are optional and can be added to an environment. To create tags in OCI see the tutorial on creating tags, at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

Use the Tagging section to associate a tag with a template or node in the template.

Note: If default tags are created in a compartment in OCI, the tags are automatically applied to all resources in that compartment at the time of creation, regardless of the permissions of the user who creates the resource.

This example illustrates the fields and controls in the Tagging section for a template or node in a template. You can find definitions for the fields and controls later on this page.

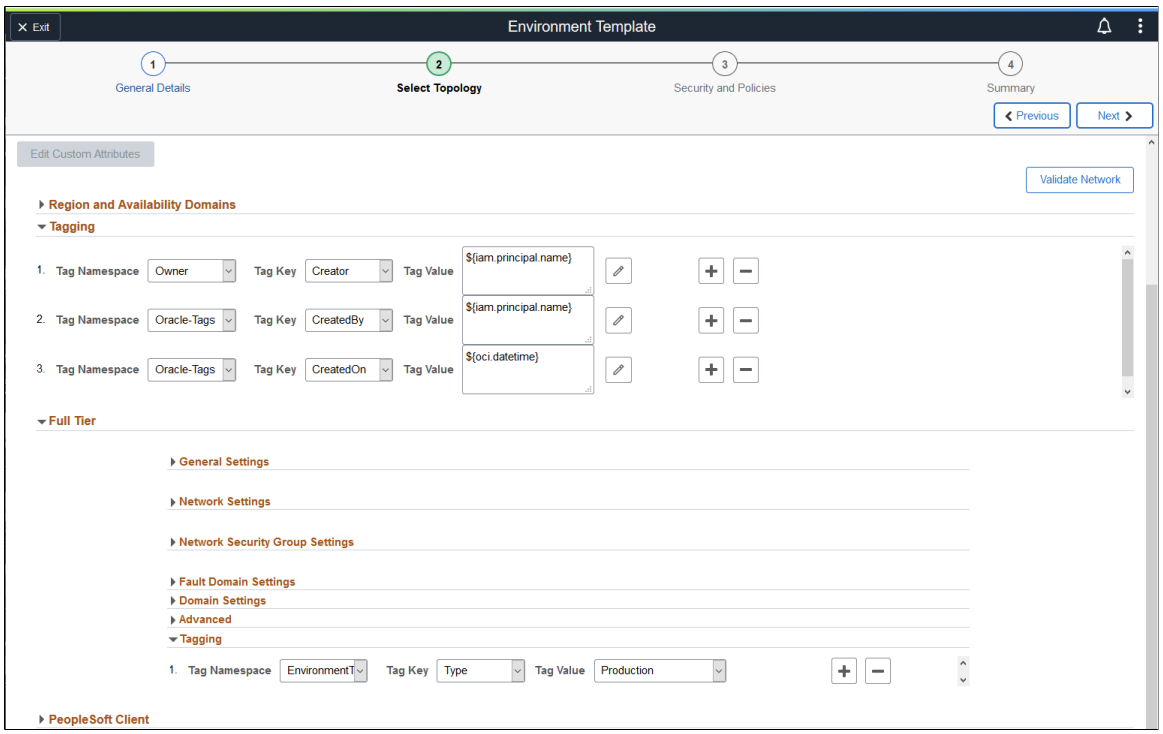

Field or Control	*Description*
Tag Namespace	Created in OCI as a container for the tag keys. This is a natural grouping that can be used to apply a policy.
Tag Key	Tag keys are created within a Tag Namespace. This is the name that refers to the tag.

Field or Control	*Description*
Tag Value	The tag value is the value that the user applying the tag adds to the tag key. This may be a list of values or user may be allowed to enter any value when the tag is applied. Tag values are listed below.

OCI Tag Variables

This table lists OCI tag variables.

Variable	*Description*
${iam.principal.name}	The name of the principal that tagged resource.
${iam.principal.type}	Type of principal that tagged resource.
${oci.datetime}	The date and time that the tag was created.

Cloud Manager Tag Variables

The following variables are supported for Cloud Manager.

Variable	*Description*
${cloudmanager.env_name}	Name of provisioned Cloud Manager environment.
${cloudmanager.app_type}	Application type (HRMS, FSCM, and so on).
${cloudmanager.tools_ver}	PeopleTools version.
${cloudmanager.app_ver}	Application Version.

Configuring Network Settings

Each node in the template must be configured for subnet where it will run. The compartment and VCN default to the values entered in the Region and Availability Domain section, these values can be changed to run in a different compartment or VCN. For more information on VCN peering see the tutorial *Use Custom or Private Network Resources with PeopleSoft Cloud Manager* at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

This example illustrates the fields and controls on the Network Setting section. You can find definitions for the fields and controls later on this page.

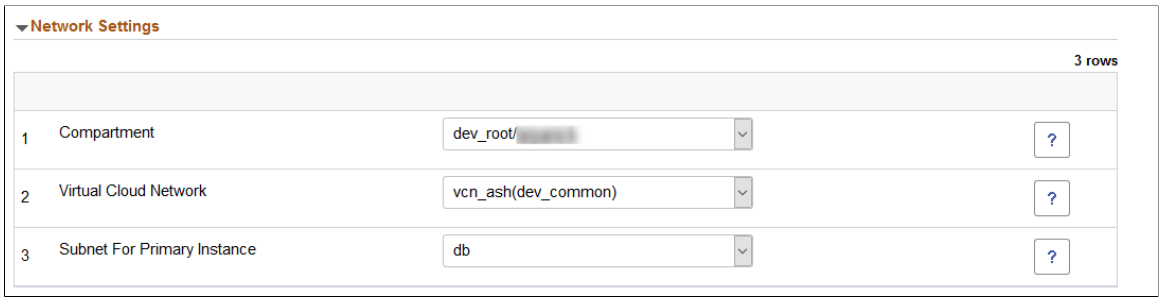

Field or Control	*Description*
Compartment	The default compartment is displayed, you can select a different compartment for the network.
Virtual Cloud Network	The default VCN is displayed, you can select a different VCN for this node.
Subnet For Primary Instance	Select the subnet where this node will run.

Configuring Network Security Group Settings

Cloud Manager supports Network Security Groups (NSG). Network Security Groups are set up outside of Cloud Manager. Up to 5 Network Security Groups can be assigned to a node.

See tutorial *Use Custom or Private Network Resources with PeopleSoft Cloud Manager*.

This example illustrates the fields and controls on the Network Security Group Settings section.

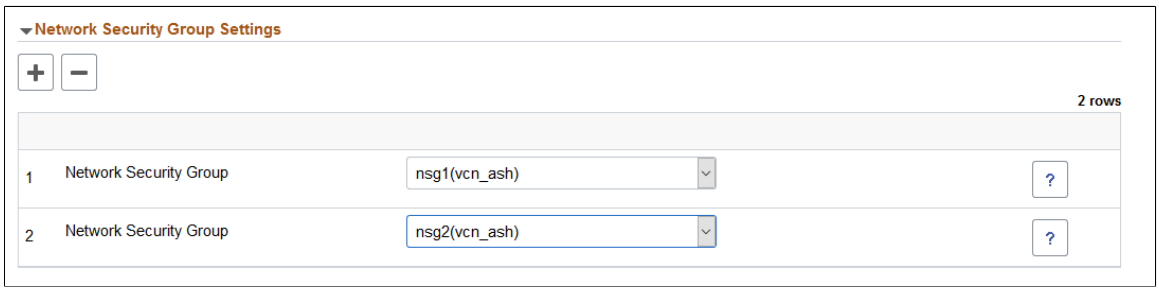

Configuring the Fault Domain

Each availability domain in OCI contains three fault domains for high availability. OCI randomizes the availability domains by tenancy to help balance capacity in data centers.

Use the Fault Domain Settings section to select which fault domain to use for a specific node. Select the fault domain from the drop down list, only fault domains in the availability domain for the node are listed and available.

If OCI Default is selected, OCI randomly selects the fault domain to use and customer is not aware of which fault domain is being used.

Note: DB Systems node does not have fault domain settings.

For more information on faults domains see Regions and Availability Domains.

This example illustrates the fields and controls on the Fault Domain section.

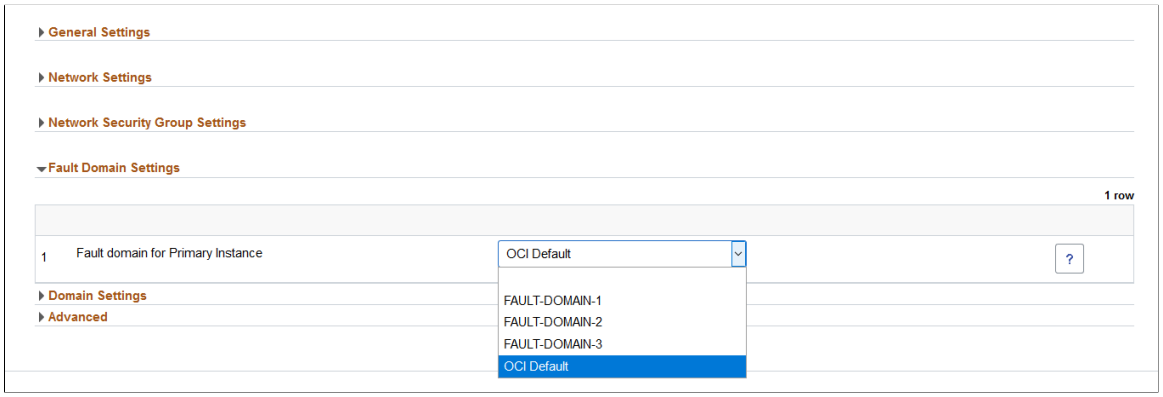

Configuring Advanced Section

The Advanced section contains kernel parameters and Ulimit values. Expert administrators have the option to change these values.

Important! No validation is performed on these fields. Only operating system experts should change these values.

Note: Add Node operation and Search nodes do not include an Advanced section.

This example illustrates the fields and controls on the Advanced section.

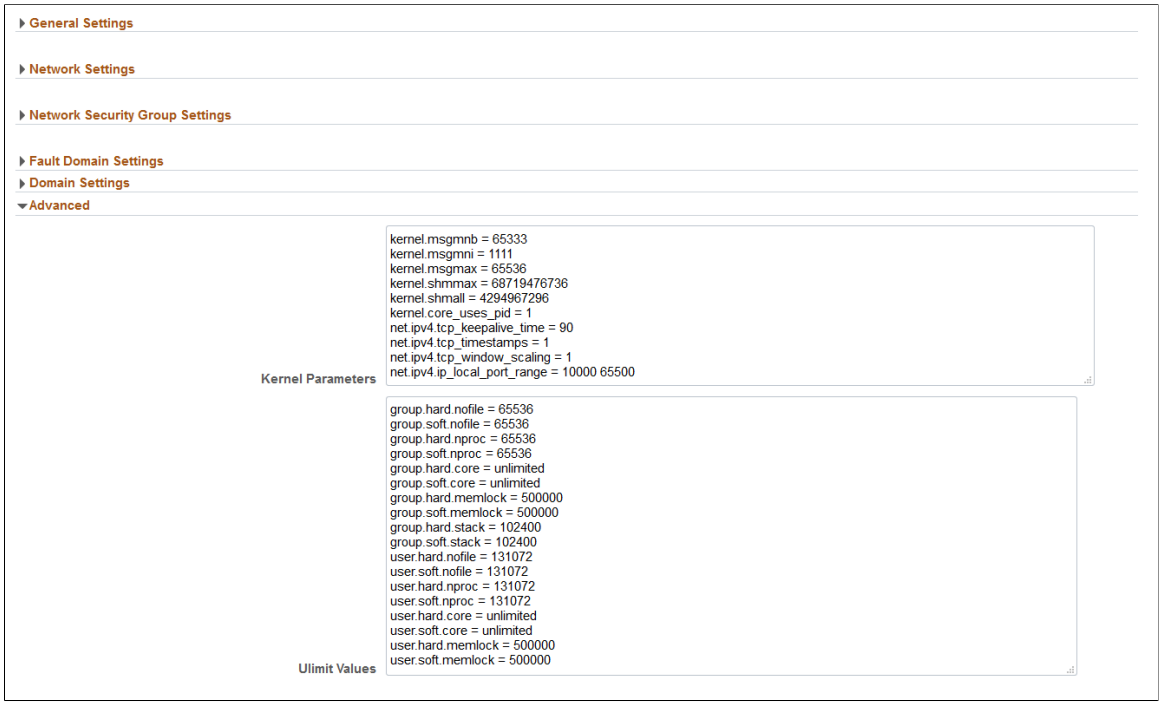

Ulimits

This input field allows you to set ulimit values for the Linux node. Each line of the input represents one ulimit entry. The format to be used for each line is as follows:

```
[user|group].[soft|hard].<limit name> = <value>
```

The left side of the expression represents a ulimit, and the right side of the expression is a value of that ulimit. They are separated by an equal sign (=).

Field or Control	***Description***
user or group	The key word “user” represents a limit for a Linux user, and the key word “group” represents a limit for a user group. Exactly which users and groups are added to the Linux configuration files is internally determined by the environment deployment process. You cannot set those actual user/group names here.
soft or hard	The keyword “soft” denotes the limit a process can use. It can later be increased by the corresponding user or group. The keyword “hard” denotes the maximum limit to which the soft limit can be raised to.
limit name	This is the name of the limit. Refer to the table below for the list of names

Field or Control	*Description*
value	This is either a numeric value or the string “unlimited”. Refer to table below for a list of values and the units in which they have to be expressed

This table lists the ulimit names.

Name	*Unit*
cpu	Seconds
fsize	Blocks
data	Kilobytes
stack	Kilobytes
core	Blocks
rss	Kilobytes
nofile	Number of file descriptors
as	Kilobytes
nproc	Number of processes
memlock	Kilobytes
lock	Number of locks
sigpending	Number of queued signals
nice	Nice level (an integer)
rtprio	Realtime priority (an integer)
rttime	Number

Kernel Parameters

The kernel parameters input field can be used for setting Linux kernel parameters. Each line represents one Linux parameter. The format of each line should be as follows:

```
<kernel parameter name> = <value>
```

The list of all kernel parameters can be found from the main page of the Linux sysctl command.

Configuring Full Tier Template Settings

Example Full Tier Node with Search Stack and COBOL

Based on the topology selected, the full tier contains the appropriate settings for the node. The examples in this section display full tier topology that includes Search Stack and COBOL.

This example illustrates the fields and controls on the Full Tier - General Settings section (1 of 2). You can find definitions for the fields and controls later on this page.

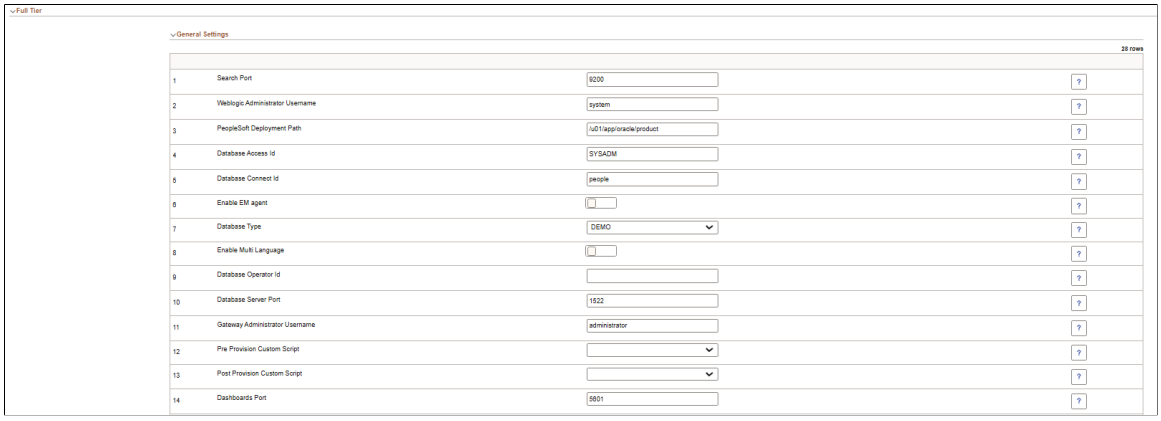

This example illustrates the fields and controls on the Full Tier - General Settings section (2 of 2). You can find definitions for the fields and controls later on this page.

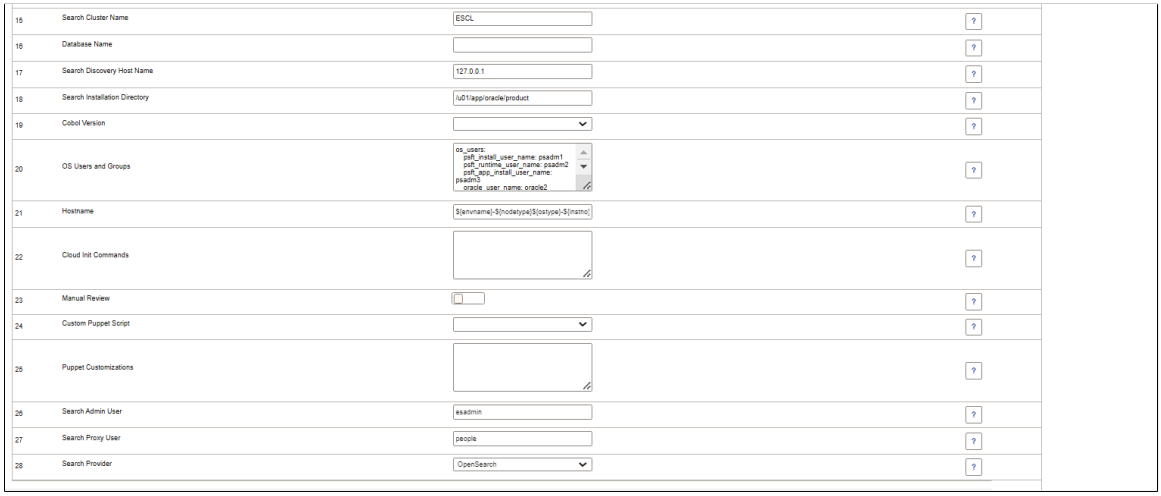

Field or Control	*Description*
Search Port	Elasticsearch or OpenSearch port.
Weblogic Administrator Username	User name of the Oracle WebLogic administrator. This is used for accessing WebLogic console.
PeopleSoft Deployment Path	Location where the PeopleSoft application is deployed. **Note:** PeopleSoft Deployment Path must not end with a slash.
Database Access Id	Access ID for the database. The default is SYSADM.
Database Connect Id	Connect ID for the database. The default is people.
Enable EM agent	Select either Yes to enable Environment Management agent, or No to disable the EM agent.
Database Type	Select the required database type. Available database types are DEMO or SYS.
Enable Multi Language	Select either Yes to enable, or No to disable, multi language support.
Database Operator Id	Database Operator ID. For HCM, CS, and ELM, the default is PS. For FSCM, IH, and CRM, the default is VP1.
Database Server Port	Listener port number.
Gateway Administrator Username	User name of the Integration Gateway administrator.
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page.

Field or Control	*Description*
Post Provision Custom Script	Select an uploaded script to run post provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page.
Dashboards Port	Search dashboards port.
Search Cluster Name	Name of the search cluster.
Database Name	Name of the database.
Search Discovery Host Name	Host names for any nodes that are already members of a cluster. Host names (or IP/DNS) are required for letting each search server (Elasticsearch or OpenSearch) know where it can ping and find other search servers during booting up.
Search Installation Directory	Installation directory for search components.
Cobol Version	If COBOL is enabled in the selected topology, you must select the Cobol version. **Note:** The COBOL license must be configured on the Cloud Manager Settings page. See Cloud Manager Settings Page

Field or Control	*Description*

OS Users and Groups	Use this field to specify custom users and groups to set up the full-tier instance.
	
	`os_users:`
	`  psft_install_user_name: psadm1`
	`  psft_runtime_user_name: psadm2`
	`  psft_app_install_user_name: psadm3`
	`  oracle_user_name: oracle2`
	`  psft_search_user_name: esuser`
	`  remote_login_user: opc`
	`os_user_groups:`
	`  psft_runtime_group_name: psft`
	`  psft_app_install_group_name: appinst`
	`  oracle_install_group_name: oinstall`
	`  oracle_runtime_group_name: dba`
	
	The field is populated with the default Linux users and groups in the form of a YAML customization field. All entries are mandatory. Do not modify the format. The format (for example, spacing and punctuation) and values are validated when you move to another field on the page.
	
	User names or groups should be between 1 and 32 characters long and may contain only lower and upper case letters, digits (1-9), periods (.), underscores (_), or dashes (-). They can end with a dollar sign ($). Dashes and periods are not allowed at the beginning of the user name. Fully numeric names are not allowed.
	
	The remote_login_user is required to access the instance. If you want to specify a remote_login_user other than opc, you must first create a custom Linux image with the custom user. The custom user must have root (sudo) privileges.

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames are validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.

Cloud Init Commands	Enter the OS level settings like package installation or OS upgrade/ update that you need to customize when you create the infrastructure. The commands/statements must be separated by semicolon (;).
	
	Some sample commands:
	
	`sudo yum -y install nfs-utils;`
	`sudo yum -y install rpcbind;`
	`sudo yum -y install glibc.i686;`

Manual Review	Select to enable manual review. On enabling the Manual Review field, new manual stop steps are added to Customer DPK processing activity as the final step. This field is disabled by default. When the field is enabled, the activity execution pauses on reaching this step. See Manually Reviewing Steps During Processing.

Field or Control	*Description*
Custom Puppet Script	Select a script that includes the custom DPK. The scripts are shown as Zip files. See Upload Custom Scripts Page.

Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML.
	
	Sample YAML input data for puppet customization:
	
	`"custom_input":`
	`"env_variables":`
	`"CS_BIN_PATH": "//10.1.1.6/CloudManager⇒`
	
	`Utils"`
	`"CS_BINARY": "crowdstrike-6.54.16808”`
	`"crafted_dpk_customization":`
	`"peoplesoft_site_name": "peopesoft_forge⇒`
	
	`t_password_site"`
	
	To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page

Search Admin User	Enter the search admin user name.
Search Proxy User	Enter the search proxy user name.
Search Provider	Select the search provider. You can select either Elasticsearch or OpenSearch.

Domain Settings Section

Full tier includes Web Server, Appserver and Process Scheduler.

Field or Control	*Description*
Web Server Settings	For Web Server setting see Configuring Web Server Tier Settings.
Appserver Setting	For App Server Setting see Configuring AppServer Tier Domain Settings.

Field or Control	*Description*
Process Scheduler Settings	For Process Scheduler Setting see Configuring Process Scheduler Domain Settings.
Process Scheduler Server Definition Parameters	For Process Scheduler Setting see Configuring Process Scheduler General Setting.

Configuring DB Systems Settings

DB Systems General Settings

The DB System options differ depending on whether the database system is on a VM or Exadata.

This example illustrates the fields and controls on the DB Systems – General Settings. You can find definitions for the fields and controls later on this page.

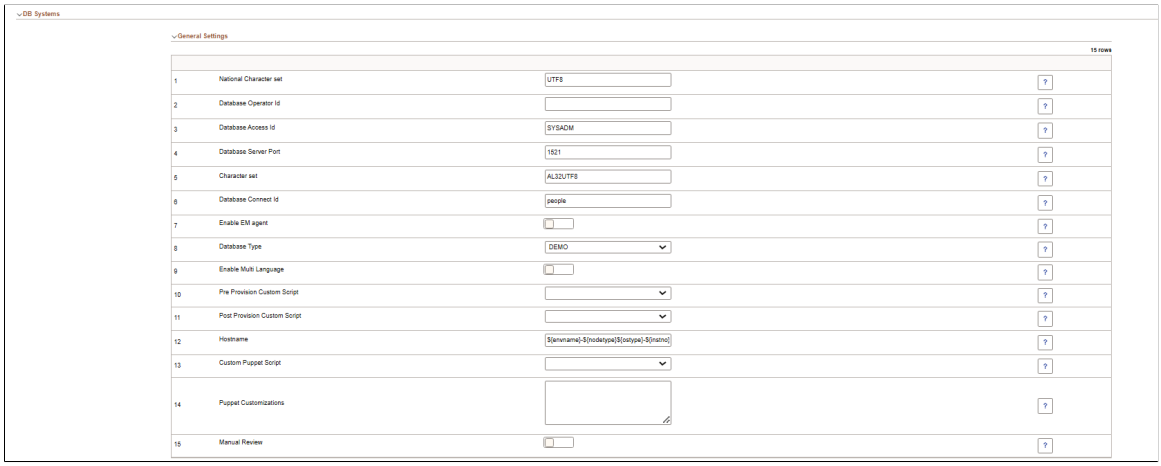

Field or Control	*Description*
National Character Set	The national character set for the database.
Database Operator ID	Default database operator ID. For HCM, CS, and ELM, the default is PS. For FSCM, IH, and CRM, the default is VP1.
Database AccessID	Access ID for the database. The default is SYSADM.
Database Server Port	Listener port number.

Field or Control	*Description*
Character Set	The character set for the database.
Database Connect Id	Connect ID for the database. The default is people.
Enable EM Agent	Select Yes to enable Environment Management agent for creating the infrastructure that is required to deploy an EM agent.
Database Type	Select the required database type. Available database types are DEMO or SYS.
Enable Multi Language	Select either Yes or No to enable multi language support.
Post Provision Custom Script	Select an uploaded script to run post provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page.
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page.

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames are validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or dbaas. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.
Custom Puppet Script	Select a script that includes the custom DPK. The scripts are shown as Zip files. See Upload Custom Scripts Page.

Field or Control	*Description*
Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML. Sample YAML input data for puppet customization: `"custom_input": "env_variables": "CS_BIN_PATH": "//10.1.1.6/CloudManager⇒ Utils" "CS_BINARY": "crowdstrike-6.54.16808” "crafted_dpk_customization": "peoplesoft_site_name": "peopesoft_forge⇒ t_password_site"` To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page.
Manual Review	Select to enable manual review. On enabling the Manual Review field, new manual stop steps are added to Customer DPK processing activity as the final step. This field is disabled by default. When the field is enabled, the activity execution pauses on reaching this step. See Manually Reviewing Steps During Processing.

DB System Options

This example illustrates the fields and controls on the DB System Options section. You can find definitions for the fields and controls later on this page.

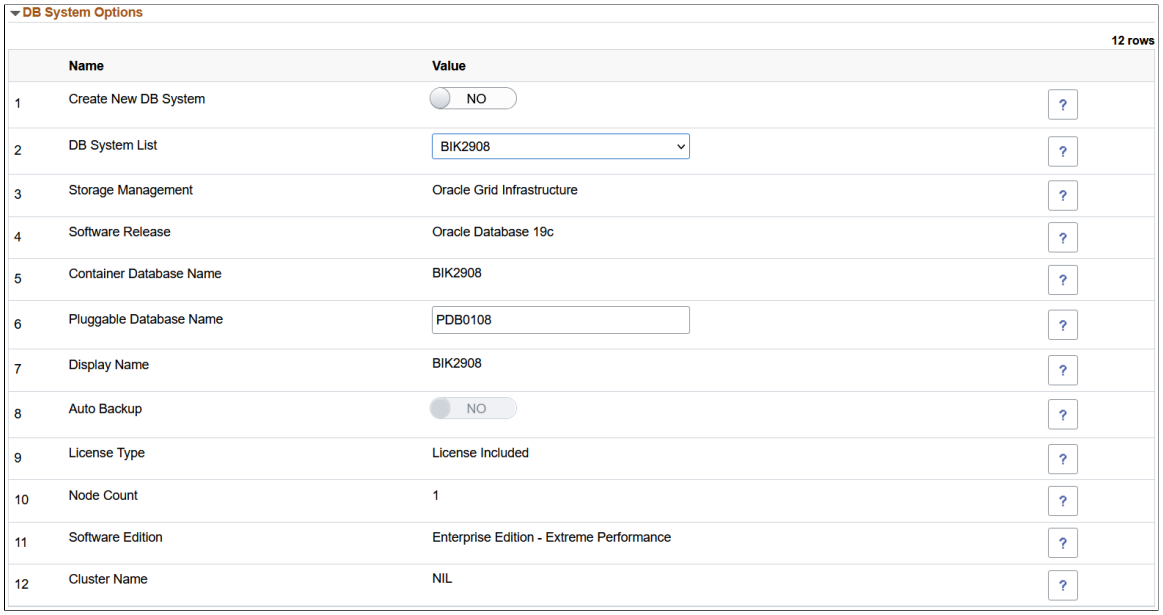

Field or Control	*Description*
Create New DB System	Select **Yes** to create a new DB system. Select **No** to use one of the available DB systems, within which Pluggable Databases (PDB) can be attached to the existing Container Database (CDB). **Note:** This feature is supported for VM-based DB systems only and not Exadata. If you use a DB system that is not created by Cloud Manager for PUM provisioning or Remote Clone, you must update Cloud Manager SSH Key to the target DB system.
DB System List	Displays a list of databases if you select **No** in the **Create New DB System** field. When you select an existing DB system, all the fields except **Pluggable Database Name** are automatically populated with corresponding values.
Storage Management	Select the Storage Management System. • Oracle Grid Infrastructure (default) Select to use Oracle Automatic Storage Management. Recommended for production workloads. • Logical Volume Manager Select to quickly provision your DB system using Logical Volume Manager storage management software.

Field or Control	*Description*
Software Release	Oracle database release version. Select the software release from the drop-down list. The following software releases are currently supported for DB systems. • Oracle Database 12c Release 1 • Oracle Database 12c Release 2 • Oracle Database 18c • Oracle Database 19c • Oracle Database 21c • Oracle Database 23ai **Note:** Oracle 23ai is currently not supported for full-tier nodes and database tier nodes on Compute. It is supported for DB System, Exadata, and Autonomous Database, or ADB. The database release version must be chosen based on the database version in the PeopleSoft Update Image or customer's lifted database version. **Note:** Oracle Database versions are certified by PeopleTools release. The Oracle Database version and PeopleTools release are listed in the PeopleSoft Update Image manifest. See the My Oracle Support Certifications area and the PeopleSoft Cloud Manager Home Page, My Oracle Support, Doc ID 2231255.2 for support information.
Container Database Name	Name of the container database in the DB system. The name of CDB is automatically filled if you choose an existing DB system. You can edit the name of CDB if you opt to create a new DB system.
Pluggable Database Name	Name of the pluggable database within the CDB. This name must be unique within a CDB.
Display Name	Display name for the DB system. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system.
Auto Backup	Displays whether automatic incremental backups for this database are enabled or disabled.

Field or Control	*Description*
License Type	The type of license you want to use for the DB system. Your choice affects metering for billing. License included means the cost of the cloud service includes a license for the Database service. Bring Your Own License (BYOL) means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
Node Count	The number of nodes in the DB system. The number depends on the shape you select. You can specify 1 or 2 nodes for virtual machine DB systems, except for VM.Standard2.1 and VM.Standard1.1, which are single-node DB systems. **Note:** Some shapes may not be available in new tenancies. **Note:** Except 1.1 and 2.1, all other shapes seem to be supported for RAC (2-node DB system). **Note:** Multi-node Virtual Machine DB systems require Oracle Automatic Storage Management and cannot be created using Logical Volume Manager option.
Software Edition	The database edition supported by the DB system.
Cluster Name	A unique cluster name for a multi-node DB system. The name must begin with a letter and contain only letters (a-z and A-Z), numbers (0-9) and hyphens (-). The cluster name can be no longer than 11 characters and is not case sensitive.
Fault domain	Select the fault domain for the database node, if the DB System contains 2 nodes, you will be able to configure the fault domain for each node.

Example Database System on Exadata

This example illustrates the fields and controls on the database system options on Exadata. You can find definitions for the fields and controls later on this page.

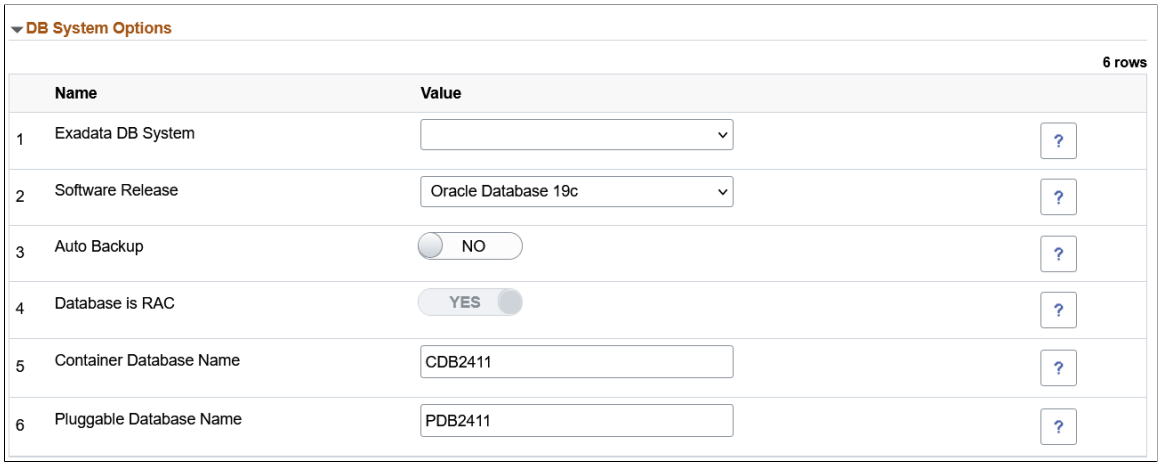

Field or Control	*Description*
Exadata DB System	Displays a list of Exadata DB systems that are available in your availability domain. If there is only one Exadata DB System, the value is auto-populated.
Software Release	Oracle database release version. The database release version must be chosen based on the database version in the PeopleSoft Update Image or customer's lifted database version.
Auto Backup	Displays whether automatic incremental backups for this database is enabled or disabled.
Database is RAC	Displays whether database is RAC.
Container Database Name	Name of the container database.
Pluggable Database Name	Name of the pluggable database.

Configuring Distributed Middle Tier Environment Template Settings

For templates using a topology with multiple middle tiers, you will configure the custom attributes for each middle tier.

This example illustrates the fields and controls on the Custom Attributes section showing multiple middle tiers.

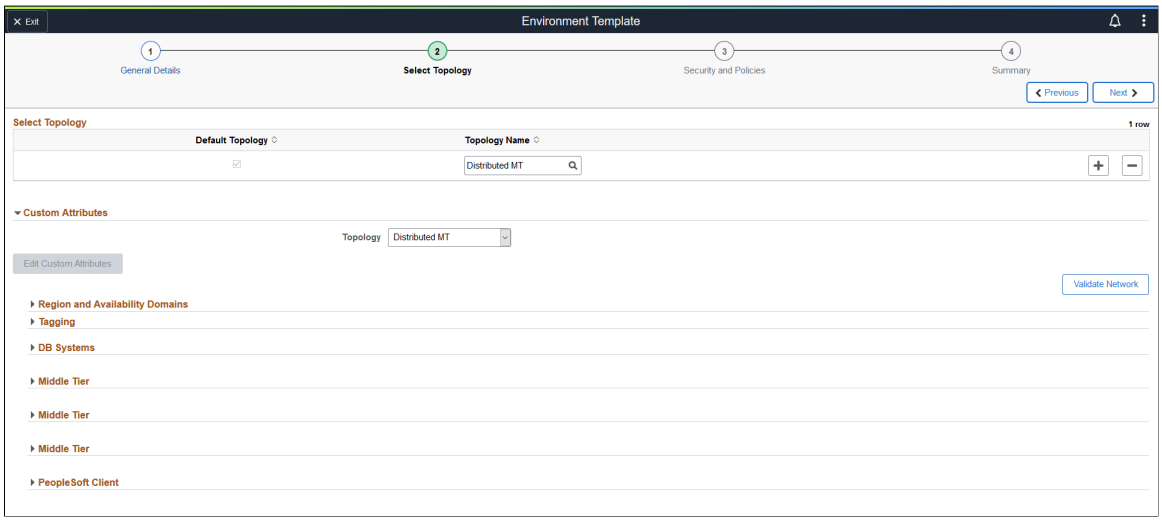

Configuring Web Server Tier Settings

General Settings

This example illustrates the fields and controls on Middle Tier - General Settings for Web Server.

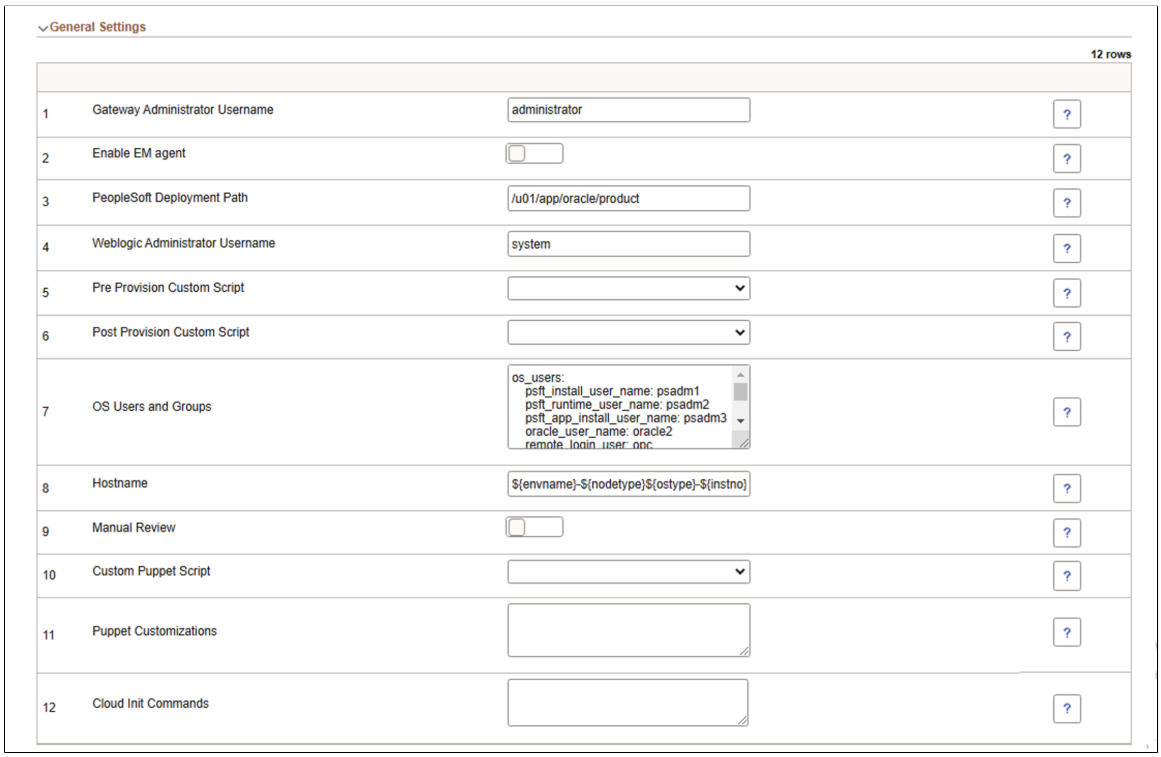

Field or Control	*Description*
Gateway Administrator Username	User ID of the Integration Gateway administrator.
Enable EM Agent	Select either Yes or No to enable or disable EM agent.
PeopleSoft Deployment Path	Location where the PeopleSoft application is deployed. **Note:** PeopleSoft Deployment Path must not end with a slash.
Weblogic Administrator Username	User name of the WebLogic administrator. This is used for accessing the Oracle WebLogic console.
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page
Post Provision Custom Script	Select an uploaded script to run after the environment provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page

Field or Control	*Description*

OS Users and Groups	Use this field to specify custom users and groups to set up the middle tier instance.
	
	`os_users:`
	`  psft_install_user_name: psadm1`
	`  psft_runtime_user_name: psadm2`
	`  psft_app_install_user_name: psadm3`
	`  oracle_user_name: oracle2`
	`  remote_login_user: opc`
	`os_user_groups:`
	`  psft_runtime_group_name: psft`
	`  psft_app_install_group_name: appinst`
	`  oracle_install_group_name: oinstall`
	`  oracle_runtime_group_name: dba`
	
	The field is populated with the default Linux users and groups in the form of a YAML customization field. All entries are mandatory. Do not modify the format. The format (for example, spacing and punctuation) and values are validated when you move to another field on the page.
	
	User names or groups should be between 1 and 32 characters long and may contain only lower and upper case letters (a-z and A-z), digits (1-9), periods (.), underscores (_), or dashes (-). They can end with a dollar sign ($). Dashes and periods are not allowed at the beginning of the user name. Fully numeric names are not allowed.
	
	The remote_login_user is required to access the instance. If you want to specify a remote_login_user other than opc, you must first create a custom Linux image with the custom user. The custom user must have root (sudo) privileges.

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames will be validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or midtier. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.
Manual Review	Select to enable manual review. On enabling the Manual Review field, new manual stop steps are added to Customer DPK processing activity as the final step. This field is disabled by default. When the field is enabled, the activity execution pauses on reaching this step. See Manually Reviewing Steps During Processing.
Custom Puppet Script	Select a script that includes the custom DPK. The scripts are shown as Zip files. See Upload Custom Scripts Page.

Field or Control	*Description*
Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML. Sample YAML input data for puppet customization: `"custom_input": "env_variables": "CS_BIN_PATH": "//10.1.1.6/CloudManager⇒ Utils" "CS_BINARY": "crowdstrike-6.54.16808” "crafted_dpk_customization": "peoplesoft_site_name": "peopesoft_forge⇒ t_password_site"` To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page
Cloud Init Commands	Enter the OS level settings like package installation or OS upgrade/ update that you need to customize when you create the infrastructure. The commands/statements must be separated by semicolon.

Domain Settings

This example illustrates the fields and controls on the Middle Tier - Domain Settings for Web Server. You can find definitions for the fields and controls later on this page.

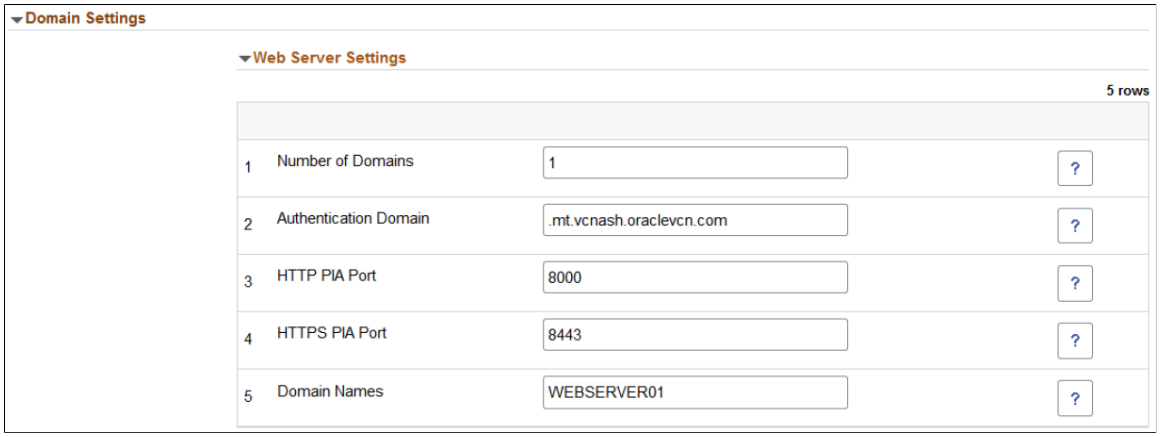

Field or Control	*Description*
Number of Domains	Enter the number of web server domains. Number of domains can be 1 to 5.

Field or Control	*Description*
Authentication Domain	The domain in which the portal is running and across which the single sign-on authentication token is valid. **Note:** The PIA URL must be modified appropriately to access the environment if you have entered a custom authentication token domain value.
HTTP PIA Port	There will be as many ports, equal to the given number of domains, separated by comma.
HTTPS PIA Port	There will be as many ports, equal to the given number of domains, separated by comma.
Domain Names	Accept the default or enter a custom domain name.

Configuring AppServer Tier Domain Settings

The General Settings fields for AppServer tier are the same as those described in Configuring Web Server Tier Settings.

This example illustrates the fields and controls on the Domain Settings for AppServer. You can find definitions for the fields and controls later on this page.

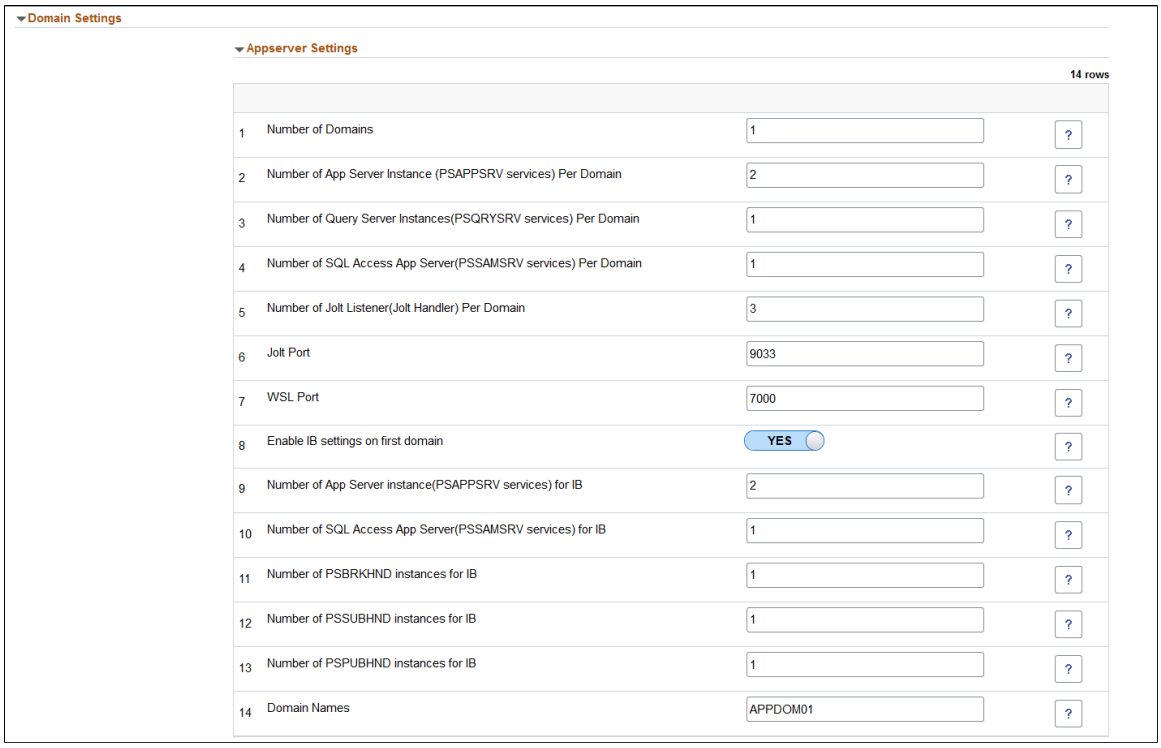

Field or Control	*Description*
Number of Domains	Number of application server domains. Number of domains can be 1 to 5.
Number of App Server Instance Per Domain	Number of PSAPPSRV instances required. This configuration is applied to all App Server domains.
Number of Query Server Instances Per Domain	Number of PSQRYSRV instances required. This configuration is applied to all App Server domains.
Number of SQL Access App Server(PSSAMSRV) Per Domain	Number of PSSAMSRV instances required. This configuration is applied to all App Server domains.
Number of Jolt Listener per Domain	Number of Jolt Listener per Domain.
Jolt Port	There will be as many ports, equal to the given number of domains, separated by comma.
WSL Port	There will be as many ports, equal to the given number of domains, separated by comma.
Enable IB Domain on first Domain	If Yes is selected IB will be enabled in the first App Domain.
Number of App Server Instance (PSAPPSRV services) for IB	Number of App Server Instance (PSAPPSRV services) for IB.
Number of SQL Access App Server (PSSAMRSRV services) for IB	Number of SQL Access App Server (PSSAMRSRV services) for IB
Number of PSBRKHND instances for IB	Number of PSBRKHND instances for IB
Number of PSSUBHND instances for IB	Number of PSSUBHND instances for IB
Number of PSPUBHND instances for IB	Number of PSPUBHND instances for IB
Domain Names	Accept the default or enter a custom domain name. Separate multiple domain names with a comma.

Configuring Process Scheduler General Setting

If the middle tier is Process Scheduler and COBOL is enabled in the topology, the General Settings will include Cobol Version. The other General Settings fields for Process Scheduler are essentially the same as those described in Configuring Web Server Tier Settings.

Field or Control	*Description*
Cobol Version	If COBOL is enabled in the selected topology, you must select the Cobol version. **Note:** The COBOL license must be configured on the Cloud Manager Settings page. See Cloud Manager Settings Page

Configuring Process Scheduler Domain Settings

This example illustrates the fields and controls on the Domain Settings for Process Scheduler Settings and Process Scheduler Server Definition Parameters. You can find definitions for the fields and controls later on this page.

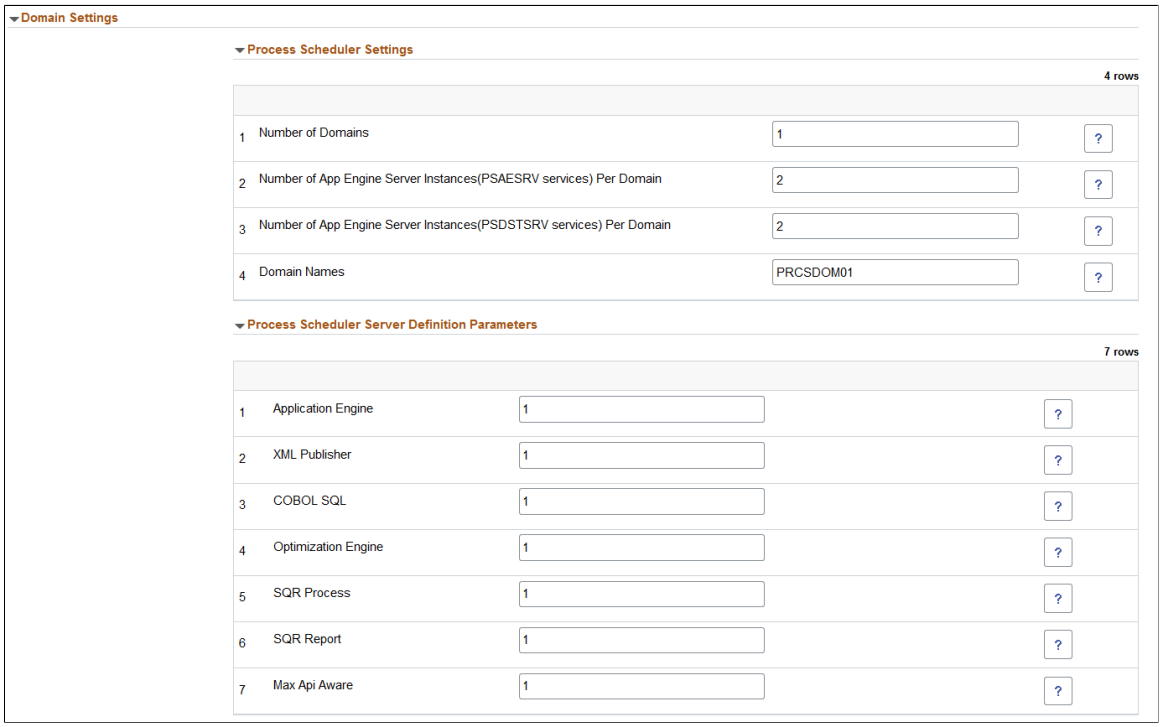

Process Scheduler Settings

Field or Control	*Description*
Number of Domains	Number of process scheduler domains.
Number of App Engine Server Instances(PSAESRV services) Per Domain	Number of application engines required.
Number of App Engine Server Instances(PSDSTSRV services) Per Domain	Number of application servers required.

Field or Control	*Description*
Domain Names	Accept the default or enter a custom domain name.

Process Scheduler Server Definition Parameters

Field or Control	*Description*
Application Engine	Number of application engine processes.
XML Publisher	Number of XML publishers.
COBOL SQL	Number of COBOL SQL processes.
Optimization Engine	Number of optimization engines.
SQR Process	Number of SQR processes.
SQR Report	Number of SQR reports.
Max Api Aware	Number of Max Api Aware.

Configuring Windows Middle Tier General Settings

This example illustrates the fields and controls on the Middle Tier - General Settings section for Windows Middle Tier.

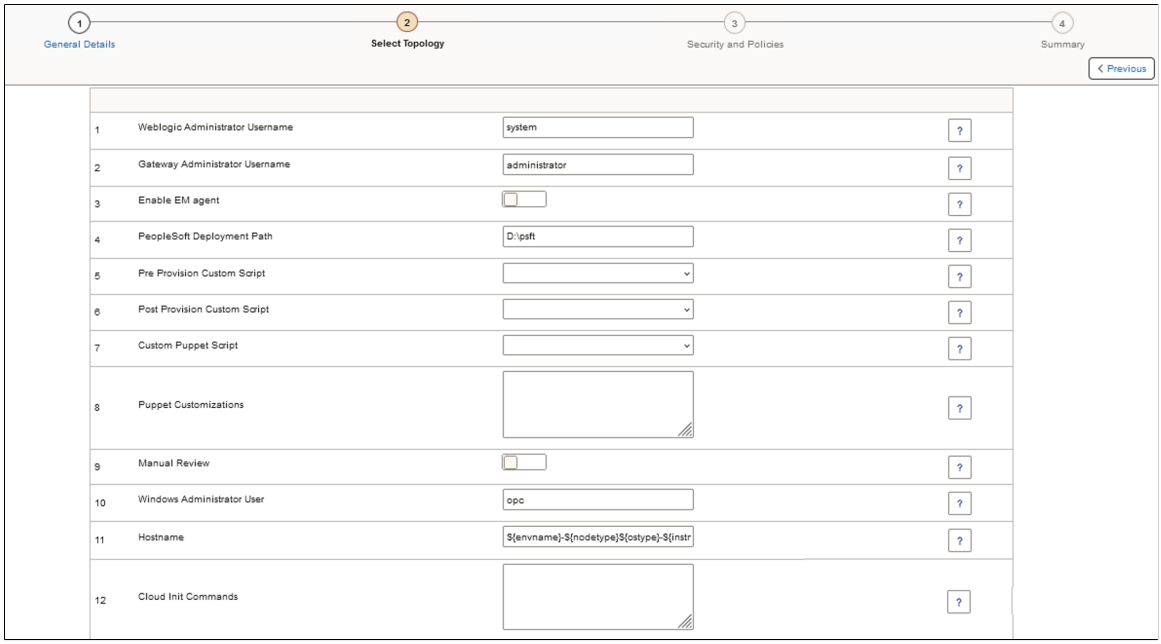

Field or Control	*Description*
WebLogic Administrator Username	User name of the WebLogic administrator. This is used for accessing Oracle WebLogic console. The default is system.
Gateway Administrator Username	User ID of Integration Broker Gateway.
Enable EM Agent	Select Yes to enable or No to disable Environment Management agent.
PeopleSoft Deployment Path	Location where the PeopleSoft application is deployed. **Note:** PeopleSoft Deployment Path must not end with a slash.
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page
Post Provision Custom Script	Select an uploaded script to run after environment provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page
Custom Puppet Script	Select a script that includes the custom DPK. The scripts are shown as Zip files. See Upload Custom Scripts Page.

Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML.
	Sample YAML input data for puppet customization:
	`"custom_input":`
	` "env_variables":`
	` "CS_BIN_PATH": "//10.1.1.6/CloudManager⇒`
	`Utils"`
	` "CS_BINARY": "crowdstrike-6.54.16808”`
	` "crafted_dpk_customization":`
	`"peoplesoft_site_name": "peopesoft_forge⇒`
	`t_password_site"`
	To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page

Field or Control	*Description*
Manual Review	Select to enable manual review. On enabling the Manual Review field, new manual stop steps are added to Customer DPK processing activity as the final step. This field is disabled by default. When the field is enabled, the activity execution pauses on reaching this step. See Manually Reviewing Steps During Processing.
Windows Administrator User	The Windows Administrator User is required to access the instance. If you want to specify a user other than the default, opc, you must first create a custom Windows image and add the custom user. The custom user must have administrative privileges. See the tutorial Create a Custom Windows Image for PeopleSoft Cloud Manager in Oracle Cloud Infrastructure (Optional) at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames will be validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or midtier. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.
Cloud Init Commands	Enter the operating system-level settings, like package installation or OS upgrade or update, that you need to customize when you create the infrastructure. The commands and statements must be separated by semicolons.

Configuring PeopleSoft Client General Settings

The PeopleSoft Client uses a Microsoft Windows operating system.

This example illustrates the fields and controls on the PeopleSoft Client - General Settings section.

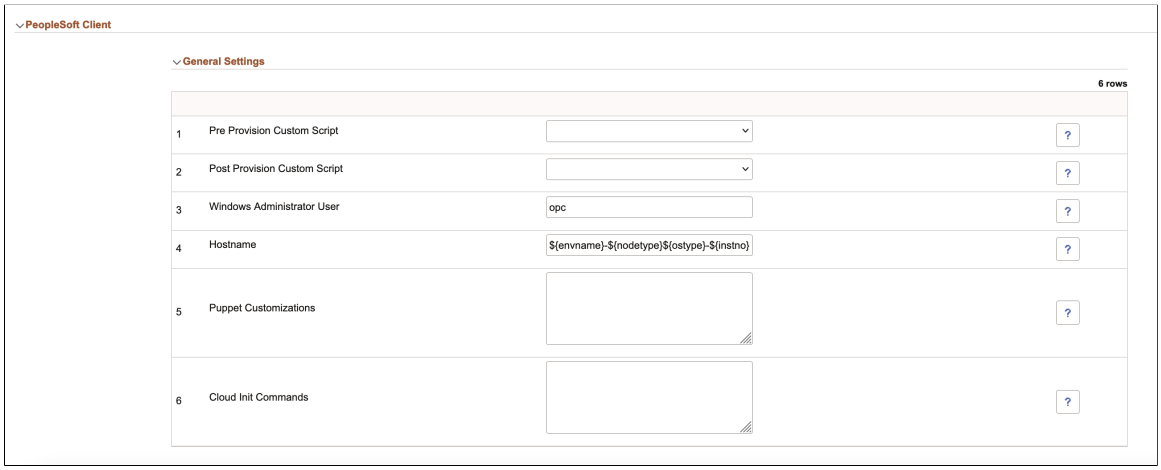

Field or Control	*Description*
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page
Post Provision Custom Script	Select an uploaded script to run after the environment provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page
Windows Administrator User	The Windows Administrator User is required to access the instance. If you want to specify a user other than opc, you must first create a custom Windows image and add the custom user. The custom user must have administrative privileges. See the tutorial Create a Custom Windows Image for PeopleSoft Cloud Manager in Oracle Cloud Infrastructure (Optional) at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames will be validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or psftclient. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.

Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML.
	
	Sample YAML input data for puppet customization:
	
	`"custom_input":`
	`"env_variables":`
	`"CS_BIN_PATH": "//10.1.1.6/CloudManager⇒`
	
	`Utils"`
	`"CS_BINARY": "crowdstrike-6.54.16808”`
	`"crafted_dpk_customization":`
	`"peoplesoft_site_name": "peopesoft_forge⇒`
	
	`t_password_site"`
	
	To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page

Cloud Init Commands	Enter the operating system-level settings, like package installation or OS upgrade or update, that you need to customize when you create the infrastructure. The commands and statements must be separated by semicolons.

Configuring Database Tier

The database tier includes general settings and subnet settings.

This example illustrates the fields and controls on the Database Tier – General Settings page (1 of 2). You can find definitions for the fields and controls later on this page.

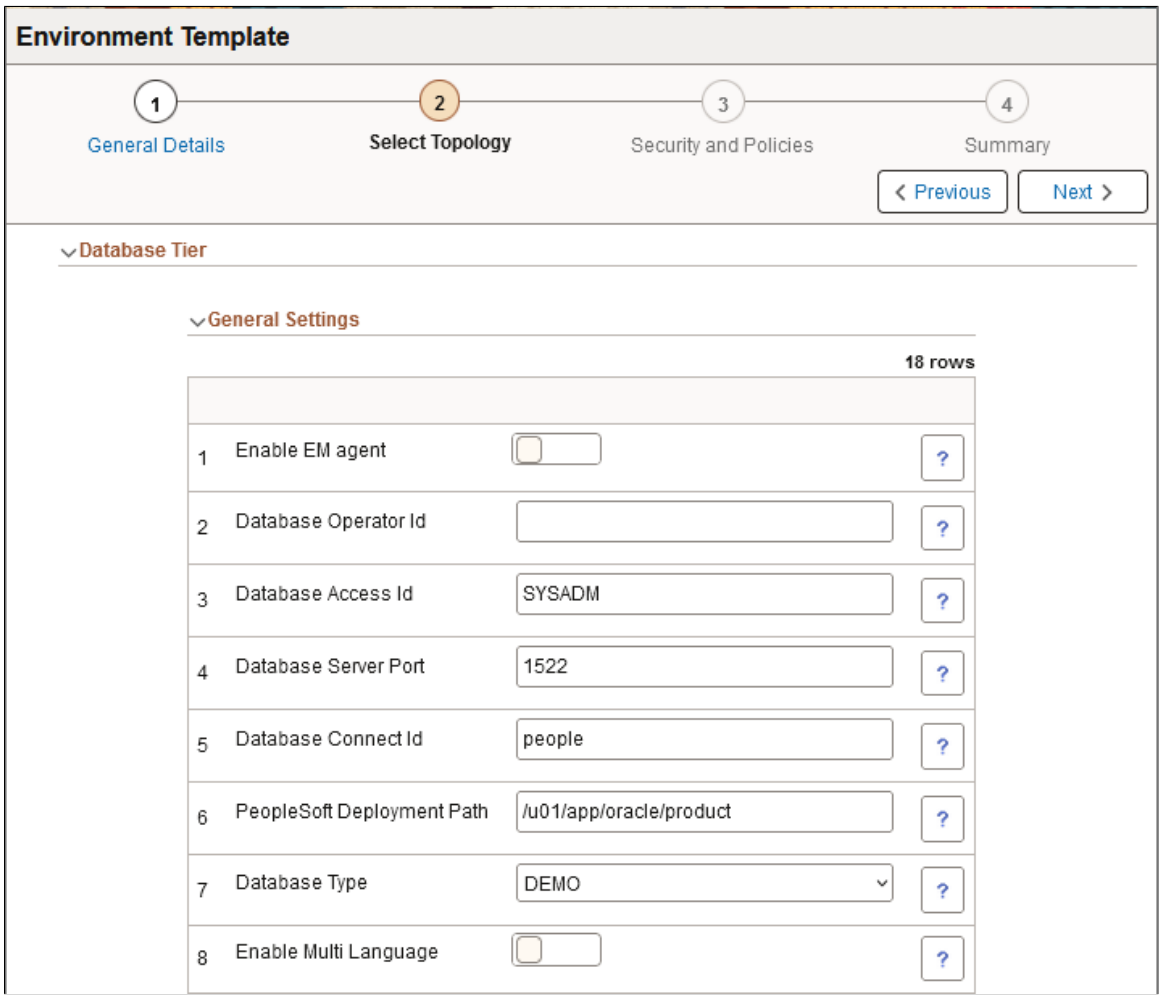

This example illustrates the fields and controls on the Database Tier – General Settings page (2 of 2). You can find definitions for the fields and controls later on this page

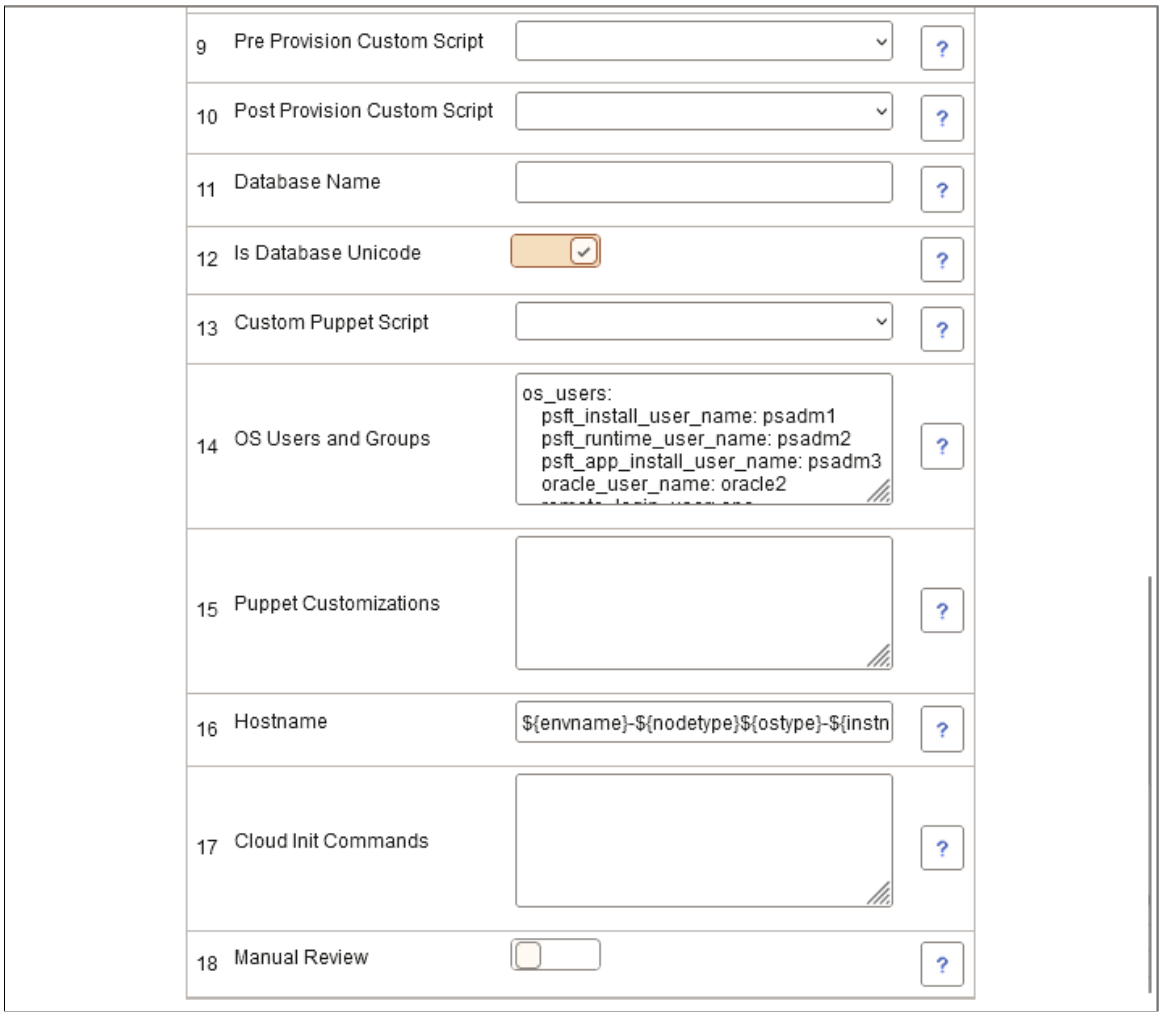

Field or Control	*Description*
Enable EM agent	Select Yes to enable Environment Management agent for creating the infrastructure that is required to deploy an EM agent.
Database Operator Id	Default database operator ID. For HCM, CS, and ELM, the default is PS. For FSCM, IH, and CRM, the default is VP1.
Database Access Id	Access ID for the database. The default is SYSADM.
Database Server Port	Listener port number. The default is 1522.
Database Connect Id	Connect ID for the database. The default is people.

Field or Control	*Description*
PeopleSoft Deployment Path	Location where the PeopleSoft application is deployed. The default is /u01/app/oracle/product. **Note:** PeopleSoft Deployment Path must not end with a slash.
Database Type	Select the required database type. Available database types are DEMO or SYS.
Enable Multi Language	Select either Yes or No to enable multi language support.
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page
Post Provision Custom Script	Select an uploaded script to run post provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page
Database Name	Name of the database.
Is Database Unicode	Select either Yes or No.
Custom Puppet Script	Select a script that includes the custom DPK. The scripts are shown as Zip files. See Upload Custom Scripts Page.

Field or Control	*Description*

OS Users and Groups	Use this field to specify custom users and groups to set up the full-tier instance.
	
	`os_users:`
	`  psft_install_user_name: psadm1`
	`  psft_runtime_user_name: psadm2`
	`  psft_app_install_user_name: psadm3`
	`  oracle_user_name: oracle2`
	`  psft_search_user_name: esuser`
	`  remote_login_user: opc`
	`os_user_groups:`
	`  psft_runtime_group_name: psft`
	`  psft_app_install_group_name: appinst`
	`  oracle_install_group_name: oinstall`
	`  oracle_runtime_group_name: dba`
	
	The field is populated with the default Linux users and groups in the form of a YAML customization field. All entries are mandatory. Do not modify the format. The format (for example, spacing and punctuation) and values are validated when you move to another field on the page.
	
	User names or groups should be between 1 and 32 characters long and may contain only lower and upper case letters, digits (1-9), periods (.), underscores (_), or dashes (-). They can end with a dollar sign ($). Dashes and periods are not allowed at the beginning of the user name. Fully numeric names are not allowed.
	
	The remote_login_user is required to access the instance. If you want to specify a remote_login_user other than opc, you must first create a custom Linux image with the custom user. The custom user must have root (sudo) privileges.

Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post Provision and Pre Provision Custom Scripts. The values are provided in the form of YAML.
	
	Sample YAML input data for puppet customization:
	
	`"custom_input":`
	`  "env_variables":`
	`  "CS_BIN_PATH": "//10.1.1.6/CloudManager⇒`
	
	`Utils"`
	`  "CS_BINARY": "crowdstrike-6.54.16808”`
	`  "crafted_dpk_customization":`
	`"peoplesoft_site_name": "peopesoft_forge⇒`
	
	`t_password_site"`
	
	To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames will be validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or dbtier. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.
Cloud Init Commands	Enter the operating system-level settings, like package installation or OS upgrade or update, that you need to customize when you create the infrastructure. The commands and statements must be separated by semicolons.
Manual Review	Select to enable manual review. On enabling the Manual Review field, new manual stop steps are added to Customer DPK processing activity as the final step. This field is disabled by default. When the field is enabled, the activity execution pauses on reaching this step. See Manually Reviewing Steps During Processing.

Configuring Search Stack General Settings

See Provisioning and Sharing Search Clusters for information on setting up search clusters.

This example illustrates the fields and controls on the General Settings for Search Stack tier (1 of 2). You can find definitions for the fields and controls later on this page.

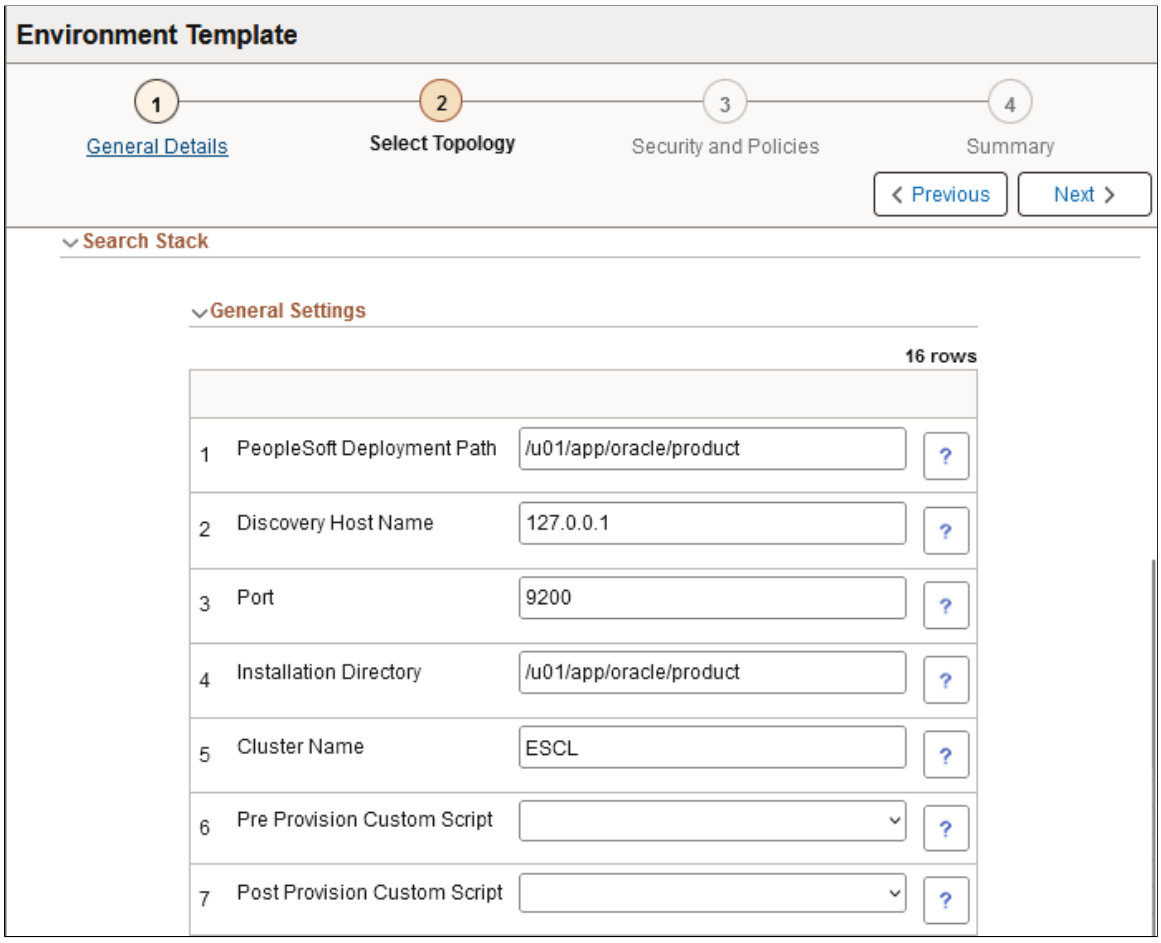

This example illustrates the fields and controls on the General Settings for Search Stack tier (2 of 2). You can find definitions for the fields and controls later on this page.

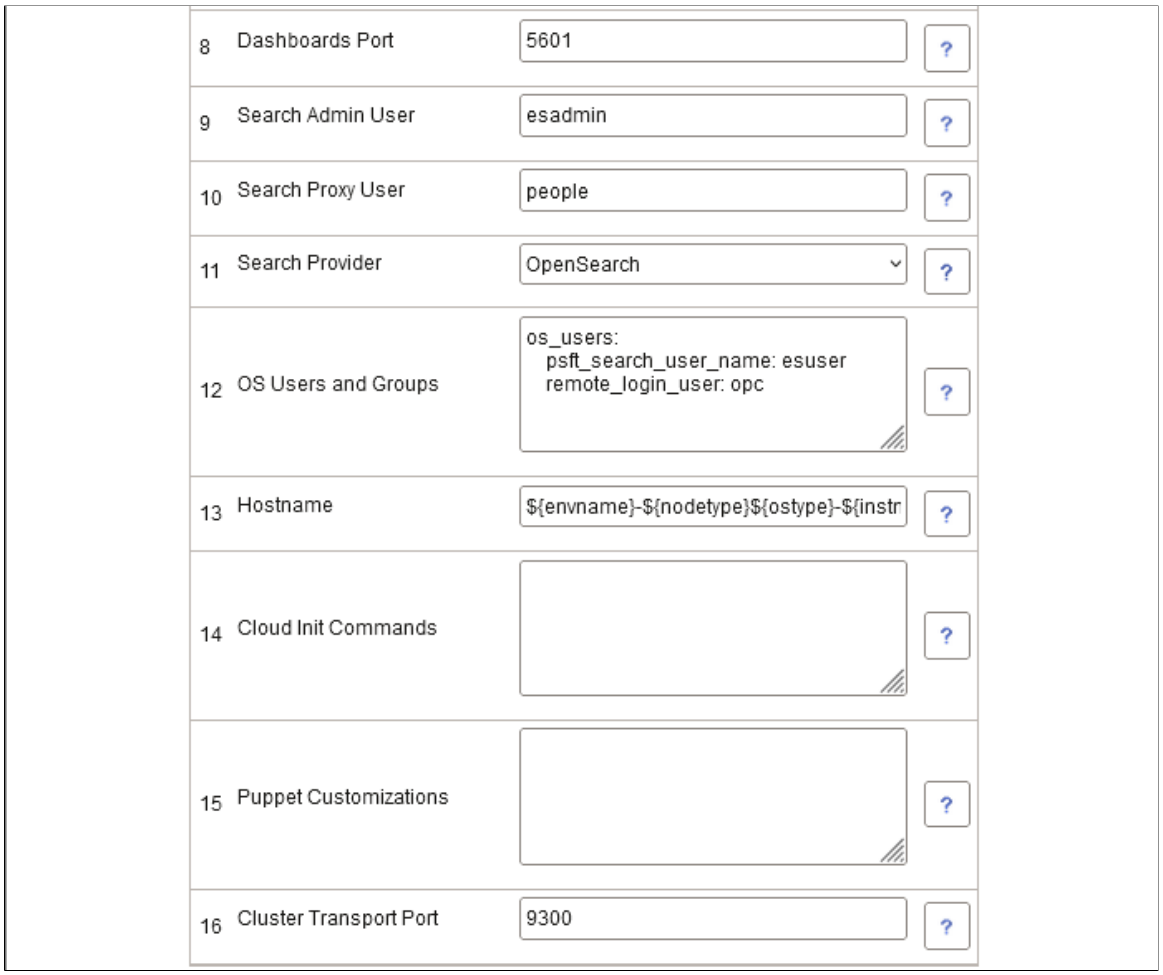

Field or Control	*Description*
PeopleSoft Deployment Path	Location where the PeopleSoft application is deployed. **Note:** PeopleSoft Deployment Path must not end with a slash.
Discovery Host Name	The host name for any nodes that are already members of a cluster.
Port	Elasticsearch or OpenSearch port. The default is 9200.
Installation Directory	The path to install Elasticsearch (or OpenSearch) and/or Kibana (or OpenSearch Dashboards).
Cluster Name	The name of the OpenSearch cluster. The default is ESCL.

Field or Control	*Description*
Pre Provision Custom Script	Select an uploaded script to run prior to provisioning the environment. The environment variable defined in the puppet customization can be accessed in Pre Provision Custom Script. See Upload Custom Scripts Page
Post Provision Custom Script	Select an uploaded script to run post provisioning. The environment variable defined in the puppet customization can be accessed in Post Provision Custom Script. See Upload Custom Scripts Page
Dashboards Port	Port for OpenSearch Dashboards or Kibana. The default is 5601.
Search Admin User	Enter the admin user for search stack. The default is esadmin.
Search Proxy User	Enter the proxy user for search stack. The default is people.
Search Provider	Select the search provider. You can select either Elasticsearch or OpenSearch.

OS Users and Groups	Use this field to specify custom users to set up the search stack instance.
	
	`os_users:`
	`  psft_search_user_name: esuser`
	`  remote_login_user: opc`
	
	The field is populated with the default Linux users in the form of a YAML customization field. All entries are mandatory. Do not modify the format. The format (for example, spacing and punctuation) and values are validated when you move to another field on the page.
	
	User names should be between 1 and 32 characters long and may contain only lower and upper case letters, digits (1-9), periods (.), underscores (_), or dashes (-). They can end with a dollar sign ($). Dashes and periods are not allowed at the beginning of the user name. Fully numeric names are not allowed.
	
	The remote_login_user is required to access the instance. If you want to specify a remote_login_user other than opc, you must first create a custom Linux image and add the custom user. The custom user must have root (sudo) privileges.
	
	See the tutorial Create a Custom Linux Image for PeopleSoft Cloud Manager (Optional) at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

Field or Control	*Description*
Hostname	Accept the default or enter a custom Hostname. **Note:** Hostnames will be validated when template is submitted. The text ${envname}-${nodetype}${ostype}-${instno} represents macros that Cloud Manager expands. • ${envname} is replaced by lowercase environment name. • ${nodetype} is replaced by node type; for example, fulltier or elasticsearchtier. • ${ostype} is replaced by operating system; for example, linux. • ${instno} is replaced by a numeric value. For example, suppose you are provisioning an environment with the name HCM03, with a single full-tier node on Linux. If you do not change the default text in the **Hostname** field, the value of hostname would be hcm03-fulltierlinux-1. If you want to enter a custom hostname, you can replace one or more of the macros. For example, replace ${envname} with Test02, but retain the other values: Test02-${nodetype}${ostype}-${instno} This will result in the hostname value: Test02-fulltierlinux-1.
Cloud Init Commands	Enter the operating system-level settings, like package installation or OS upgrade or update, that you need to customize when you create the infrastructure. The commands and statements must be separated by semicolons.

Puppet Customizations	Enter DPK customization values or custom environment variables that can be accessed in Custom Puppet Script as well as the Post and Pre Provision Custom Scripts. The values are provided in the form of YAML.
	
	Sample YAML input data for puppet customization:
	
	`"custom_input":`
	`"env_variables":`
	`"CS_BIN_PATH": "//10.1.1.6/CloudManager⇒`
	
	`Utils"`
	`"CS_BINARY": "crowdstrike-6.54.16808”`
	`"crafted_dpk_customization":`
	`"peoplesoft_site_name": "peopesoft_forge⇒`
	
	`t_password_site"`
	
	To refer the usage of environment variables and DPK YAML values, see Upload Custom Scripts Page

Field or Control	*Description*
Cluster Transport Port	The port for communication for shared OpenSearch search clusters. The default is 9300.

Environment Template – Security and Policies Page

Use the Environment Template – Security and Policies page to associate the zone in which the environment is created, the role that will have access to the template and policies associated with the template.

Navigation:

- Click Next on the Select Topology step.
- Click Step 3, Security and Policies, at the top of the page to navigate to the Environment Template – Security and Policies page in the guided process.

This example illustrates the fields and controls on the Environment Template — Security and Policies page. You can find definitions for the fields and controls later on this page.

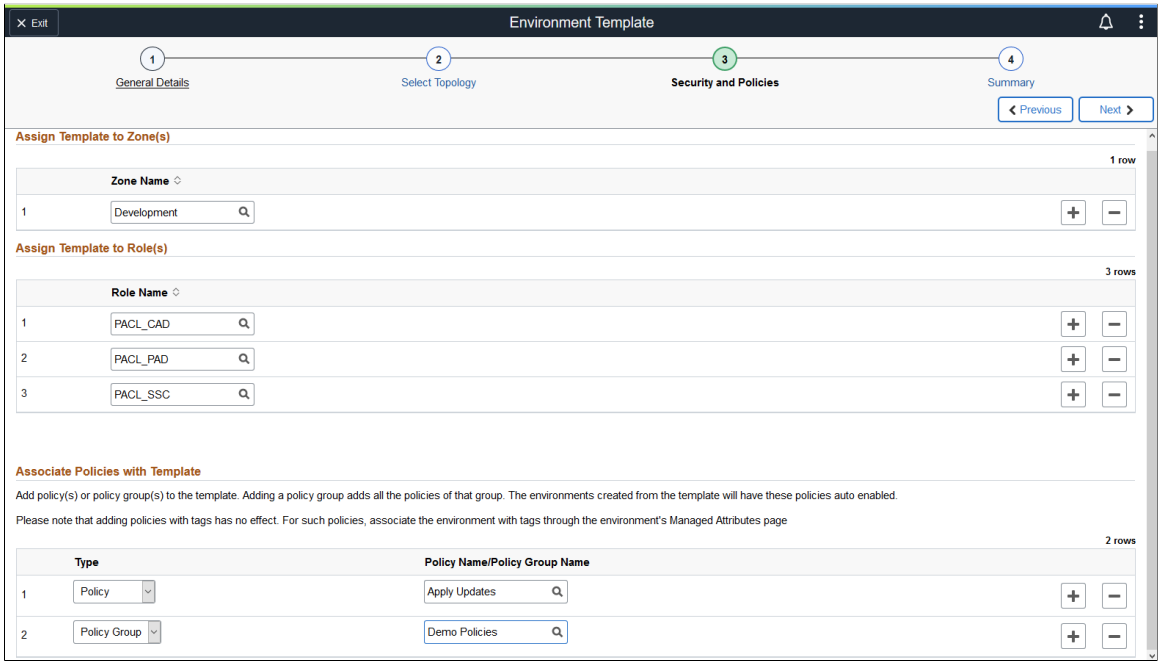

Field or Control	*Description*
Zone Name	Indicates the zone in which the environment is created.

Field or Control	*Description*
Role Name	Indicates the roles that have access to the template for creating environments. Only the users belonging to the role specified will be able to access the template while creating environment. The delivered Cloud Manager roles are: • Cloud Administrator (PACL_CAD) • Cloud PeopleSoft Administrator (PACL_PAD) • Self-Service User (PACL_SSC)
Type	Select the policy type: • Policy • Policy Group **Note:** Adding a policy group adds all the policies of that group.
Policy Name/Policy Group Name	Select policy name or policy group name for this template. The environments created from the template will have these policies auto enabled. See Using Policy Editor.

Environment Template – Summary Page

Use the Environment Template – Summary page (ECL_TEMPL_REV_FL) to review and submit the template details.

Navigation:

- Click Next on the Security and Policies step.
- Click step 4, Summary, at the top of the page to navigate to the Environment Template – Summary page in the guided process.

This example illustrates the fields and controls on the Environment Template – Summary page. You can find definitions for the fields and controls later on this page.

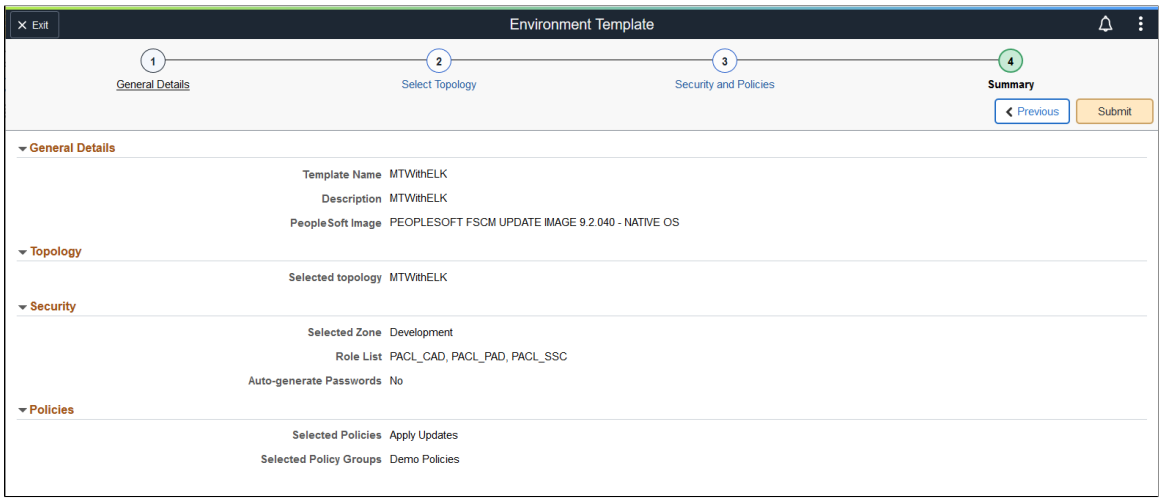

The details provided in all the pages in the Environment Template wizard is displayed here.

Field or Control	***Description***
Submit	Click this button to submit the details for template creation.

Chapter 6

Managing Environments

Environments Overview

Cloud Manager provisions PeopleSoft environments on-demand with just a few clicks. The entire provisioning process is automated. At the end of provisioning, a ready-to-use environment is available within a short time. The environments can be created by a three step process:

1. Create Topology
2. Create Template
3. Create Environment

Note: Prior to creating an environment, ensure that the required DPKs are already downloaded in the Repository.

An administrator defines a template for creating an environment. The topology is encapsulated inside the template. Users can select a template, override topologies, change any attributes, if needed and provision PeopleSoft environments on demand.

Users are allowed to perform actions on a running environment, such as stop, view details, create new template from it, and so on. For details, see the Actions on the Environment section under the Environments Page.

Note: Also, you must ensure to tune the servers, database, and PeopleSoft system for optimum performance once the deployment is completed.

Pages Used to Manage Environments as an Administrator

Page Name	***Definition Name***	***Usage***
Environments Tile	ECL_ENVPROV_FL_GBL (Content reference for the tile.)	Access the Environments landing page.
Environments Page	ECL_ENVPRO_FL	Access the Environments landing page.
Create Environment Creating an Environment	ECL_ENV_ADD_SCF	Create a new environment.

Page Name	*Definition Name*	*Usage*
Environment Details Accessing Environment Details	ECL_ENV_DET_FL	Access more details of the environment from one location.
Manage Attributes Managing Environment Attributes	ECL_ENV_ATTR_FL	Centrally administer the configuration of all the managed instances using Cloud Manager.
Manage Tags Managing Tags	ECL_ENV_TAGS_FL	Add, delete or update tags for each node in the environment.
Manage Passwords Managing Passwords	ECL_ENV_RESET_FL	Update Cloud Manager with environment attributes, if a user modifies it outside Cloud Manager.
Manage PUM Connections Managing PUM Connections	ECL_SA_MANAGEPM_FL	Manage PUM connections.
Infrastructure CPU Patches Applying Infrastructure CPU Patches	ECL_ENV_INFRUPD_FL	Display currently applied third party component details.
Apply PeopleTools Patch Applying PeopleTools Patch	ECL_ENV_PTCHUPD_FL	Apply latest patches.
Upgrade PeopleTools Upgrading PeopleTools	ECL_ENV_UPGD_FL	Update PeopleTools version (major version changes).
Provision Task Status Viewing Provision Task Status	ECL_PTS_PROV_TASK	Use to check provisioning status and retry or resume.
Policies Associating Policies with Environment	ECL_POLICY_ENVS	Associate policies with the environment name.
Logs Viewing Environment Logs	ECL_ESEARCH_FL	View logs of all operations such as create, delete, actions performed on the environment, and the like.
Monitoring Monitoring Environments	ECL_ML_ALERTS_FL	Enable or disable monitoring for the environment.

Page Name	*Definition Name*	*Usage*
Load Balancing Configuring Load Balancer Settings	ECL_LB_BACKENDS_FL	Set up load balancing.
Advisory Settings Configuring and Reviewing Advisories	ECL_ENVRECM_SET_FL	Configure or update advisory settings for the environment.
Sparse Hierarchy Details Configuring Sparse Hierarchy Details	ECL_EXA_SPARSE_FL	Create test masters and manage their sparse clones. This page is available only for environments with databases running on Exadata.
Cluster Settings Setting Up Unified Navigation Clusters	ECL_CLUST_SETNG_FL	Create and manage unified navigation clusters.

Environments Tile

Use the Environments tile (ECL_ENVPROV_FL_GBL) to access the Environments landing page.

Navigation:

The Environments tile is delivered as part of the Cloud Manager home page.

This example illustrates the Environments tile.

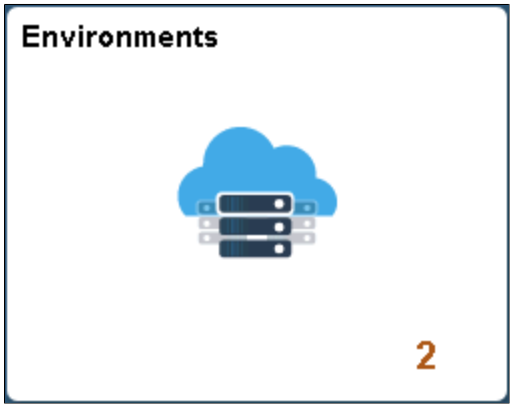

Environments Page

Use the Environments page (ECL_ENVPRO_FL) to manage, create, and import provisioned environments. You can view the environments in card or grid (list) format. This documentation typically refers to the card format.

Navigation:

Click the Environment tile on the delivered Cloud Manager Fluid home page. The Environments page is displayed.

This example illustrates the fields and controls on the Environments page. You can find definitions for the fields and controls later on this page.

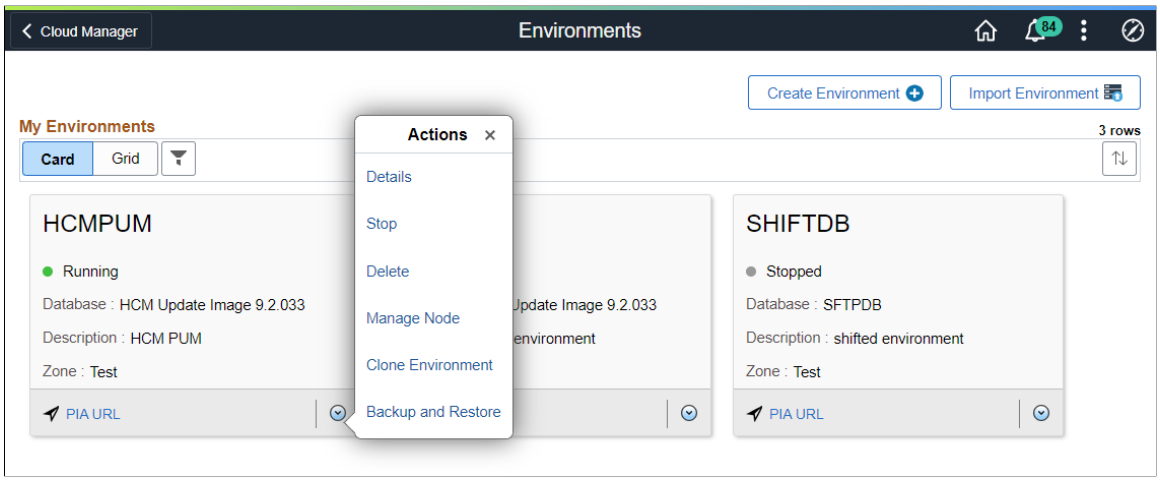

Field or Control	***Description***
Name	Name of the environment. **Note:** Length of Environment name and identity domain name should not exceed 20 characters in OCI.

Field or Control	*Description*
Status	Status of the environment provisioned through Cloud Manager. The main statuses associated with the environment are: • Initiating • Provisioning • Failed • Stopping • Starting • Running • Deleting • Applying PeopleTools Patch. • Upgrading PeopleTools. • Refreshing • Restoring • Infra Creation Complete You may see other statuses associated with specific tasks performed on environments.
Description	Description of the environment.
Zone	Zone in which the environment is deployed.
PIA URL	Indicates the URL used to connect to the provisioned environment.
Create Environment button	Click this button to access the Create Environment page, where you can create new environments.
Related Actions button 	Click this button to perform different actions for managing the environment as a whole.
Import Environment button	Click this button to import a database system environment. See Importing Environment

Field or Control	*Description*
Filter button 	Click the Filter icon to refine search results based on these criteria: • Environment Name • Environment Status • App version • Cluster Name
Sort button 	Select a criteria to sort the list of environments. You can sort by: • Environment Name • Status • Environment Status • App version • Environment Description • Cluster • PIA URL

Actions on the Environment

You can perform a variety of actions on the environment by using the Related Actions button corresponding to each environment. The actions can be:

- **Details**: Select this option to view environment details and to perform additional actions on the environment such as performing a health check, applying a PeopleTools patch, viewing logs, and managing PUM connections. Alternatively, you can access the Environment Details page by clicking anywhere on the tile of an environment.
- **Start**: Select this option to start all the instances and then all the domains within them.
- **Stop**: Select this option to stop all domains and shutdown all the instances. In case of database, only compute database instances are shutdown. You cannot stop the DB system if it is used as a database tier and has multiple PDBs.
- **Delete**: Select this option to remove the environment.

  Note: If DB system is used as a database tier and has multiple PDBs, only the PDB is deleted and the DB system is not terminated.

- **Manage Node**: Select this option to scale an environment up or down, add nodes, remove nodes, restart, stop, start, and work with search clusters.
- **Clone Environment**: Select this option to clone an existing environment.

- **Refresh**: Select this option to refresh the database or the database, Application Home (PS_APP_HOME) and Custom Home (PS_CUST_HOME).

  This option is only available for DBaaS environments. This option is disabled for environments that use DB system as database tier and have multiple PDBs.

- **Backup and Restore**: Select this option to backup or restore an environment.

  This option is disabled for environments that use DB system as database tier and have multiple PDBs.

- **Deploy**: Select this option to continue deploying an environment that has paused after infrastructure creation.

- **Delegate Access**: This functionality is no longer supported. See Role Based Security Page for the currently supported action.

- **Disaster Recovery**: Select this option to set up an environment as a standby in case of failure.

Creating an Environment

Use the Create Environment page (ECL_ENV_ADD_SCF) to create a new environment.

Important! Before creating an environment in OCI, ensure that the template is updated with OCI-specific Infrastructure Settings such as region, compartment, VCN and subnet settings.
When you use the existing DB system during PUM provisioning, you must ensure that the database versions in the DB system and DPK are the same and that the same DB administrator password that was provided while creating the DB system is used.

Note: Deploying a PeopleSoft Update Image requires a Microsoft Windows platform image that is updated with the latest Windows updates and patches. If the Windows image is not on the latest updates and patches the provisioning of PeopleSoft Client will fail. Refer to the Cloud Manager installation tutorials on Oracle Help Center https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html.

Navigation:

Click the **Create Environment** button on the Environments landing page.

This example illustrates the fields and controls on the Create Environment page. You can find definitions for the fields and controls later on this page.

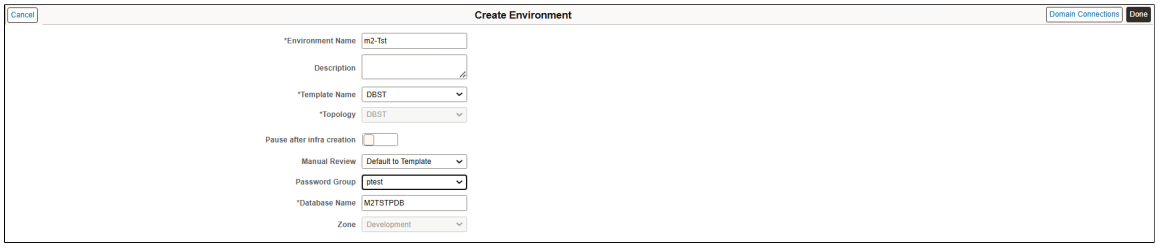

Field or Control	*Description*
Environment Name	Name of the environment that you want to create. It is a required field. **Note:** Length of environment name must not exceed 20 characters in OCI.
Description	Description for the environment that you want to create.
Template Name	Select a template and the zone. On selecting the template, zone options are automatically displayed. It is a required field. For details on templates, see the Creating a Template section under Environments Page.
Pause after infra creation	Select Yes for the environment provisioning to pause after completion of the Infrastructure task. This provides the user the opportunity to do additional setup, actions, or operations on the newly created environment outside of Cloud Manager before proceeding with the PeopleSoft deployment. **Note:** When you are ready to proceed to the PeopleSoft deployment, select Deploy from the related actions menu for the environment. Select No (default) to continue provisioning the environment when the infrastructure layer is complete.
Manual Review	Select the global option to enable or disable all the manual reviews selected in the chosen template in all the instances. The default option is 'Default to Template' to use the Manual Review values in chosen environment template.
Database Name	Enter the name of the pluggable database (PDB) in the provisioned environment. If you did not set the database name in the environment template, Cloud Manager generates a default name based on the environment name during PUM provisioning. The default name of PDB gets generated using the following rule: • If the length of the remaining characters in the environment name after removing special characters is less than 6, Database Name converts that into upper case and appends 'PDB' at the end of it. • If the length of the remaining characters in the environment name after removing special characters is less than 9, Database Name converts it into upper case. • If the length of the remaining characters in the environment name after removing special characters is greater than 9, Database Name converts it into upper case and truncates it to the first 8 characters.

Field or Control	*Description*
Zone	Select the zone for the environment. If only one zone was defined on the environment template used for this environment, it will be displayed as read-only.
Password Group	Select the password group that contains the passwords that are already created as secret OCIDs in OCI Vault. The credentials are automatically filled, once you select a password group.
Domain Connections	Click the button to configure connections between application server domains and web server domains using Domain Connections page.

This example illustrates the actions on the Environments page after Infrastructure creation is complete.

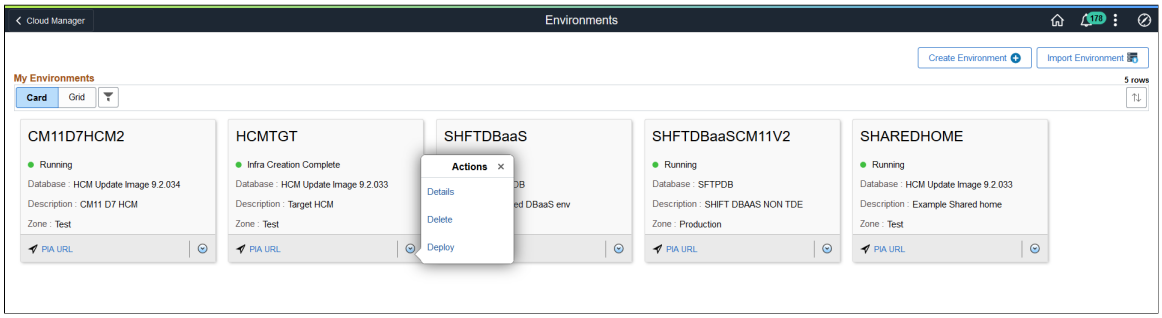

Creating an Environment

After creating topology and template, you can create an environment.

Important! Before creating an environment in OCI, ensure that the template is updated with OCI-specific Infrastructure Settings such as region, primary availability domain, default compartment, default VCN, network settings, and network security group.

To create an environment:

1. Enter the required environment attributes.

   Note: Region and Availability Domains, Network Settings and Network Security sections are read-only.

   You can add multiple web server, application server and process scheduler server domains with custom configurations for nodes in the Domain Settings section.

This example illustrates the fields and controls in the Domain Settings section of the Create Environment page.

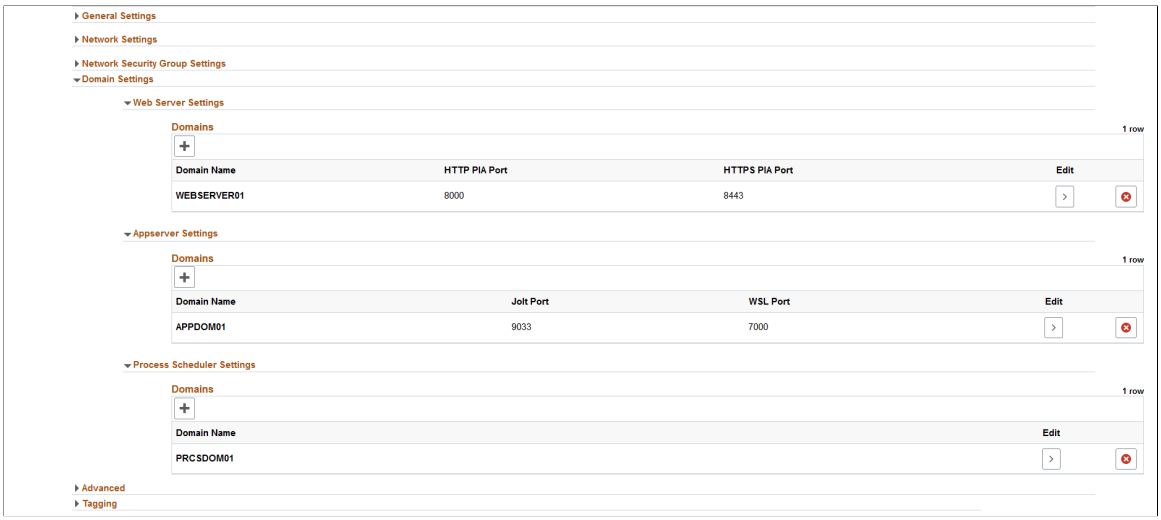

The grid-like structure enables you to customize attribute values at each domain level.

2. Click > in the Edit column in Web Server Settings section to edit the web server domain configurations.

   This example illustrates the fields and controls in the Web Server Domain Configuration. You can find definitions for the fields and controls later on this page.

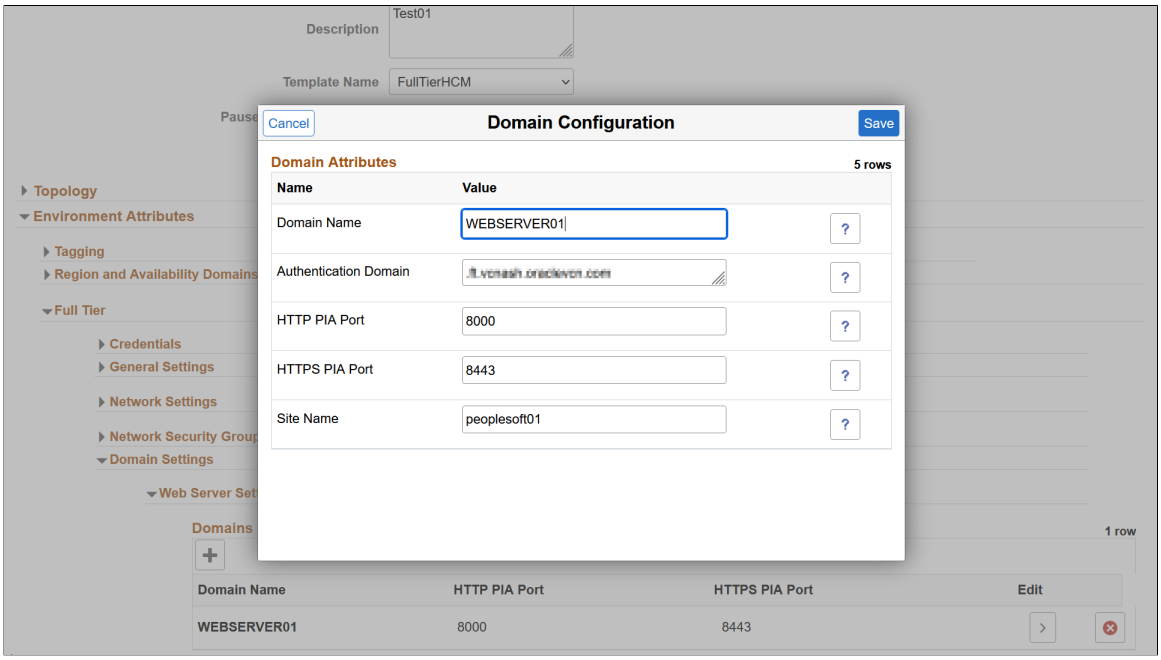

Field or Control	*Description*
Domain Name	Enter a custom name for web server domain.

Field or Control	*Description*
Authentication Domain	The domain in which the portal is running and across which the single sign-on authentication token is valid. **Note:** The PIA URL must be modified appropriately to access the environment if you have entered a custom authentication token domain value.
HTTP PIA Port	HTTP port for the web server domain.
HTTPS PIA Port	HTTPS port for the web server domain.
Site Name	Name of the site to be created with web server domain.

Click > in the Edit column in App Server Settings section to edit the app server domain configurations.

This example illustrates the fields and controls in the App Server Domain Configuration. You can find definitions for the fields and controls later on this page.

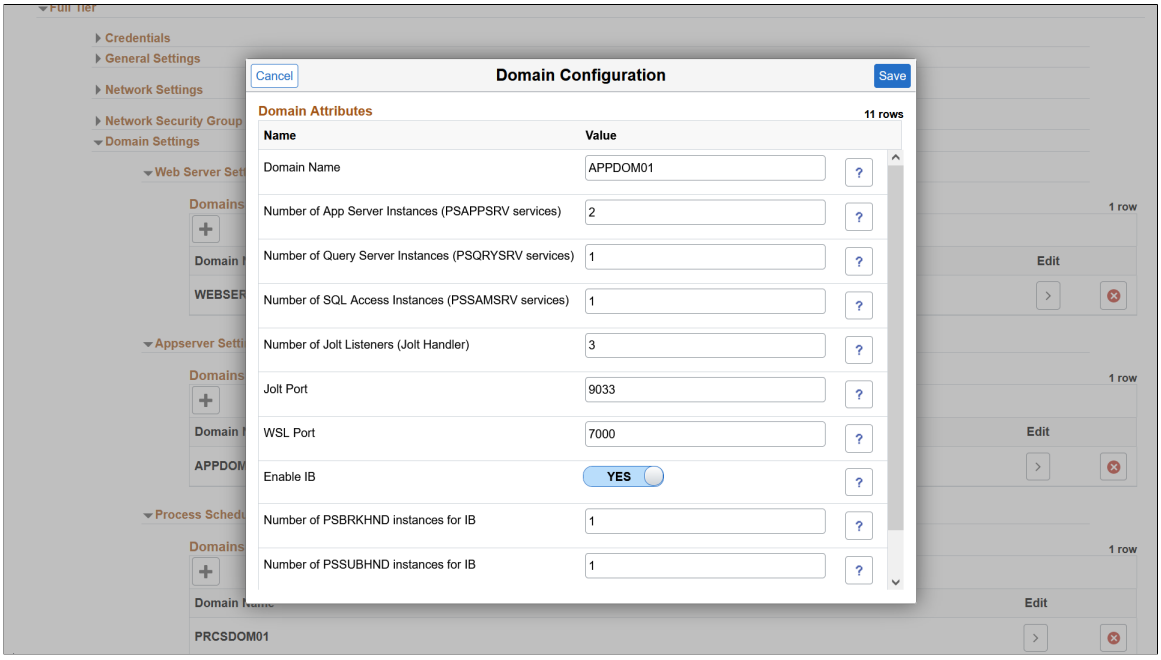

Field or Control	*Description*
Domain Name	Enter a custom name for application server domain.
Number of App Server Instances(PSAPPSRV services)	Number of PSAPPSRV instances required.

Field or Control	*Description*
Number of Query Server Instances(PSQRYSRV services)	Number of PSQRYSRV instances required.
Number of SQL Access App Server(PSSAMSRV services)	Number of PSSAMSRV instances required.
Number of Jolt Listeners (Jolt Handler)	Number of Jolt Listener per Domain.
Jolt Port	Jolt port for the app domain.
WSL Port	WSL port for the app domain.
Enable IB	Select Yes to enable IB in the App Domain.
Number of PSBRKHND instances for IB	Number of PSBRKHND instances for IB.
Number of PSSUBHND instances for IB	Number of PSSUBHND instances for IB.
Number of PSPUBHND instances for IB	Number of PSPUBHND instances for IB.

Click > in the Edit column in Process Scheduler Server Settings section to edit the process scheduler server domain configurations.

This example illustrates the fields and controls in the Process Scheduler Server Domain Configuration. You can find definitions for the fields and controls later on this page.

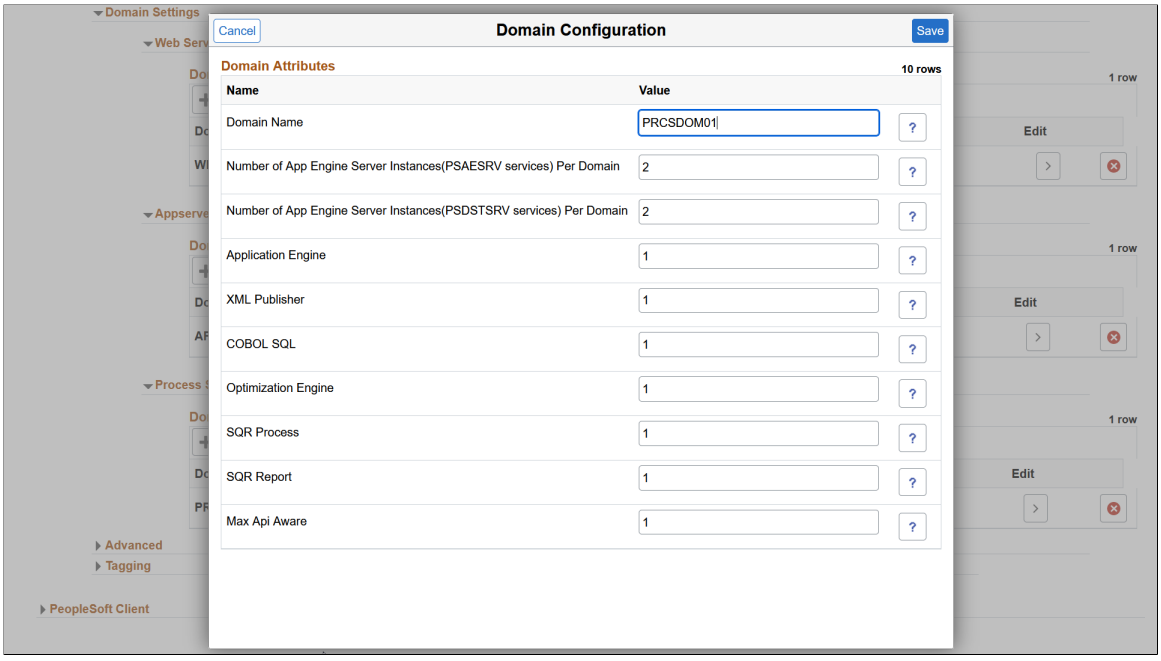

Field or Control	*Description*
Domain Name	Accept the default or enter a custom domain name.
Number of App Engine Server Instances(PSAESRV services) Per Domain	Number of application engines required.
Number of App Engine Server Instances(PSDSTSRV services) Per Domain	Number of application servers required.
Application Engine	Number of application engine processes.
XML Publisher	Number of XML publishers.
COBOL SQL	Number of COBOL SQL processes.
Optimization Engine	Number of optimization engines.
SQR Process	Number of SQR processes.
SQR Report	Number of SQR reports.
Max Api Aware	Number of Max Api Aware. Indicates the total number of tasks that a Process Scheduler can initiate concurrently.

Click **Save** to save the changes.

Click + on the respective domain grid to add a new domain and click **X** to delete a domain.

3. Click **Done** to start environment provisioning.

Note: The system validates available resources before starting the provisioning process. See Validating Resources.

---

4. Select the Provision Task Status link to view the progress of the environment creation. See Viewing Provision Task Status. If a failure occurs, you can retry and resume the operation. See Retrying and Resuming Provisioning.

Alternately, you can override the default topology and environment attributes while environment provisioning.

The default database operator id for each PeopleSoft PUM instance is listed below:

- For HCM, default database operator id is PS.
- For FSCM, default database operator id is VP1.
- For CRM, default database operator id is VP1.

- For ELM, default database operator id is PS.
- For IH, default database operator id is VP1.
- For CS, default database operator id is PS.

Configuring Domain Connections

Configure connections between application server domains and web server domains using Domain Connections page.

To select the required configuration for an environment or node:

1. Click Domain Connections on the Create Environment page.
2. Select the application server domains to be connected with each web server node. Multiple Integration Broker-enabled application server domains can be connected to web servers.

   This example illustrates the fields and controls on the Domain Connections page.

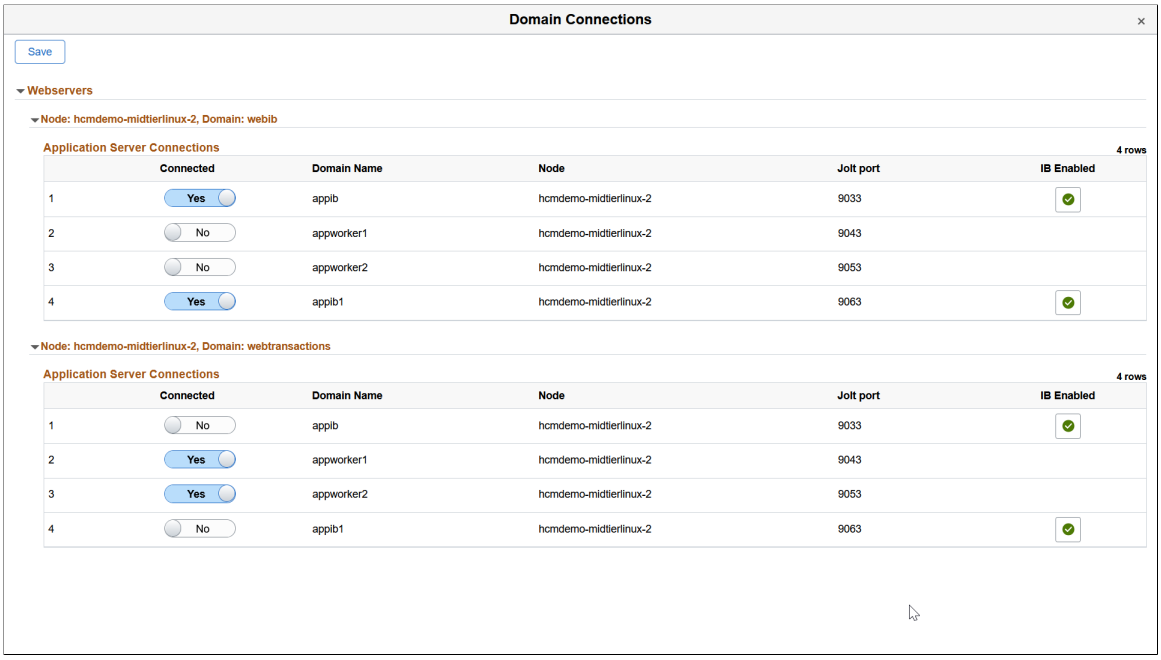

Field or Control	*Description*
Connected	Select Yes to choose the application server domains to be connected with each web server node.
Domain Name	Displays the domain associated with the application server.
Node	Displays the node containing application server domain.
Jolt port	Displays the Jolt port for the application server domain.

Field or Control	*Description*
IB Enabled	Flagged rows represent the application server domains that are enabled for Integration Broker.

3. Click Save. The node is provisioned with the selected configuration.

Note: At least one application server domain must be connected to each web server.

Overriding Default Topology and Attributes

If the template contains multiple topologies, you can override default topology and attributes.

To override the topology:

1. Select Yes in Override Topology field.

   This example illustrates the fields and controls in the Topology section of the Create Environment page.

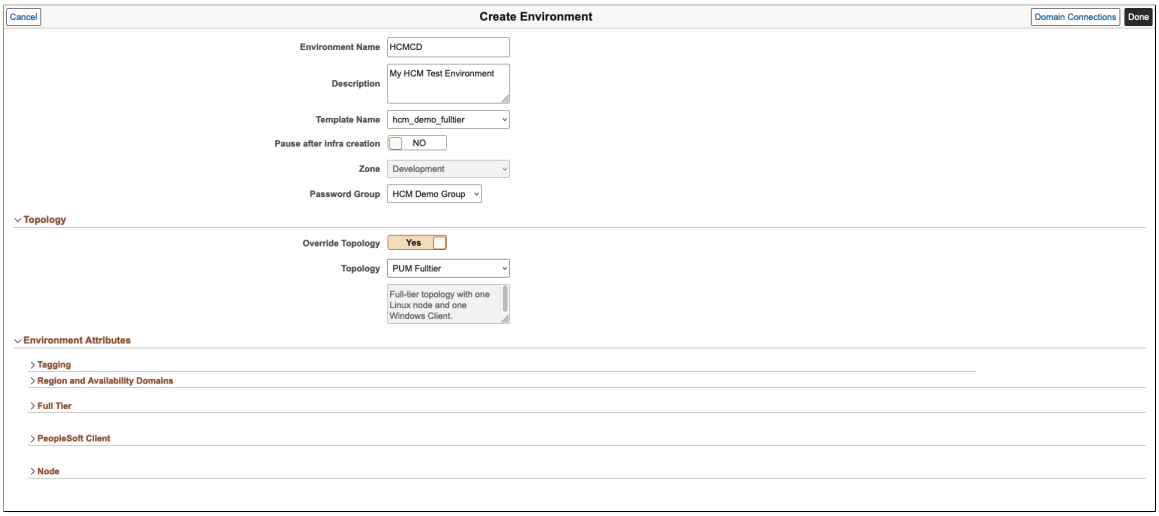

2. Select an appropriate Topology. Corresponding description is displayed in the below text area.
3. Input the required environment attributes. The different attributes are:

   - Full Tier: Full Tier is the VM where application server domain, process scheduler domain, and the web server domain are installed.
   - Middle Tier: Middle Tier Node can be created in either Linux or Windows. The Linux Middle Tier is the VM where application server domain, process scheduler domain, and the web server domain are installed. The Windows Middle Tier is the VM where Windows process scheduler is installed.
   - Database Tier: Database tier is the VM where the database (non-DbaaS) is installed for the new PSFT system.
   - PeopleSoft Client: PeopleSoft Client is the VM where PeopleTools client (for example, PeopleSoft Application Designer, or pside) and Change Assistant are pre-installed.

- Database as a Service: PeopleSoft database is deployed on DBaaS.
- Search Stack: Search Stack Tier is the VM where Elasticsearch server and Kibana (or OpenSearch and OpenSearch Dashboards) are installed.

4. Enter the PeopleSoft Client credentials and other required attributes.

   Note: In case of OCI, the password for the PeopleSoft Client instance should meet the password complexity as per the OCI requirement.

   Some custom attributes are displayed based on the selected topology nodes. If you select a Search Stack node, then you need to provide a couple of input parameters and passwords. Password must be at least 9 characters long and contain a numeric and one uppercase letter. Special characters are not accepted.

5. Click Done to start environment provisioning.

   Note: The system validates available resources before starting the provisioning process. See Validating Resources.

6. Select the Provision Task Status link to view the progress of the environment creation. See Viewing Provision Task Status. If a failure occurs, you can retry and resume the operation. See Retrying and Resuming Provisioning.

Note: Ensure to tune the servers, database, and PeopleSoft system for optimum performance once the deployment is completed.

Accepting Licensing Agreement

If you selected to use the Linux Image from Marketplace and this is the first time you are provisioning an environment in a specific compartment, a License agreement will be displayed and must be accepted to continue.

Each compartment needs to accept the license if it has not already been accepted for that compartment.

See Infrastructure Settings Page.

Using Shared File System for Linux Middle Tier using File Storage Service

Multiple middle tiers in an environment can share PS_HOME, PS_APP_HOME and PS_CUST_HOME.

Note: Shared File System is only supported for Linux middle Tier. Windows middle tier is not supported.

User must first create a new File Storage Service (FSS) in the OCI. It is recommended that this new FSS be in the same availability domain (AD) where the middle tier of the environment is provisioned.

Important! You should not use the File Storage System that you created for the Cloud Manager repository as the file system for a middle tier node in a provisioned environment.

When creating the FSS, keep the following in mind:

- All VMs in the subnet should have read/write access.

  See the tutorial Plan the Virtual Cloud Network for PeopleSoft Cloud Manager (Optional).

- File storage mount target (TCP ports 111, 2048, 2049, 2050; UDP ports 111 and 2048) is specific to FSS ports, this has to be opened in Linux MT machines.

- FSS export path requires full read/write access.

- Network access (ports and security rules) must be configured to Mount Targets from the middle tier nodes.

- Mount Target must be a minimum of 1024 GB.

  This example illustrates the fields and controls on the Mount Target Details page.

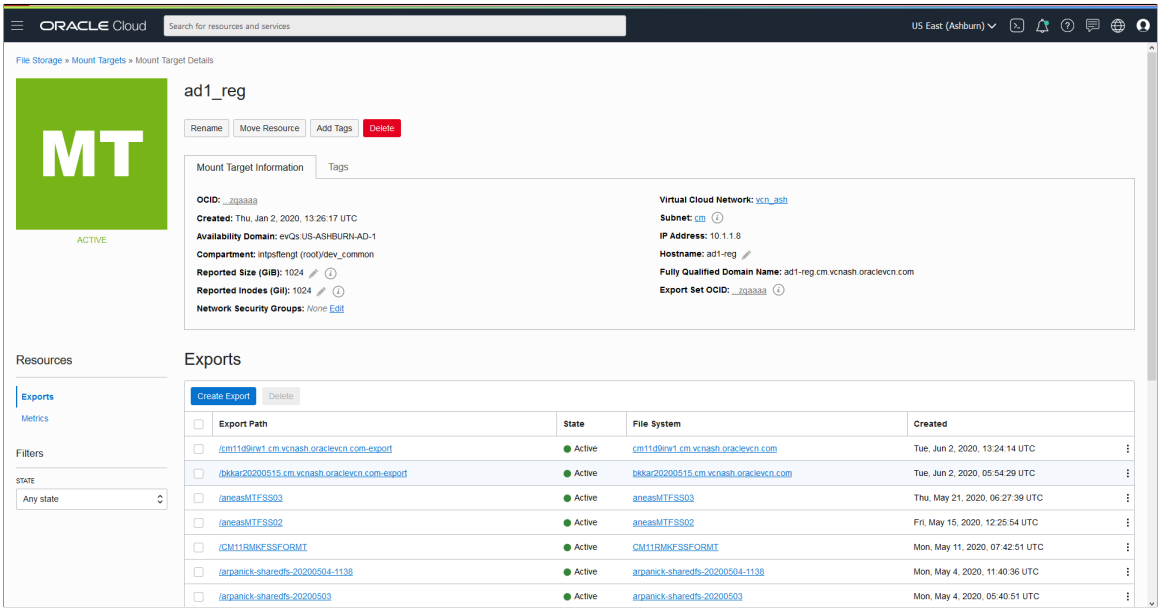

- Once the environment is running, user can update the FSS export path read/write permission to Linux Middle Tier.

  See Managing File Systems

This example illustrates the fields and controls on the File System (Linux Midtier) section. You can find definitions for the fields and controls later on this page.

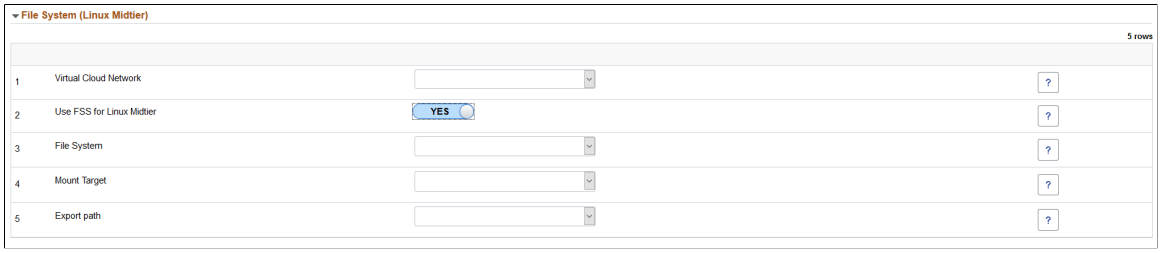

This example illustrates the File System warning message.

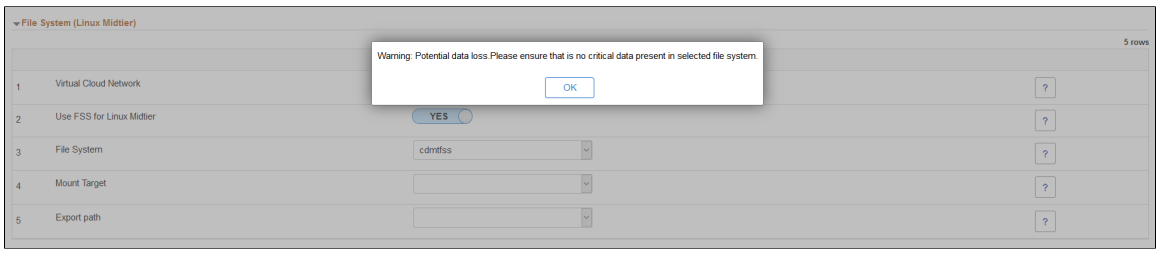

Field or Control	*Description*
Virtual Cloud Network	Select the Virtual Cloud Network.
Use FSS for Linux Midtier	Select Yes to use FSS for Linux Midtier
File System	Select the File System from the available file systems in OCI. **Note:** When you tab off the field, a warning is displayed for potential loss of data. The file system should not contain any critical data.
Mount Target	Select the Mount Target from the drop down list.
Export path	Select the Export path from the drop down list.

Adding a New Middle Tier

After creating an environment that includes a shared file system, use the Manage Node action to add a new middle tier that will share the same PS_HOME, PS_APP_HOME and PS_CUST_HOME.

See Managing Nodes

This example illustrates the fields and controls on the Manage Node page - Add Middle Tier with FSS.

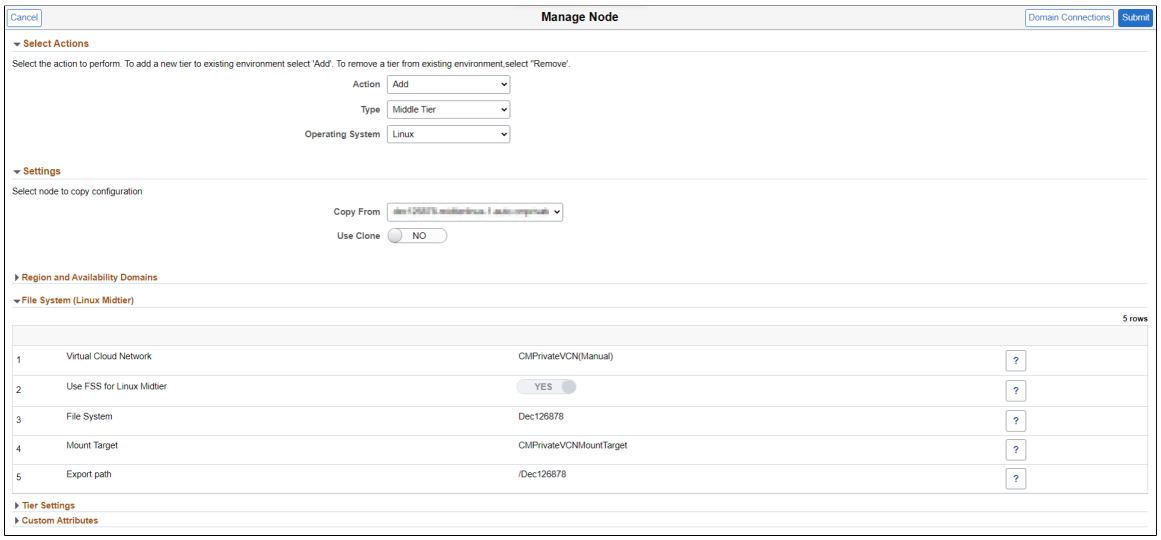

Manually Reviewing Steps During Processing

When you create an environment, you have the global option to enable all the manual reviews selected in the chosen template, or alternatively enable or disable all the manual review steps in the template.

If the status of an action step is "Manual Review", the Retry action is enabled in all the activities and action steps in that task. When you select Retry for any activity or action step, all the activities or action steps after that are run again.

By default, the option to go with the selection made during template creation is enabled. You can enable or disable all manual review steps using this field.

This example illustrates the fields and controls on the Create Environment page with Manual Review.

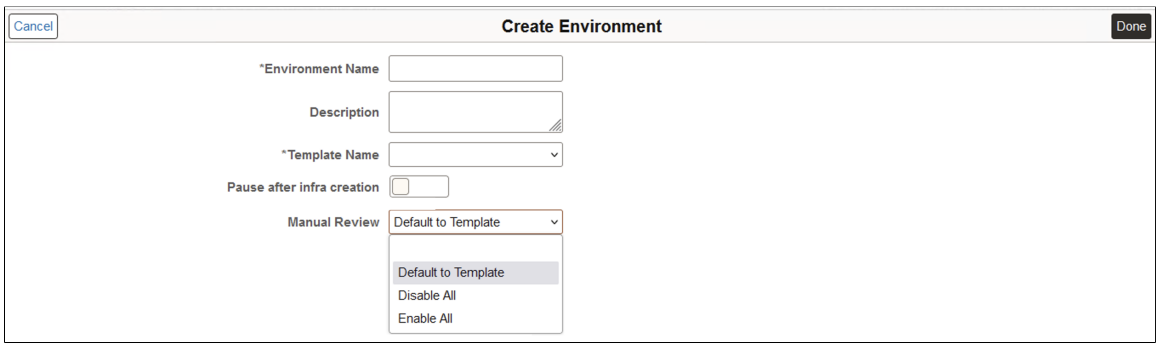

Field or Control	*Description*
Default to Template	Select this option to use the Manual Review values in chosen environment template as the default value in environment provisioning. You can override this selection in individual instances.

Field or Control	*Description*
Disable All	Select this option to disable Manual Review in all the instances.
Enable All	Select this option to enable Manual Review in all the instances.

You can modify the selection for individual instances in General Settings. By default, the selection for Manual Review is disabled.

This example illustrates the fields and controls on the General Settings tab.

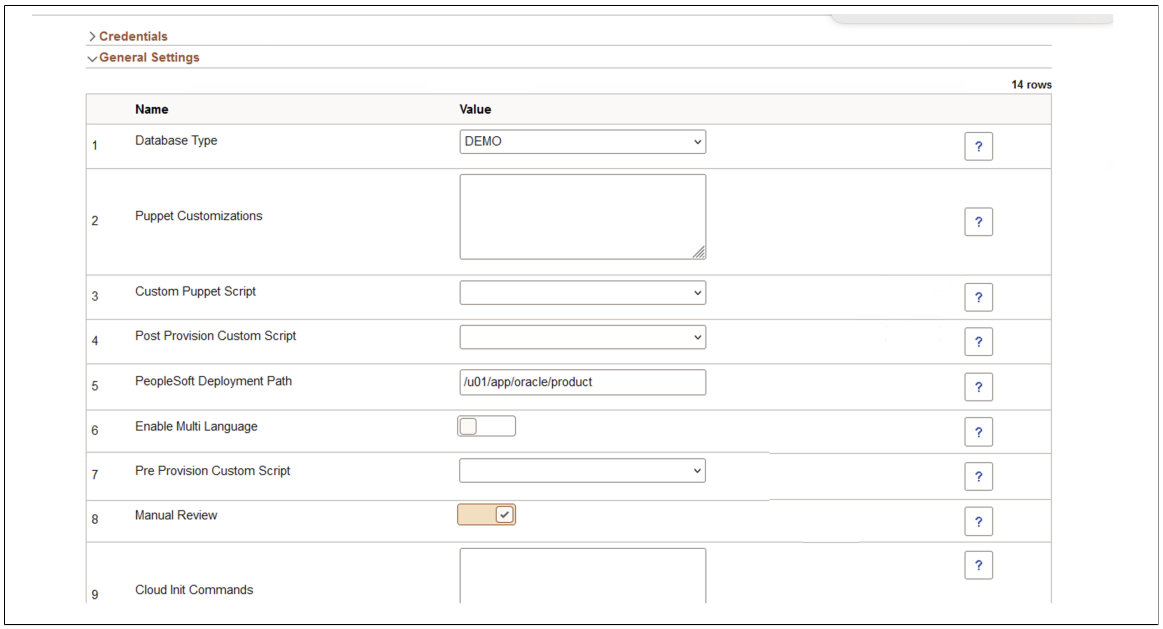

The environment tile shows the status as “Provisioning – Manual Review” when the manual review is in progress.

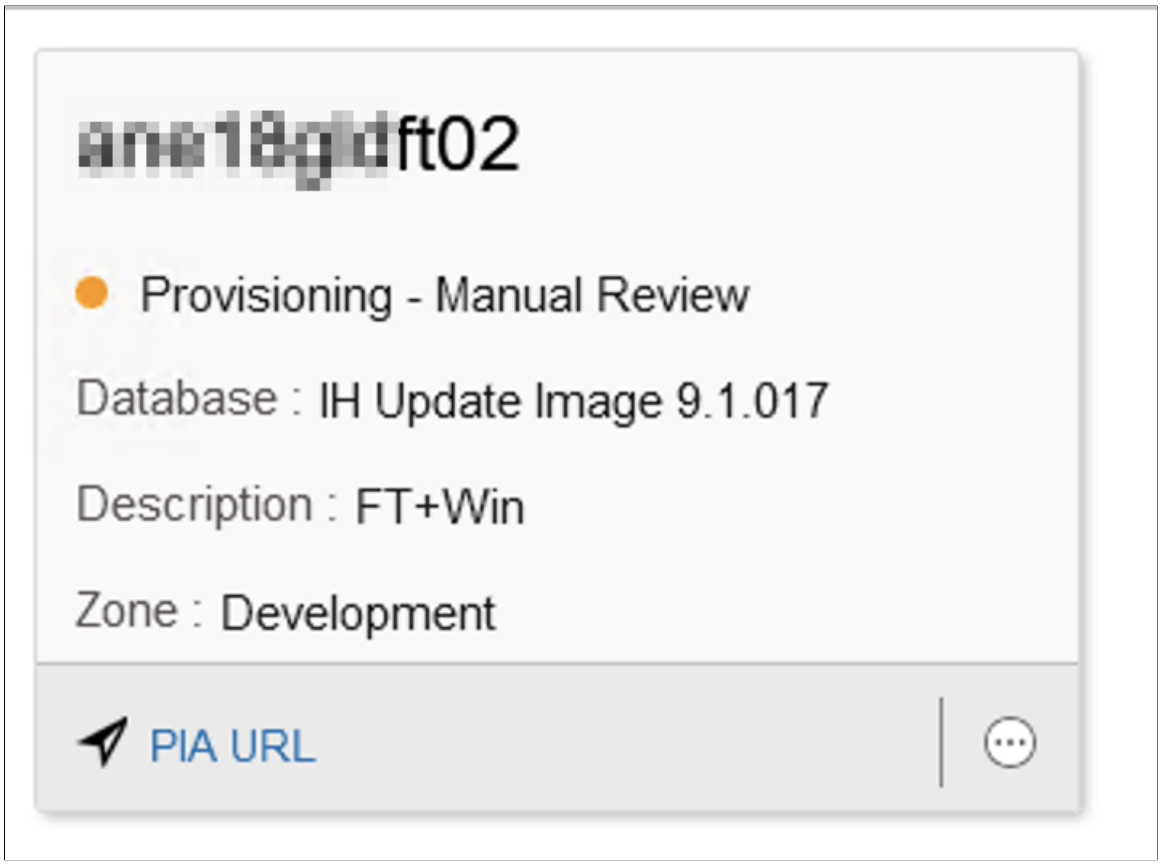

Manual reviewing of steps is currently available only during provisioning use case and is not applicable for shift provisioning and scaling up use cases. Manual Review is enabled for following environment types:

- Full Tier
- Middle Tier
- Database Tier
- Database Systems
- Windows MT

Note: The Manual Review step is included for reviewing the processing of Customer DPK. This capability will be extended to other actions and activities during future releases of Cloud Manager.

Validating Resources

When creating a new environment using Create Environment, Clone Environment, Shift Environment or adding nodes to an existing environment, the system will validate the resources before proceeding with the request.

If resources are not available for provisioning, the Resource Availability Validation page is displayed.

This example illustrates the Resource Availability Validation page when resources are not available.

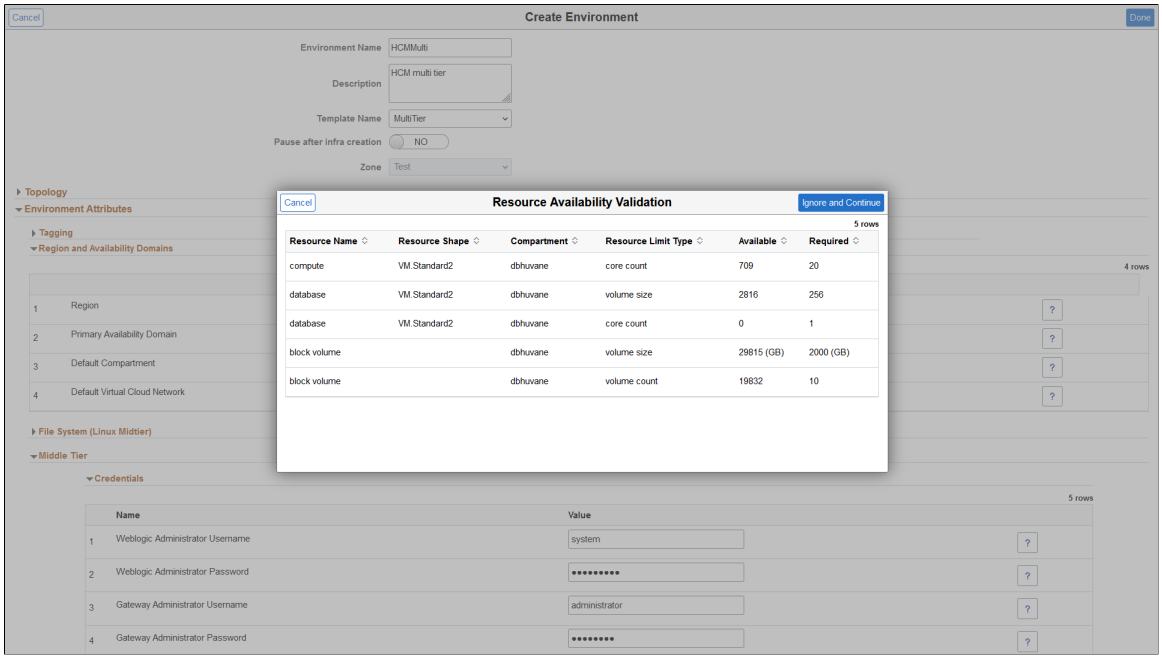

Field or Control	***Description***
Cancel	Select to cancel the request.
Ignore and Continue	If you know the resource will be available, you can select Ignore and Continue to proceed with the provisioning.

Accessing Environment Details

The Environment Details page (ECL_ENV_DET_FL) is a navigation collection that enables administrators to access more details of the environment from one location. It also enables the user to perform additional actions that can be performed on the environment such as performing applying a PeopleTools patch, viewing logs, and managing PUM connections.

Navigation:

Click the tile corresponding to an environment. Alternatively, click the Related Actions button corresponding to the environment and then select Details. The Environment Details page is displayed.

This example illustrates the fields and controls on the Environments Details Page. You can find definitions for the fields and controls later on this page.

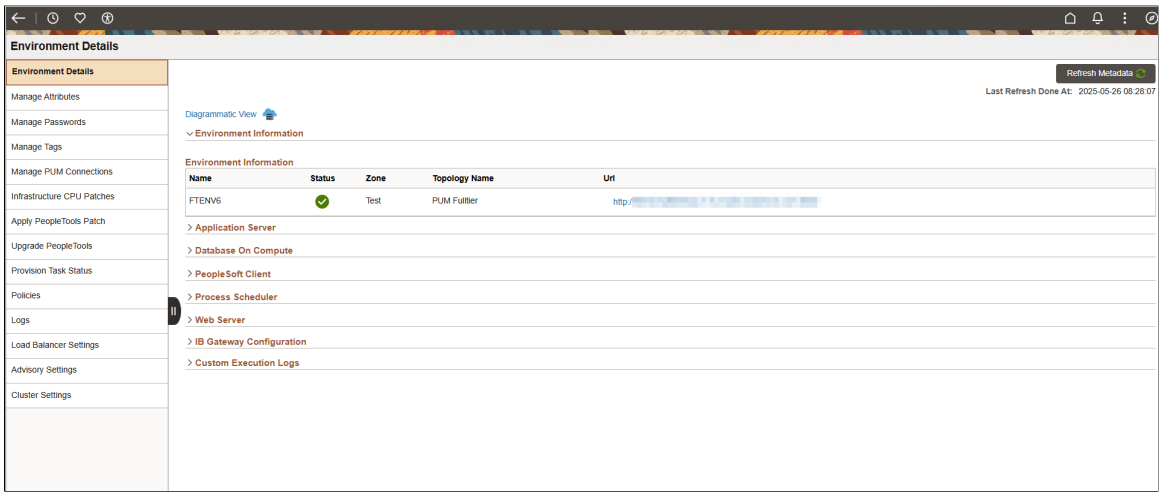

Field or Control	*Description*
Refresh Metadata	Click the Refresh Metadata button, at the upper-right corner of the page, to fetch the current status of the nodes in the environment and the PeopleSoft domains within the nodes. This also updates the diagrammatic view and configured gateways for the environment.
Diagrammatic View	Click to view a diagrammatic representation of all the instances and domains running inside the VMs.
Process Scheduler	This section provides details of the process scheduler component of the deployed PeopleSoft application environment. The Process Scheduler is responsible for processing scheduled tasks or jobs that typically do not happen during the course of a user's browser request.
Application Server	This section provides details of the application server component of the deployed PeopleSoft application environment. The application server acts as the business logic engine of the PeopleSoft system.
Database on: Compute	This section provides details of the database server of the deployed PeopleSoft application environment. The PeopleSoft applications refers to Oracle PeopleSoft products such as PeopleSoft Customer Relationship Management (CRM), PeopleSoft Enterprise Learning Management (ELM), PeopleSoft Financials and Supply Chain Management (FSCM), PeopleSoft Human Capital Management (HCM), and PeopleSoft Interaction Hub.
Web Server	This section provides details of the web server component of the deployed PeopleSoft application environment.

Field or Control	*Description*
IB Gateway Configuration	This section displays the ID and URL of existing Integration Broker gateways for the managed environment, which aid in environment discovery.
Database on: DBaaS	This section provides details of the database server of the deployed PeopleSoft application environment. The PeopleSoft applications refers to Oracle PeopleSoft products such as PeopleSoft Customer Relationship Management (CRM), PeopleSoft Enterprise Learning Management (ELM), PeopleSoft Financials and Supply Chain Management (FSCM), PeopleSoft Human Capital Management (HCM), and PeopleSoft Interaction Hub. **Note:** The 'Database on: DBaaS' section is displayed only when a user selects 'Database as a Service' node in topology.
PeopleSoft Client	This section provides details of the Windows Client of the deployed PeopleSoft application environment. This is the Microsoft Windows virtual machine on which PeopleSoft Application Designer and PeopleSoft Change Assistant will be installed.

Note: To access PSIDE (PeopleSoft Application Designer) and Change Assistant applications for this environment, connect with Remote Desktop (RDP) to the Windows VM using the IP address or hostnames provided under the PeopleSoft Client section.

This example illustrates the fields and controls on the Diagrammatic View page that can be viewed on clicking Diagrammatic View link.

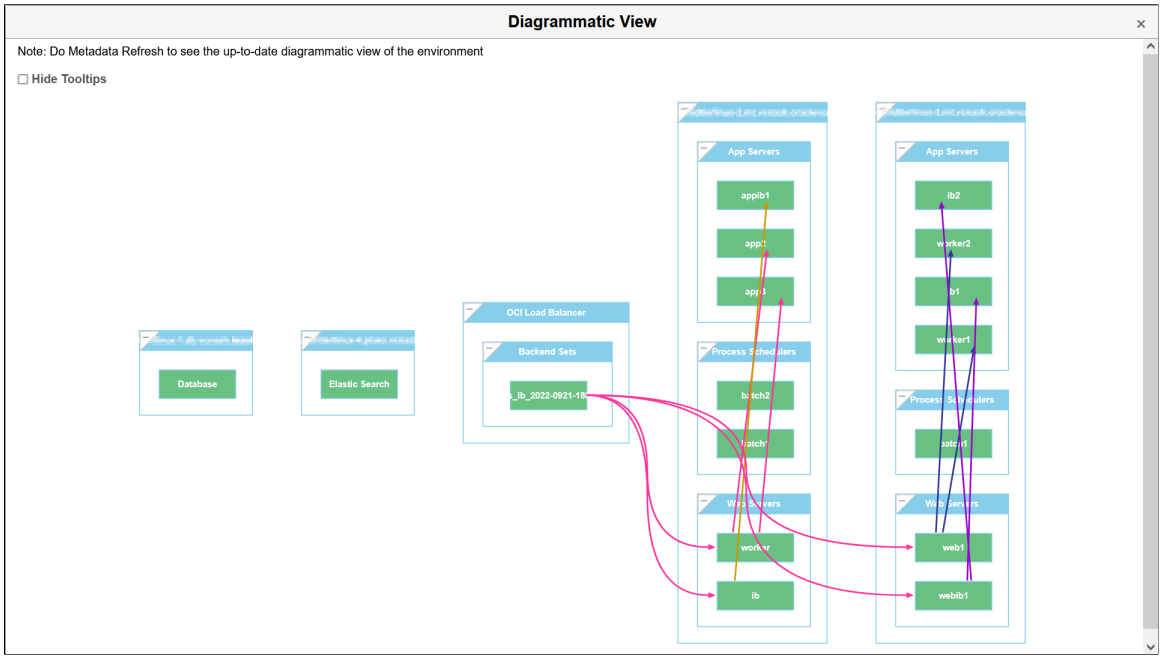

Hover the mouse over each instance for viewing the details.

You can view the status of different PeopleSoft services running within the VMs (application servers domains, process scheduler domains, web server domain, and the like) as shown:

This example illustrates the fields and controls on the Diagrammatic View Instance Details page.

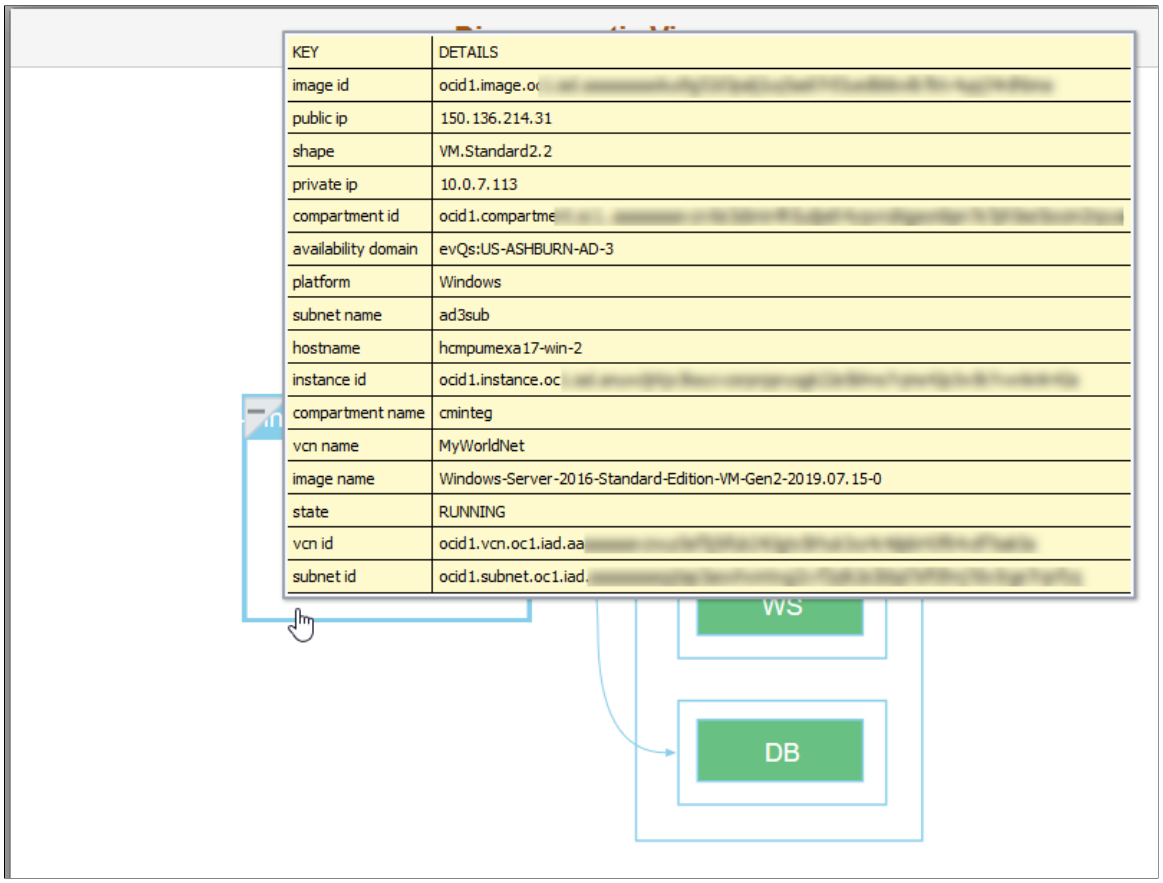

As an illustration, Process Scheduler domain details are described in the following section.

Process Scheduler Domain

This section provides details of the Process Scheduler component of the deployed PeopleSoft application environment. The Process Scheduler is responsible for processing scheduled tasks or jobs that typically do not happen during the course of a user's browser request.

Navigation:

Expand Process Scheduler available on the Environment Details page.

This example illustrates the fields and controls on the Process Scheduler section for Process Scheduler middle tier on Linux.

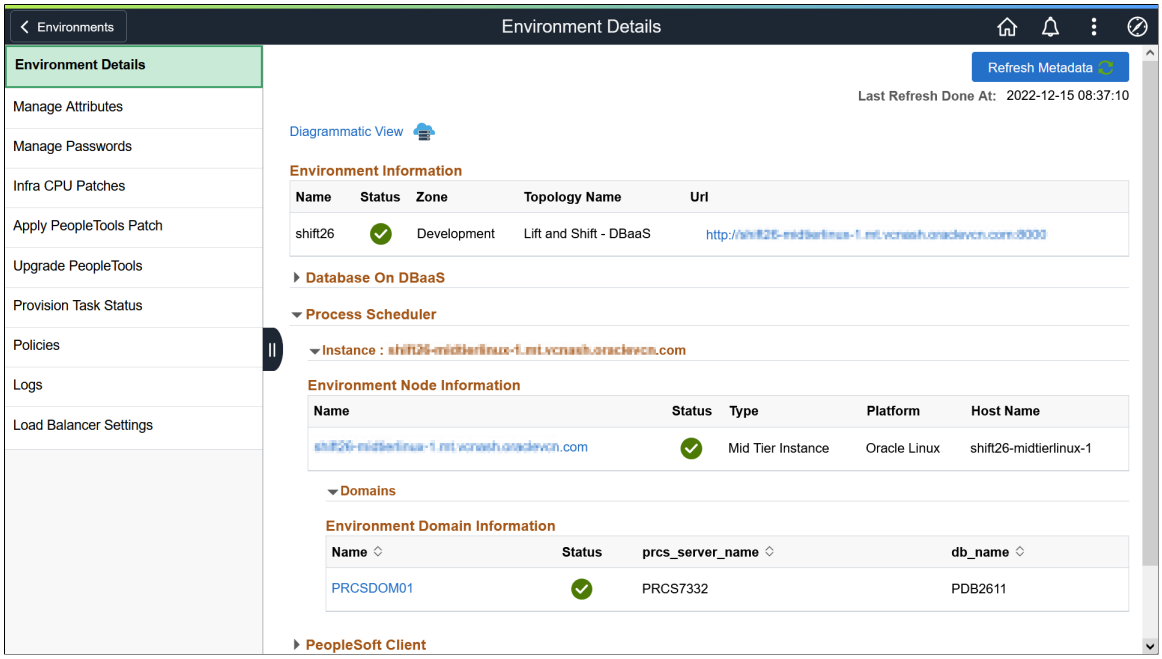

This is an example of where Process Scheduler was configured on a Windows instance.

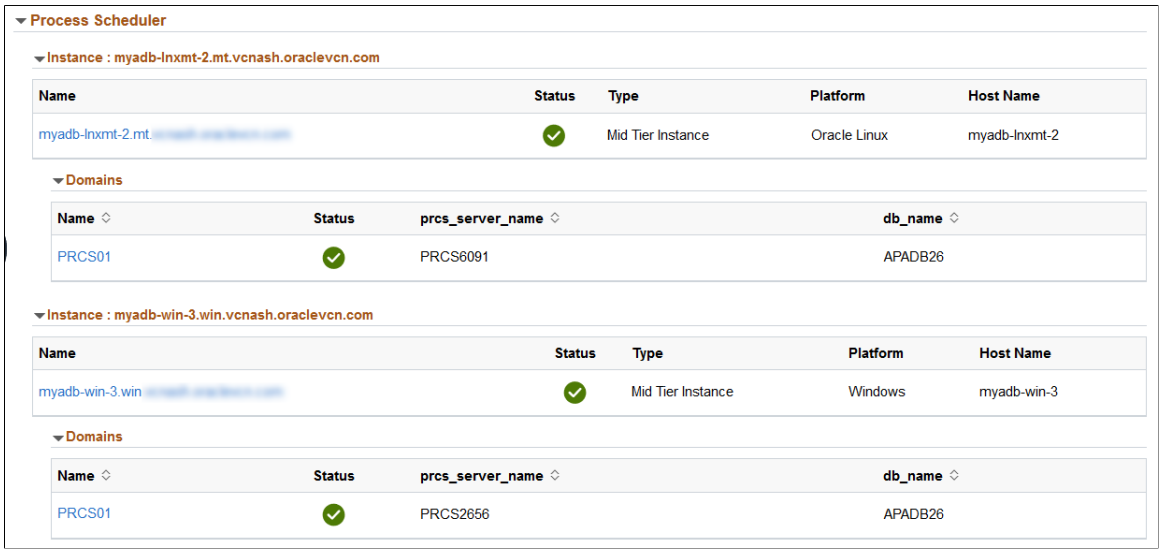

Instance Details Modal Window

Use Instance Details modal window to view more details about the virtual machine.

Navigation:

Click on the instance name.

This example illustrates the instance details for PeopleSoft Client.

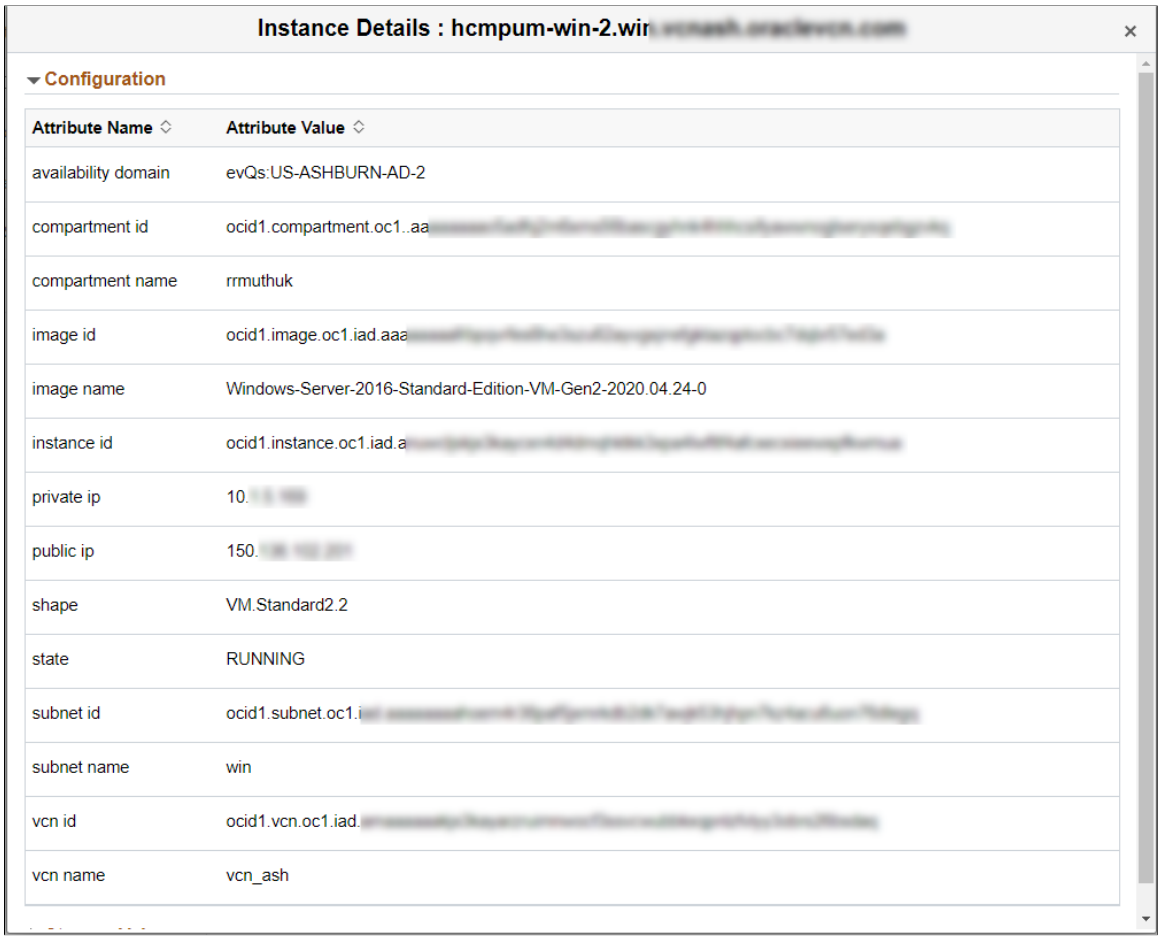

Domain Details Modal Window

Use the Domain Details modal window to view domain details.

Navigation:

Click on the domain name.

This example illustrates the configuration domain details for the Process Scheduler PRCS01.

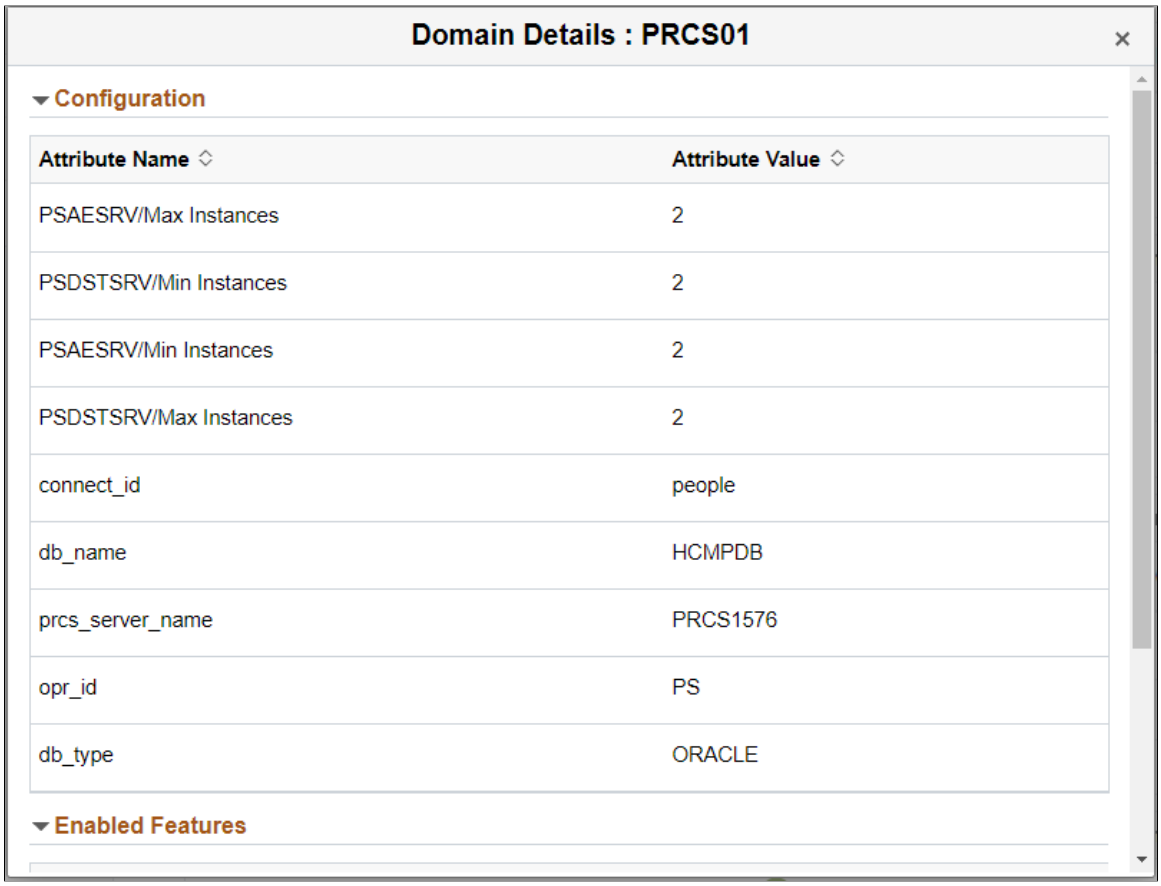

This example illustrates the enabled features for the Process Scheduler domain PRCS01.

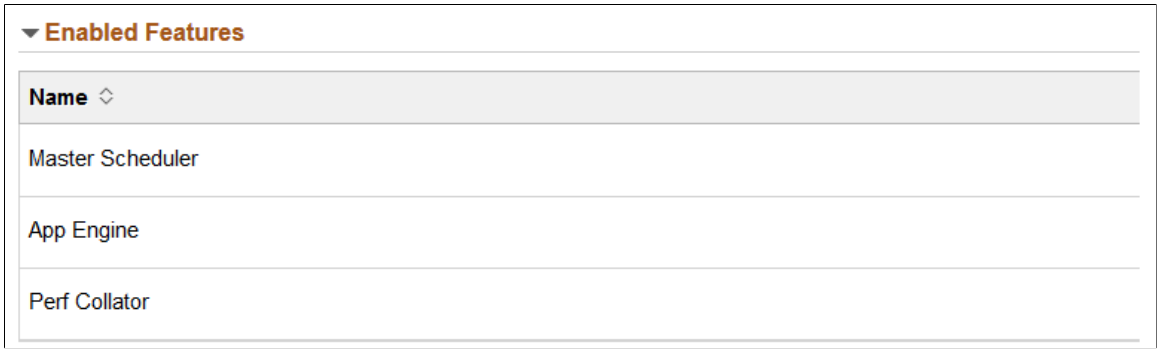

Master Scheduler, Application Engine, and Performance Monitor features can be enabled for the instance.

Configuring IB Gateway

The IB Gateway Configuration section can be used for configuring IB gateways on managed environments.

You can use the IB Gateway Configuration section on the Environment Details page to support the following IB configurations from Cloud Manager:

- Reconfiguring local IB gateway for an environment

- Configuring multiple environments using a single shared gateway
- Configuring remote gateway for an environment
- Configuring load balancer on local gateway for an environment
- Renaming IB nodes

You can utilize pre-defined policies to configure IB on an environment according to your dynamic requirements. The IB Gateway Configuration section on Environment Details page displays the existing IB gateways for a managed environment.

This example illustrates the IB Gateway Configuration section on Environment Details page.

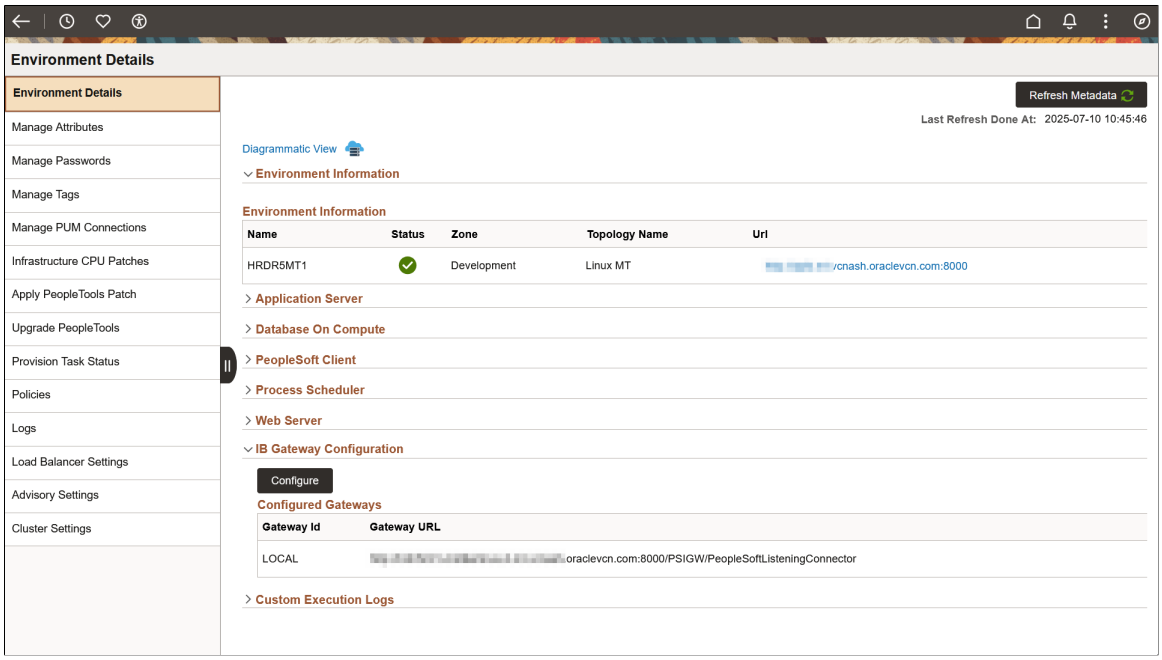

Click Configure to open the Configure IB Gateway page that lets you edit the parameters of the pre-defined internal policy action Configure IB Gateway.

Reconfiguring Local IB Gateway for Environments

To configure local IB gateway on an existing environment:

1. Click Configure on IB Gateway Configuration section. The IB Gateway Configuration page is displayed.

   Alternatively, navigate to Orchestration Manager > Policy Editor and click the View/Edit button corresponding to Configure IB Gateway internal policy.

2. Click the Parameters button on Policy Action Parameters field, corresponding to the policy action Configure IB Gateway.

This example illustrates the fields on Policy Action Parameters dialog box.

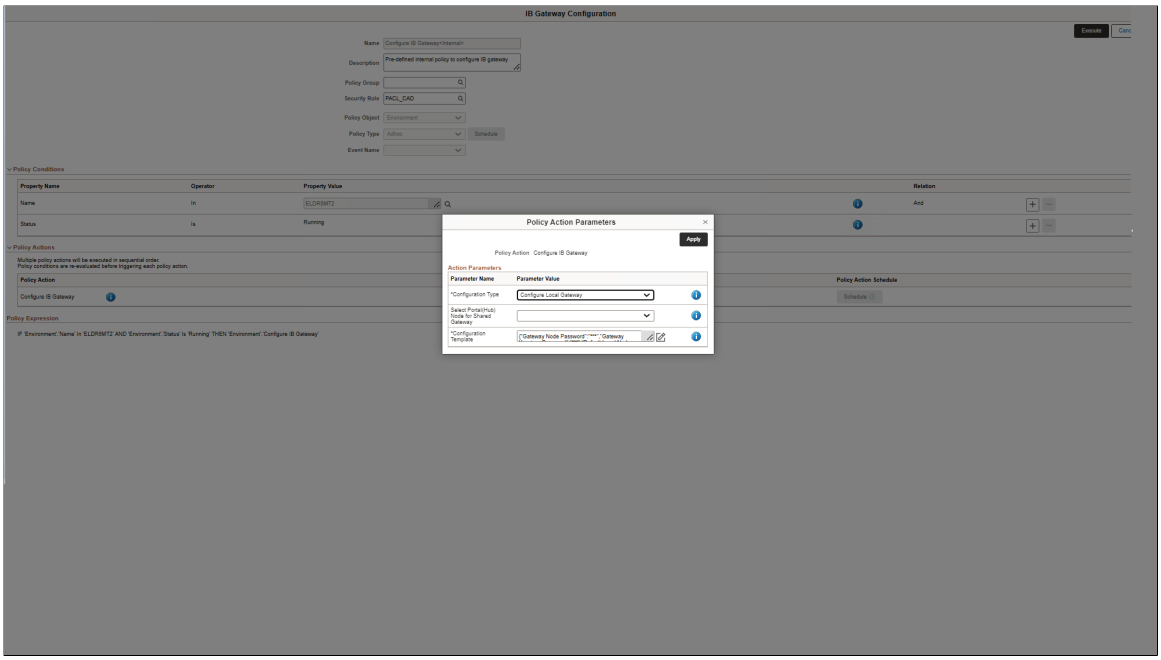

3. Select Configure Local Gateway as the Configuration Type. You can ignore the field for selecting portal node because it applies only for shared gateways.

4. Click the Activate button next to Configuration Template field, which is auto-populated when you select the configuration type. The Configuration Template page is displayed.

This example illustrates the fields on Configuration Template for Local Gateway.

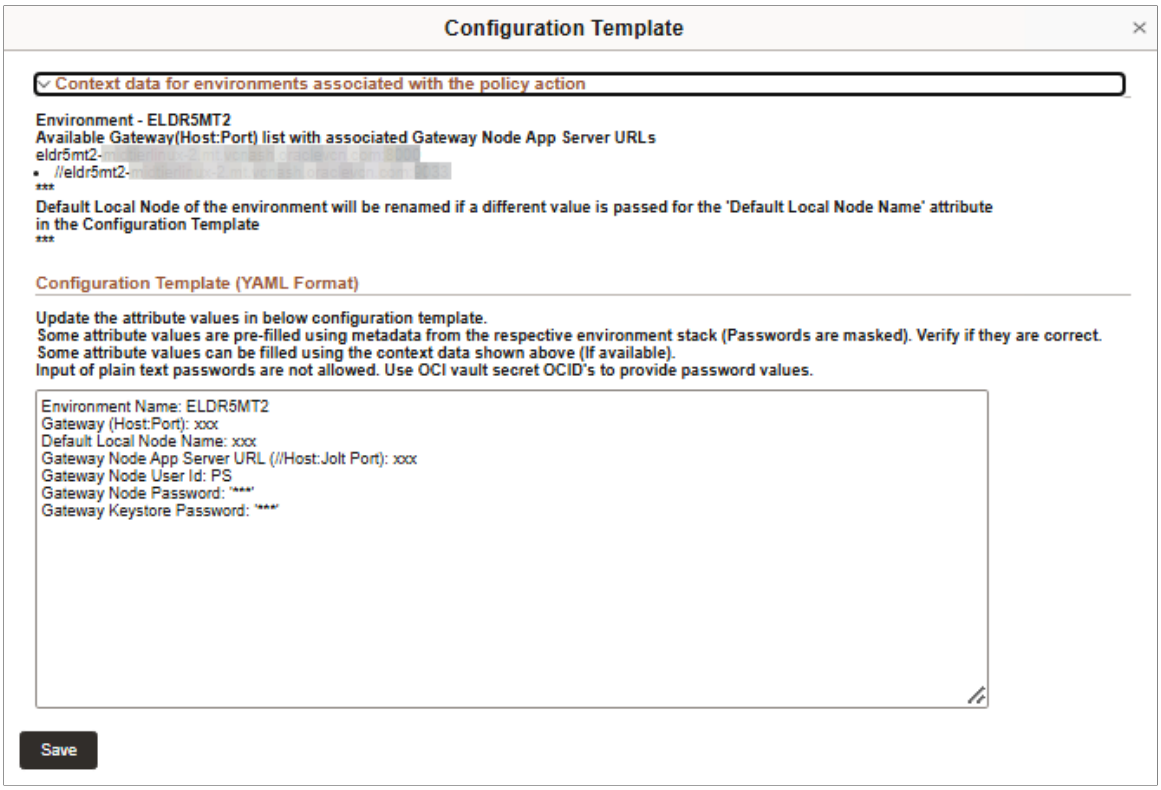

Use the attribute values provided in the Context Data section to replace the variables given in attributes in the YAML format.

Field or Control	***Description***
Context data for environments associated with policy action	This section provides context help for filling up the configuration template. The available environments and gateways are listed.

Field or Control	*Description*
Configuration Template (YAML Format)	The following attribute values are listed in YAML format: • **Environment Name:** This value is pre-filled. • **Gateway (Host:Port):**Enter an available Gateway (Host:Port) from the context data section. • **Default Local Node Name:**Enter the name of the local node in the selected environment. • **Gateway Node App Server URL:**Enter an available application server URL associated with the selected web domain from the context data section. • **Gateway Node User Id:**The user ID associated with gateway node. This value is pre-filled. • **Gateway Node Password:**The password associated with gateway node. This value is pre-filled. • **Gateway Keystore Password:**This value is pre-filled for new environments provisioned through Cloud Manager. Enter the OCID vault secret if the value is not pre-filled. Plain text passwords do not work in this field.

5. Click Save.
6. Click Apply. The policy action parameters are validated.
7. Click Execute on the IB Gateway Configuration page. A confirmation message is displayed. When you confirm, Cloud Manager runs the policy and the Policy Monitor page is displayed. When you run the policy from the Policy Editor page in Orchestration Manager, click the Run Policy button corresponding to the Configure IB Gateway policy to run the policy.

This example illustrates the fields on Policy Editor page.

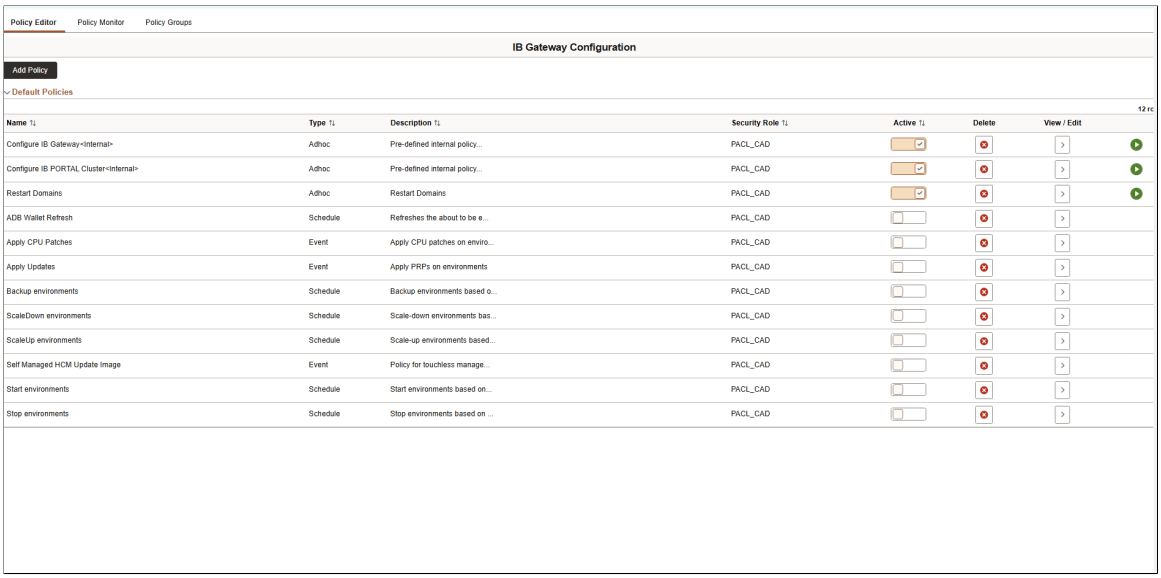

You can also configure the local IB gateway for a newly created environment while selecting the required configuration for an environment or node using Domain Connections while provisioning an environment with multiple middle tier nodes or while adding a new node.

Configure connections between application server domains and web server domains using Domain Connections page.

This example illustrates the fields and controls on the Domain Connections page.

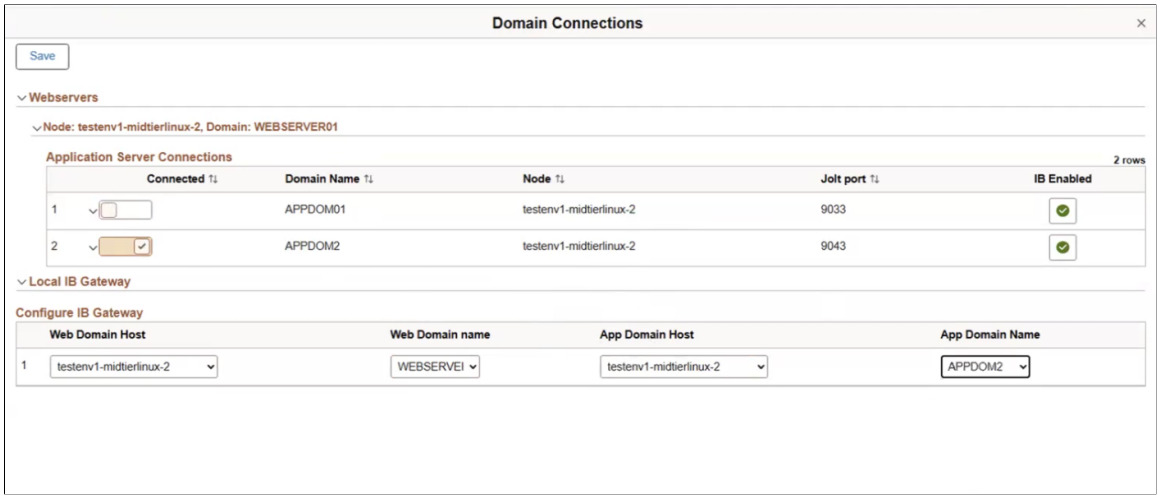

Field or Control	***Description***
Web Domain Host	Select one of the available middle tier nodes, which is IB enabled.

Field or Control	*Description*
Web Domain Name	Select the name of the web server associated with the web domain host.
App Domain Host	Select the app server node that is connected to the selected web server.
App Domain Name	Select the domain that you want the web domain to be connect with.

Renaming IB Nodes

You can use the IB Gateway Configuration page to rename the IB nodes before you refresh the existing IB node definitions.

To rename an IB node on an existing gateway:

1. Click Configure on IB Gateway Configuration section. The IB Gateway Configuration page is displayed.

   Alternatively, navigate to Orchestration Manager > Policy Editor and click the View/Edit button corresponding to Configure IB policy.

2. Click the Parameters button on Policy Action Parameters field, corresponding to the policy action Configure IB Gateway.

3. Select Rename Node as the Configuration Type. You can ignore the field for selecting portal node.

4. Click the Activate button next to Configuration Template field, which is auto-populated when you select the configuration type. The Configuration Template page is displayed.

This example illustrates the fields on Configuration Template for Rename Node Configuration Type.

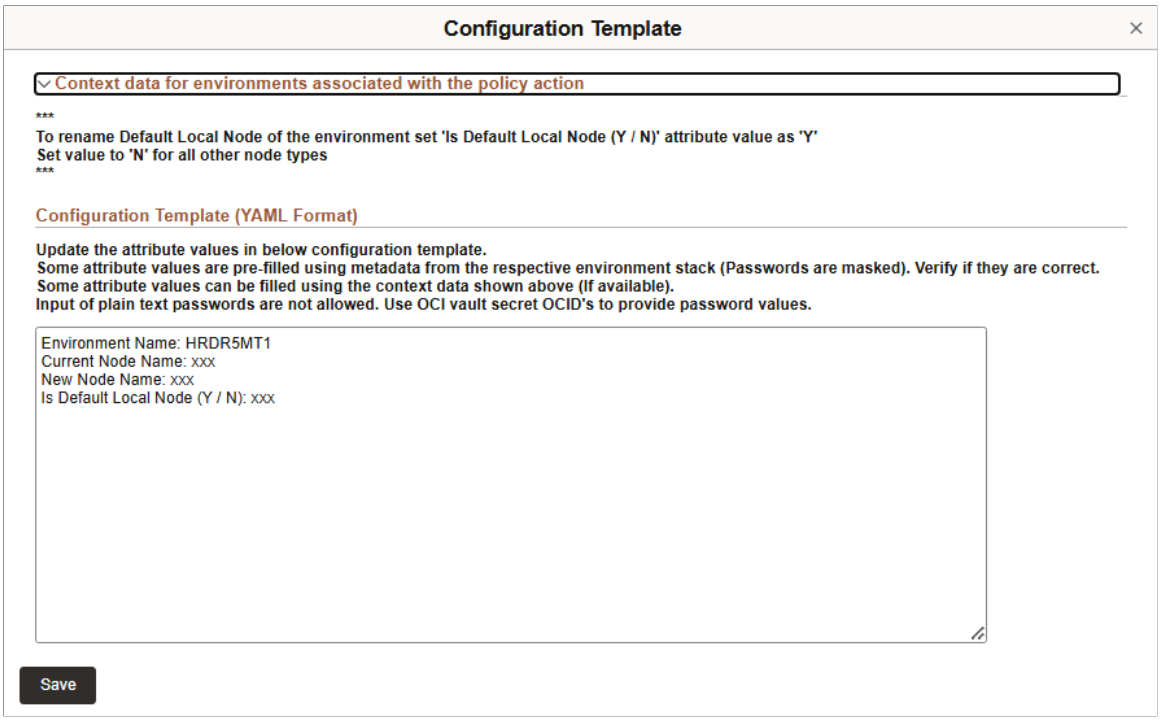

Use the attribute values provided in the Context Data section to replace the variables given in attributes in the YAML format.

Field or Control	***Description***
Context data for environments associated with policy action	This section provides context help for filling up the configuration template. The available environments and gateways are listed.
Configuration Template (YAML Format)	The following attribute values are listed in YAML format: • **Environment Name:** This value is pre-filled. • **Current Node Name:**Enter the name of the node you want to rename. • **New Node Name:**Enter the new name of the node. • **Is Default Local Node (Y/N):**Enter 'Y' if you want to rename the default local node associated with the selected environment. Enter 'N' if you want to rename any other node.

5. Click Save.
6. Click Apply. The policy action parameters are validated.

7. Click Execute on the IB Gateway Configuration page. The policy to rename an IB node on an existing gateway is triggered.

Configuring Shared Gateway

You can use the IB Gateway Configuration page to configure a setup where multiple PeopleSoft applications share the same local IB gateway.

The following are the prerequisites for configuring environments with a shared gateway:

- The shared gateway node must be at the same or higher release of PeopleTools.This restriction is not applicable for content nodes.
- The environments must not be separated through firewalls.
- All the environments with a shared gateway must be in Running state.

To configure a shared gateway:

1. Click Configure on IB Gateway Configuration section. The IB Gateway Configuration page is displayed.

   Alternatively, navigate to Orchestration Manager > Policy Editor and click the View/Edit button corresponding to Configure IB policy.

2. Click the search icon next to the default environment name that is displayed in the Policy Conditions section and associate all the environments for the shared cluster to the policy definition.
3. Click the Parameters button on Policy Action Parameters field, corresponding to the policy action Configure IB Gateway.
4. Select Configure Shared Gateway as the Configuration Type.
5. Select the gateway environment as the Portal (Hub) node for Shared Gateway.
6. Click the Activate button next to Configuration Template field, which is auto-populated when you select the configuration type. The Configuration Template page is displayed.

This example illustrates the fields on Configuration Template for Shared Gateway.

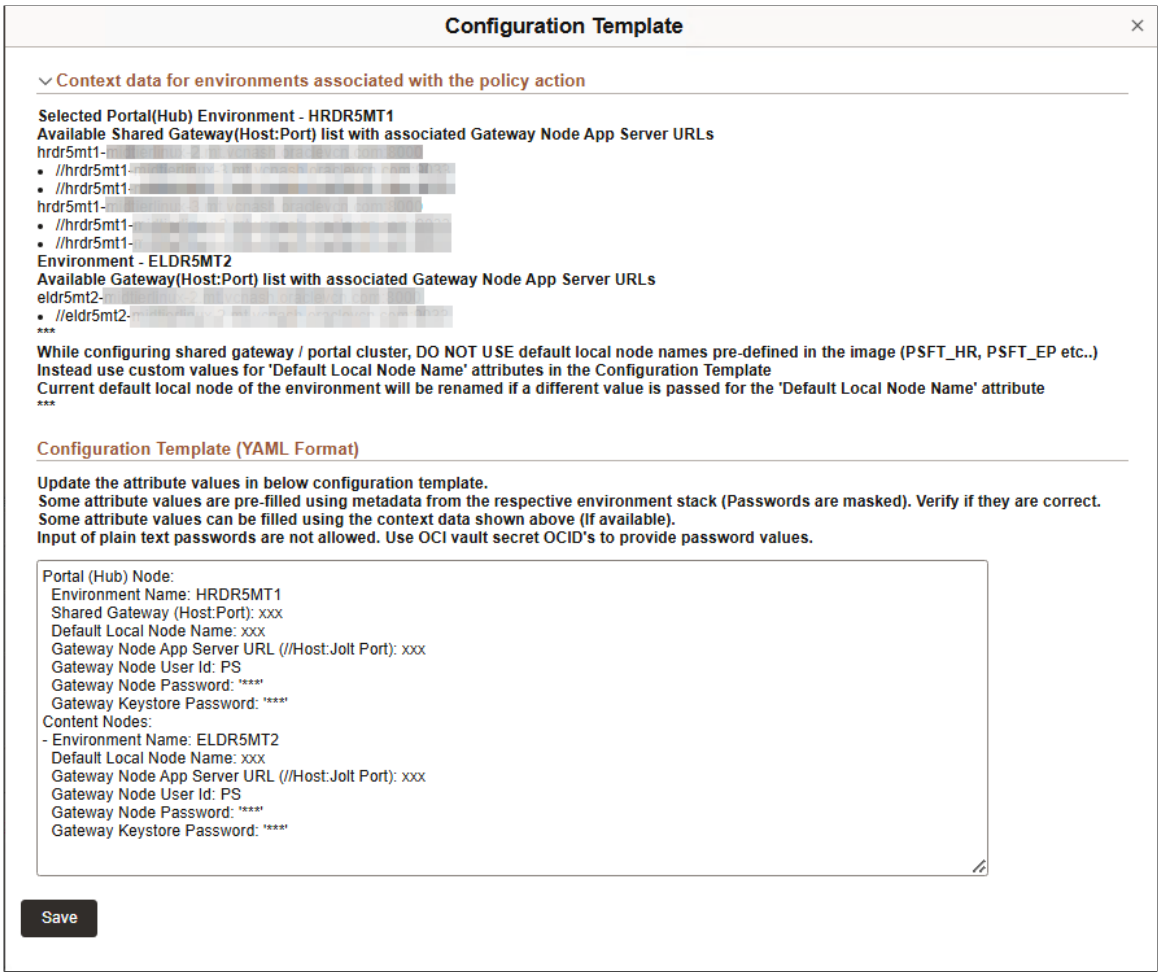

Use the attribute values provided in the Context Data section to replace the variables given in attributes in the YAML format.

Field or Control	***Description***
Context data for environments associated with policy action	This section provides context help for filling up the configuration template. The available environments and gateways are listed. There must be at least two environments listed here.

Field or Control	*Description*
Configuration Template (YAML Format)	The attributes for portal and content nodes are listed separately. The following attribute values are listed in YAML format: • **Environment Name:** This value is pre-filled. • **Gateway (Host:Port):**Enter an available gateway (host:port) from the context data section. • **Default Local Node Name:**Enter the name of the local node in the selected environment. • **Gateway Node App Server URL:**Enter an application server URL associated with the selected gateway from the context data section. • **Gateway Node User Id:**The user ID associated with gateway node. This value is pre-filled. • **Gateway Node Password:**The password associated with gateway node. This value is pre-filled. • **Gateway Keystore Password:** This value is pre-filled for new environments provisioned through Cloud Manager. Enter the OCID vault secret if the value is not pre-filled. Plain text passwords do not work in this field.

7. Click Save.
8. Click Apply. The policy action parameters are validated.
9. Click Execute on the IB Gateway Configuration page. The policy to configure a shared gateway is triggered.

Configuring Remote Gateway

Remote gateways enable two-way communication between integration participants where the local gateway for one application serves as the remote gateway for the other application. You can configure a remote gateway when a direct connection between integration participants is not possible through the internet. IB uses the default remote gateway connector (PeopleSoft Listening connector) on the local gateway to send messages to the PeopleSoft listening connector on the remote gateway.

Note: Both the environments must be in Running state for configuring the remote gateway.

To configure a remote gateway:

1. Click Configure on IB Gateway Configuration section. The IB Gateway Configuration page is displayed.

Alternatively, navigate to Orchestration Manager > Policy Editor and click the View/Edit button corresponding to Configure IB policy.

2. Click the search icon next to the default environment name that is displayed in the Policy Conditions section and associate both the environments participating in the remote gateway with the policy definition.
3. Click the Parameters button on Policy Action Parameters field, corresponding to the policy action Configure IB Gateway.
4. Select Configure Remote Gateway as the Configuration Type.
5. Click the Activate button next to Configuration Template field, which is auto-populated when you select the configuration type. The Configuration Template page is displayed.

   This example illustrates the fields on Configuration Template for Remote Gateway.

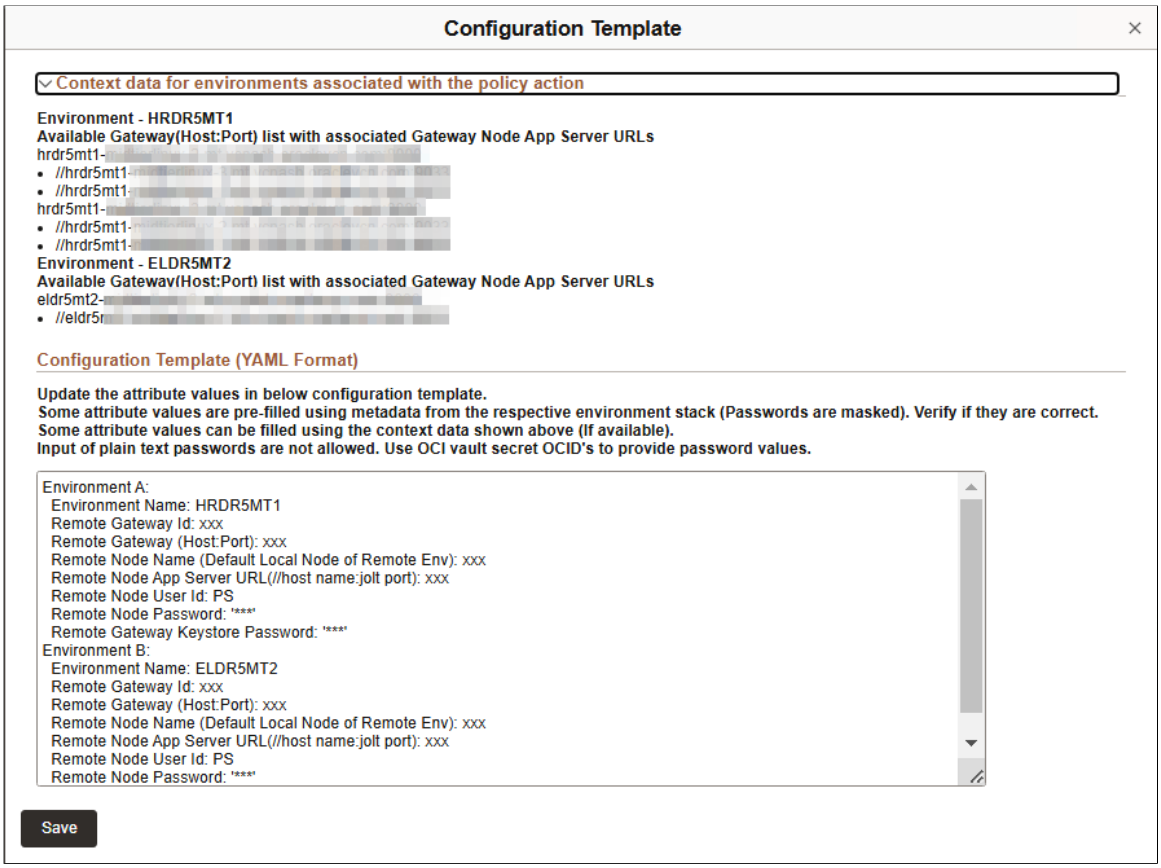

Use the attribute values provided in the Context Data section to replace the variables given in attributes in the YAML format.

Field or Control	*Description*
Context data for environments associated with policy action	This section provides context help for filling up the configuration template. The available environments and gateways are listed. There must be at least two participating environments listed for remote gateways.
Configuration Template (YAML Format)	The attributes for both the environments participating in the remote gateway are listed separately. The following attribute values are listed in YAML format: • **Environment Name:** This value is pre-filled. • **Remote Gateway Id:**The user ID associated with remote gateway node. This value is pre-filled. • **Remote Gateway (Host:Port):** Enter an available gateway (host:port) from the context data section. • **Remote Node Name (Default Local Node of Remote Env):**Enter the name of the local node in the selected environment. • **Remote Node App Server URL (//host name:jolt name):** Enter an application server URL associated with the selected gateway from the context data section. • **Remote Node User Id:**The user ID associated with remote node. This value is pre-filled. • **Remote Node Password:**The password associated with remote node. This value is pre-filled. • **Remote Gateway Keystore Password:**This value is pre-filled for new environments provisioned through Cloud Manager. Enter the OCID vault secret if the value is not pre-filled. Plain text passwords do not work in this field.

6. Click Save.
7. Click Apply. The policy action parameters are validated.
8. Click Execute on the IB Gateway Configuration page. The policy to configure a remote gateway is triggered.

Configuring Load Balancer on Local Gateway

You can configure a Load Balancer(LB) as the local IB gateway URL for inbound calls to an environment by specifying an LB URL and then defining and configuring one or more physical gateways. Inbound calls that are made on the LB are delegated to one of the configured physical gateways.

To configure a load balancer on local gateway:

1. Click Configure on IB Gateway Configuration section. The IB Gateway Configuration page is displayed.

   Alternatively, navigate to Orchestration Manager > Policy Editor and click the View/Edit button corresponding to Configure IB policy.

2. Ensure that the Environment Name field displays the environment where LB URL must be configured on the local IB gateway.

3. Click the Parameters button on Policy Action Parameters field, corresponding to the policy action Configure IB Gateway.

4. Select Configure Load Balancer as the Configuration Type.

5. Click the Activate button next to Configuration Template field, which is auto-populated when you select the configuration type. The Configuration Template page is displayed.

   This example illustrates the fields on Configuration Template for Load Balancer.

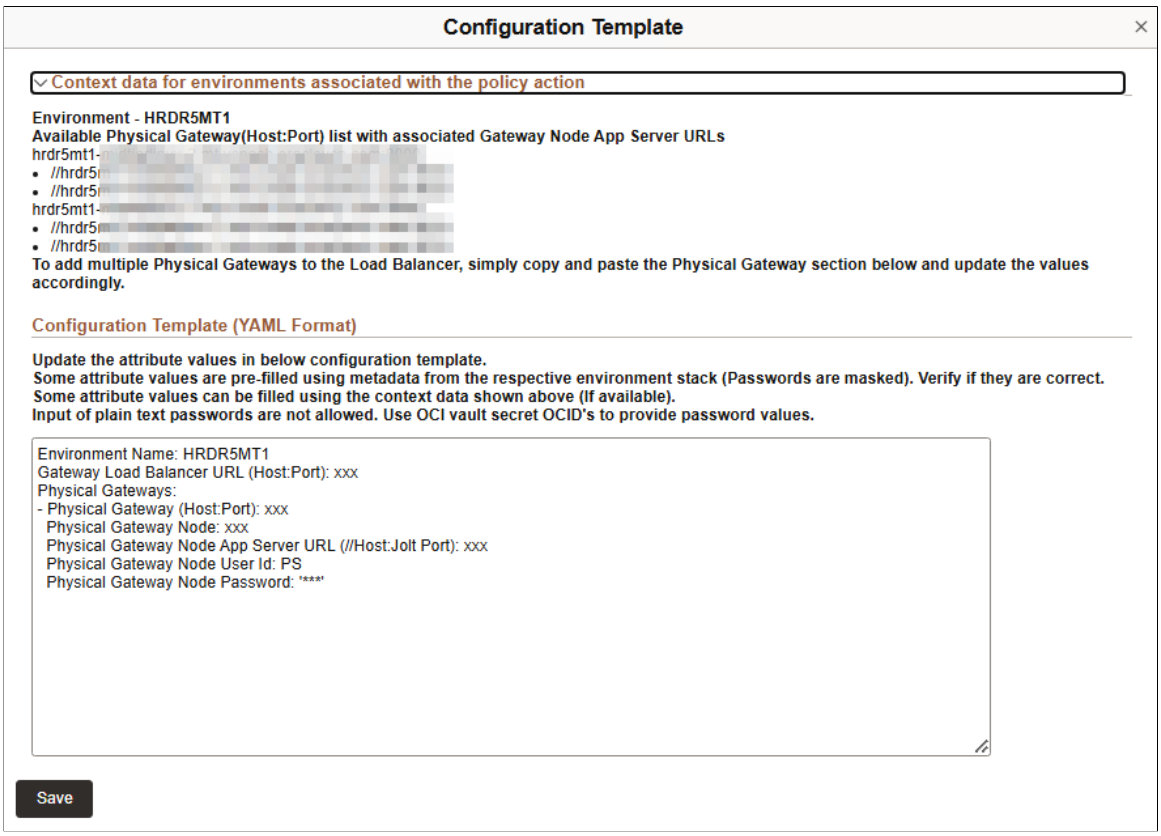

   Use the attribute values provided in the Context Data section to replace the variables given in attributes in the YAML format.

Field or Control	*Description*
Context data for environments associated with policy action	This section provides context help for filling up the configuration template. The available environments and gateways are listed.
Configuration Template (YAML Format)	The following attribute values are listed in YAML format: • **Environment Name:** This value is pre-filled. • **Gateway Load Balancer URL (Host:Port):**Enter a web domain that is displayed in the context data section. • **Physical Gateway (Host:Port):**Enter the name of a local node in the selected environment. • **Physical Gateway Node**Enter an application server URL associated with the selected web domain. • **Physical Gateway Node App Server URL (//Host: Jolt Port):**The user ID associated with gateway node. This value is pre-filled. • **Physical Gateway Node User Id:**The user ID associated with physical gateway node. This value is pre-filled. The password associated with gateway node. This value is pre-filled. • **Physical Gateway Node Password:**The password associated with physical gateway node. This value is pre-filled.

6. Click Save.
7. Click Apply. The policy action parameters are validated.
8. Click Execute on the IB Gateway Configuration page. The policy to configure load balancer is triggered.

Accessing Provisioned Environments

To access PIA of provisioned environment, click on the PIA URL link on the environment card. This will launch PIA of the newly created environment. To view more details about the environment, refer to Accessing Environment Details.

Note: If environments are provisioned on private subnets, then use a Bastion server or a Windows instance as a jump host. The bastion or the Windows instance must be set up on a public subnet accessible from internet.
The PIA URL must be modified appropriately to access the environment if you have entered a custom authentication token domain value in web server domain.

Linux instances can be accessed using SSH. SSH private keys are required to connect to the provisioned instances. There are two private keys that can be used:

1. Cloud Manager SSH Keys for Administration — This is the SSH key pair that is created by Cloud Manager and the public key is automatically injected into the newly provisioned instances. The SSH private key file cm_adm_pvt_key is available under /home/psadm2/psft/data/cloud/ocihome/keys/.

2. User SSH key — A Cloud Manager user can create a set of SSH key pair and configure the public key in My Settings page. Using the private key, user can connect to the provisioned instances. For more details refer Configuring My Settings .

Windows instances can be accessed using Remote Desktop (RDP). Ensure to enable RDP ports in OCI security lists, as well as in the client-side firewalls.

The IP addresses for Linux and Windows instances can be determined from the Environment Details Page. See Accessing Environment Details.

Important! It is user's responsibility to back up SSH keys for Administration and User SSH Keys to avoid losing access to provisioned instances due to loss of Cloud Manager instance or any fatal failures.

Updating SSH Keys

In Oracle Cloud Infrastructure, SSH keys are used to provide secure access to all Linux instances. It is user's responsibility to manage and secure the SSH keys that are used in OCI. Cloud Manager also uses SSH keys for managing environment nodes. Cloud Manager injects two SSH public keys into any node that it provisions. The SSH keys are:

1. SSH keys for Administration

   This key pair is generated by Cloud Manager at the time of installation (bootstrap). This key pair is used to connect and manage Linux instances provisioned as PeopleSoft environments. The key pair is available under the path /home/psadm2/psft/data/cloud/ocihome/keys. The two files for this key pair are:

   - Private key: cm_adm_pvt_key
   - Public key: cm_adm_pvt_key.pub

2. User SSH keys

   Users can create their own personal SSH key pair and configure an additional SSH key under My SSH Public Key. This gets automatically configured in a newly provisioned node, enabling users to use their own keys to access PeopleSoft instances. This key is optional and will be injected into provisioned instances only if it is configured.

Note: SSH keys for Administration will be injected into all provisioned instances, irrespective of User SSH key configuration.

Updating SSH Keys for Administration

Cloud Manager uses SSH keys to connect to Linux instances deployed and managed by Cloud Manager. The public and private SSH key pair used by Cloud Manager to manage instances are located under the path /home/psadm2/psft/data/cloud/ocihome/keys. The public key file is cm_adm_pvt_key.pub and the private key file is cm_adm_pvt_key. From time to time, an organization will want to update or rotate SSH keys. For example:

1. A employee who was a Cloud Administrator or Cloud Manager Administrator has left the organization.
2. As a company policy, it is mandated to update keys periodically.

In such situations, an administrator must ensure to update SSH keys on both Cloud Manager instance as well as on all the managed instances that were created by Cloud Manager. The administrator must create a new pair of Administration keys and update in two locations:

- On Cloud Manager

  Back up the existing keys and replace the keys cm_adm_pvt_key.pub and cm_adm_pvt_key under /home/psadm2/psft/data/cloud/ocihome/keys. The file names should be retained as they are.

- Managed instances

  Using the old private key, SSH into each of the instances provisioned by Cloud Manager as 'opc' user. On the managed instance, update the /home/opc/.ssh/authorized_keys. Remove the previous Administration public key entry and add the new public key.

Updating User SSH Keys

To update any user SSH keys that were injected by Cloud Manager:

1. Generate a new pair of user SSH keys.
2. Log in to the managed instance using either the existing User SSH key or the Cloud Manager's SSH key for Administration.
3. Update the file /home/opc/.ssh/authorized_keys with new key and remove the existing key. Ensure to remove the correct entry.

Generating New SSH Keys

Guidelines for generating new SSH keys:

1. New SSH key pair must be generated using the openssh ssh-keygen utility. If the key pair is generated using any other utility, then it must be converted to openssh format before using them in Cloud Manager.
2. Cloud Manager does not support encrypted ssh key. That is, ssh keys should not be protected by a passphrase.

3. When new **SSH keys for Administration** are generated, ensure to retain the same names for the private and public key files. The permissions of these files should be as shown below.

```
-r-x------. 1 psadm2 oinstall 1675 Jan 21 08:08 cm_adm_pvt_key

-r-x------. 1 psadm2 oinstall  382 Jan 21 08:08 cm_adm_pvt_key.pub
```

4. When new **User SSH Keys** are generated, the file names can be user defined but the permissions must be same as above.

Managing PUM Connections

Use the Manage PUM Connections page (ECL_SA_MANAGEPM_FL) for setting up environments for selective adoption. This page appears only for environments that were deployed using a PeopleSoft Update Image and has a PeopleSoft Client (Windows Client) as part of the environment. This environment can act as a PUM Source environment. You can manage target databases for the PUM Source from this page, which will add or remove specified target databases to the PUM source environments. After adding target databases, administrators can use the PIA URL shown on this page to access PUM Dashboard to define change packages. To create and apply change packages, access Change Assistant that is installed on the PeopleSoft client. To access Change Assistant, use remote desktop (RDP) to Windows Client.

Navigation:

Click the Manage PUM Connections link available on the left panel of the Environment Details page. The Manage PUM Connections page is displayed in the right panel.

This example illustrates the fields and controls on the Manage PUM Connections page.

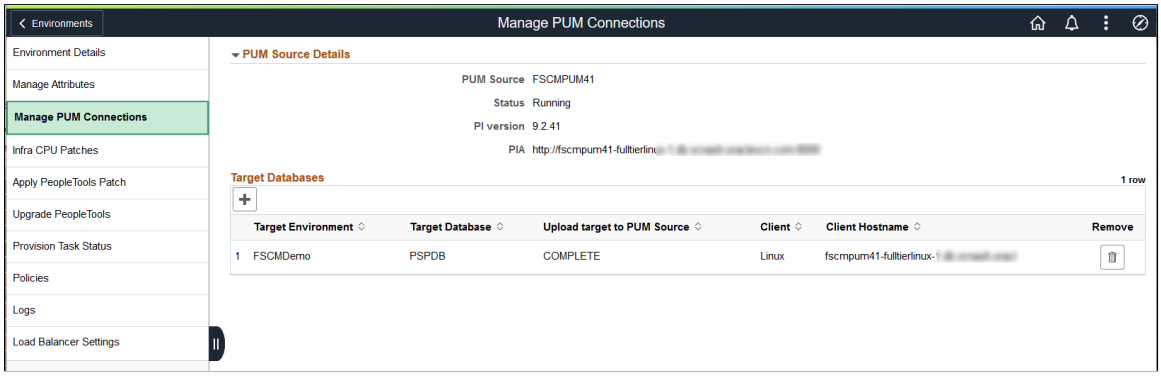

The Upload target to PUM Source status is displayed as either In progress, Complete or Failed.

Adding Target Databases

To add a target database which you want to update, perform the following:

1. Click Add target button available in the Target Databases section.
2. Select the required target environment.

This example illustrates the fields and controls on the Select Target modal window.

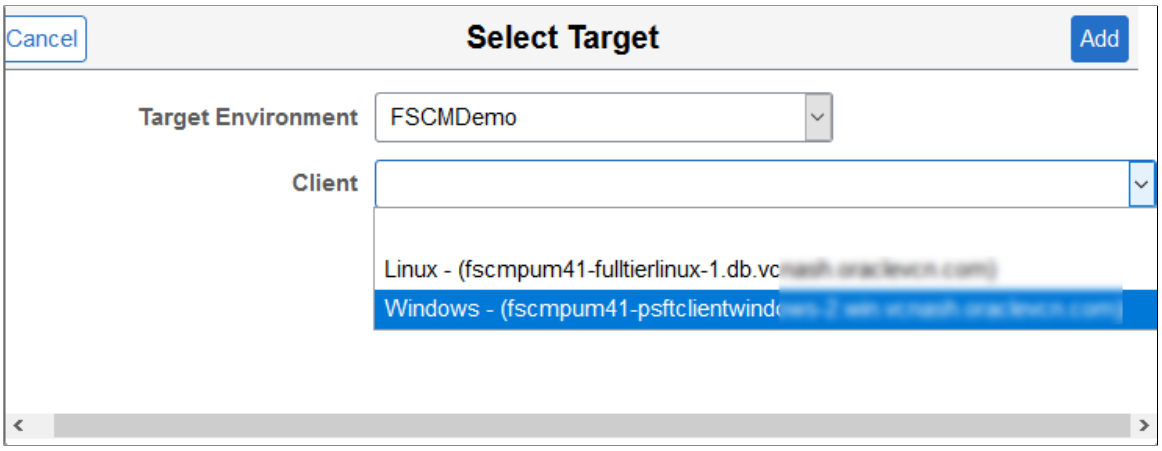

3. Select the client.
4. Click Add. This action starts the ‘Add Target’ and ‘Upload to PUM Source’ functionality. The status is displayed as either In Progress, Complete, or Failed.

Applying Infrastructure CPU Patches

Use the Infrastructure CPU Patches (ECL_ENV_INFRUPD_FL) page to apply critical patch updates to the infrastructure components such as Java, Tuxedo, WebLogic, and Oracle Database client.

The Infrastructure CPU Patches page shows the currently applied third party component version details.

If a valid Infra DPK is available, a message "Latest CPU Available” is displayed.

This example illustrates the fields and controls on the Infrastructure CPU Patches page, showing that patches are available.

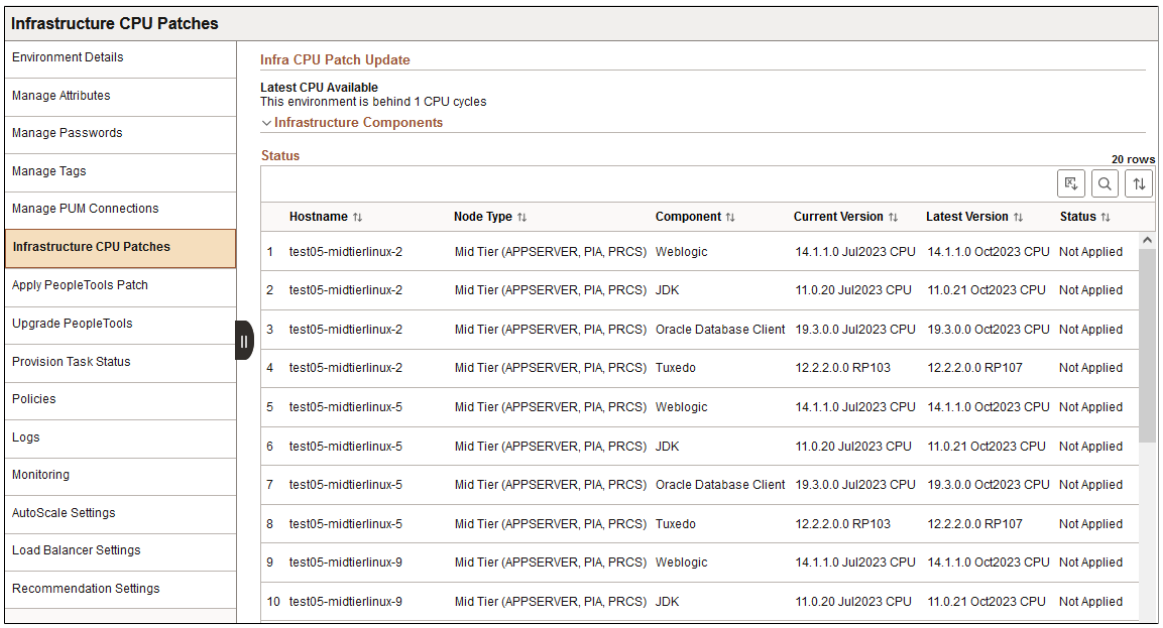

The Infrastructure Components version information lists the available updates.

If there are components that have not been applied, use the Apply button to apply the components.

For information on working with advisories for Infrastructure CPUs, see Configuring and Reviewing Advisories.

Applying PeopleTools Patch

Use the Apply PeopleTools Patch page (ECL_ENV_PTCHUPD_FL) for applying latest PeopleTools patches.

It is recommended to take a backup of the environment prior to applying a PeopleTools patch.

Note: The Apply PeopleTools Patch link is available only if a Windows client node (PeopleSoft Client or Windows middle tier) is associated with the selected environment.

For environments on PeopleTools 8.57 or higher, customizations made to Application Server (psappsrv.cfg) and Process Scheduler Server (psprcs.cfg) are preserved during a PeopleTools Patch. The following Web Server files are also preserved.

- $PS_CFG_HOME/webserv/WEBSERVER01/config/config.xml and any custom folders
- $PS_CFG_HOME/webserv/WEBSERVER01/applications/peoplesoft/PORTAL.war/WEB-INF/psftdocs

Note: User has to manually update other customization files in the web server.

The PeopleTools patch process saves the configuration files during the unprovisioning task. The files are then imported using PSADMIN.

Note: If Auto Scaling is enabled for the environment, the JSON files are not recreated and the data collection/prediction stops after applying the patch. To recreate the files and re-enable monitoring access the Monitoring page and select *Regenerate Logstash configuration*, then disable and re-enable monitoring. See Monitoring Environments.

Navigation:

Click the Apply PeopleTools Patch link available on the left panel of the Environment Details page. The Apply PeopleTools Patch page is displayed in the right panel.

Note: Ensure that the latest PeopleTools patch is already downloaded and available in the Repository.

This example illustrates the fields and controls on the Apply PeopleTools Patch page. You can find definitions for the fields and controls later on this page.

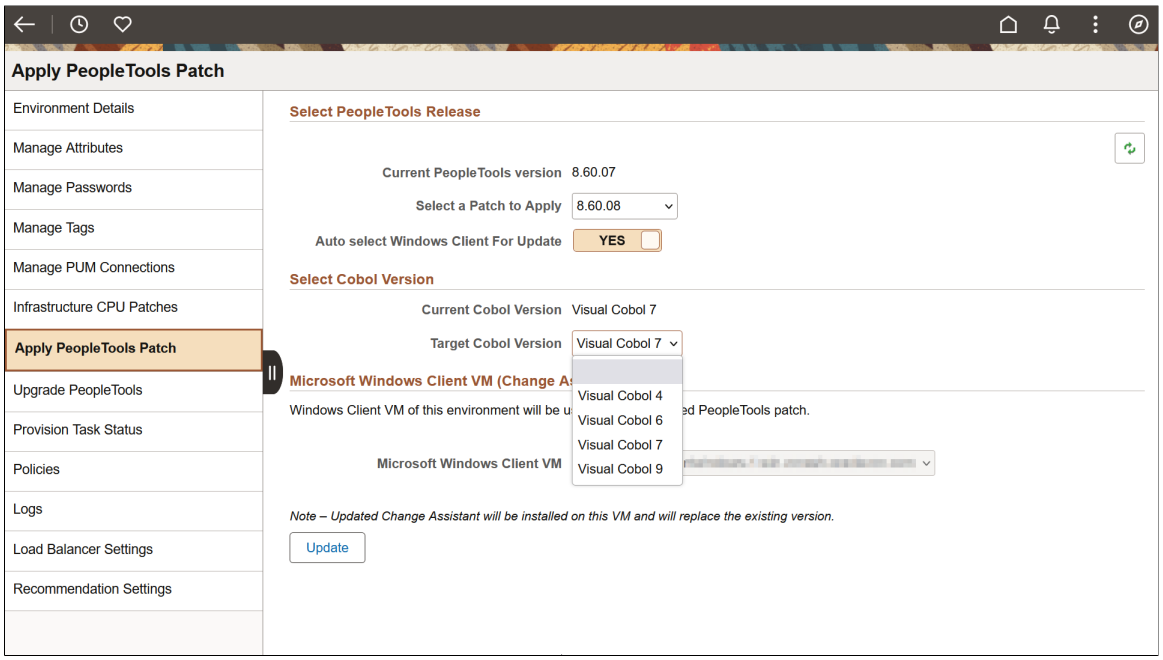

Field or Control	*Description*
Select a Patch to Apply	Select an appropriate PeopleTools patch to be applied on the target environment.
Auto select Windows Client for Update	Select Yes to auto select. Select No to select the Windows Client to use for the upgrade from the Microsoft Windows Client VM drop-down list box.
Target Cobol Version	Select the Visual COBOL version that must be applied on upgrading PeopleTools to the selected version. The default value is the current COBOL version. The available values are Visual Cobol 4, Visual Cobol 6, Visual Cobol 7 and Visual Cobol 9. **Note:** The Select Cobol Version section is enabled only when COBOL is already installed in the selected environment. You must add the Visual COBOL license details on the Cloud Manager Settings page before selecting a COBOL version on this page. In order to be applied, the selected Visual COBOL version must be supported by the PeopleTools version. Visual COBOL versions 6, 7, and 9 are supported for PeopleTools 8.61.

Field or Control	*Description*
Update	Click this button to apply the changes. **Note:** User can select the patch update entry from the grid to see a window which shows the tasks that were run for the patch update process and their real-time status. There is a provision to mark failed tasks as complete so as to complete the patch update use case in failure scenarios. If the task fails, you can see the error details by clicking the Error icon against the failed task on the Job Status Information page.

Depending on the PeopleTools release, the option to select target search provider is available. You can select either Elasticsearch or OpenSearch as the search provider. OpenSearch and OpenSearch Dashboards are only available with PeopleTools 8.60.07 or later, in addition to Elasticsearch and Kibana. Support is also available for PeopleTools 8.59.21 patch or later. PeopleTools 8.61 supports only OpenSearch and OpenSearch Dashboards.

This example illustrates the fields and controls on the Apply PeopleTools Patch page when OpenSearch option is enabled in the target PeopleTools release.

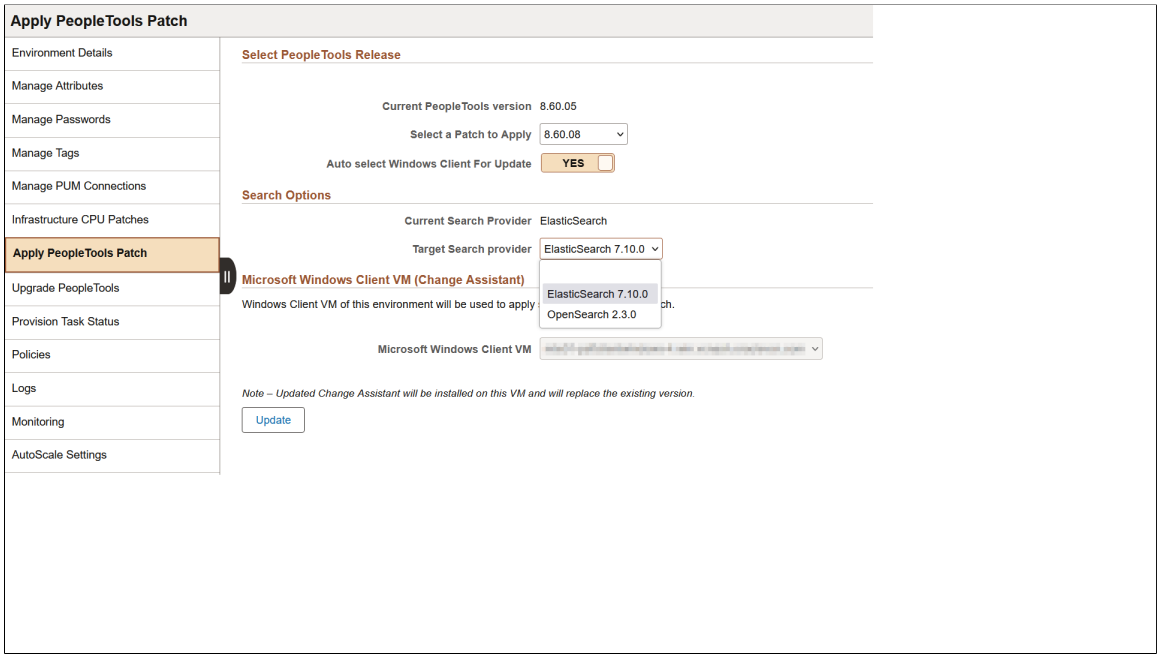

Upgrading PeopleTools

Use Upgrade PeopleTools page (ECL_ENV_UPGD_FL) to upgrade PeopleTools version (major version changes).

It is recommended to take a backup of the environment prior to performing a PeopleTools upgrade.

For environments on PeopleTools 8.57 or higher, customizations made to Application Server (psappsrv.cfg) and Process Scheduler Server (psprcs.cfg) are preserved during a PeopleTools Upgrade. Web server configurations are not preserved in a PeopleTools upgrade. The web server will be redeployed.

The PeopleTools patch process saves the configuration files during the unprovisioning task. The files are then imported using PSADMIN.

Note: The Upgrade PeopleTools link is available only if a Windows client node (PeopleSoft Client or Windows middle tier) is associated with the selected environment.

Note: If Auto Scaling is enabled for the environment, the JSON files are not recreated and the data collection/prediction stops after applying the patch. To recreate the files and re-enable monitoring access the Monitoring page and select *Regenerate Logstash configuration*, then disable and re-enable monitoring. See Monitoring Environments.

Navigation:

Click the Upgrade PeopleTools link available on the left panel of the Environment Details page. The Upgrade PeopleTools page is displayed in the right panel.

This example illustrates fields and controls on the Upgrade PeopleTools page.

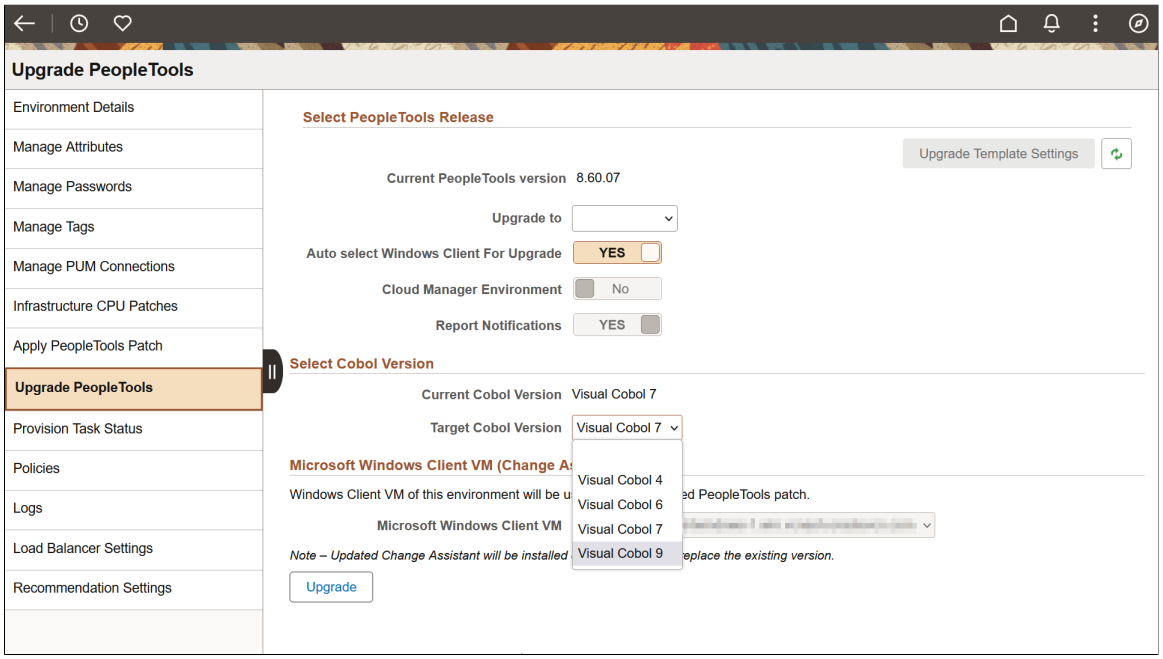

Field or Control	***Description***
Upgrade to	Select the major PeopleTools version.
Auto select Windows Client For Upgrade	Select Yes to auto select. Select No to select the Windows Client to use for the upgrade from the Microsoft Windows Client VM drop-down list box.

Field or Control	*Description*
Report Notifications	Select to receive notifications when reports are available. This applies to upgrades to PeopleTools 8.60.x. In order to enable notification, the Cloud Manager administrator must complete the prerequisites and update the Notifications section on the Infrastructure Settings page. See Infrastructure Settings Page.
Target Cobol Version	Select the Visual COBOL version that must be applied on upgrading PeopleTools to the selected version. The default value is the current COBOL version. The available values are Visual Cobol 4, Visual Cobol 6, Visual Cobol 7, and Visual Cobol 9. **Note:** The Select Cobol Version section is enabled only when COBOL is already installed in the selected environment. You must add the Visual COBOL license details on the Cloud Manager Settings page before selecting a COBOL version on this page. In order to be applied, the selected Visual COBOL version must be supported by the PeopleTools version. Visual COBOL versions 6, 7, and 9 are supported for PeopleTools 8.61.
Upgrade	Click this button to apply the changes. Before doing Upgrade, user must ensure to take a backup of the environment.
Upgrade Template Settings	Select to configure which reports to review. This applies to upgrades to PeopleTools 8.60.x. If the task fails, you can see the error details by clicking the Error icon against the failed task on the Job Status Information page. See Upgrade Template Settings Page

Depending on the PeopleTools release, the option to select target search provider is available. You can select either Elasticsearch or OpenSearch as the search provider. OpenSearch and OpenSearch Dashboards are only available with PeopleTools 8.60.07 or later, in addition to Elasticsearch and Kibana. Support is also available for PeopleTools 8.59.21 patch or later. PeopleTools 8.61 supports only OpenSearch and OpenSearch Dashboards.

This example illustrates the fields and controls on the Upgrade PeopleTools page when OpenSearch option is enabled in the target PeopleTools release.

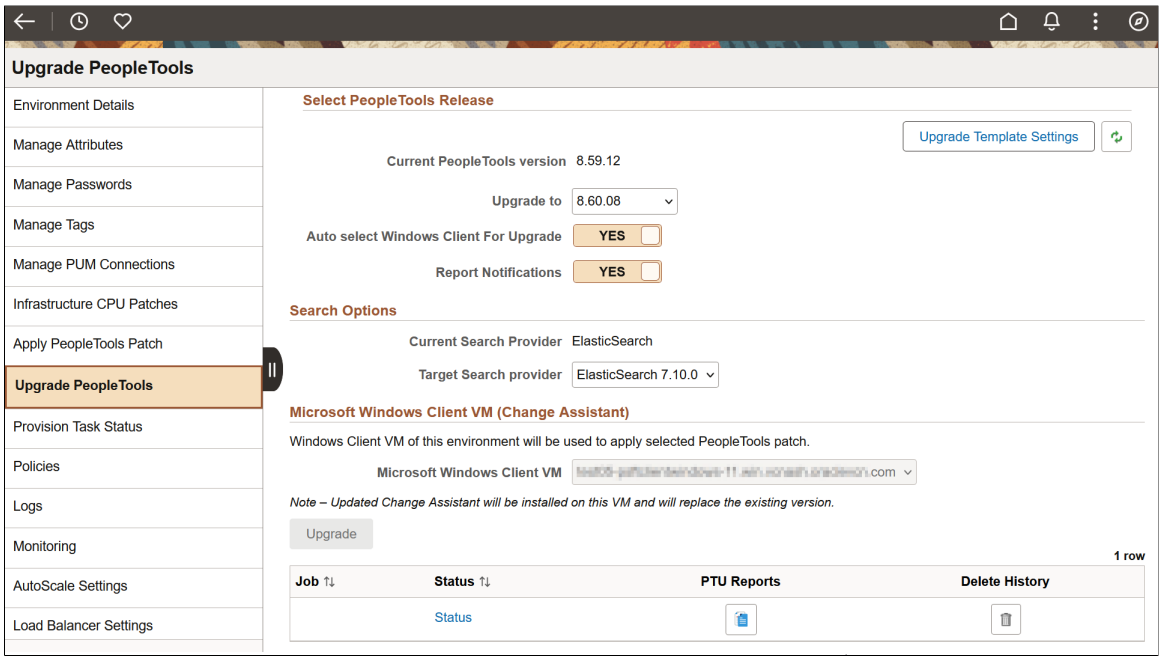

Upgrade Template Settings Page

Use the Upgrade Template Settings page to configure the upgrade template and select the reports to include. Manual stops can be added to review the selected reports.

This example illustrates the fields and controls on the Upgrade Template Settings page. You can find definitions for the fields and controls later on this page.

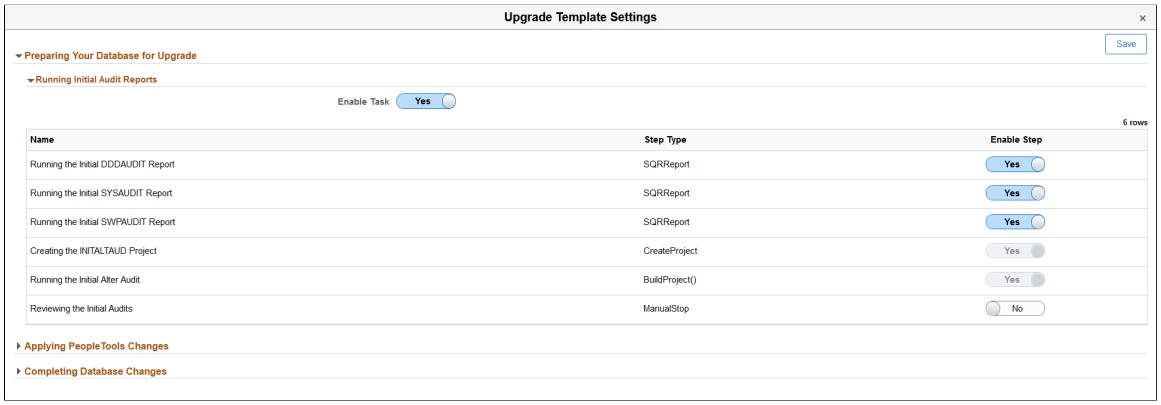

The Upgrade Settings page contains 3 chapters that can be expanded to enable tasks and select the steps to run.

Expand each chapter and select the steps to run. When you select Reviewing steps, a ManualStop is added to the template. When the manual stop step is encountered, the upgrade job will pause for you to review the reports.

Note: Reviewing steps should only be selected if reports are also selected.

Field or Control	*Description*
Enable Task	Select Yes to enable at the task level. When selected, you can then select which steps to run. When you select No, none of the steps for this task will run.
Name	The name of the step.
Step Type	The step type.
Enable Step	Select Yes for each step you want to run.

After enabling or disabling task and steps, save the page.

Job Status Information Page

The Job Status Information page displays the current job status.

After starting the PeopleTools Upgrade, select Upgrade PeopleTools in the left panel to view the current status.

This example illustrates the Upgrade PeopleTools page showing the job step and status.

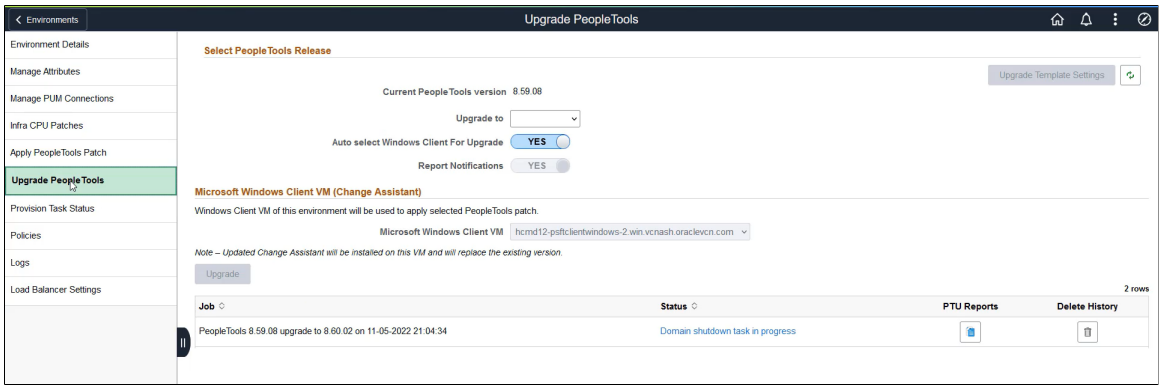

Click on the link in the Status column to view the Job Status Information.

This example illustrates the fields and controls on the Job Status Information Modal Window where the Database Upgrade is in progress.

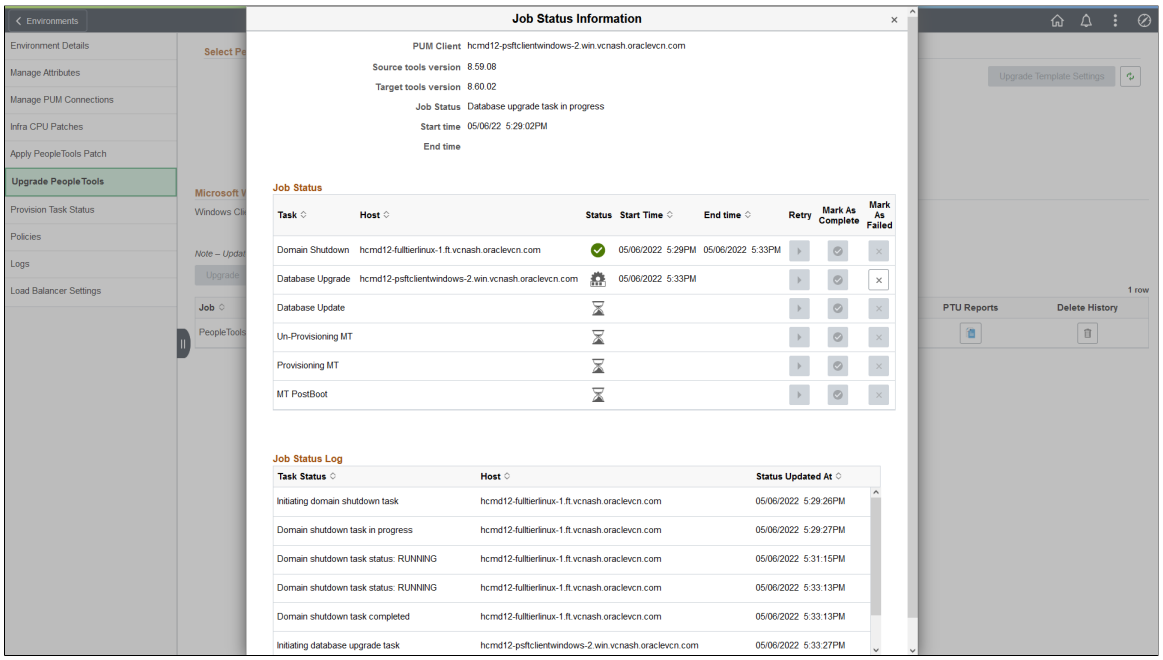

You can view upgrade process details such as jobs completed successfully, jobs which are in pending status, and failed jobs. If a task fails, you can view the details of the failure by clicking the status icon corresponding to the failed task.

Viewing Compare Reports

Use the PTU Reports icon to view compare reports.

When a manual stop is encountered in the template, the PeopleTools Upgrade will be in a paused state.

This example illustrates the PeopleTools Upgrade in a paused status.

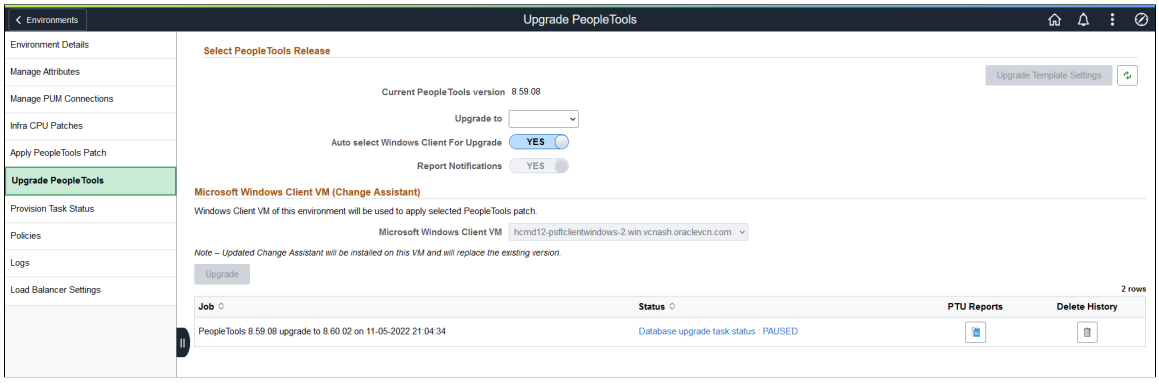

Field or Control	***Description***
Status	Select the link in the Status field to open the Job Status Information page.

Field or Control	*Description*
PTU Reports	Select the PTU Reports icon to view the reports.
Delete History	Select the Delete History icon to delete an upgrade job.

Select the PTU Reports icon to view the reports.

This example illustrates the Compare Reports page.

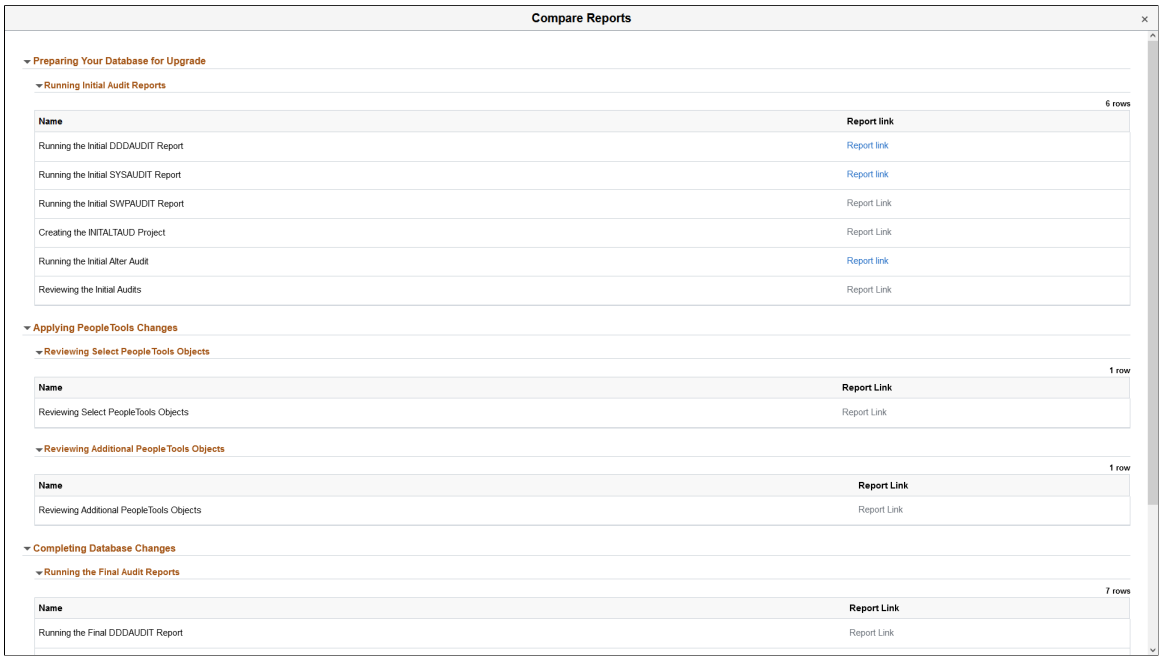

The Report link will be active when the report is available. Click the Report link to view the report.

You can open the report in a browser or save the file.

This example illustrates how to open the report.

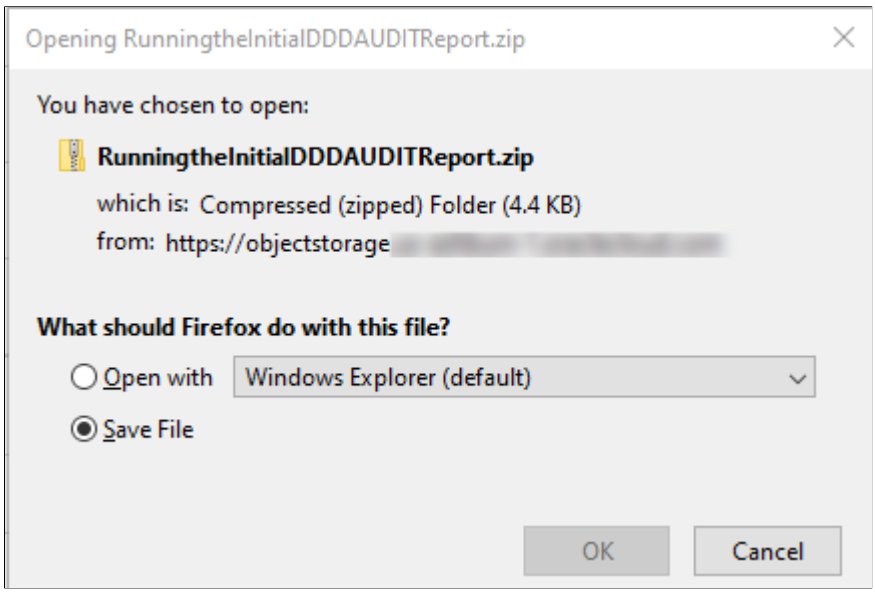

Compare Reports can be viewed in the Compare Report Viewer.

1. On the Compare Reports page, click the link for either Reviewing Select PeopleTools Objects or Reviewing Additional PeopleTools Objects.
2. Select to Open the file.
3. Select pscmviewer.exe.
4. Click Extract All.
5. Select a destination for extracting the files and click Extract.
6. Select pscmviewer.exe.

This example illustrates the PeopleSoft Compare Report Viewer.

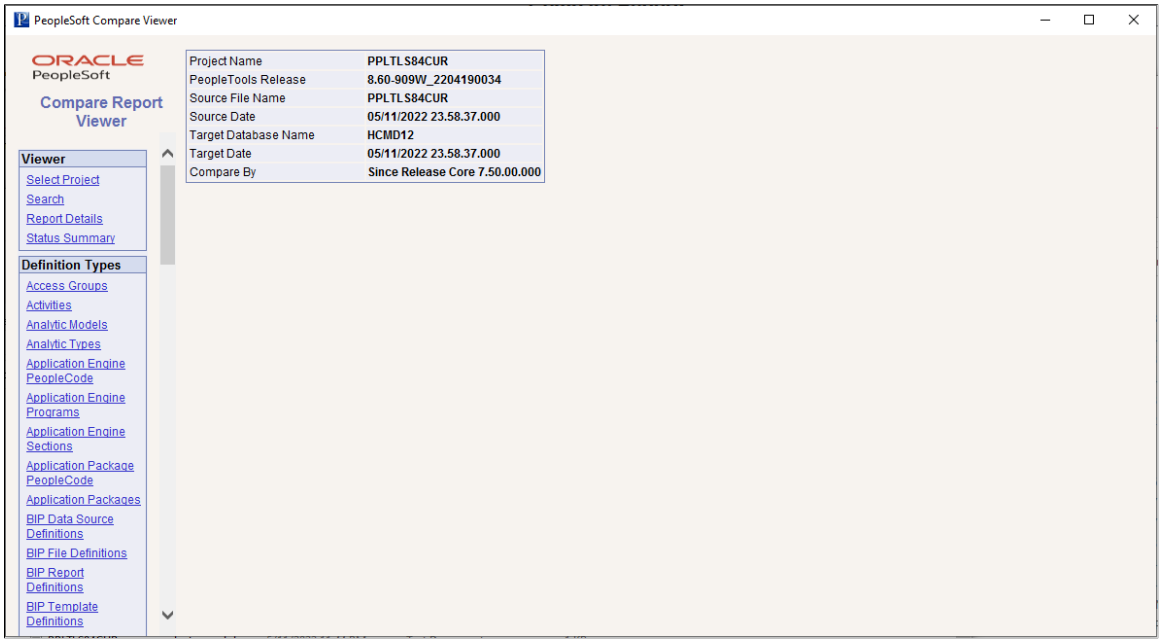

For more information on using the Compare Report Viewer, see PeopleTools Lifecycle Management, "Working with Browser Reports".

Viewing Provision Task Status

Use the Provision Task Status page to view the current provisioning status for the environment.

Provision Task Status is available for the following lifecycle operations:

- Provisioning a new environment
- Shifting an environment
- Importing an environment
- Cloning an environment
- Refreshing an environment
- Adding nodes
- Removing nodes
- Applying PRPs on an existing environment
- Applying CPU Patches on an environment
- Starting or Stopping environments or nodes

Navigation:

1. Click the Related Actions button corresponding to the environment.

2. Select Details.
3. Select the Provision Task Status link available on the left panel of the Environment Details page.

This example illustrates the Provision Task Status page for a scale up operation.

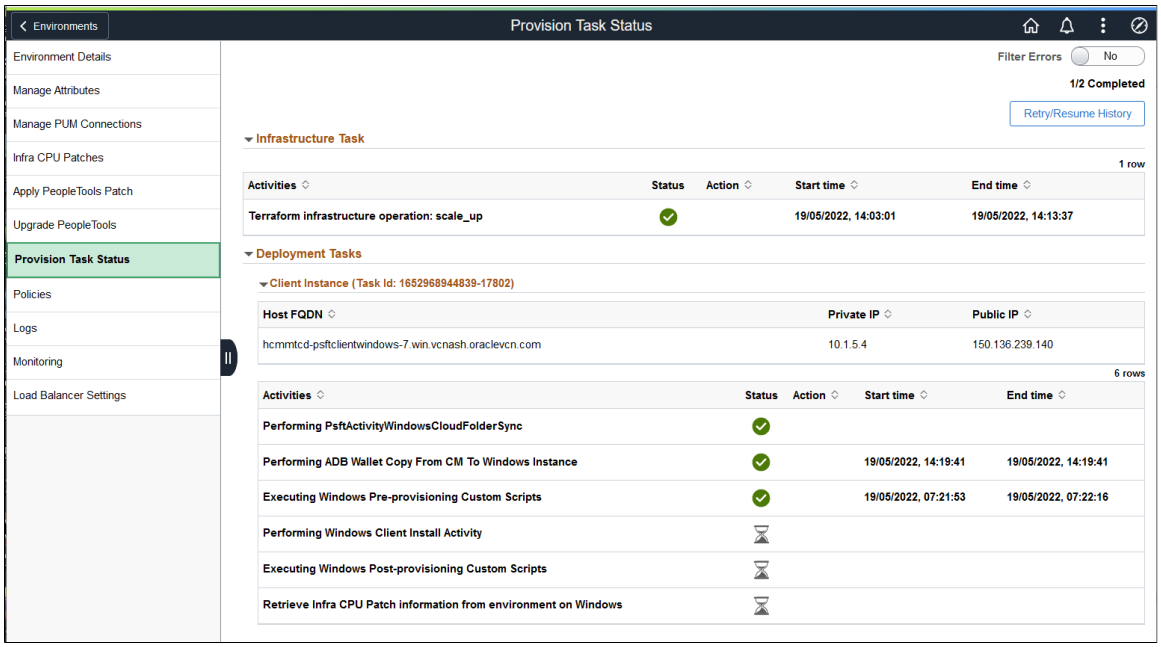

The Provision Task Status page will display the status for the most recent lifecycle operation on the environment. If a failure occurs, the status for the task status will display the Failed icon. See Retrying and Resuming Provisioning.

Retrying and Resuming Provisioning

Cloud Manager provides the ability to resume provisioning when a recoverable failure occurs. This applies to provisioning a new environment, shifting an environment or adding nodes using Manage Node action.

To view the failed task:

1. Select the Environments tile.
2. Select Details for the environment.
3. Select Provision Task Status.
4. Use the Filter Error toggle to show the failing step, without scrolling through all steps to locate the failing step.

This example illustrates the Provision Task Status page filtered for errors.

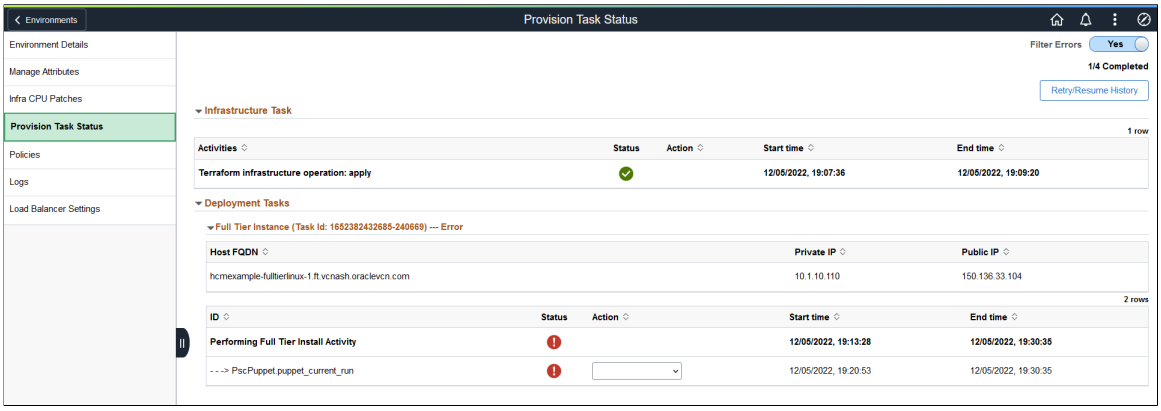

5. Click on the Failed icon to view the error. The Error Details window has two sections:

   - The *Error code details* section shows the error code and message indicating the failure.
   - The *Errors fetched from logs* section shows the log directory for the failure as well as errors fetched from log files.

   This example illustrates the fields and controls on the Error Details page.

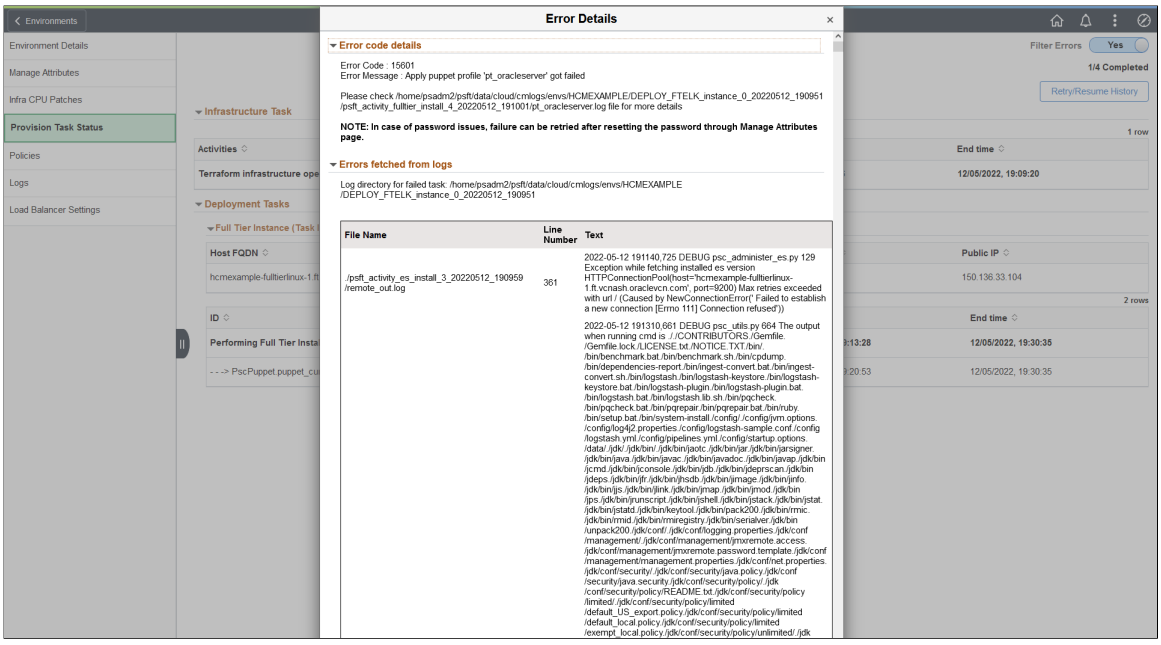

6. After correcting the error, return to the Provision Task Status.
7. Select the action from the drop down list.

Note: If any task fails, the option to retry is enabled for all the steps. You can retry any activity or action step present in the same task, including those that are successful. The option to ‘Mark As Complete’ is only available for failed activities and action steps.
Infrastructure Task (Terraform operations) as well as actions in Deployment Tasks can be retried.

This example illustrates the fields and controls on the Provision Task Status Retry/Resume.

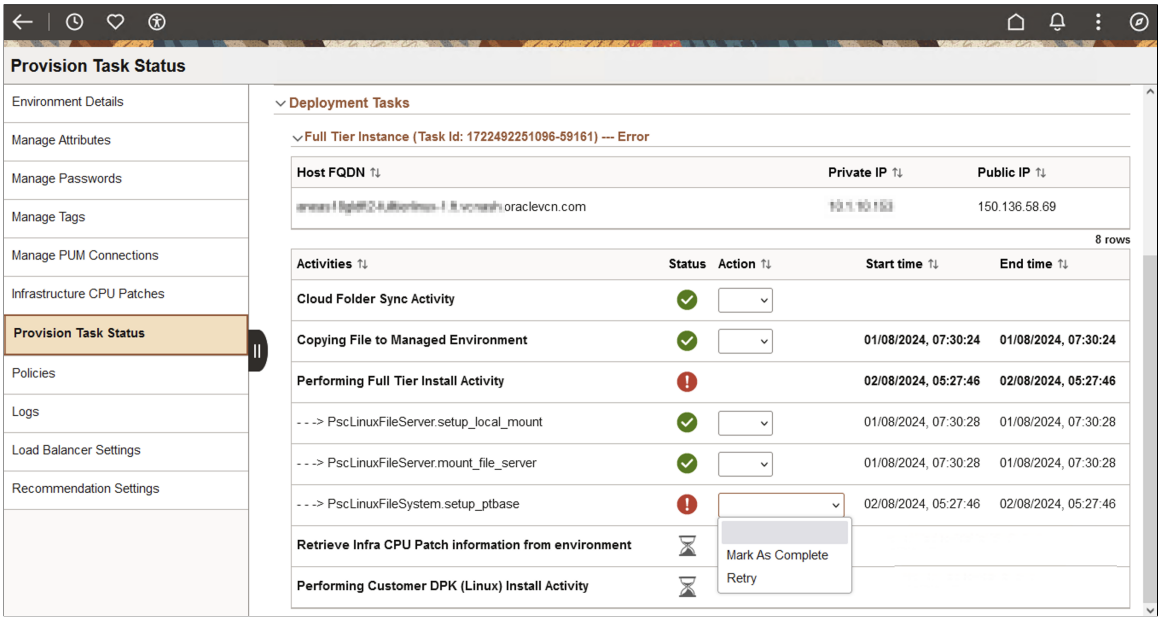

Field or Control	***Description***
Mark as Complete	If you manually corrected the error, select this action to mark the failed task as complete and continue with the next task. ***Warning!*** Manually corrected errors are not validated by Cloud Manager.
Retry	Select this action to retry a successful or failed task. If any attributes for the environment are updated (through Managed Attributes page), then the retry will be done using the updated attributes.

8. On selecting either ‘Retry’ or ‘Mark As Complete’ option for a task, the corresponding action is triggered and the menus are disabled further. The provisioning process continues.

9. Use the Retry/Resume History button to review the history of tasks for which you had triggered a retry action.

This example illustrates the fields and controls on the Retry/Resume History page where provisioning is complete.

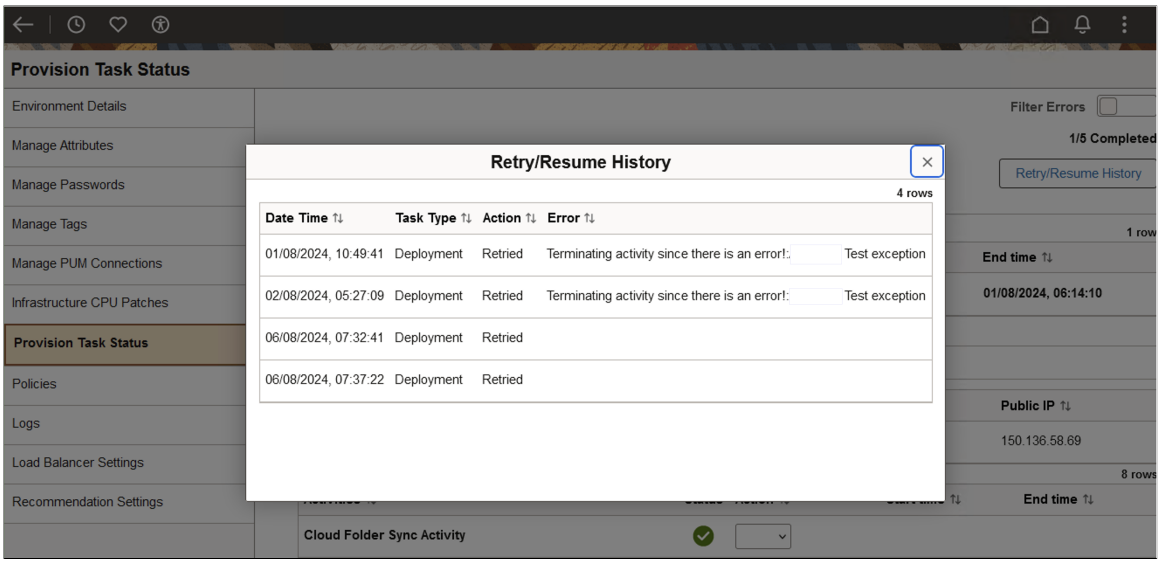

You can also add manual stop steps to Customer DPK activity to manually review the processing of steps during provisioning. See Manually Reviewing Steps During Processing.

The Manual Review step is associated only with Mark As Complete action and not Retry action. For all the other activities or action steps in the same task, only the Retry option is available. When you select the Retry action, all the activities or action steps after that step are run again. When the review process is completed, you can mark the manual review step as complete for resuming the processing of the activity.

This example illustrates the manual review step on the Provision Task Status page.

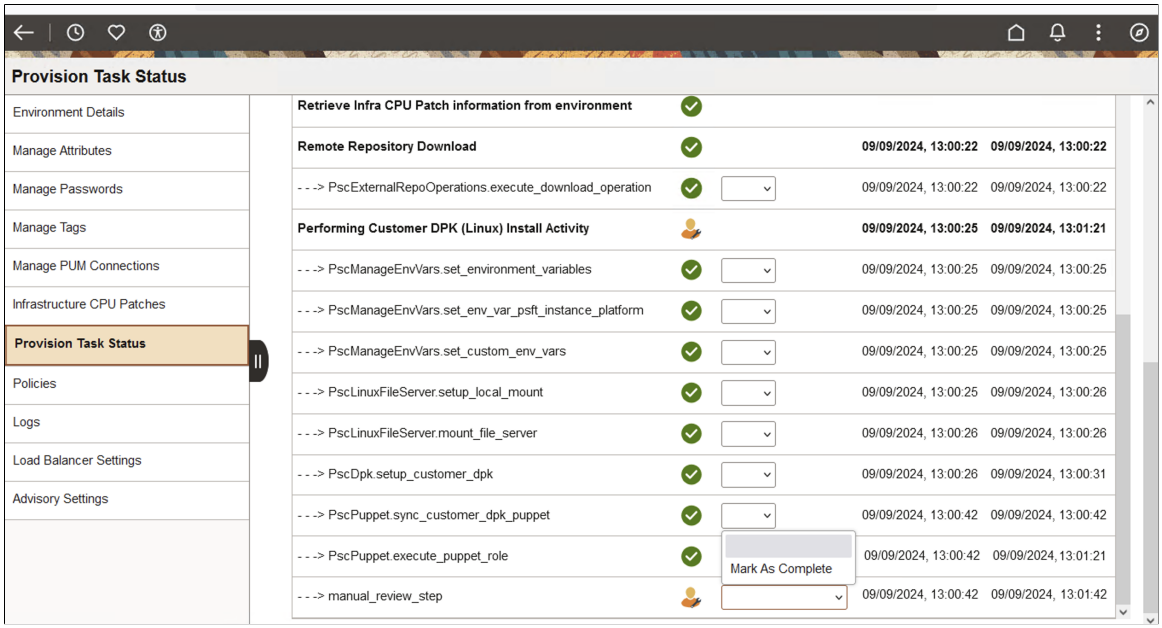

The information is then updated in Retry History.

Associating Policies with Environment

Use the Policies page (ECL_POLICY_ENVS) to associate policies with the environment name.

Navigation:

Click the Policies link available on the left panel of the Environment Details page.

This example illustrates the fields and controls on the Policies page. You can find definitions for the fields and controls later on this page.

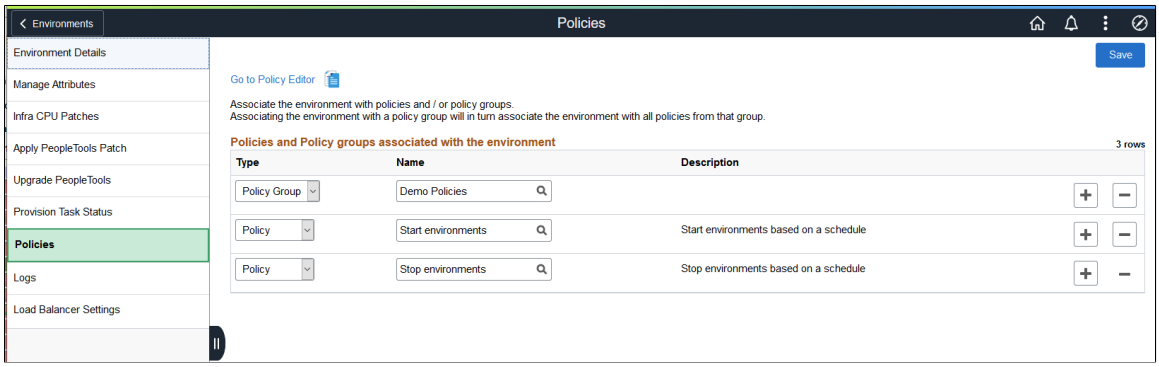

Use the Go to Policy Editor link to view the existing policies. See Using Policy Editor.

The policies and policy groups associated with the environment are shown.

You can add additional policies or policy groups or remove existing ones for the environment. When a policy is added, the environment name will be appended to the Environment Names condition property or action parameter for the Policy. Likewise, if the policy is deleted, it will be removed from the Environment Names condition property or action parameter for the Policy.

When you add a policy group, the environment will be associated with all the policies in that group.

Note: Policies associated with OCI tags can only be added using the Manage Attributes page to add or remove tags. Once the tags are added, the corresponding policies for the tags will be displayed on this page.

See Setting Policy Conditions and Actions for Environment Policy Object.

Managing Passwords

Use the Manage Password page (ECL_ENV_RESET_FL) to directly initiate a password reset of managed environment passwords or to update passwords when you modify passwords outside Cloud Manager.

The Manage Passwords page must be updated when you modify the environment passwords directly on the instance. For example, if you modify the OPRID Password on the instance directly, you must update and save the password for Cloud Manager to store in its database. Otherwise, the password stored in Cloud Manager will be stale and any operation that is dependent on this OPRID Password will fail. You can also initiate password reset directly from this page.

Click the Update Metadata button to fetch the current status of the environment.

Navigation:

Click the Manage Passwords link available on the left panel of the Environment Details page. The Manage Passwords page is displayed on the right panel.

This example illustrates the fields and controls on the Manage Passwords page.

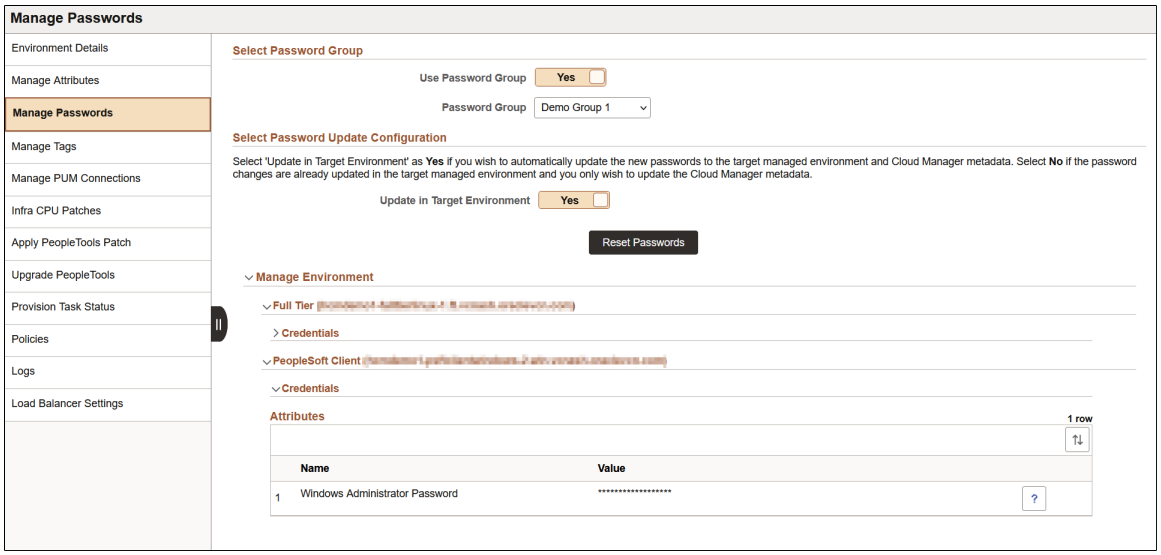

Field or Control	***Description***
Use Password Group	Select Yes to use password groups that are already created as secret OCIDs in the secure password storage mechanism called OCI Vault, so that you need not manually enter the passwords.
Password Group	Select the required password group from the drop down.
Update in Target Environment	Select Yes to automatically update the new passwords to the target managed environment and Cloud Manager metadata on clicking Reset Passwords button. When you select No, the button name changes to Update Metadata.
Reset Passwords or **Update Metadata**	Click the Reset Passwords button to automatically update the passwords to the target environment. Click the Update Metadata button if the password changes are already updated in the target environment, and you want to update only the Cloud Manager metadata.

When the environment uses a password group, the password fields are read-only. When you do not select a password group, you can edit the required passwords by expanding each node in the environment and these changes can be applied by clicking the Reset Passwords or Update Metadata button.

Resetting Passwords

You can reset passwords from the Manage Passwords page. The following passwords are supported for reset:

- Database Operator Password
- Database Connect Password
- Database Access Password
- Weblogic Administrator Password
- Gateway Administrator Password
- Windows Administrator Password
- Search Administrator Password
- Search Proxy Password

Note: You must perform the update manually on the environment and then refresh the metadata from the Manage Passwords page for Database Administrator Password, TDE Wallet Password, and Web Profile Password.

Note: To reset passwords on shared search nodes see Sharing a Search Cluster Across Multiple Environments.

To reset passwords when a password group is used:

1. Select a password group from the Password Group drop down. The custom attribute passwords are pre-filled in the Credentials section for each node, according to the selected password group.

This example illustrates the fields and controls on the Manage Passwords page when a password group is selected.

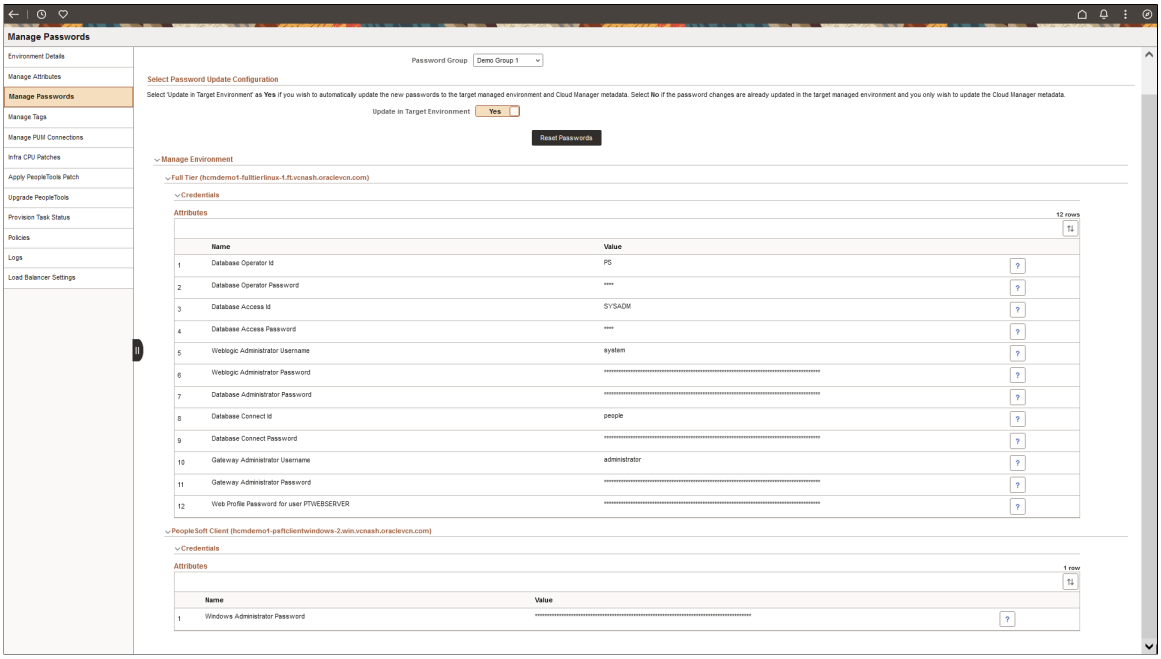

2. Modify the secret content from OCI. See Updating Secret Content.
3. Select Yes in the Update in Target Environment field if you want to update the new passwords in the target environment as well as Cloud Manager metadata. Select No if you only wish to update the Cloud Manager metadata.
4. Click the Reset Passwords or Update Metadata button. A confirmation message appears, which lists the passwords that were changed in OCI.
5. Click OK. An additional section called Messages appears. You can view the status of the task on Provision Task Status page.

   This example illustrates the Messages section on the Manage Passwords page.

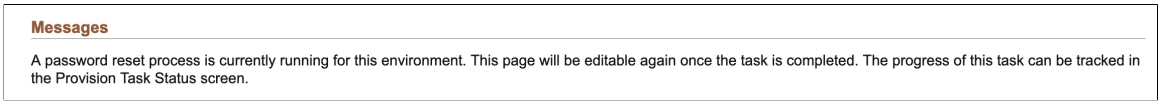

To reset passwords when a password group is not used:

1. Modify the required custom attribute passwords in the Credentials section for each node.
2. Select Yes in the Update in Target Environment field if you want to update the new passwords in the target environment as well as Cloud Manager metadata. Select No if you only wish to update the Cloud Manager metadata.
3. Click the Reset Passwords or Update Metadata button. A confirmation message appears.

This example illustrates the fields and controls on the Manage Passwords page when password reset task is initiated.

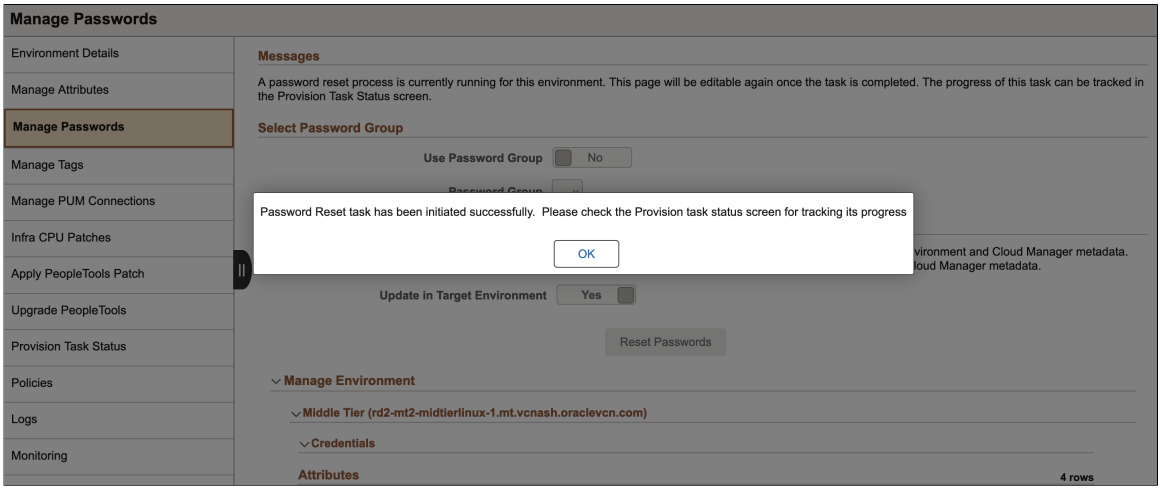

4. Click OK. An additional section called Messages appears. You can view the status of the task on Provision Task Status page.

   This example illustrates the fields and controls on the Provision Task Status page, showing the status of password reset task.

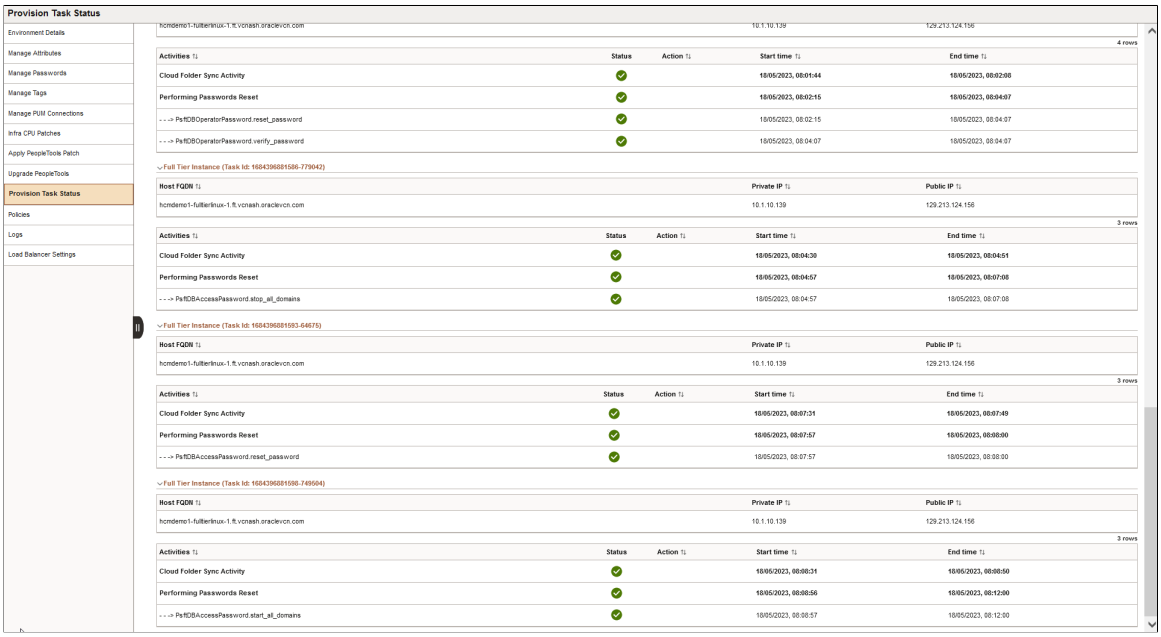

When the process of resetting passwords is in progress, the corresponding environment displays the status as 'Resetting Passwords'.

Restoring PeopleSoft Insights and Search Functionality After Modifying Passwords

If you modify the Database Operator Password in the Database Tier, or either the Proxy Password or Administrator Password in the Search Stack, the PeopleSoft Insights displayed in Remote Worker

Dashboard tile in the managed environment may be displayed incorrectly. You may also experience difficulty in using the search functionality in the managed environment. To resolve this issue:

1. Navigate to **Menu**>**PeopleTools**>**Search Framework**>**Search Admin Activity Guide**>**Search Instance**. The Search Administrator page appears.
2. Click Search.
3. Click **Update Deploy Definitions** in the Call Back Properties section. This restores the PeopleSoft Insights and search functionality in the managed environment.

Managing Tags

Use the Manage Tags page (ECL_ENV_TAGS_FL) to update managed environment tags, if user modified any parameter outside Cloud Manager.

Navigation:

Click the Manage Tags link available on the left panel of the Environment Details page. The Manage Tags page is displayed on the right panel.

Tagging

The latest values of Tag Namespace, Tag Key, and Tag Value are displayed in the Tagging section for each node in the environment. See Configuring Tagging.

You can add, delete or update tags for each node in the environment. Click the Update button to save the changes made to tags.

Tags can be used with policies to associate a group of multiple environments to a specific policy.

This example illustrates the fields and controls in the Tagging section on Manage Tags page.

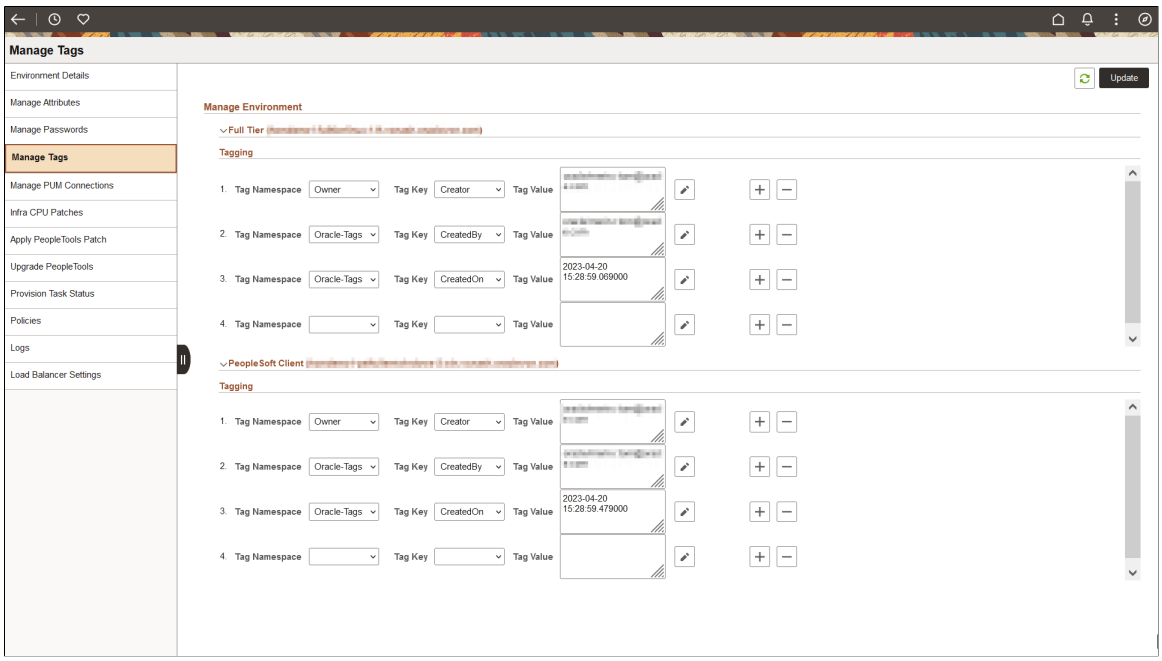

Click the Refresh button to fetch the current status of the environment node.

Policies can be associated with multiple environments using tags. See Setting Policy Conditions and Actions for Environment Policy Object.

Configuring Sparse Hierarchy Details

Use the Sparse Hierarchy Details page (ECL_EXA_SPARSE_FL) to create test masters and manage their sparse clones using Cloud Manager.

Navigation:

Click the Sparse Hierarchy Details link available on the left panel of the Environment Details page. The Sparse Hierarchy Details page is displayed on the right panel.

Cloud Manager enables you to create and manage test masters and create environments with sparse cloned databases for Exadata Database systems. Sparse clones can only be performed on Test masters created in Exadata snapshot enabled clusters. Exadata Sparse Clone is a native feature of Exadata that enables the creation of thinly provisioned databases for non-production purposes like development and testing. All read operations to a sparse cloned database points to the parent database's disk block, until a disk block is created in the sparse cloned database itself with a write operation. Sparse clones are quicker to create than full copies of databases and need much less storage because they track only the changes in data.

This example illustrates the high level architecture of sparse cloning process.

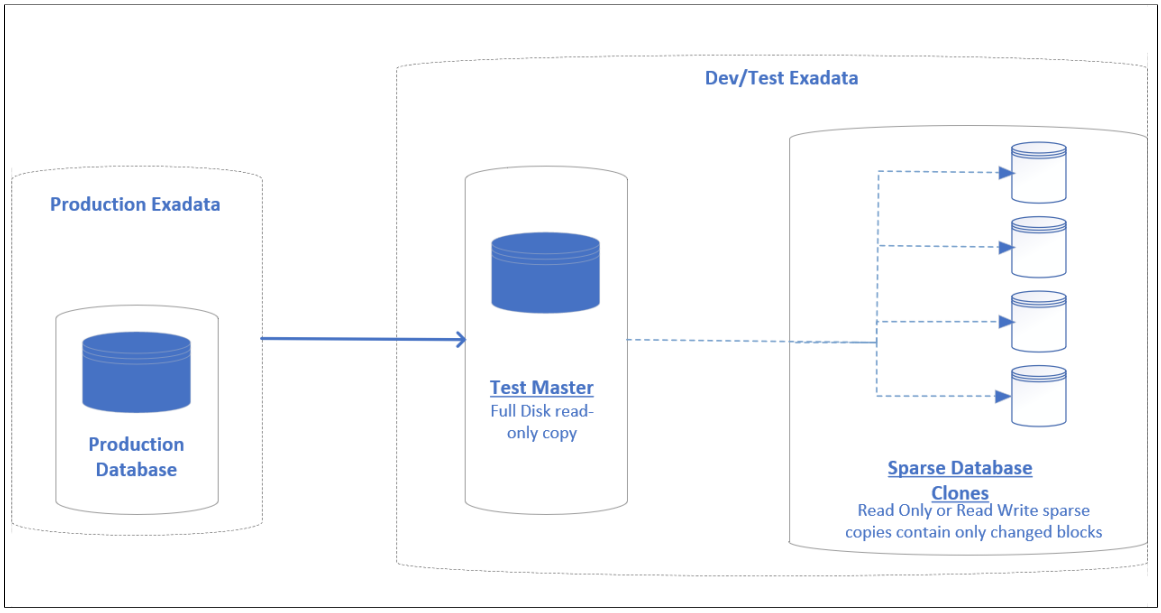

Before creating sparse clones, you must create a full disk clone of the environment from Exadata production database. Such a full read-only copy of the database is called a test master. Although the creation of a test master takes time, you can quickly create multiple space efficient development or test databases using sparse clones.

To create a test master:

1. Click the Create Test Master button. The Create New Test Master page appears.

   This example illustrates the Create Test Master button on Sparse Hierarchy Details page.

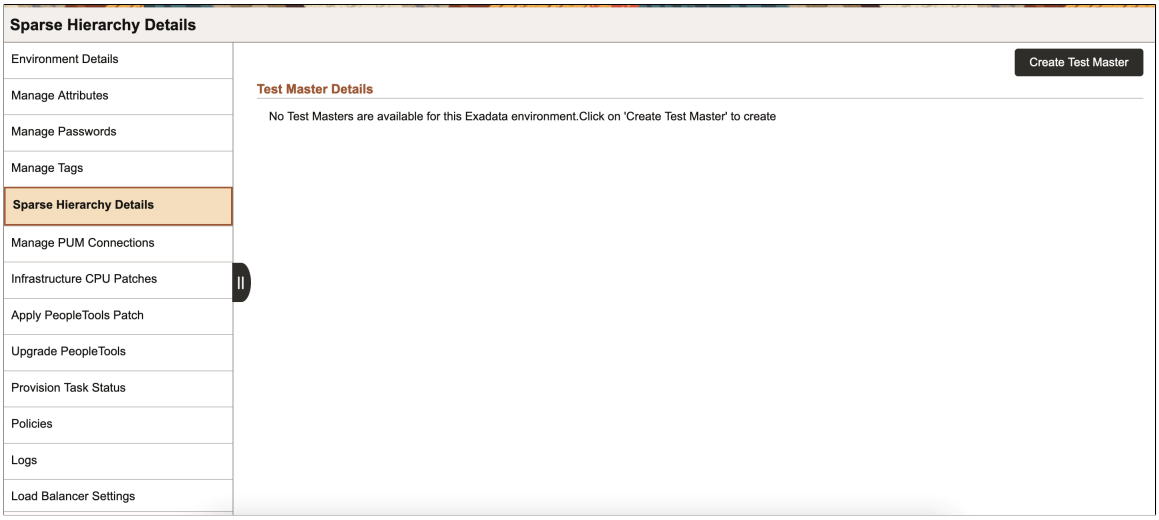

2. Make the necessary changes to the Network Configuration and General Settings for the new test master.

   This is an example of the page used to create a new test master. You can find definitions for the fields and controls later on this page.

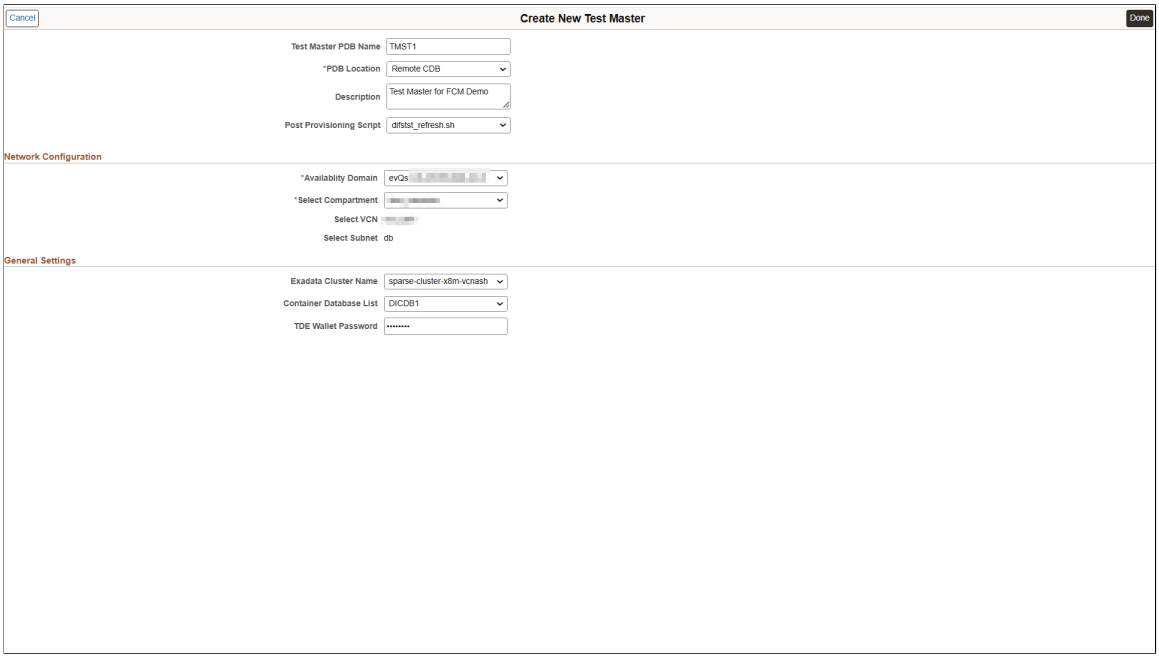

Field or Control	***Description***
Test Master PDB Name	Select the name for test master PDB, which is a clone of the source environment.

Field or Control	*Description*
PDB Location	Select the location where test master PDB is to be created. The available options are Local CDB, New CDB, Remote CDB.
Post Provisioning Script	Select a post provisioning script for the Test Master.
Select Compartment	Select the compartment containing the VCN you want to use.
Select VCN	Select the VCN for the compute node.
Select Subnet	Select the subnet within VCN for the compute node.
Shape Name	Select Exadata as the shape name.
Container Database List/ Container Database Name	Select a CDB. The list includes the CDBs in the Exadata cluster that hosts the source CDB. Container Database List is available only for Remote CDB and Container Database Name is available only for New CDB and Local CDB.
TDE Wallet Password	Specify a TDE Wallet password for the existing target CDB. This option is available only for Remote CDB.

3. Click Done to save the changes.

The Sparse Hierarchy Details page displays the test masters associated with the source environment. Sparse clones created using a specific test master are found within the collapsible panel of that test master.

This example illustrates the fields and controls on Sparse Hierarchy Details page.

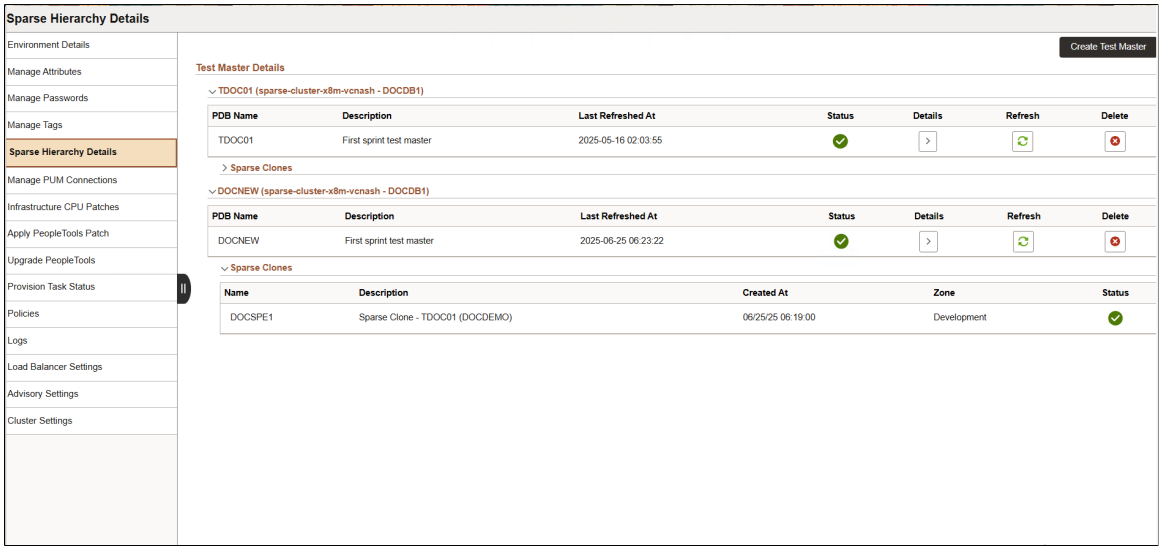

To view the details of a specific test master, click the Details button. The Test Master Details page appears.

This example illustrates the fields and controls on Test Master Details page.

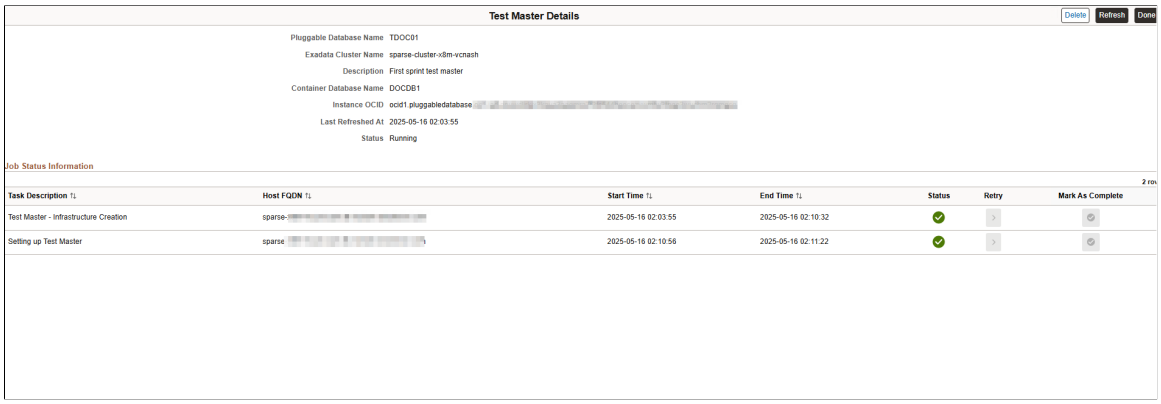

The Job Status Information section displays the latest tasks run in the lifecycle of the test master. When a task fails, the Status column displays a failure icon against the failed task. You can retry the failed task after debugging and fixing the error in the logs. Alternatively, it is possible to manually complete the step and mark the task as completed.

Click Refresh to refresh the test master with the source environment and click Delete to remove the test master. Before attempting either of these actions, ensure that all the child sparse clones of the test master are deleted.

To create sparse clones using the Clone Environment action on the Clone page, see Cloning an Environment With Database Running on Exadata and Other Nodes on Compute.

Refreshing Sparse Clones

The refresh functionality in Sparse Hierarchy Details page enables you to get the latest version of a test master.

You can either refresh all the child sparse environments or exclude some sparse environments when you refresh the test master.

This example illustrates the scenario when all the child sparse environments are refreshed.

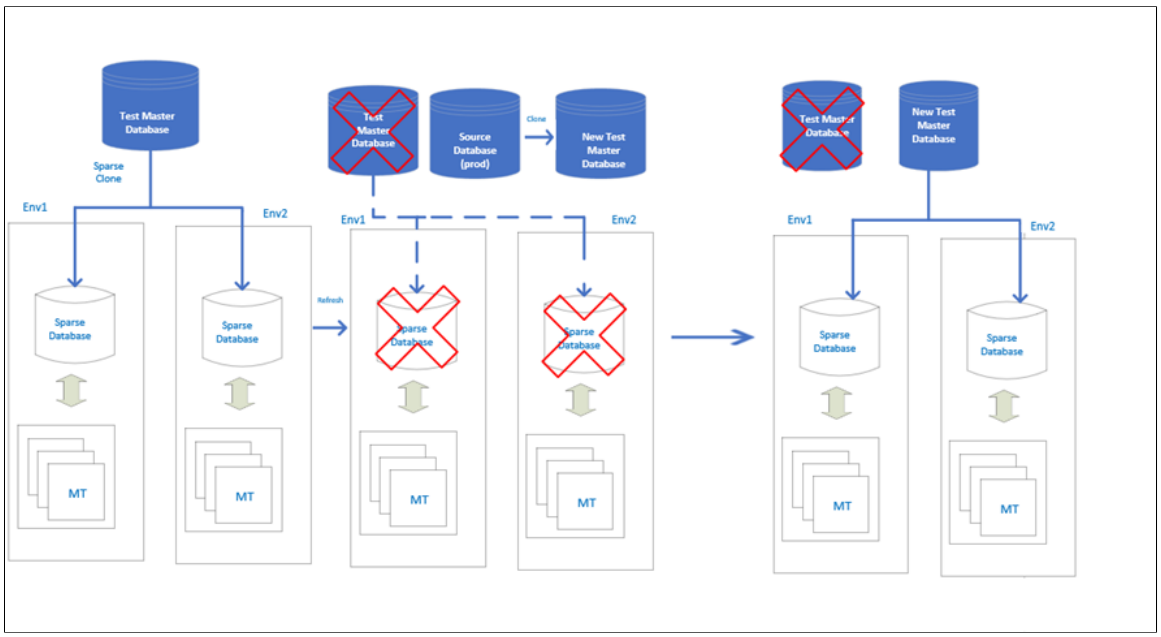

The following processes happen when you refresh the test master along with all its associated sparse environments:

1. Cloud Manager deletes all the existing sparse cloned databases and initiates the refresh of test master.
2. The existing test master is deleted.
3. Cloud Manager creates a new test master from the source database and then creates sparse clones from the new test master.
4. The middle tiers of sparse environment are reconfigured with the newly created sparse clones.

This example illustrates the scenario when certain child sparse environments are excluded from being refreshed.

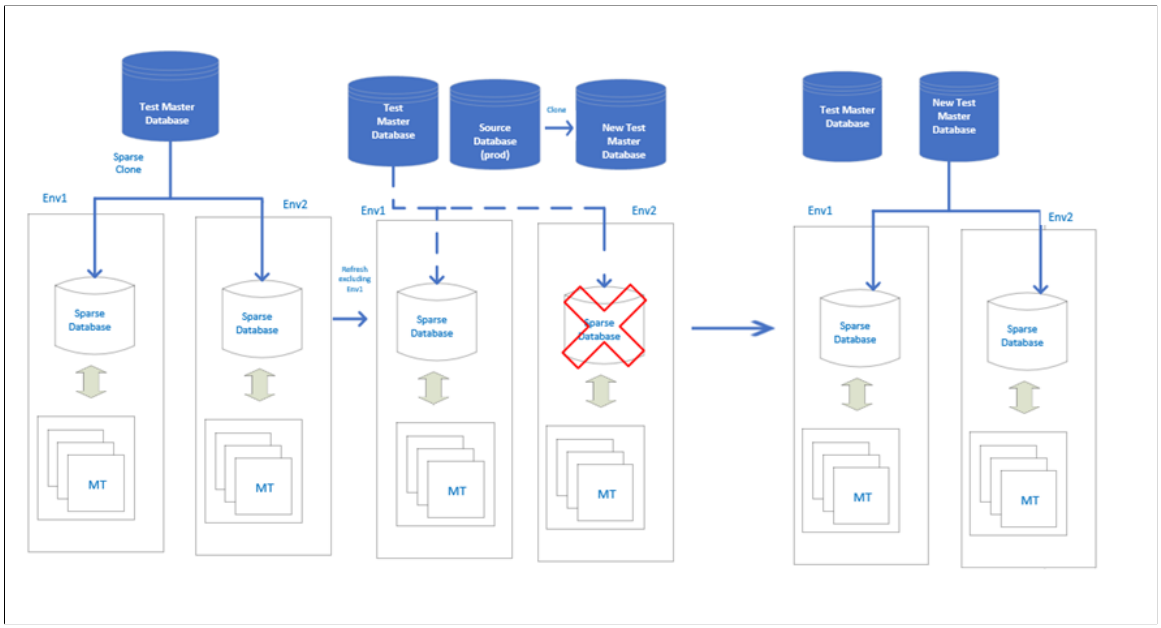

The following processes happen when you refresh the test master excluding certain sparse environments:

1. Cloud Manager deletes the sparse cloned database of the environment that is selected for refresh.
2. A new test master is created from the source database. The existing test master and its sparse environments that are excluded from refresh are not modified.
3. New sparse clones are created from the new test master for each of the sparse environments that are selected for refresh.
4. Cloud Manager reconfigures the middle tiers of sparse environments with the corresponding new sparse clone.

To refresh a test master:

1. Click Refresh on the Sparse Hierarchy Details page. Alternatively, you can click Refresh from the Test Master Details page. The Test Master Refresh page is displayed.

This example illustrates the fields and controls on Test Master Refresh page.

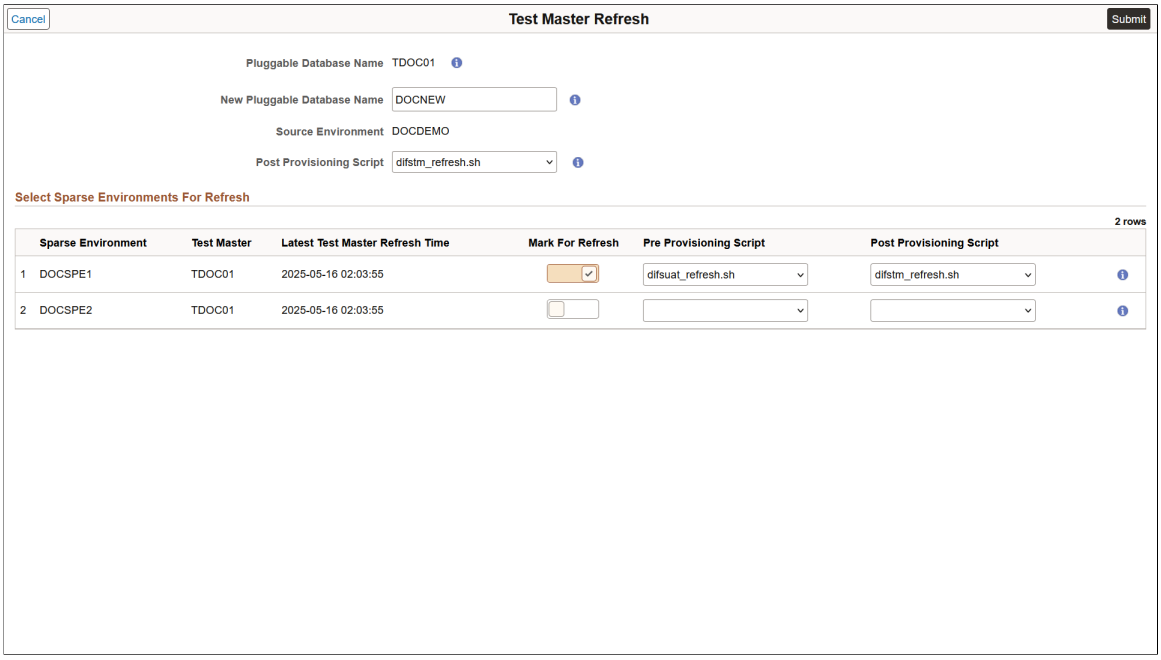

2. Mark the required sparse clones for refresh.

Note: You can also select sparse environments that are not part of the current test master, given that they were last refreshed earlier than the last refresh time of the current test master.

---

3. Select the Post Provisioning Script for the test master.
4. Select the Pre and Post Provisioning Scripts for sparse environments selected for refresh.
5. Click Submit. You must confirm the action when prompted. The refresh process gets initiated after the confirmation. You can view the status of the refresh process on the Sparse Hierarchy Details page.

This example illustrates the fields and controls on Sparse Hierarchy Details page.

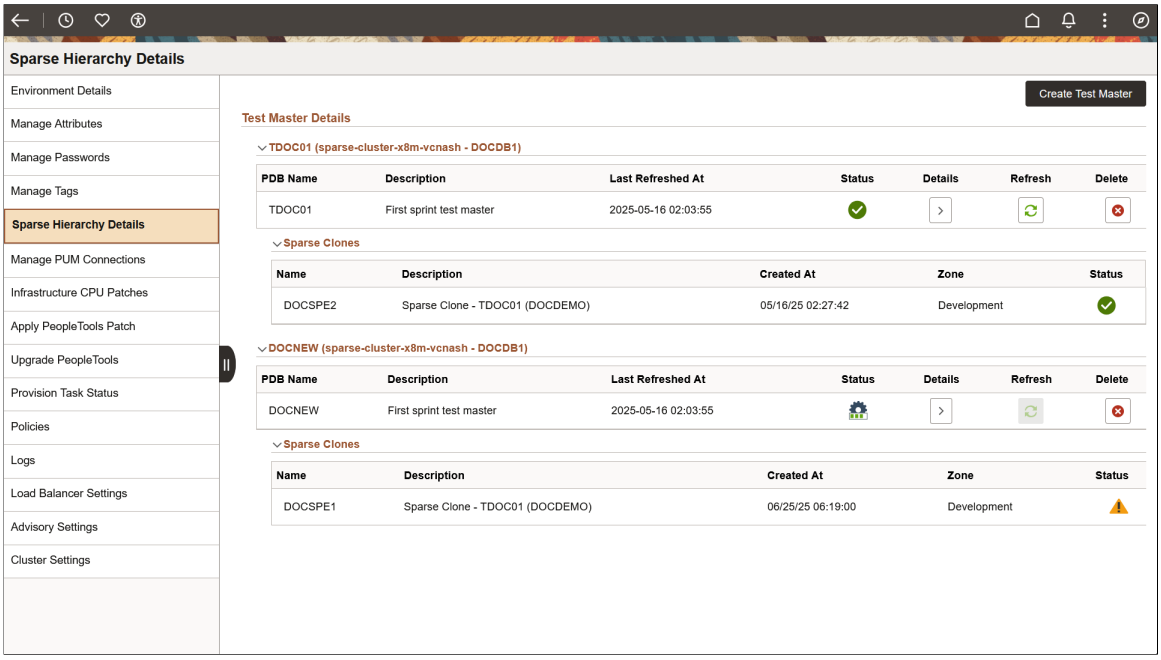

Managing Environment Attributes

Use the Manage Attributes page (ECL_ENV_ATTR_FL) to centrally administer the configuration of all the managed instances using Cloud Manager.

Navigation:

Click the Manage Attributes link available on the left panel of the Environment Details page. The Manage Attributes page is displayed on the right panel.

This example illustrates the fields and controls on the Manage Attributes page.

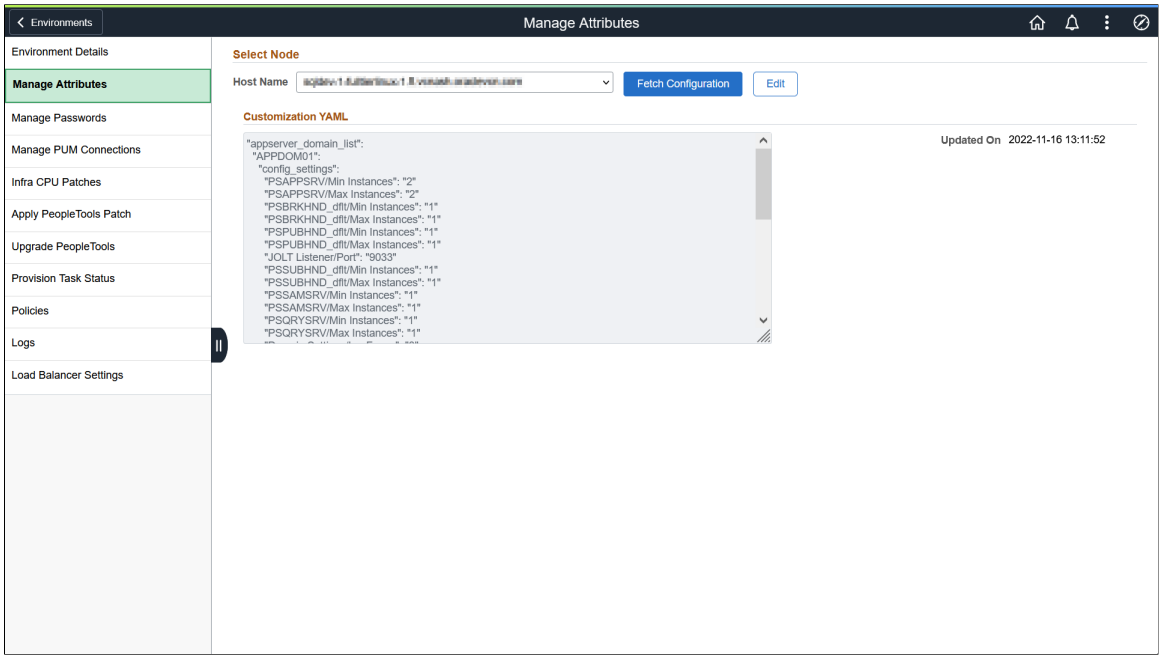

The Manage Attributes page can be used to modify the configuration parameters even after an environment is provisioned. You can select the host name and click the Fetch Configuration button to view the psft_customization.yaml file for the selected environment node. The timestamp showing when the fetch happened is displayed at the bottom of the page.

To modify the customization YAML:

1. Click the **Edit** button. The Customization YAML section is enabled.

   This example illustrates the fields and controls on the Manage Attributes page that appear on clicking the Edit button.

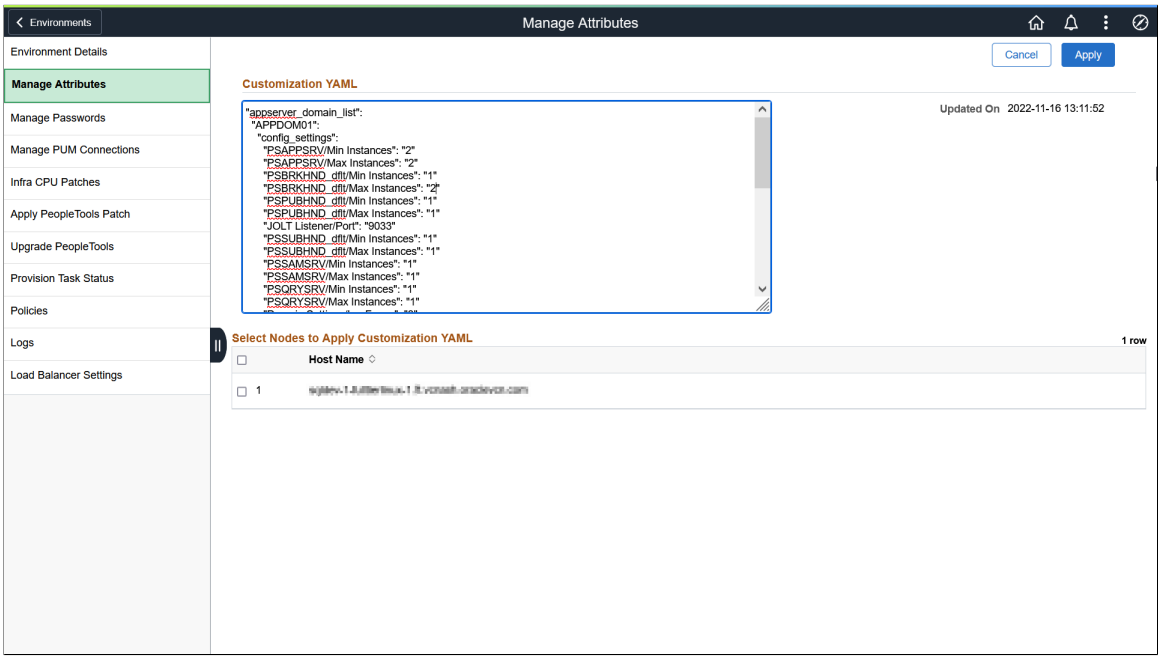

2. Modify the attributes in the application server domain, Process Scheduler domain or web server domain.

Domain	*Attribute*	*Configuration File Location*	*Description*
App Server Domain	PSAPPSRV/Min Instances	$PS_CFG_HOME/appserv/<domain_name>/psappsrv.cfg	Instances for processing functional requests, such as building and loading components.
	PSAPPSRV/Max Instances		
	PSBRKHND_dflt/Min Instances		Instances for processing the requests from dispatcher for Integration broker queues.
	PSBRKHND_dflt/Max Instances		
	PSPUBHND_dflt/Min Instances		Handler instances responsible for publications out bound messages from Integration Broker environment.
	PSPUBHND_dflt/Max Instances		
	JOLT Listener/Port		Port number used by JOLT server handlers.
	PSSUBHND_dflt/Min Instances		Handler instances responsible for subscriptions in bound messages from external system.
	PSSUBHND_dflt/Max Instances		
	PSSAMSRV/Min Instances		SQL application manager process handles the conversational SQL that is mainly associated with Application Designer.
	PSSAMSRV/Max Instances		
	PSQRYSRV/Min Instances		Handles any query run by PeopleSoft Query. This is an optional process designed to improve performance by reducing the workload of PSAPPSRV.
	PSQRYSRV/Max Instances		
	Domain Settings/LogFence		Logging detail level

Domain	*Attribute*	*Configuration File Location*	*Description*
Process Scheduler Domain	PSAESRV/Max Instances	$PS_CFG_HOME/appserv/prcs/<domain_name>/psprcs.cfg	Application Engine server process is a process that handles both Application Engine and Optimization Engine requests.
	PSDSTSRV/Min Instances	$PS_CFG_HOME/appserv/prcs/<domain_name>/psprcs.cfg	The distribution agent server process posts the reports and system log files to the report repository.
	PSDSTSRV/Max Instances		
	Domain Settings/LogFence		Logging detail level
Web server Domain	auth_token_domain	Not editable	Domain in which the portal runs. The single sign on token is valid across this domain.
	max_threads	$PS_CFG_HOME/webserv/<domain_name>/config/config.xml	Maximum number of threads WebLogic can create in the default Execute queue.
	java_heap_min	$PS_CFG_HOME/webserv/<domain_name>/bin/setEnv.sh	Web server JVM heap size
	java_heap_max		
	appserver_connections	$PS_CFG_HOME/webserv/<domain_name>/applications/peoplesoft/PORTAL.war/WEB-INF/psftdocs/<site_name>/configuration.properties	List of host and port details of the JSH process in App server. Web server forwards requests to these URLs.

3. Select the hosts where the customization must be applied on the **Select Nodes to Apply Customization YAML** section.

4. Click **Apply**. The system displays errors and warnings related to the update, if any. You can choose to ignore warnings, but the errors must be resolved.

This example illustrates the errors and warnings on the Manage Attributes page that appear on clicking the Apply button.

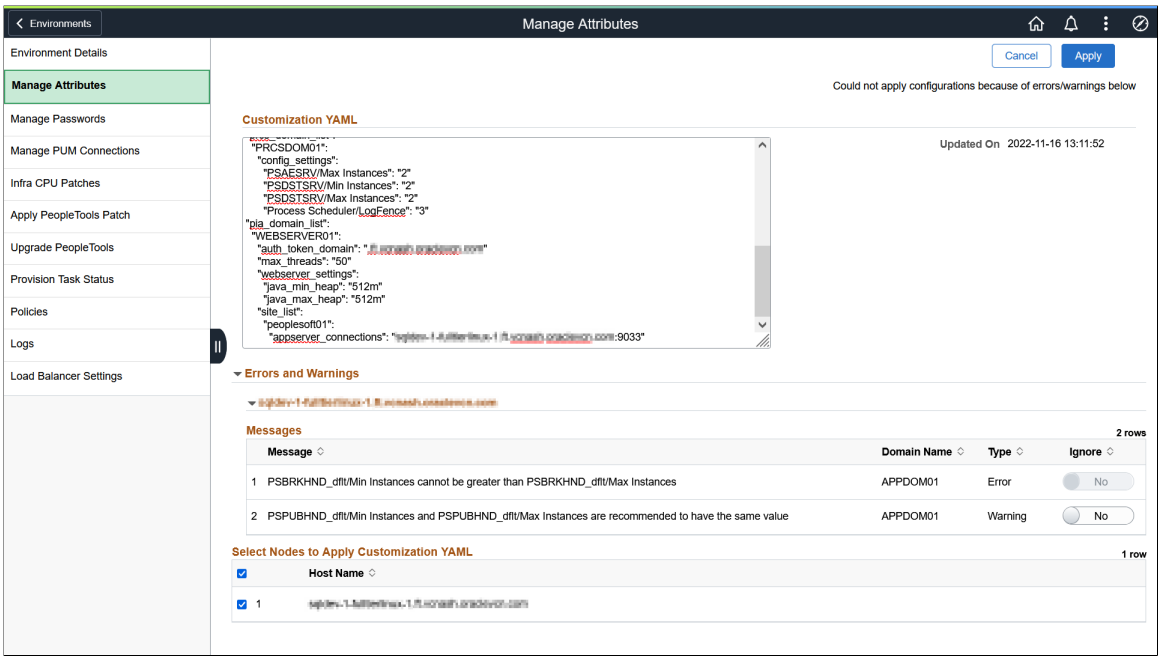

5. Click **Apply** after resolving the errors. The system displays a confirmation message.

   This example illustrates the confirmation message on the Manage Attributes page that appears after all the errors are successfully resolved and warnings are resolved or ignored.

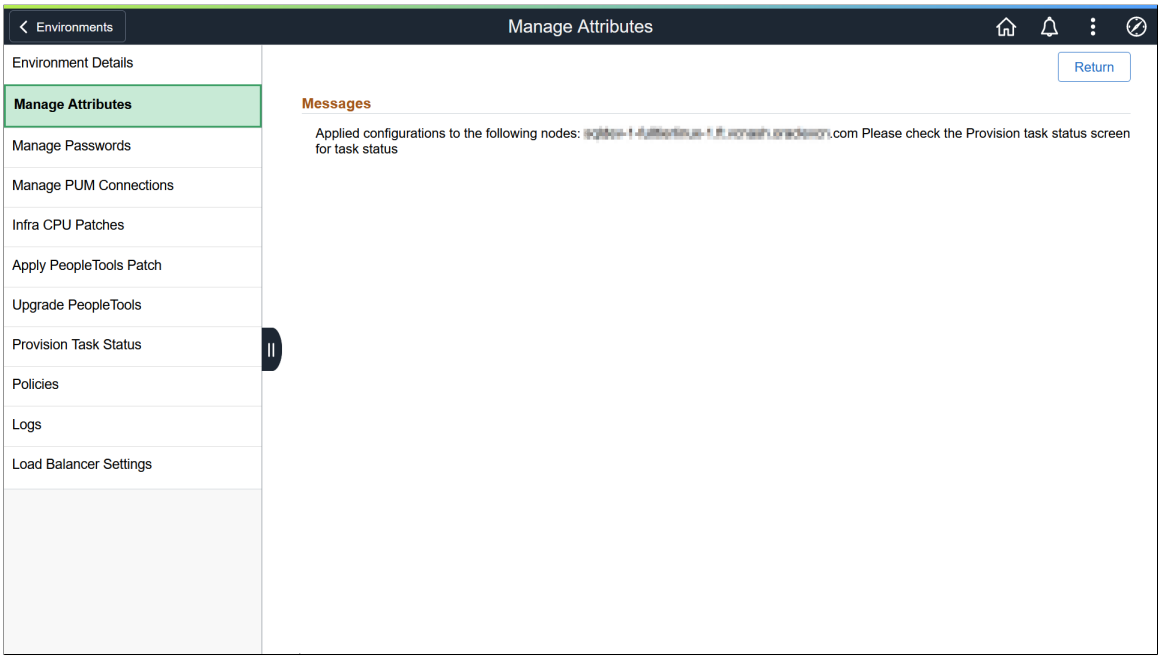

6. Check the Provision Task Status page for the status of application of custom configurations, where the status of deployment task activities is displayed along with their start and end times. See Viewing Provision Task Status.

You cannot modify the configuration of an instance until the tasks in progress for applying the configuration are completed.

Troubleshooting on Failure of Deployment Task

You can find the status of applying custom configurations on the Provision Task Status page. If the task fails, a Failed icon appears in the task **Status** column. If an environment fails due to the presence of faulty values in the Customization YAML, you must revert the configuration files corresponding to the attributes that were modified in the YAML, followed by reconfiguring the domains where attributes are modified, so that the environment can be restarted successfully. Reverting configuration files to the latest backup file is done from PS_CFG_HOME.

Locating the Default PS_CFG_HOME

When you launch PSADMIN, if a PS_CFG_HOME does not exist, the system creates the PS_CFG_HOME directory in the “user” directory of the current user (the owner of the domain). The system assumes the presence of the following environment variables:

Operating System	***Required Environment Variable***
UNIX/Linux	HOME
Windows	USERPROFILE

For example, depending on the operating system of the server, the system creates PS_CFG_HOME in the following location on the same drive as PS_HOME.

Operating System	***PS_CFG_HOME Location***
UNIX/Linux	$HOME/psft/pt/<version>
Windows	%USERPROFILE%\psft\pt\<version>

After you create a domain, the domain exists under $PS_CFG_HOME\appserv\<domain>.

With a user of *psftuser*, on UNIX/Linux this would appear as:

```
/home/psftuser/peopletools/8.58
                     peopletools.properties
                     /appserv
                       /CRM
                       /HR
                       /PRCS
                         /CRM_PRCS
                         /HR_PCRS
                       /Search
                       /webserv
                       /ver_dom
```


With a user of *psftuser*, on Windows this would appear as:

```
C:\Documents and Settings\psftuser\psft\pt\8.58
                  peopletools.properties
                     \appserv
                        \CRM
                        \HR
                        \PRCS
                           \CRM_PRCS
                           \HR_PCRS
                        \Search
                        \Temp
                        \webserv
                        pswinsrv.cfg
```

Note: The previous examples show a situation in which CRM, HR, CRM_PRCS, HR_PRCS and ver_dom are all domain directories. They are not in PS_CFG_HOME by default, and appear only after the domains are created.

To display the default PS_CFG_HOME location, you can submit the following command to PSADMIN:

```
psadmin -defaultPS_CFG_HOME
```

Note: These commands are not case sensitive.

Retrieving Configuration Files When Deployment Task Fails

To retrieve configuration files:

1. Locate the configuration backup file in the same folder as the corresponding configuration file. The backup file is found in the following format:

   ```
   [original_configuration_file_name].cfg_[timestamp_of_backup]
   ```

2. Rename the backup file to the original filename without timestamp. The following command can be used:

   ```
   mv [filename].cfg_[timestamp] [filename].cfg
   ```

   Note: Ensure that you do not alter the ownership and permissions for the configuration file when you rename the file.

3. Select *Configure this domain* on the PeopleSoft Domain Administration menu to access the Quick-Configure menu.

4. Select *Load config as shown* to apply settings to the configuration files.

5. Click the **Refresh Metadata** button on Environment Details page in Cloud Manager.

6. Click the **Fetch Configuration** button on Manage Attributes page. The backup configuration file now becomes the original configuration file.

Configuring Database Backup Settings

You can use this section to configure DB backup settings for DBS-based environments and imported environments.

Cloud Manager supports backup and restore of DBS-based environments that have single or multiple Pluggable Databases (PDB) in a Container Database (CDB). The prerequisite for this is that auto-backup must be enabled for the CDB containing the environment's PDB.

To update database backup settings:

1. Expand the DB Backup Settings section and click Edit Settings.
2. Select the appropriate settings on DB Backup Settings window.

   This example illustrates the fields and controls associated with DB Backup Settings on the Manage Attributes page. This section is enabled only for DBS-based environments. You can find definitions for the fields and controls later on this page.

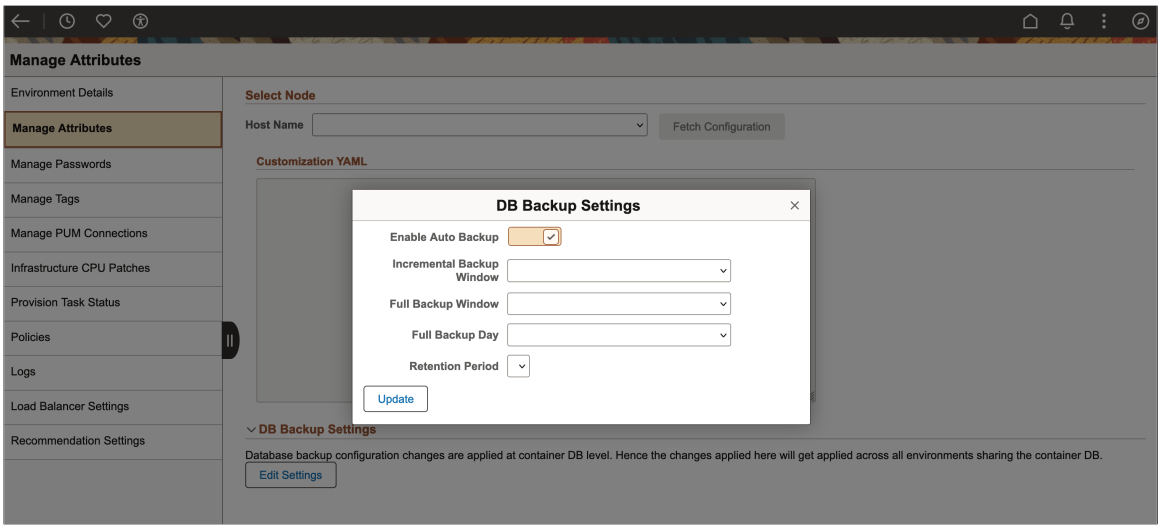

Field or Control	*Description*
Enable Auto Backup	Select to enable automatic backup. This must be selected to enable the backup of DBS-based environments that have multiple PDBs in a CDB.
Incremental Backup Window	Select the time window for performing incremental backup.
Full Backup Window	Select the time window for performing full backup.
Full Backup Day	Select the day for performing full backup.
Retention Period	Select the period for which the backup of an environment is retained. **Important!** The extent of restore is controlled by the retention period of automatic backup configuration.

3. Click Update. The settings are updated.

It may take up to 24 hours for newly deployed or refreshed PDBs to be fully available in OCI. Therefore, you must wait for up to 24 hours for the PDB to synchronize before attempting Restore action. Similarly, due to a limitation in OCI you can restore just one PDB at a time and parallel restoration of environments that share the same CDB is not possible.

Note: When you enable automatic database backup from OCI console, it is recommended to select the 'Take the first backup immediately' checkbox. If you do not want to run the first automatic backup right away, you can wait for the first automatic incremental database backup to finish before starting the Backup & Restore operations on the environment from Cloud Manager.

After a PDB is recovered to a point in time, you cannot recover the same PDB to any backup point between that point and the time in which the recovery action was initiated.

Viewing Environment Logs

Use the Logs page (ECL_ESEARCH_FL) to view the logs for all actions that are performed on the environment.

Note: The contents of the log files are displayed in reverse (latest first) order.

Navigation:

Click the Logs link available on the left panel of the Environment Details page. The Logs page is displayed in the right panel.

This example illustrates the fields and controls on the Logs page.

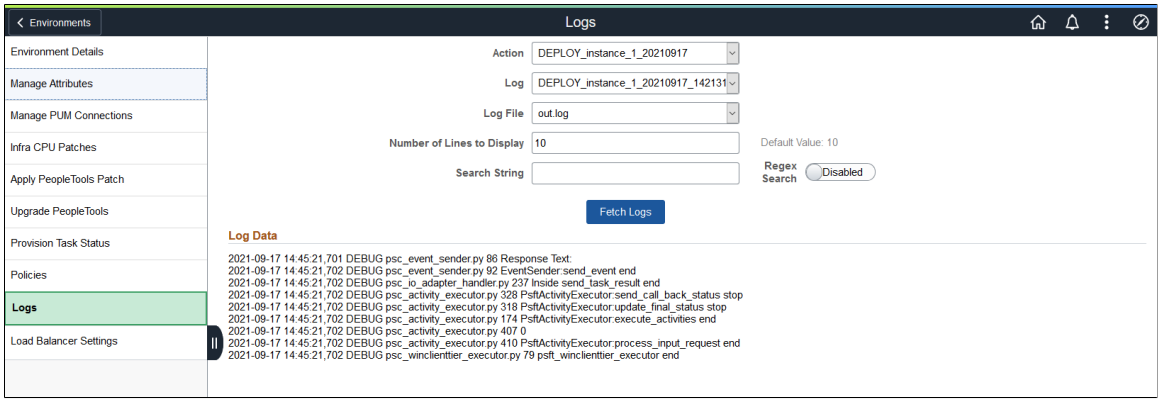

Field or Control	*Description*
Action	Select the Action from the drop down. Actions include: • ADD_NODE • CREATE_DB_BACKUP • ES_UNPROVISION • LIFT • OCI_CLONE • POSTBOOT_CONFIG • PTU_PEOPLETOOLS_PATCH_UPDATE • PTU_PEOPLETOOLS_UPGRADE • PTU_REPROV_MT • PTU_STOP • PTU_UNPROV_MT • REMOVE_NODE • SHIFT_DEPLOY_TDE • SHIFT_DEPLOY • ADD_TARGET • START • STOP
Log	Select the log associated with the action.
Log File	• console.log • out.log • tf.err Terraforma error log • tf.out Terraforma log
Number of Lines to Display	Enter number of lines to display.
Search String	Enter a search string.

Field or Control	*Description*
Fetch Logs	Click to fetch the logs.

Monitoring Environments

The Monitoring page is available on managed environments that contain a search framework, database system, and middle tier.

Monitoring must be enabled to use Auto Scaling. See Setting Up Auto Scaling.

Use the Monitoring page to enable or disable monitoring for the environment.

The monitoring graph is displayed when:

- Monitoring is enabled for the environment.
- OCI or JMX metrics are being collected.
- Data Science is used in the environment for prediction.

This example illustrates the fields and control on the Monitoring page. You can find definitions for the fields and controls later on this page.

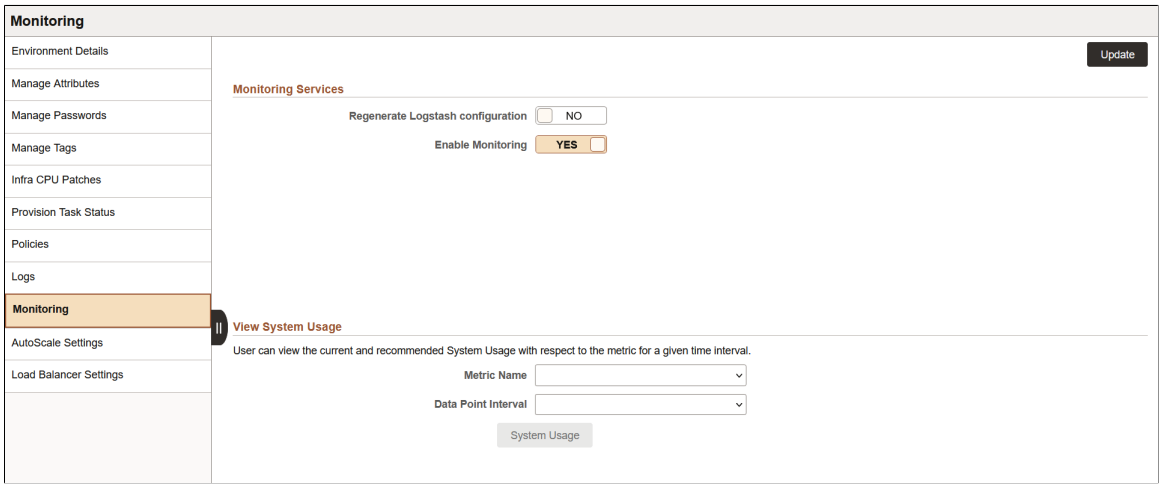

Field or Control	*Description*
Regenerate Logstash configuration	Select YES to regenerate Logstash configuration. The default value is *NO*.
Enable Monitoring	Select YES to enable monitoring. The default value is *NO*.

Field or Control	*Description*
Metric Name	Select the metric name for plotting against the System Usage. The available metrics are: • App Server Current Queue Depth • App Server Process CPU Load • App Server System CPU Load • App Server throughput • Http Requests • PIA Average Total Service Time • PIA Sessions In Use • PIA TCP Sockets in Established State • PIA TCP Sockets in wait State • Webserver Execute Thread idle Count • Webserver Jolt Throughput • Webserver Thread Queue Length • Webserver Thread throughput
Data Point Interval	Select the data point interval to plot the number of nodes in the environment. Plots are available for Public Load Balancer and Private Load Balancer. Data Point intervals available are: • Bi-Hourly – data points plotted on a 2 hour basis • Day – data points plotted per day • Half-day – data points plotted on a half-day basis • Hourly – data points plotted on an hourly basis • Weekly – data points plotted weekly
System Usage	Click this button to view the jet chart graph showing the system usage. You can view the current and recommended system usage with respect to a metric at the selected data point interval in the monitoring graph.

Monitoring System Usage

The monitoring graph displays the following information:

- Maximum mid tier nodes currently present in the environment

- Maximum predicted mid tier nodes that should be present in the environment based on model predictions

This data is displayed for private as well as public load balancers. The y-axis in the upper section represents the metric selected. The y-axis in the lower section represents the number of current, recommended, and forecasted mid-tier nodes in the system. The x-axis represents the date time points. The size of the circles indicating data points is directly proportional to the severity of the anomalous situation.

Note: The system usage chart displays daily and weekly forecasting data only for Http Requests metric and Middle Tier nodes. Forecasting is not applicable for any other metric.

This example illustrates the jet chart showing the number of incoming Http requests and the number of current and recommended mid tier nodes for private load balancer.

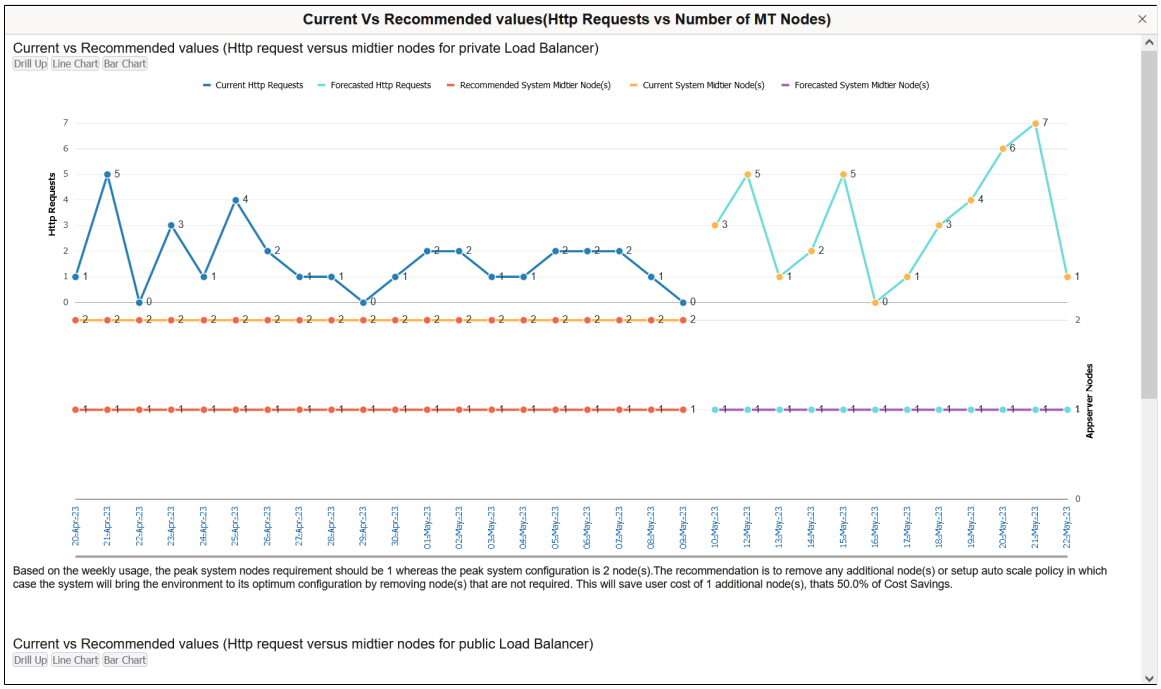

Upward and downward mouse scroll actions can be used respectively to zoom in and out the data values.

When the hourly or bi-hourly data point interval is selected, you can drill down the chart by clicking the data point interval link below the graph. In this view, the Non Anomalous nodes, Anomalous Remove nodes, and Anomalous Add nodes are represented in green, blue and red respectively. A line chart is displayed by default. You can also view it as a bar chart. On hovering above the bars on a bar chart for a particular metric, the chart displays values of the other metrics. You can go back to the original view by selecting the Drill Up option.

This example illustrates the hourly drill down chart App Server System CPU Load metric as a line chart.

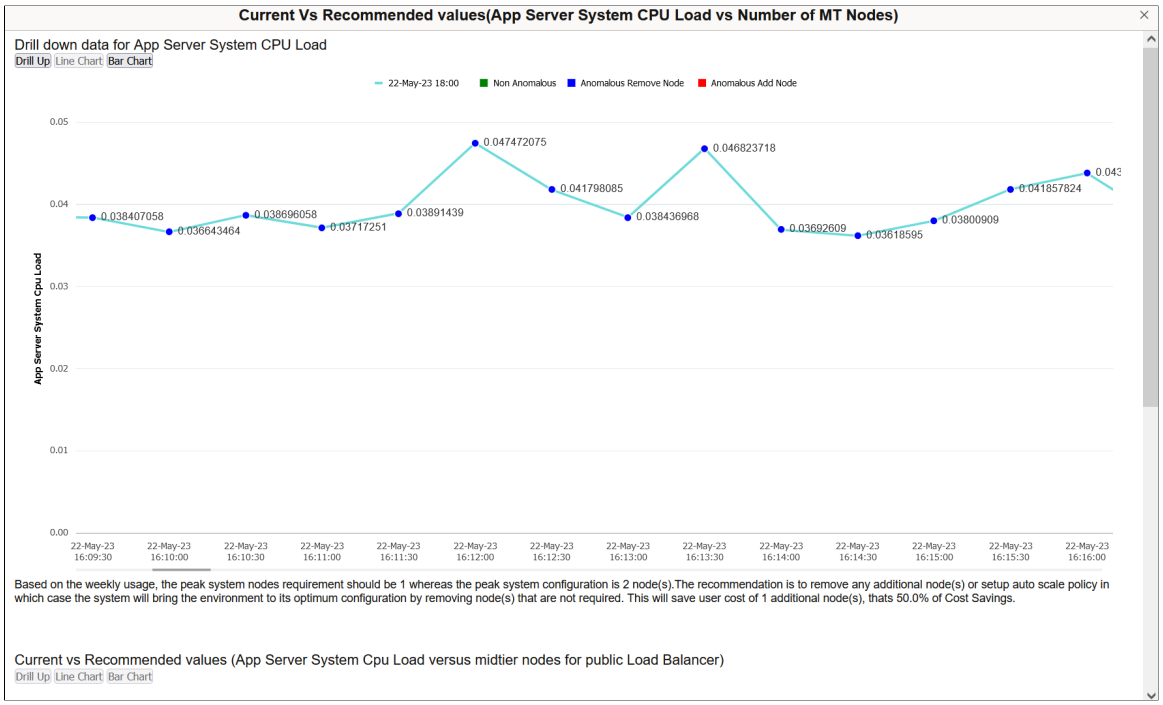

This example illustrates the hourly drill down chart App Server System CPU Load metric as a bar chart.

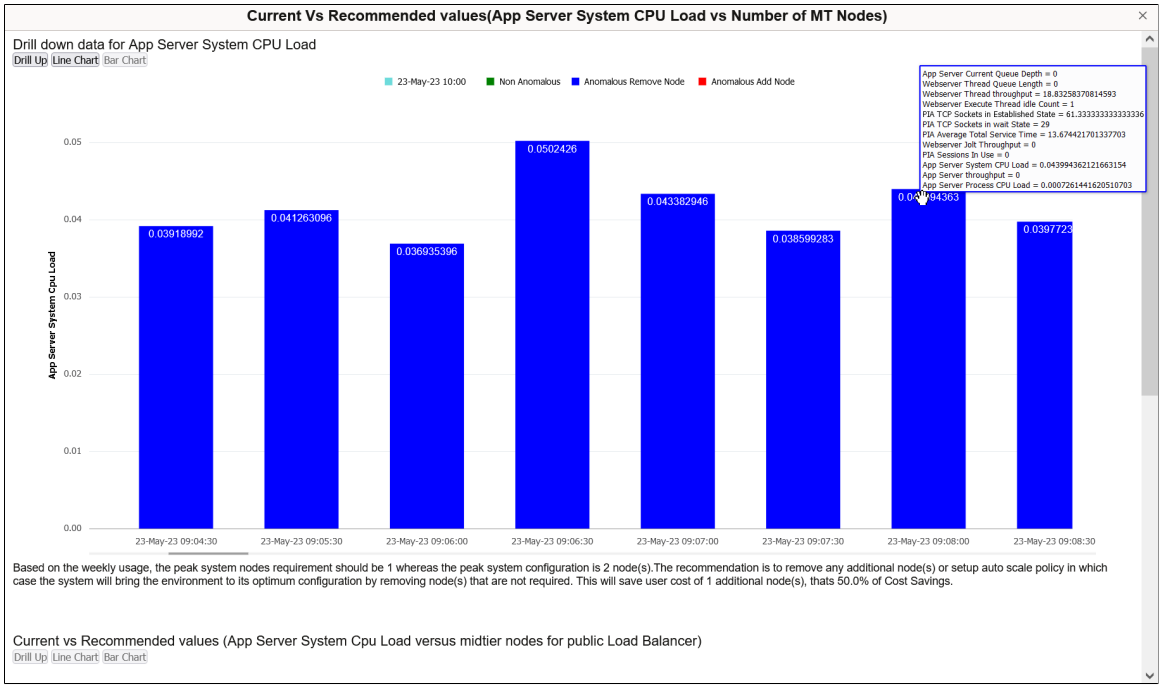

You can also find recommendations at the end of the graph to optimize the usage of nodes based on weekly statistics. Data Interval is used to divide the time into multiple date time points.

Weekly Notifications for Environment Settings

Cloud Manager sends you recommendations for environments on weekly basis, based on their weekly usage. These recommendations are sent as push notifications. You can also get this notification through email by configuring the Notification Topic OCID. See Configuring Cloud Manager Settings for OCI.

The following snippet contains sample text from an email with the recommendations for environment settings.

```
Please find list of Environment(s) and their setting recommendation based on Usage:⇒


 - <Environmentname>:
        public_lb:
        https://objectstorage.us-ashburn-1.oraclecloud.com/p/xxxxx/n/tenancyABC/b/x⇒

xxxx/o/1656048898406-330880_20221029_061223_1667023943.jpg

        private_lb: Based on the weekly usage, the peak system nodes requirement sh⇒

ould be 1 whereas the peak system configuration is 3 node(s).
        The recommendation is to remove any additional node(s) or setup auto scale ⇒

policy in which case the system will bring the environment to its optimum configura⇒

tion by removing node(s) that are not required.
        This will save user cost of 2 additional node(s), thats 66.66667% of Cost S⇒

avings.
        https://objectstorage.us-ashburn-1.oraclecloud.com/p/xxxxx/n/tenancyABC/b/x⇒

xxxx/o/1656048898406-330880_20221029_061221_1667023941.jpg
```

You can also find push notifications by clicking on the Alerts section under Notification panel on the home page. On clicking the push notification for a particular environment, you can find the Jet Chart containing details related to the weekly node usage, weekly node requirements, and cost comparison for current and predicted node configuration.

This example illustrates the weekly recommendation chart that appears on clicking the notification specific to an environment on the Alerts section under Notification panel on the home page.

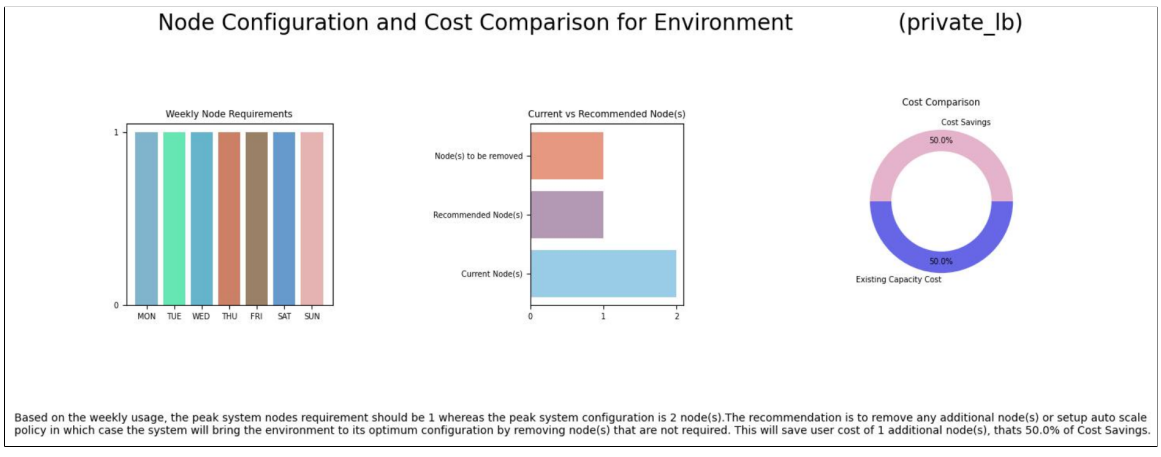

Configuring Load Balancer Settings

Use the Load Balancer page (ECL_LB_BACKENDS_FL) to configure or update load balancer settings.

Load Balancer can be configured for the following:

- Fulltier environment with or without Kibana.
- Midtier environment.
- ES node with Kibana.

Note: If multiple web server domains are included for the load balancer, you must ensure that each web server has a valid IB configuration. Verify that the integrationGateway properties file for each PIA domain has a valid IB configuration. It may be necessary to copy the IB configuration from another web server domain integrationGateway properties file.

The Oracle Cloud Infrastructure Load Balancing service provides automated traffic distribution from one entry point to multiple servers reachable from your virtual cloud network (VCN). To set up the Load Balancer in OCI refer to the Cloud Manager installation tutorials at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials

To configure load balancer:

1. Click the Environment tile, and select the PeopleSoft environment for which to use the load balancer.
2. It is necessary to access the PIA URL for the environment at least once before configuring the load balancer.

   Note: This only needs to be done one time after provisioning the environment.

3. On the environments page, select Details from the actions menu for the environment.
4. Select Load Balancer Setting from the menu on the left.
5. Expand FQDN for end URLs and enter the values.
6. Configure Backend Set.

   In the Configure Backend Set section, you can set up multiple backend sets. Backend sets can be configured for PIA or Kibana. For each backend complete the following sections.

   a. Choose Backend Set
   b. Choose Backend
   c. Configure Listener

7. Select Update after entering all the values.

Note: The Load Balancer page must be updated for cloned environments or after an environment is restored.

FQDN for end URLs

This example illustrates the fields and controls on the Load Balancer Settings page for FQDN for end URLs. You can find definitions for the fields and controls later on this page.

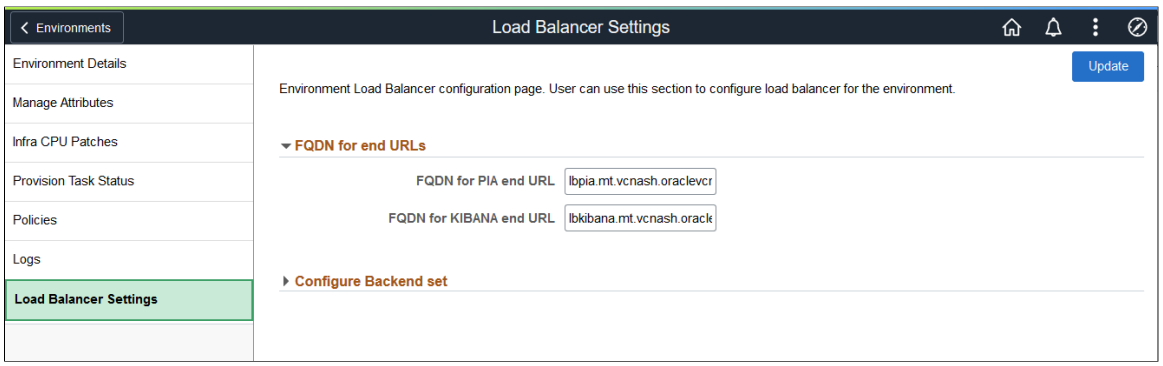

Field or Control	*Description*
FQDN for PIA end URL	The PIA FQDN is the fully-qualified domain name that is resolved for the PIA backend sets. Enter a descriptive name for the first part of the FQDN. Before deploying the environment plan for the load balancer FQDN. When deploying the environment, set the Authentication Domain to match the load balancer FQDN. The last portion of the FQDN must match the authentication domain for the web server. **Note:** Ensure that the FQDN of the load balancer is resolved correctly from the end-user machines. This is typically done by adding DNS entries.
FQDN for KIBANA end URL	The KIBANA FQDN is the fully-qualified domain name that is resolved for the Kibana backend sets. Enter a descriptive name for the first part of the FQDN.

Configure Backend Set for PIA nodes

To configure the backend set for PIA nodes, expand each section and enter the values.

Note: When configuring the backend set, review that the listener and backend set are properly defined in OCI and listener is configured to use the specific backend set.

This example illustrates the fields and controls on the Load Balance Settings page for backend set for PIA. You can find definitions for the fields and controls later on this page.

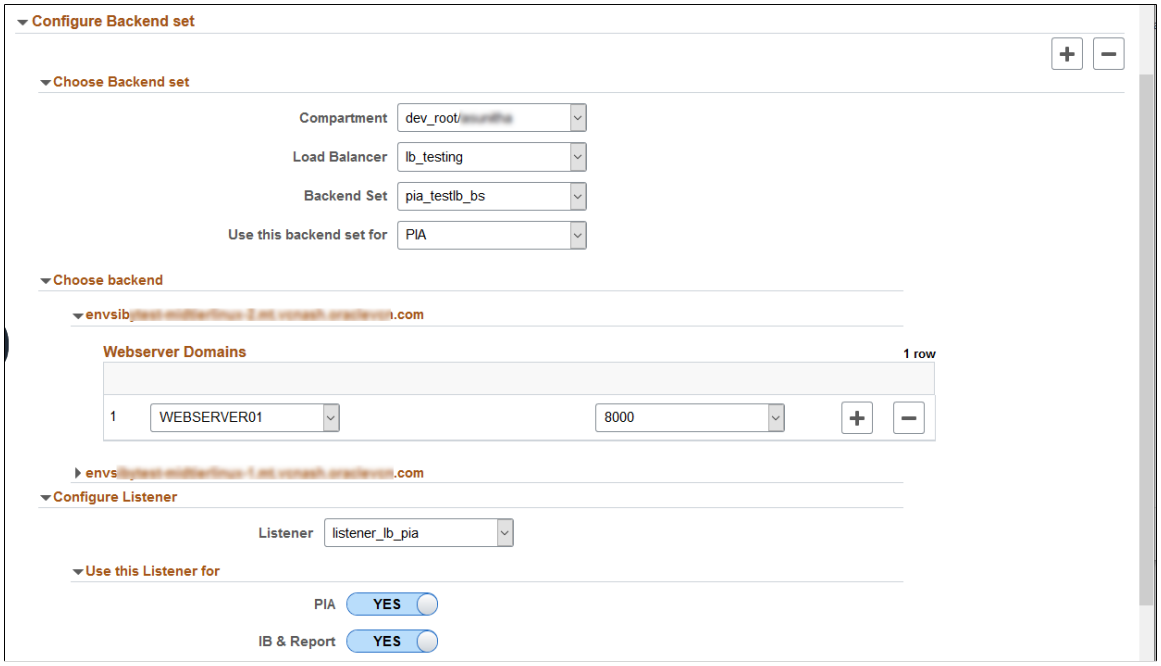

Field or Control	***Description***
Compartment	Compartment where Load Balancer is configured.
Load Balancer	Select the Load Balancer that was set up for PIA in OCI.
Backend Set	Select the backend set name that was set up in OCI.
Use this backend set for	Select PIA when configuring the PIA backend set.

Select the Domain Name and Port for participating nodes. You can add or delete rows as necessary.

Field or Control	***Description***
Listener	The listener name is specified when setting up the Load Balancer in OCI.

Field or Control	*Description*
Use this Listener for	Select the nodes to use this Listener for. • PIA • IB & Report It is mandatory to select a listener for both PIA and IB & Report. Listener for PIA and IB & Report can be selected in a single backend set or two different backend sets, one for PIA and one for IB & Report. When you select Update, the system will verify that a Listener is set for PIA and IB & Report.

Configure Backend Set for Kibana

This is an example of the configuration for Kibana.

This example illustrates the fields and controls on the Load Balance Settings page for Kibana backend.

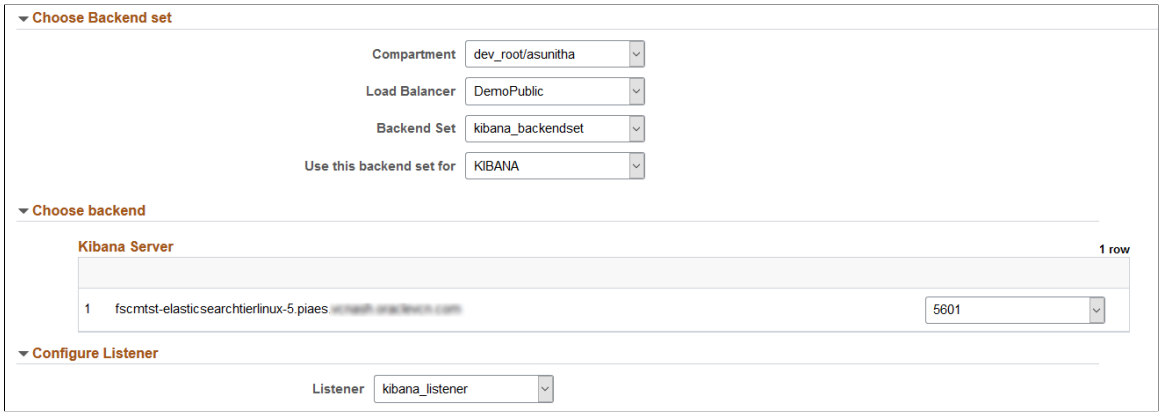

Field or Control	*Description*
Kibana Server	Port for Kibana
Listener	The listener name is specified when setting up the Load Balancer in OCI.

Cloning Environment

Use the Clone Environment action to duplicate an existing PeopleSoft environment running in Cloud Manager.

The new environment need not be an exact copy of the original environment, since you can choose the instances you wish to copy. The new environment is built by reconfiguring the disks from deep clone, saving installation and deployment time.

Reviewing Requirements for Cloning

The following restrictions apply to cloning environments:

- Parallel cloning for same source environment is not supported.
- Clone feature is not supported for Oracle database versions prior to 19.0.0.0.
- Clone feature will not be able to trigger a clone if a previous clone for the same source environment is in a failed state.

  Delete the failed cloned environment using Cloud Manager. Cloud Manager will clean up all the resources relating to the failed Oracle Cloud Infrastructure (OCI) clone environment.
- The source Pluggable Database (PDB) should be available in OCI to be cloned.
- Clone feature does not support bare metal (BM) instance cloning.
- If DB system is created manually and TDE Wallet Password is not the same as Database Administrator Password, TDE Wallet Password on the Manage Attributes page must be updated before initiating Local clone.
- Local cloning and Remote cloning features are supported for DB Systems of VM shape and Exadata.
- To create a Remote clone:
  - The source and destination databases must be in the same availability domain.
  - The source and destination databases must use the same Oracle database software version and be in the same software edition.

Updating the Active Web Profile on the Cloned Environment

Clone feature does not update the active web profile on the cloned environment. The active web profile for the cloned environment retains the load balancer information in the virtual addressing section.

See the information on configuring web profiles in the product documentation *PeopleTools: Portal Technology*.

To set up the load balancer on the cloned environment, select Details for the environment and configure the Load Balancer Settings page.

To use the cloned environment without a load balancer, you need to manually update the host, port in virtual addressing settings for the active web profile. To update the web profile:

1. Switch to a different active web profile by editing the WebProfile property of configuration.properties file in the web server domain and restarting the web server domain.

   To use PSADMIN to administer a PIA site, see the product documentation, *PeopleTools: System and Server Administration*, Using the Web (PIA) Server Menu.

2. Correct host, port in virtual addressing section of webprofile profile configuration for the intended web profile.

3. Switch back to the intended web profile by editing the WebProfile property of configuration.properties file in the web server domain and restarting the web server domain.

Reviewing Cloning Scenarios

Here are examples of use cases for cloning environments.

- Duplicate environment with Database, Middle Tier, Search Stack Node including OpenSearch Dashboards or Kibana, Web Server, and Windows client running in distributed nodes with the exact configuration and data as the source environment.
- A scaled down copy with Database, Middle Tier and Web Server for a test environment.
- All environment components are running on compute instances.

  See Cloning Compute Instances later in this section.

- Database is running on DBS (Oracle Base Database Service) and other components are running on compute instances (Infrastructure as a Service).

  See Cloning an Environment With Database Running on DBS and Other Nodes on Compute later in this section.

- Database is running on Autonomous Database (ADB-Dedicated or ADB-Shared) and other components are running on compute instances (Infrastructure as a Service).

  See Cloning an Environment With Database Running on ADB and Other Nodes on Compute later in this section.

- Database is running on Exadata Database Service on Dedicated Infrastructure (ExaCS) and other components are running on compute instances (Infrastructure as a Service).

  See Cloning an Environment With Database Running on Exadata and Other Nodes on Compute later in this section.

Note: The cloned environment need not be an exact copy of the existing environment, because you can select the tiers you want to clone. Once the cloned environment is running, you can perform scaling and Lifecycle Management actions.
See Managing Nodes.

Cloning Compute Instances

When all of the components for the PeopleSoft environment are running on Compute, the clone process uses OCI Clone APIs to clone the boot volume and secondary block volume into a volume group. The volume group is then used to create a new instance of the environment.

This diagram illustrates the cloning process.

The cloning operation will copy both the boot volume and the secondary block volume into a volume group, then perform a fast deep copy to create a new environment that is the exact copy of the source.

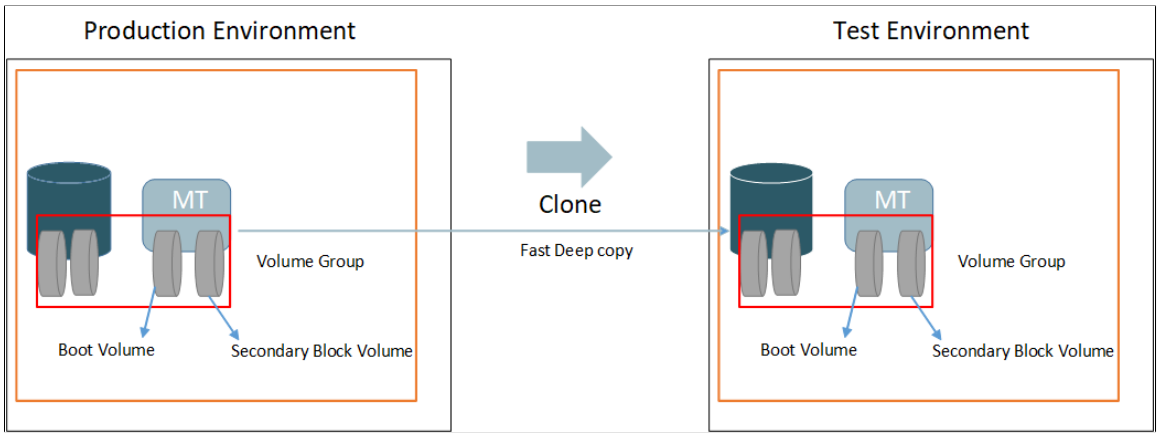

To clone an environment with all components running on compute instances:

1. From the Cloud Manager Homepage, select the Environments tile.
2. Click on the Related Actions button corresponding to the environment to be cloned and select Clone Environment.
3. The Clone Window displays all the components for the environment.

   This is an example of the Clone page for a full-tier environment, where all components are running on compute. You can find definitions for the fields and controls later on this page.

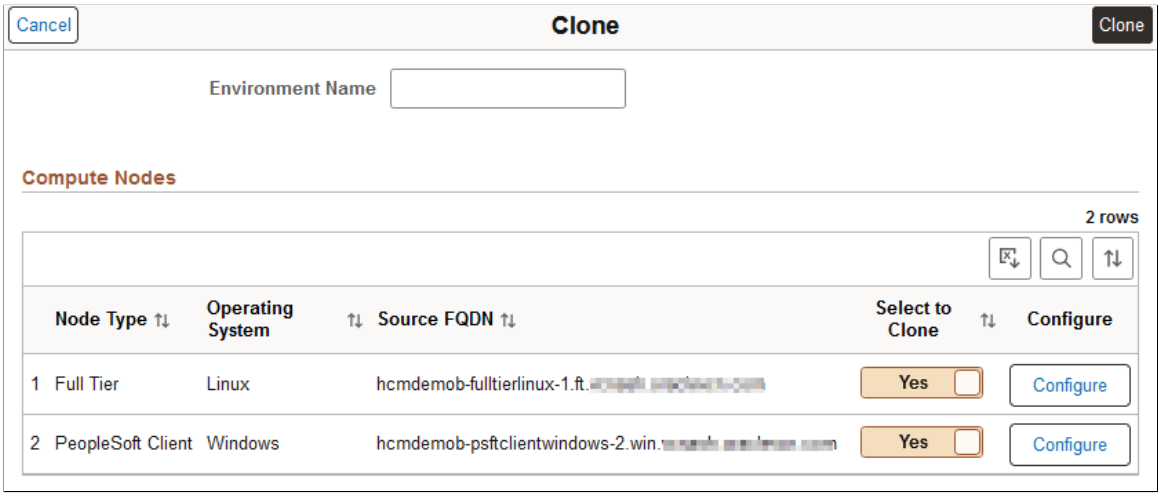

4. Enter a new Environment Name.
5. Remove any instance of your choice from the target clone environment by selecting **No** in the Select to Clone field corresponding to that environment.

   Note: The option to select the field to be cloned is disabled for database instances and full-tier instances, because they are essential to create the target environment.

6. Click the **Configure** button.

The Clone window displays the Network Configuration and General Settings associated with the selected instance for cloning.

This is an example of the page used to configure compute instances for cloning. You can find definitions for the fields and controls later on this page.

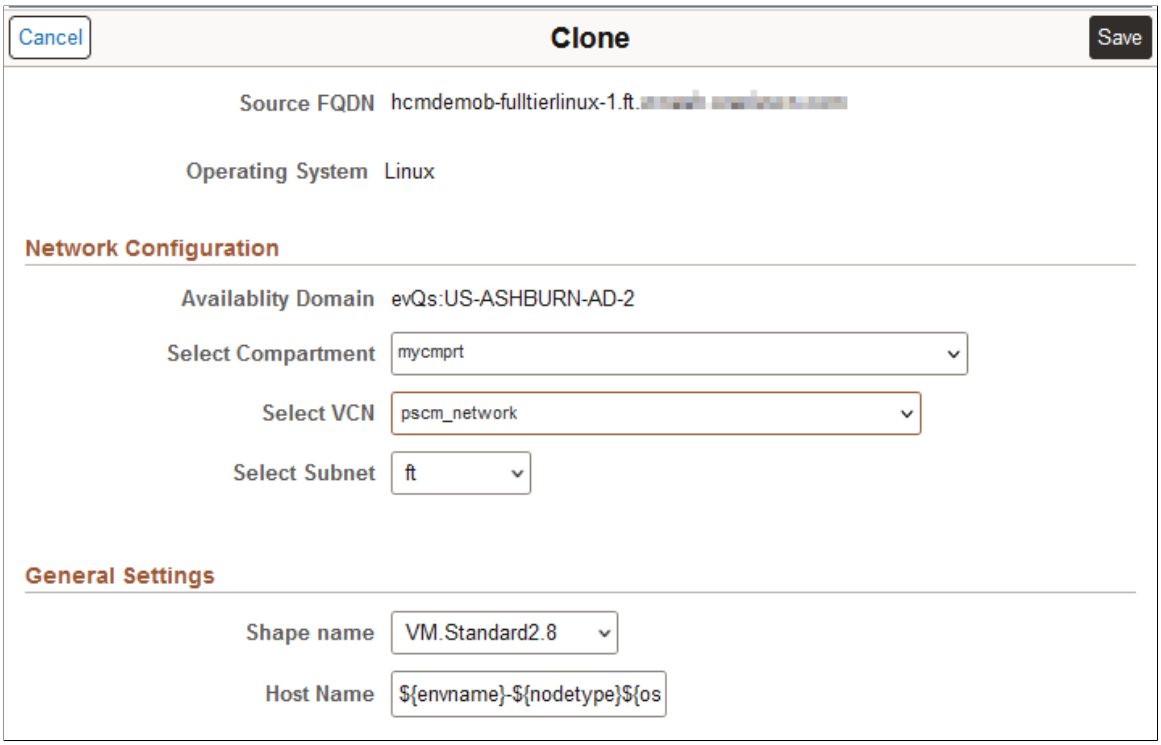

7. Make the necessary changes to the Network Configuration and General Settings for the cloned environment.

Field or Control	***Description***
Select Compartment	Select the compartment containing the VCN you want to use.
Select VCN	Select the VCN for the compute node.
Select Subnet	Select the subnet within VCN for the compute node.
Shape Name	Select the name of the standard/ flexible shape that you want to use for the compute node.
Host Name	Enter a new host name if necessary. By default, the host name will have the following format: ${envname}-${nodetype}${ostype}-${instno}, where envname stands for Environment Name and instno stands for Instance Number.

8. Click **Save** to save the changes.
9. Click the **Clone** button.
10. Select Yes to confirm. Cloning will initiate.

    The cloning process will take 10 to 15 minutes.

    Note: The system validates available resources before starting the cloning process. See Validating Resources

    ---

11. Use the Refresh button to view the status. Status will change from Initiating to Provisioning and then to Running.

Cloning an Environment With Database Running on DBS and Other Nodes on Compute

For environments with database on DBS, the Clone operation requires OCI clone APIs to create the corresponding database node in the cloned environment.

There are three options for cloning environments with database on DBS: local cloning, remote cloning, and new DB System.

Local Clone Type for DB System

Perform local cloning by creating a clone of the PDBs within the existing DB system and CDB. Using this operation, you can clone the PDB alone, without creating a new DB system.

This diagram illustrates cloning the PDB within the existing DB system.

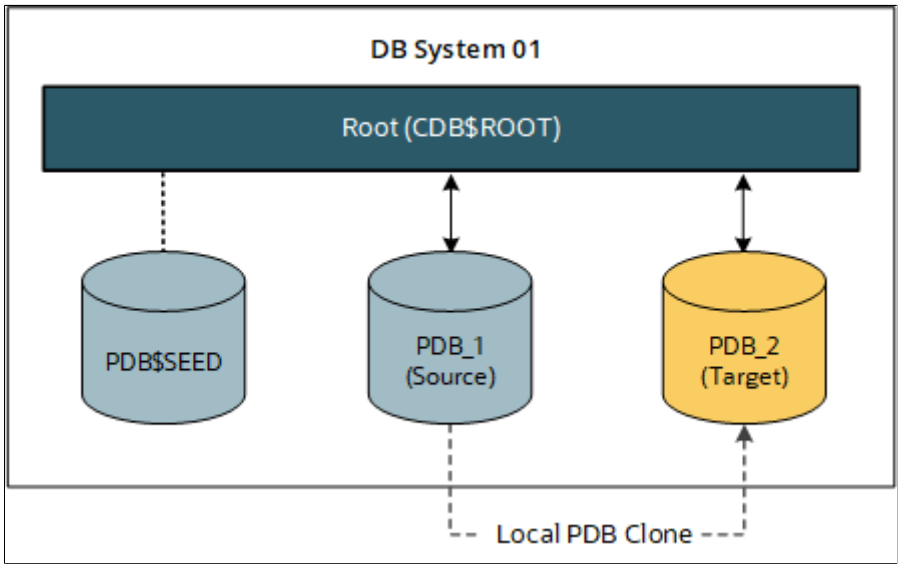

To perform local cloning for an environment where the database is running in DBS and other nodes are on compute:

1. From the Cloud Manager Homepage, select the Environments tile.

2. Click on the Related Actions button corresponding to the DBS environment to be cloned and select Clone Environment.
3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

   This example illustrates the fields and controls on the Clone page for an environment where the database is running as a service. You can find definitions for the fields and controls later on this page.

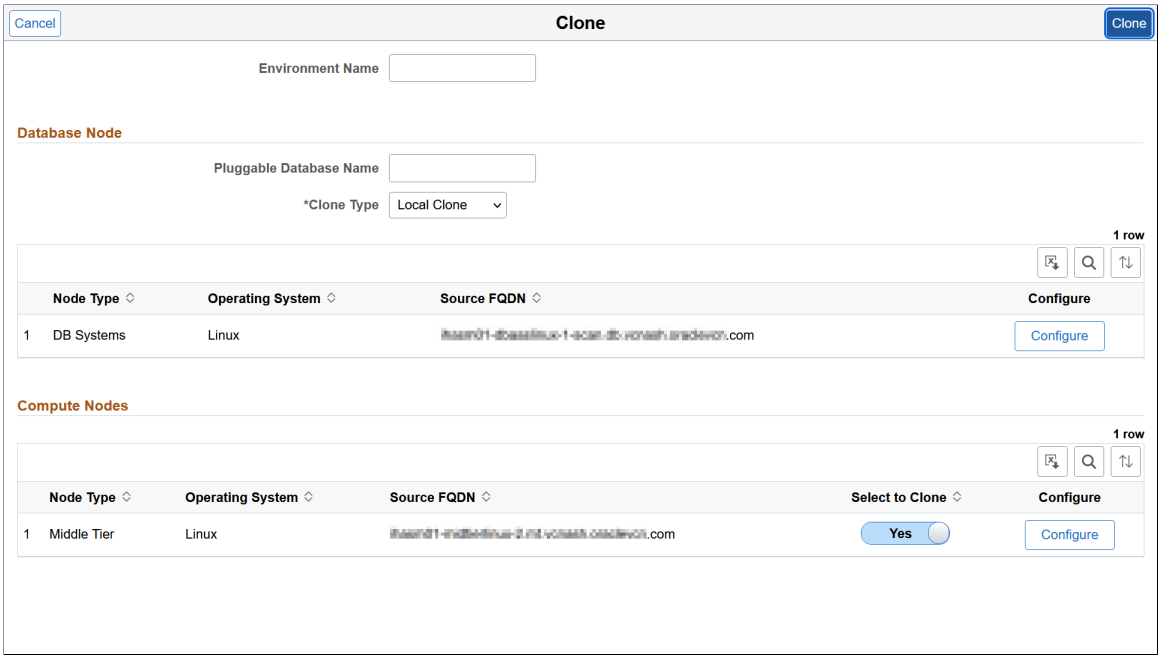

Field or Control	*Description*
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name of the pluggable database to be cloned.
Clone Type	Select *Local Clone* type. The other available options are *Remote Clone* and *New DBS*.
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment.

This example illustrates the fields and controls on the Clone page for an environment where the database is running as a service, with File System details included as part of the Compute Nodes section and Middle Tier using the File System service. You can find definitions for the fields and controls later on this page.

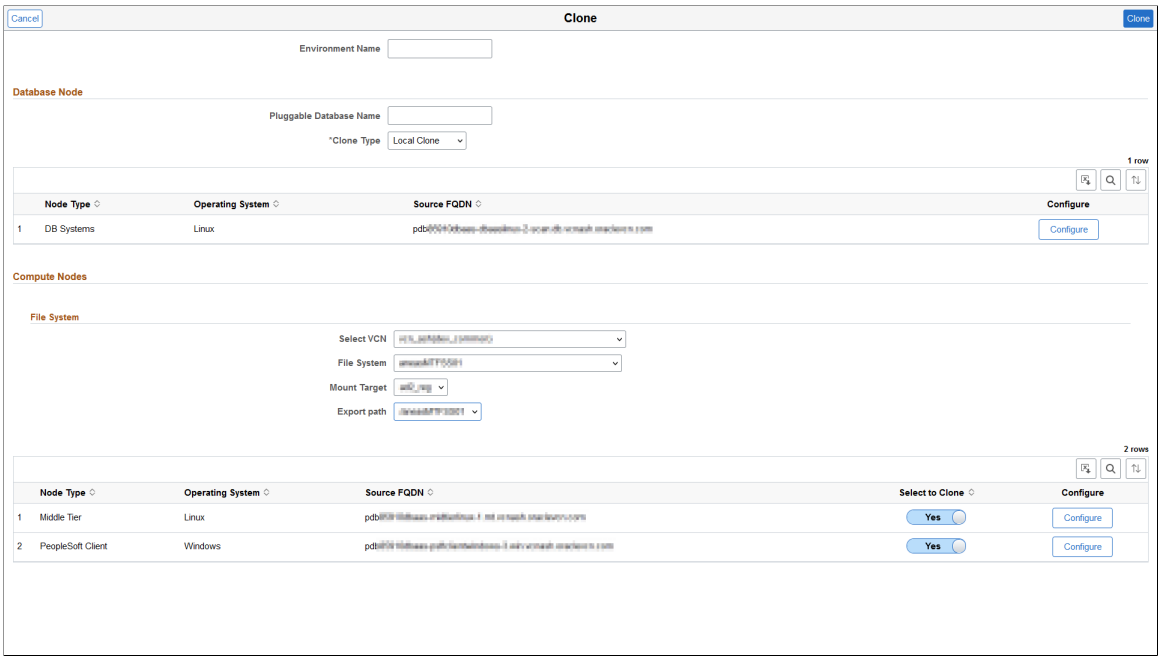

Field or Control	***Description***
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name of the pluggable database to be cloned.
Clone Type	Select the type of cloning you wish to perform on the PDB. The default option is *Local Clone*. The other available options are *Remote Clone* and *New DBS*.
Select VCN	Select the Virtual Cloud Network that contains the desired Mount Target for the file system.
File System	Select the name of the file system.
Mount Target	Select the Mount Target from the drop down list.
Export path	Select the Export path from the drop down list.

Field or Control	*Description*
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment.

4. Enter a new Environment Name.
5. Enter a name for the Pluggable Database.
6. Select a Clone Type from the drop down menu.
7. Click the **Configure** button corresponding to the database instance.

   For the Local Clone type, the Network Configuration and General Settings fields are not editable.

   This example illustrates the fields and controls on the Clone page for a DB system instance when the Clone Type is selected as Local Clone. You can find definitions for the fields and controls later on this page.

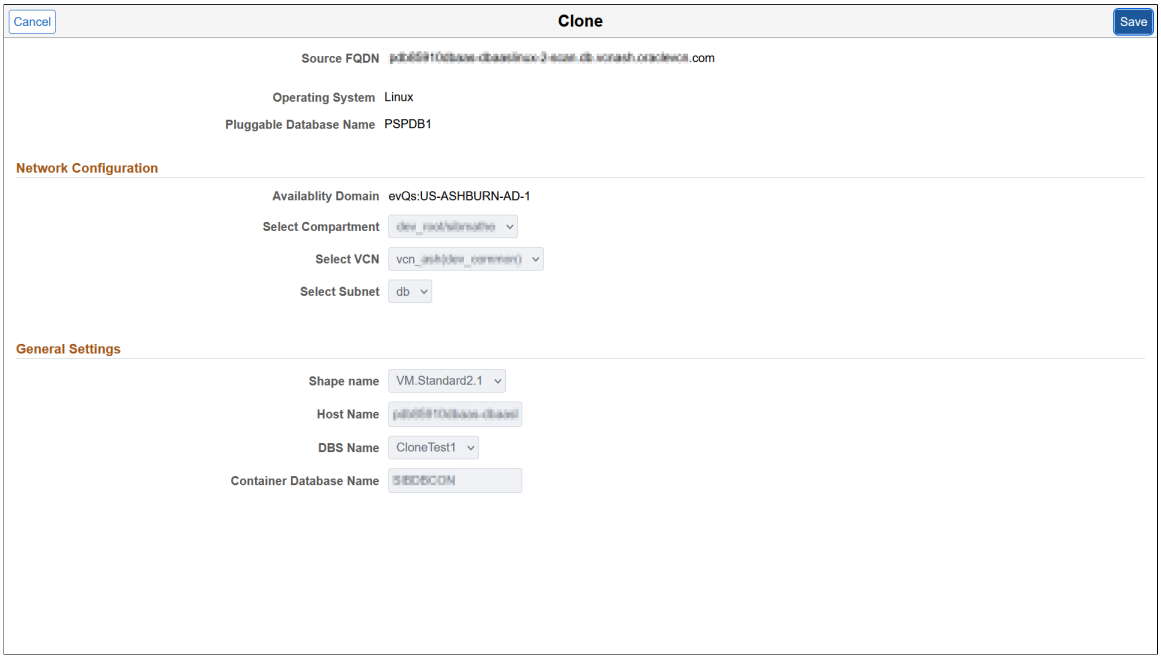

8. Click **Save** to continue.
9. Remove any instance of your choice from the target clone environment by selecting **No** in the Select to Clone instance field corresponding to that environment.

   Note: The option to select the field to be cloned is disabled for database instances, because they are essential to create the target environment.

   ---

10. Click the **Configure** button corresponding to the compute instance to configure the Network Configuration and General Settings. This option is enabled only on the compute instance that is selected to be cloned. The Clone page displays the same configuration details regardless of the Clone Type.

See the example in Cloning Compute Instances.

11. Click **Save** to save the changes.
12. Click the **Clone** button.
13. Select Yes to confirm. Cloning will initiate.

Note: The system validates available resources before starting the cloning process. See Validating Resources

---

14. Use the Refresh button to view the status. Status will change from Initiating to Provisioning and then to Running.

Remote Clone Type for DB System

Perform remote cloning by cloning a PDB to a remote CDB, which is in another DB system. This operation too does not create a new DB system.

Note: For remote cloning, the source and destination databases can be in different compartments and in different Virtual Cloud Networks (VCN). However, there must be a peer connection between the VCNs before you remotely clone a PDB across databases in different VCNs.

This diagram illustrates cloning the PDB to an existing CDB in a remote DB system.

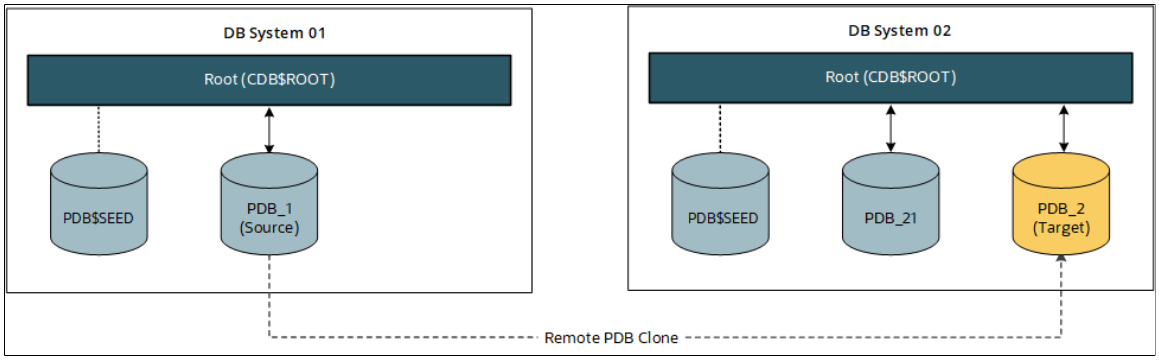

To clone an environment, with Remote Clone type, where the database is running in DBS and other nodes are on compute:

1. From the Cloud Manager Homepage, select the Environments tile.
2. Click on the Related Actions button corresponding to the DBS environment to be cloned and select Clone Environment.
3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

This example illustrates the fields and controls on the Clone page for an environment where the database is running as a service. You can find definitions for the fields and controls later on this page.

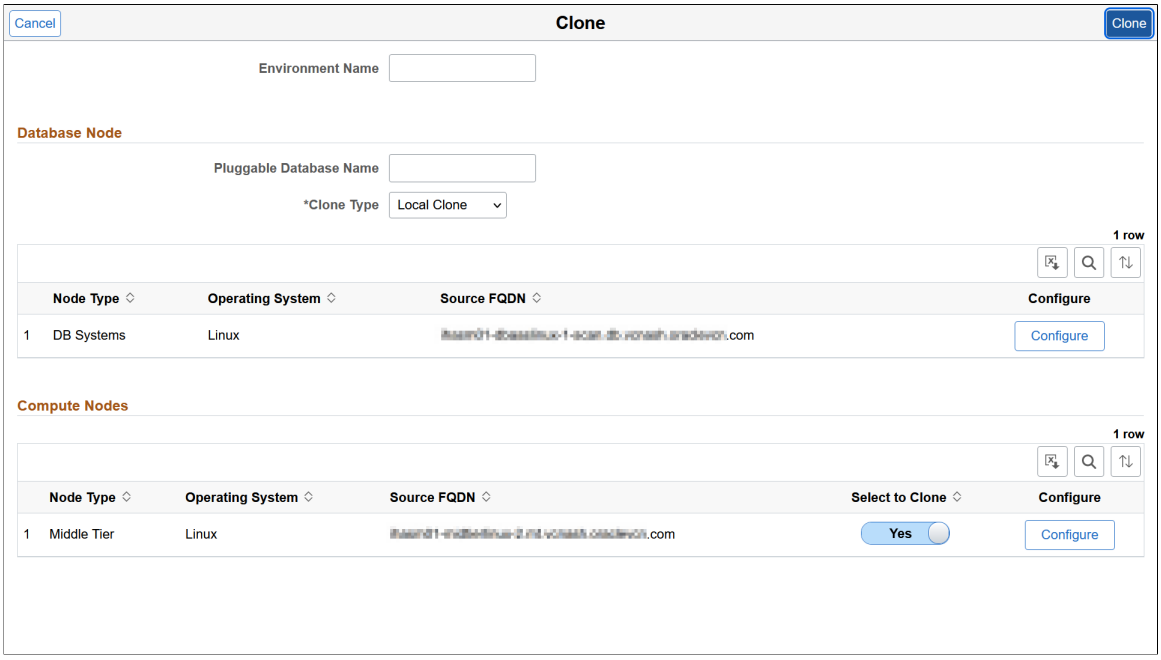

Field or Control	***Description***
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name of the pluggable database to be cloned.
Clone Type	Select *Remote Clone* type. The other available options are *Local Clone* and *New DBS*.
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment.

4. Enter a new Environment Name.
5. Enter a name for the Pluggable Database.
6. Select a Clone Type from the drop down menu.
7. Click the **Configure** button corresponding to the database instance.

This example illustrates the fields and controls on the Clone page for a DB system instance when the Clone Type is selected as Remote Clone. You can find definitions for the fields and controls later on this page.

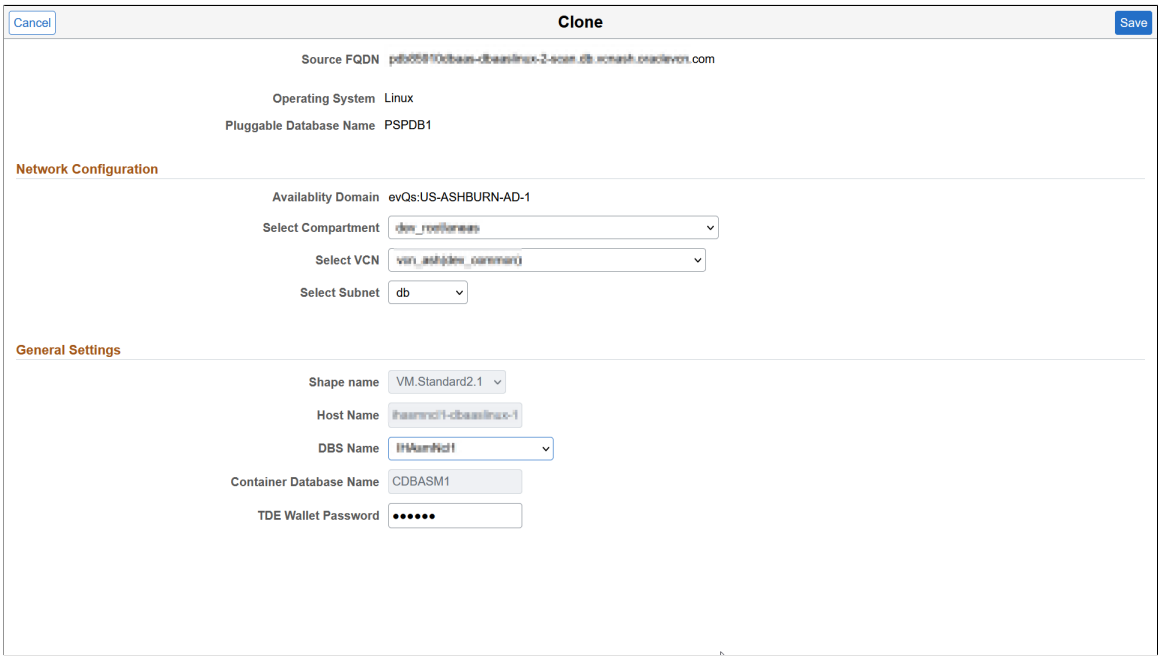

8. Make the necessary changes for the cloned environment and click **Save**.

Field or Control	***Description***
Select Compartment	Select a compartment that contains a DB system to perform remote cloning.
Select VCN	Select a VCN in the selected compartment to use when creating the remote clone.
Select Subnet	Select a subnet within the VCN to use when creating the remote clone.
DBS Name	Select the name for the DB System to create by remote cloning.
TDE Wallet Password	Enter the TDE wallet password for the specified DB system to be created.

9. Remove any instance of your choice from the target clone environment by selecting **No** in the Select to Clone instance field corresponding to that environment.

Note: The option to select the field to be cloned is disabled for database instances, because they are essential to create the target environment.

10. Click the **Configure** button corresponding to the compute instance to configure the Network Configuration and General Settings. This option is enabled only on the compute instance that is selected to be cloned. The Clone page displays the same configuration details regardless of the Clone Type.

    See the example in Cloning Compute Instances.

11. Click **Save** to save the changes.

12. Click the **Clone** button.

13. Select Yes to confirm. Cloning will initiate.

Note: The system validates available resources before starting the cloning process. See Validating Resources

14. Use the Refresh button to view the status. Status will change from Initiating to Provisioning and then to Running.

New DBS Clone Type for DB System

Create a new database environment as part of the clone operation and do a remote PDB clone in the new DB system.

This diagram illustrates creating a new DB system and a new CDB for cloning the PDB.

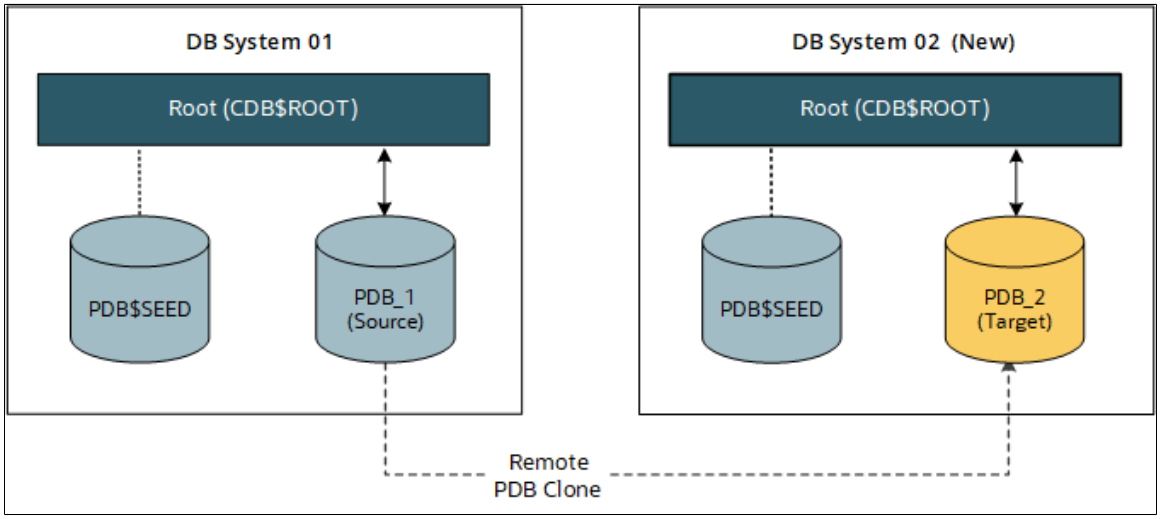

To clone an environment with New CDB clone type, where the database is running in DBS and other nodes are on compute:

1. From the Cloud Manager Homepage, select the Environments tile.

2. Click on the Related Actions button corresponding to the DBS environment to be cloned and select Clone Environment.

3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

   This example illustrates the fields and controls on the Clone page for an environment where the database is running as a service. You can find definitions for the fields and controls later on this page.

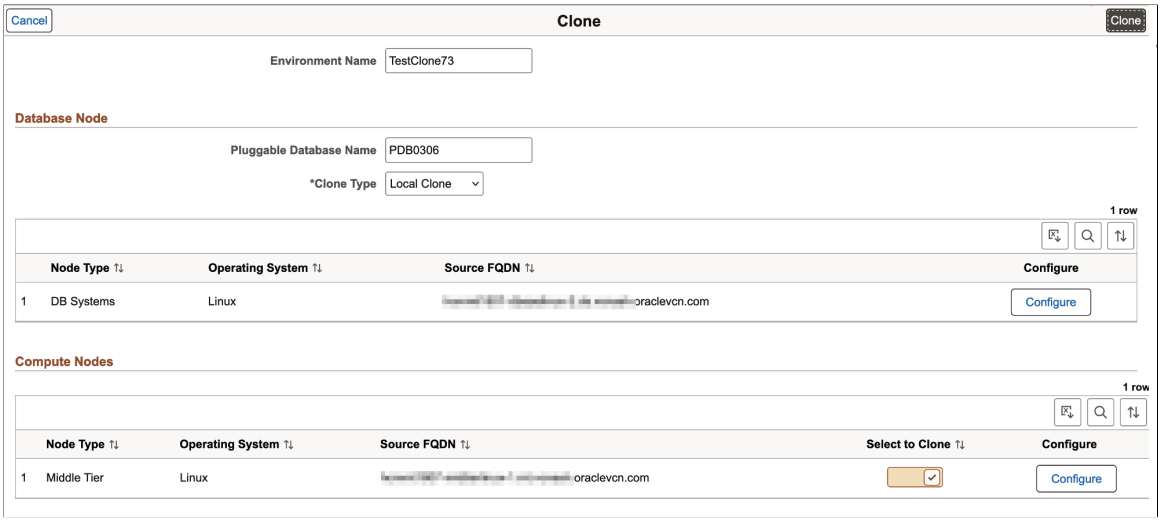

Field or Control	***Description***
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name of the pluggable database to be cloned.
Clone Type	Select *New DBS* clone type. The other available options are *Local Clone* and *Remote Clone*.
Select to Clone	Select to include an instance in the target cloned environment. Deselect the option to exclude an instance in the source environment from being cloned to the target environment.

4. Enter a new Environment Name.
5. Enter a name for the Pluggable Database.
6. Select a Clone Type from the drop down menu.
7. Click the **Configure** button corresponding to the database instance.

This example illustrates the fields and controls on the Clone page for a DB system instance when a new DB system is created for cloning. You can find definitions for the fields and controls later on this page.

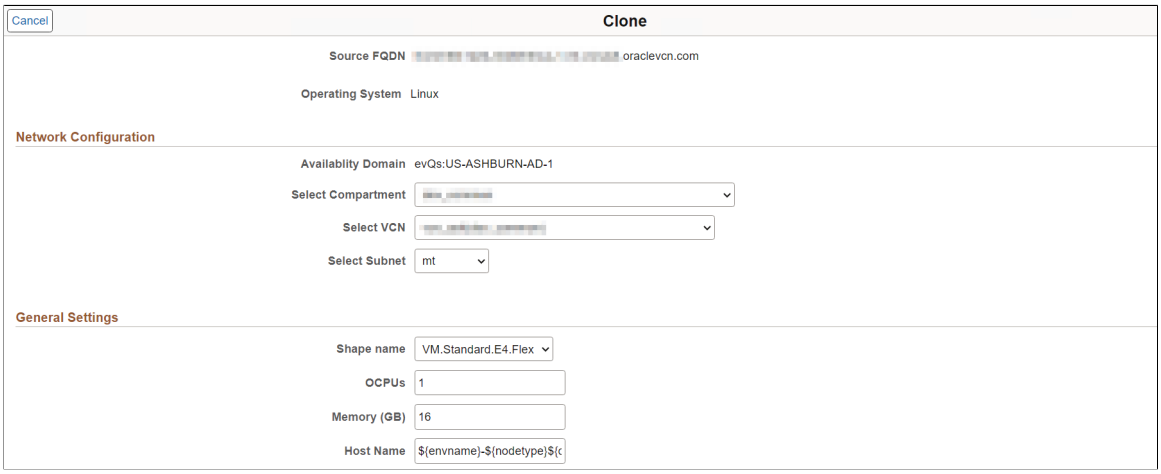

8. Make any necessary changes to the Network Configuration and General Settings, and then click **Save**.

Field or Control	***Description***
Select Compartment	Select a compartment that contains a DB system to perform remote cloning.
Select VCN	Select a VCN in the selected compartment to use when creating the remote clone.
Select Subnet	Select the subnet within the VCN to use when creating the remote clone.
Shape Name	Select the desired shape of the instance. You can select a standard or flexible shape.
OCPUs	Enter the number of Oracle CPUs required for the cloned environment. This field supports the creation of instances with flexible shapes during clone.
Memory (GB)	Enter the memory in GB required for the cloned environment. This field supports the creation of instances with flexible shapes during clone.
Host Name	Accept the default host name or enter a new host name.
Container Database Name	Enter the name of the new Container Database.

9. Remove any instance of your choice from the target clone environment by selecting **No** in the Select to Clone instance field corresponding to that environment.

   Note: The option to select the field to be cloned is disabled for database instances, because they are essential to create the target environment.

10. Click the **Configure** button corresponding to the compute instance to configure the values of Shape Name and Host Name. This option is enabled only on the compute instance that is selected to be cloned. The Clone page displays the same configuration details regardless of the Clone Type.

    See the example in Cloning Compute Instances.

11. Click **Save** to save the changes.

12. Click the **Clone** button.

13. Select Yes to confirm. Cloning will initiate.

    Note: The system validates available resources before starting the cloning process. See Validating Resources

14. Use the Refresh button to view the status. Status will change from Initiating to Provisioning and then to Running.

Cloning an Environment With Database Running on ADB and Other Nodes on Compute

To clone an environment where the database is running on Autonomous Database (ADB-Shared or ADB-Dedicated) and other nodes are on compute:

1. From the Cloud Manager Homepage, select the Environments tile.

2. Click on the Related Actions button corresponding to the DBS environment to be cloned and select Clone Environment.

3. The Clone Window displays two sections, one for the Compute Nodes and one for the Autonomous Database.

This example illustrates the fields and controls on the Clone page for an environment where the database is running in Autonomous Database-Dedicated. You can find definitions for the fields and controls later on this page.

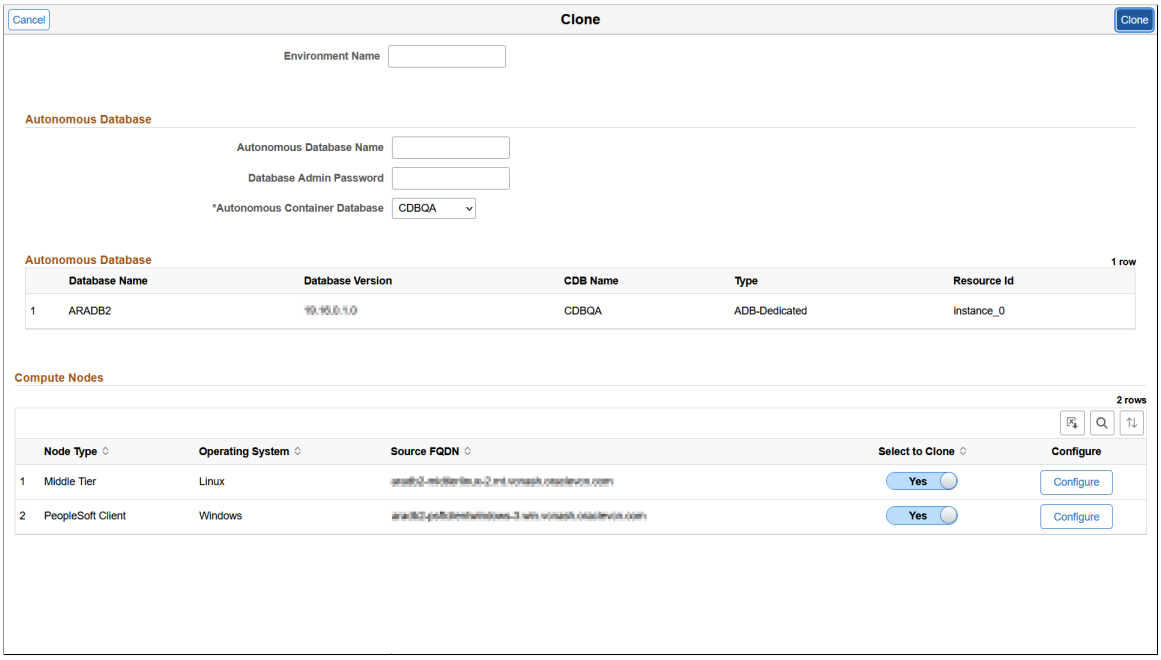

Field or Control	***Description***
Environment Name	Enter the name for the new environment.
Autonomous Database Name	Enter the name of the autonomous database to be cloned.
Database Admin Password	Enter the password for the admin user in autonomous database.
Autonomous Container Database	Select the name of the CDB that contains the autonomous database.
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment. **Note:** You can exclude any number of instances from the cloning process.

This example illustrates the fields and controls on the Clone page for an environment where the database is running in Autonomous Database-Shared. You can find definitions for the fields and controls later on this page.

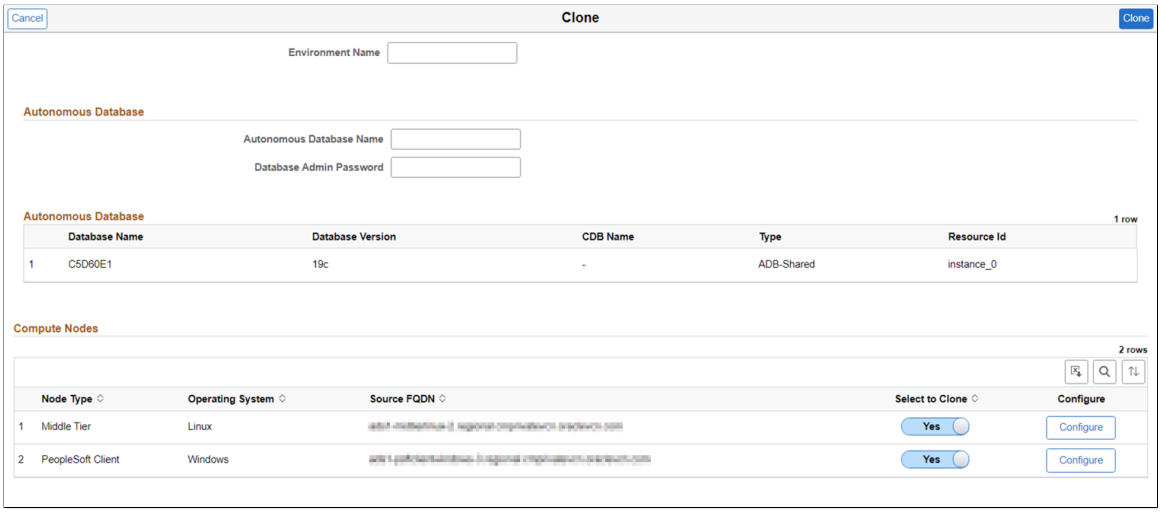

Field or Control	***Description***
Environment Name	Enter the name for the new environment.
Autonomous Database Name	Enter the name of the autonomous database to be cloned.
Database Admin Password	Enter the password for the admin user in autonomous database.
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment. **Note:** You can exclude any number of instances from the cloning process.

4. Enter a new Environment Name.
5. Enter a name for the Autonomous Database.
6. Enter the password for the admin user in autonomous database.
7. If the database type is ADB-Dedicated, select the name of the Autonomous CDB.
8. Click the **Configure** button corresponding to the compute instance to configure the values of Shape Name and Host Name.

   This option is enabled only on the compute instance that is selected to be cloned. The Clone page displays the same configuration details regardless of the database type.

   See the example in Cloning Compute Instances.

9. Click **Save** to save the changes.
10. Click the **Clone** button.
11. Select **Yes** to confirm. Cloning will initiate.

    Note: The system validates available resources before starting the cloning process. See Validating Resources

12. Use the **Refresh** button to view the status. Status will change from Initiating to Provisioning and then to Running.

Cloning an Environment With Database Running on Exadata and Other Nodes on Compute

This section describes the three options available to clone an environment where the database is running on Exadata Database Service on Dedicated Infrastructure and other nodes are on compute. You can perform a local, remote, or new CDB clone.

You may also see Exadata Database Service on Dedicated Infrastructure referred to in this and other documentation as ExaCS, or Exadata Cloud Service. See Oracle Exadata Database Service on Dedicated Infrastructure Overview.

Briefly, the Exadata Database Service hosts Exadata clusters. An Exadata cluster is a collection of container databases (CDBs) on multiple VMs (nodes). Each CDB can have multiple pluggable databases (PDBs). The OCI requirements for Exadata PDBs are as follows:

- PDB operations are supported only for Oracle Database 19c and later.
- PDBs are backed up at the CDB level, and each backup includes all the PDBs in the database. OCI does not support the creation of backups for individual PDBs.
- Restore operations are performed at the CDB level. OCI does not support restoring individual PDBs.

This section assumes that you have set up the Exadata cluster in OCI.

Local Clone Type for Exadata

When you choose the Local Clone type, the selected PDB will be cloned within the same CDB. Because you are cloning within the CDB, you do not have to provide the TDE wallet password.

This diagram illustrates cloning a PDB within the same CDB in an Exadata cluster.

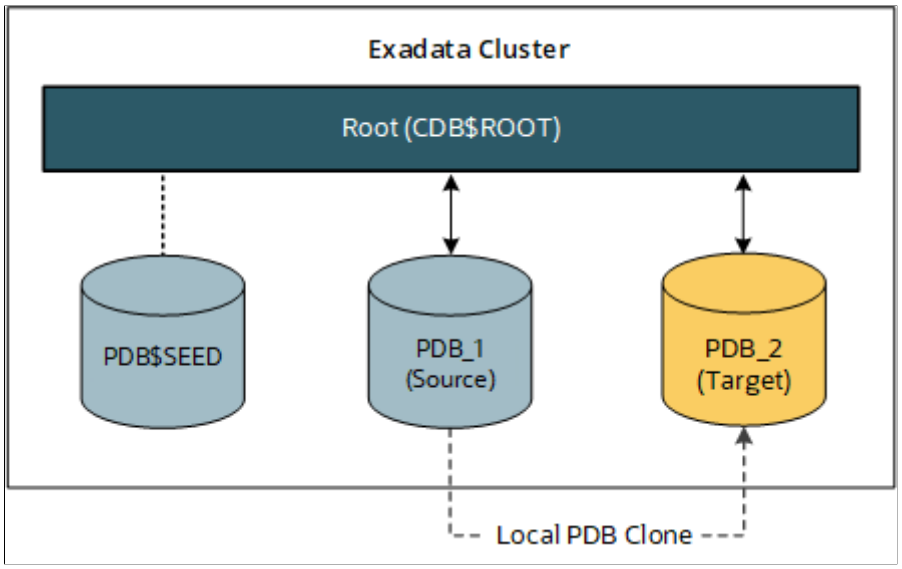

1. From the Cloud Manager Homepage, select the Environments tile.
2. Click on the Related Actions button corresponding to the Exadata environment to be cloned and select Clone Environment.
3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

   This example illustrates the fields and controls on the Clone page for an environment where the database is running in Exadata. You can find definitions for the fields and controls later on this page.

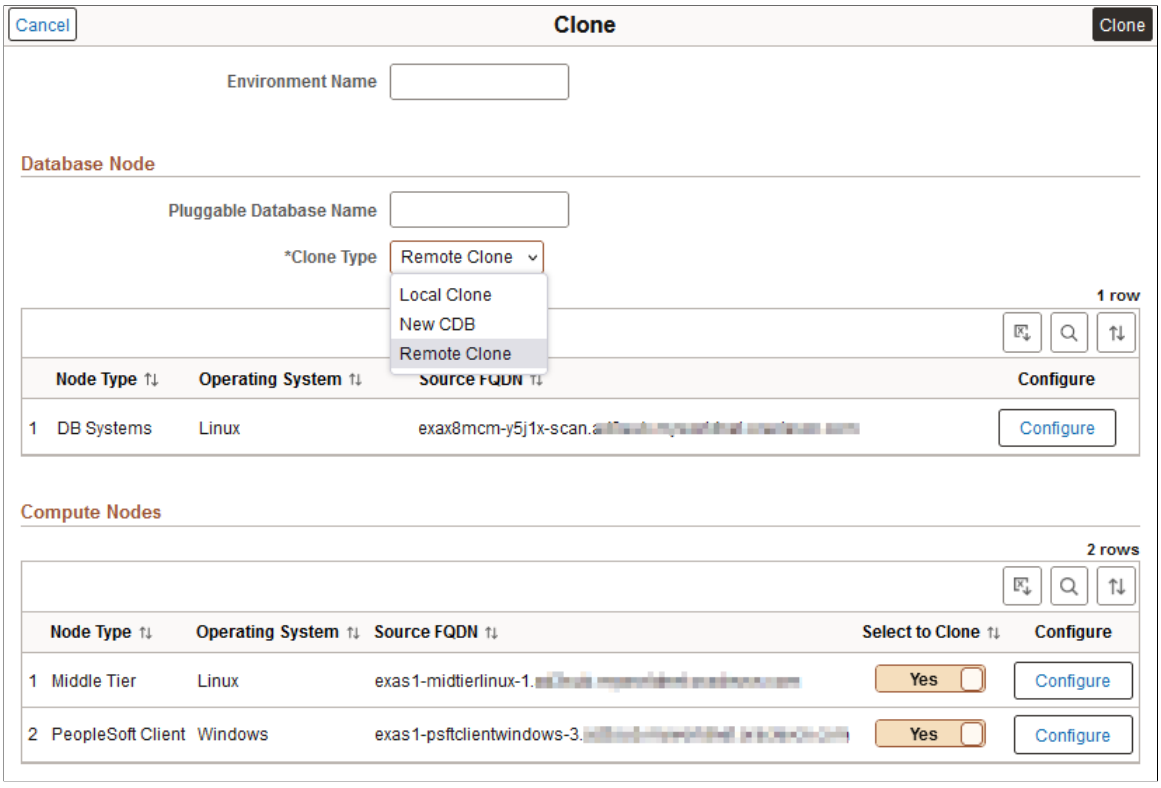

Field or Control	*Description*
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name for the new PDB in the cloned environment.
Clone Type	Select *Local Clone* as the cloning type. The other options are *Remote Clone* and *New CDB*.
Select to Clone	Select *No* to exclude an instance in the source environment from being cloned to the target environment. Select *Yes* to include the instance in the target cloned environment. **Note:** You can exclude any number of instances from the cloning process.
Configure	Click the Configure button to review or modify settings for the node.

4. Enter a new Environment Name.
5. Click the **Configure** button corresponding to the DB System node to review the Network Configuration and General Settings for the new PDB.

   When you select the Local Clone type, none of the fields are editable.

   This example illustrates the fields and controls on the Clone page for configuration for Local Clone type. You can find definitions for the fields and controls later on this page.

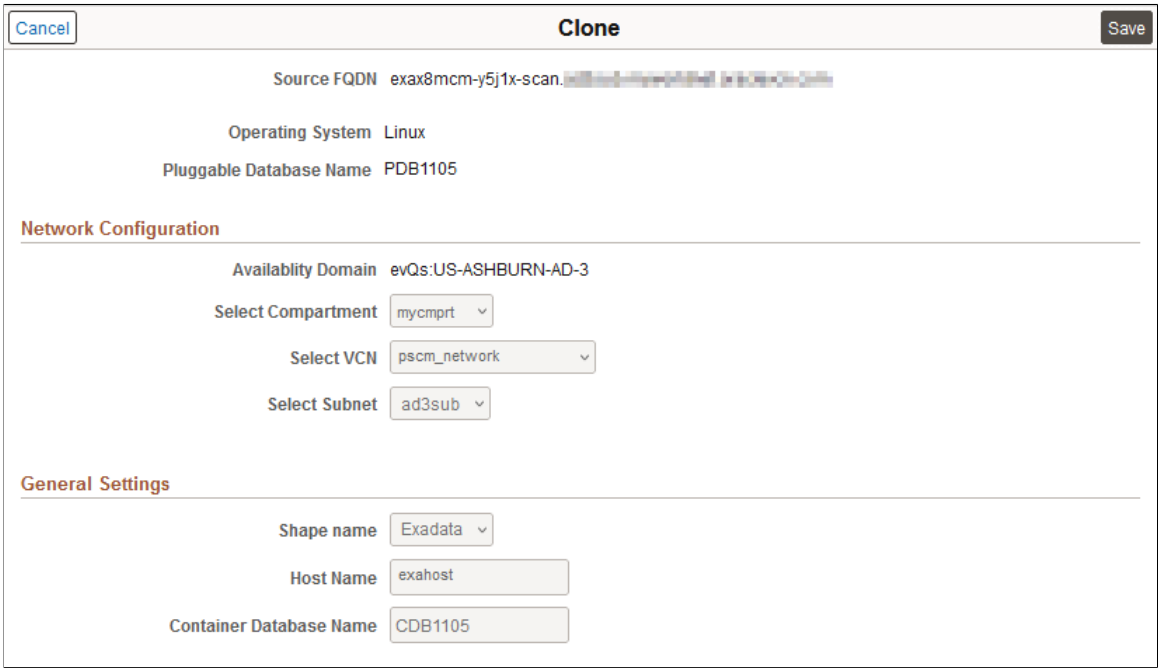

Field or Control	*Description*
Pluggable Database Name	The PDB that will be used as the source for cloning.
Network Configuration	This section lists the network objects used by the source and target PDBs. • Availability Domain • Compartment • VCN • Subnet
General Settings	This section lists the shape, host, and CDB name for the source and target PDBs.

6. Click **Save**, or Cancel to exit without saving.
7. Click the **Configure** button corresponding to the compute instance to configure networking and general settings.

   This option is enabled only on the compute instance that is selected to be cloned. The Clone page displays the same configuration details regardless of the database type.

   This example illustrates the fields and controls on the Clone page for configuration for a compute node. You can find definitions for the fields and controls later on this page.

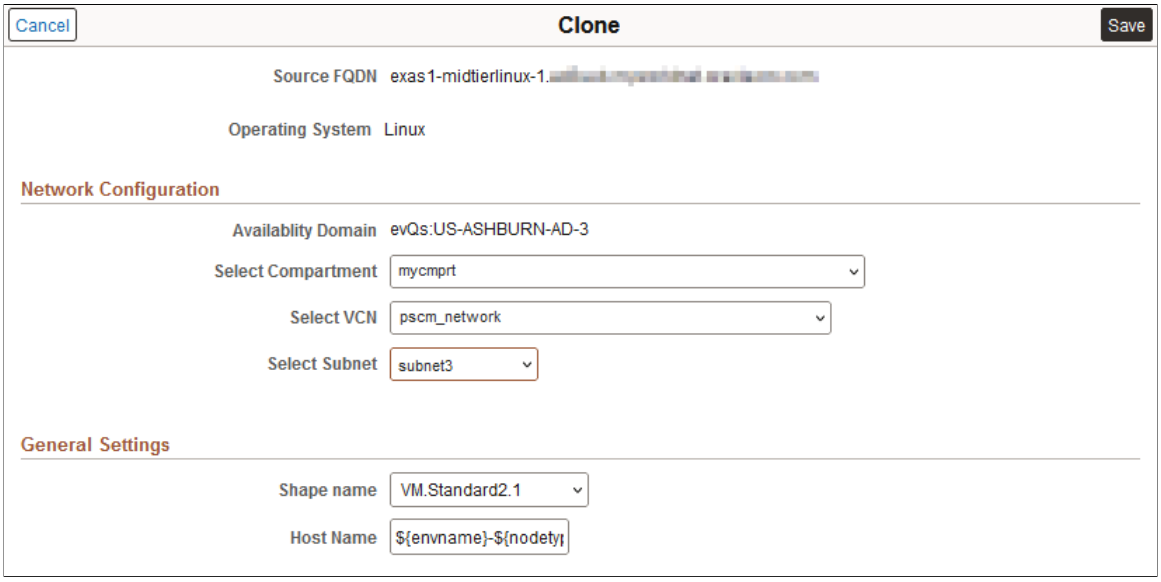

8. Make any necessary changes for the Network Configuration and General Settings for the cloned environment.

Field or Control	*Description*
Select Compartment	Select the compartment where the Virtual Cloud Network (VCN) resides.
Select VCN	Select a VCN in the selected compartment for the cloned compute instance.
Select Subnet	Select a subnet in the selected VCN for the cloned compute instance.
Shape Name	Select the desired shape of the instance.
Host Name	Accept the default host name or enter a new host name. By default, the host name will have the following format: ${envname}-${nodetype}${ostype}-${instno}, where envname stands for Environment Name and instno stands for Instance Number.

9. Click **Save**
10. Click **Clone**.
11. Select **Yes** to confirm. Cloning will initiate.

    Note: The system validates available resources before starting the cloning process. See Validating Resources

12. Use the Refresh button to view the status. Status will change from Initiating to Provisioning and then to Running.

Remote Clone Type for Exadata

To clone a PDB to a different CDB, choose the Remote Clone type and select an existing CDB within the same Exadata cluster. You will need to provide the TDE Wallet password for the existing target CDB.

This diagram illustrates cloning a PDB to a different, existing CDB within an Exadata cluster.

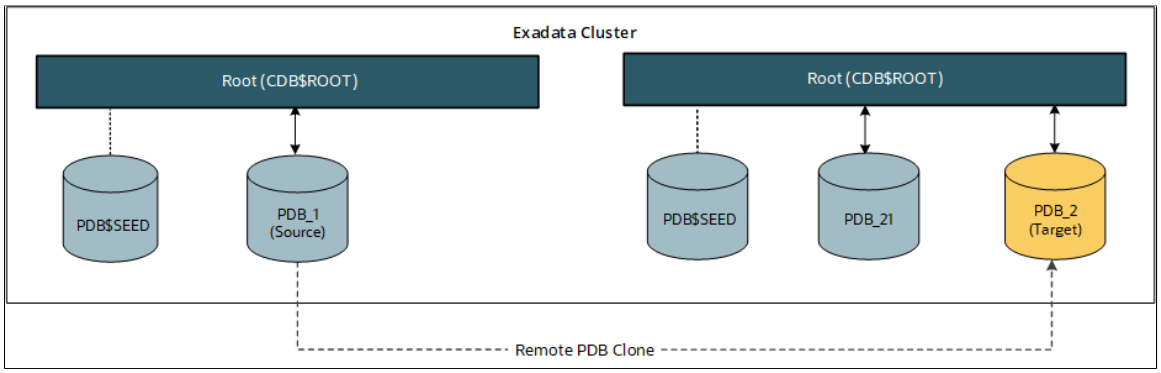

Before cloning a PDB from a source CDB to a target CDB, access the target CDB and set the open_links parameter to a non-zero value with the following commands:

1. alter system set open_links=4 scope=spfile;
2. srvctl stop database -d <CDB Unique Name> -o immediate
3. srvctl start database -d <CDB Unique Name>

See the Oracle Database documentation for details on initialization parameters and using these commands.

To clone an environment with the Remote Clone type:

1. From the Cloud Manager Homepage, select the Environments tile.
2. Click on the Related Actions button corresponding to the Exadata environment to be cloned and select Clone Environment.
3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

   See the illustration in Local Clone Type for Exadata.
4. Enter a new Environment Name.
5. Select *Remote Clone* as the cloning type.
6. Click the **Configure** button corresponding to the DB System node.

   The Network Configuration settings (compartment, VCN, and subnet) and shape are not editable for remote cloning type.

This example illustrates the fields and controls on the Clone page for configuration for Remote cloning type. You can find definitions for the fields and controls later on this page.

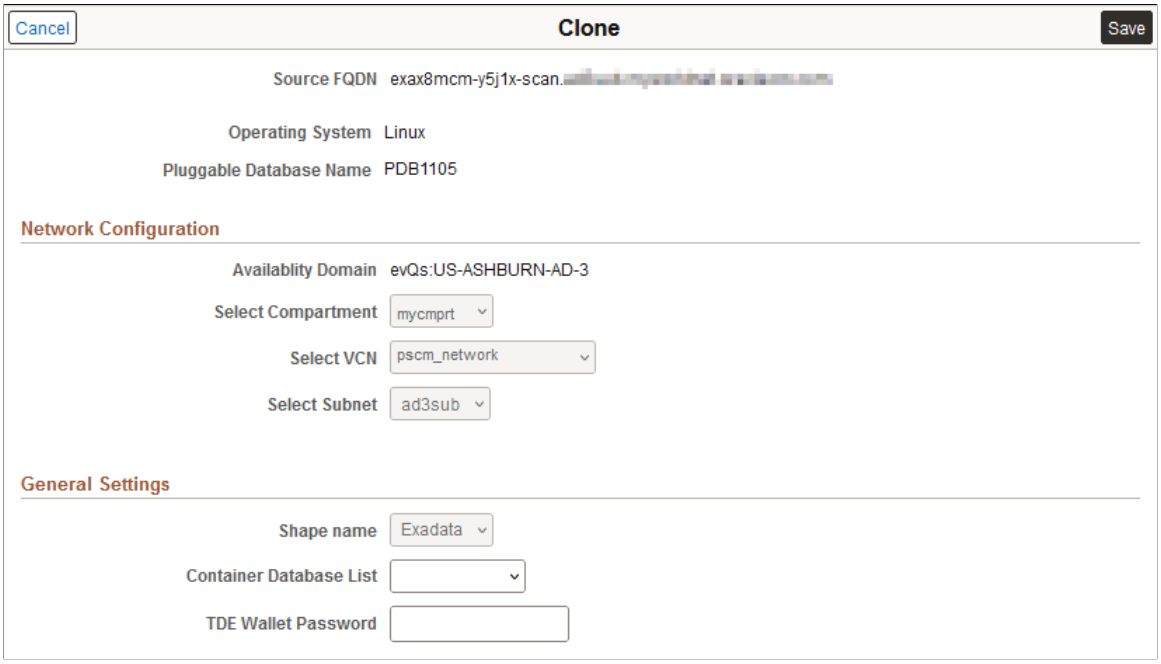

7. Make the desired changes to the General Settings and click **Save**.

Field or Control	***Description***
Container Database List	Select a CDB. The list includes the CDBs in the Exadata cluster that hosts the source CDB.
TDE Wallet Password	Specify a TDE Wallet password for the existing, target CDB. This is mandatory.

8. Click the **Configure** button corresponding to the compute instance to configure networking and general settings.

   This option is enabled only on the compute instance that is selected to be cloned. See the example in Local Clone Type for Exadata.

9. Make any desired changes to the networking configuration and general settings for the compute node and click **Save**.

10. Click the **Clone** button.

11. Select Yes to confirm. Cloning will initiate.

Note: The system validates available resources before starting the cloning process. See Validating Resources

---

12. Use the **Refresh** button to view the status. Status will change from Initiating to Provisioning and then to Running.

New CDB Clone Type for Exadata

When you choose the New CDB for Exadata, a new CDB will be created within the same Exadata cluster, and the selected PDB will be copied to it. The name you specify for the new CDB must be unique within the Exadata cluster. The new CDB will use the same CDB Administrator password as the source CDB.

This diagram illustrates cloning a PDB to a newly created CDB in an Exadata cluster.

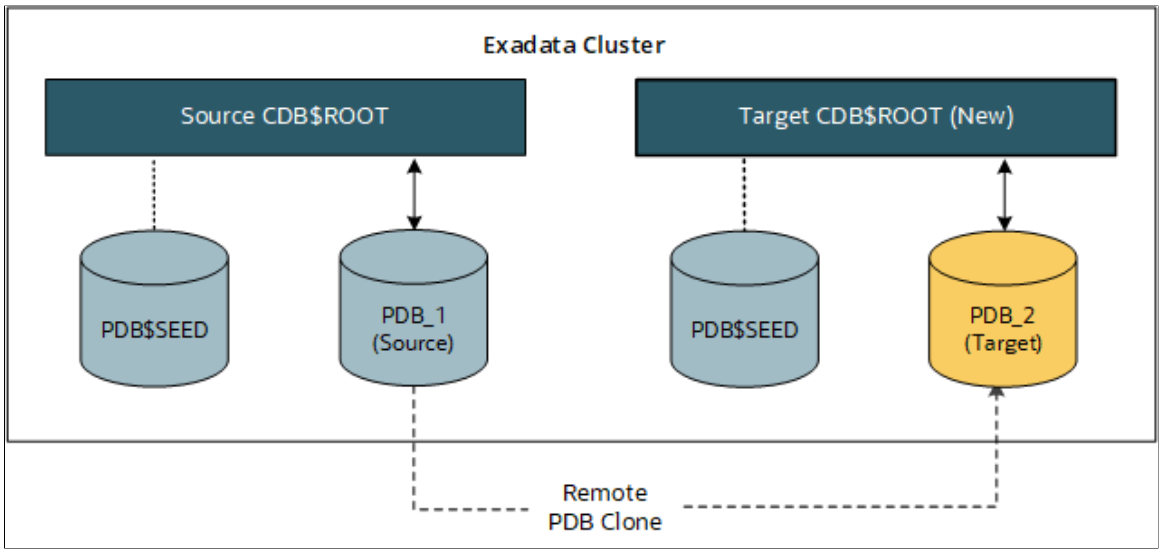

1. From the Cloud Manager Homepage, select the Environments tile.
2. Click on the Related Actions button corresponding to the Exadata environment to be cloned and select Clone Environment.
3. The Clone window displays two sections, one for the Compute Nodes and one for the Database Node.

   See the illustration in Local Cloning Type.
4. Enter a new Environment Name.
5. Select *New CDB* as the cloning type.
6. Click the **Configure** button corresponding to the DB System node.

   The Network Configuration settings are not editable for the New CDB clone type.

This example illustrates the fields and controls on the Clone page for configuration for New CDB clone type.

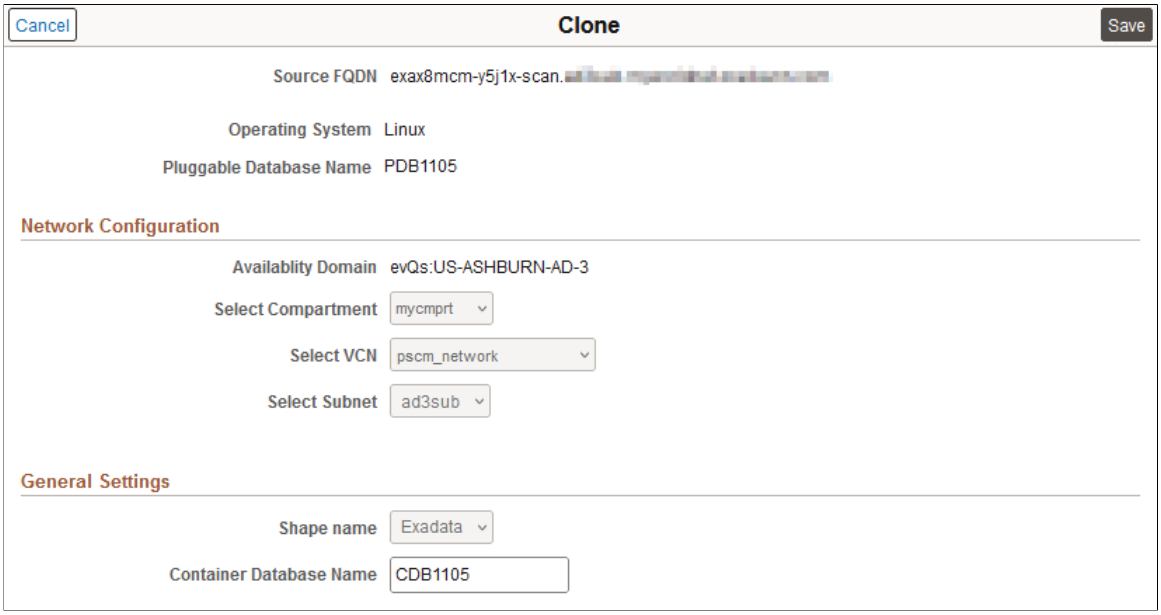

7. Enter a name for the new CDB that is unique within the Exadata cluster.
8. Click **Save**.
9. Click the **Configure** button corresponding to the compute instance to configure networking and general settings.

   This option is enabled only on the compute instance that is selected to be cloned. See the example in Local Clone Type for Exadata.
10. Make any desired changes to the networking and general settings for the compute node and click **Save**.
11. Click the **Clone** button.
12. Select Yes to confirm. Cloning will initiate.

    Note: The system validates available resources before starting the cloning process. See Validating Resources
13. Use the **Refresh** button to view the status. Status will change from Initiating to Provisioning and then to Running.

Creating Sparse Clones on Exadata Database

Exadata Sparse Clone is a native feature of Exadata that enables the creation of thinly provisioned databases for non-production purposes like development and testing. See Configuring Sparse Hierarchy Details.

After creating a test master on an Exadata database, you can create a sparse clone from the test master by enabling the Sparse Clone option on the Clone page. You can do this in the local CDB of the source environment, a remote CDB, or a new CDB.

This example illustrates the fields and controls on the Clone page for creating sparse clones.

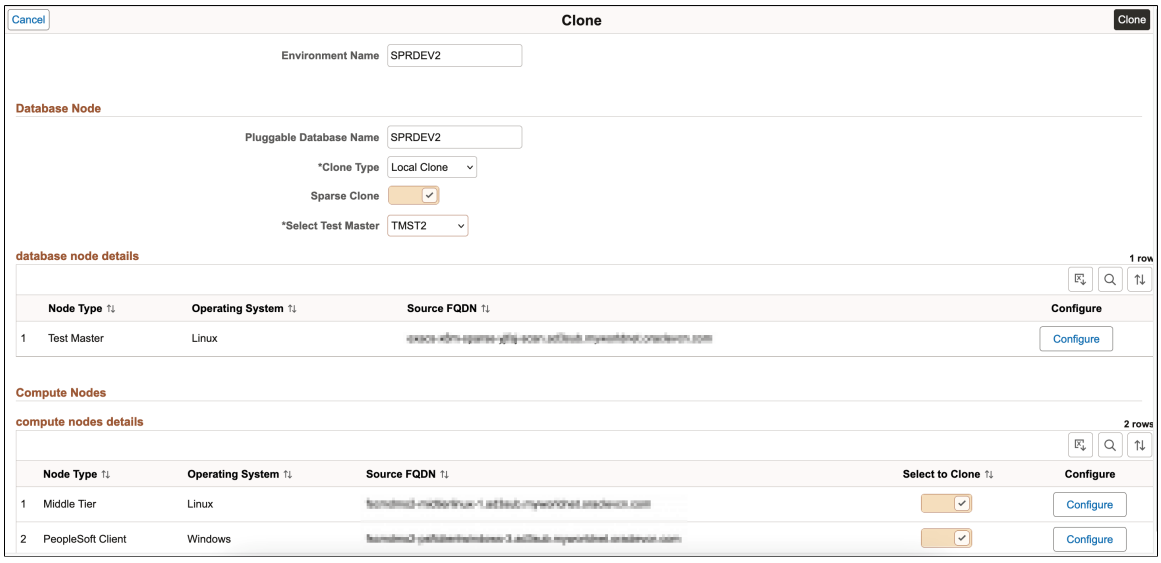

Field or Control	*Description*
Environment Name	Enter the name for the new environment.
Pluggable Database Name	Enter the name of the pluggable Exadata database to be cloned.
Clone Type	Select *Local Clone* type. The other available options are *Remote Clone* and *New DBS*.
Sparse Clone	Select to enable sparse cloning of the Exadata database.
Select Test Master	Select a test master created on the Sparse Hierarchy Details page.

Importing Environment

Cloud Manager supports importing PeopleSoft environments that are running on Oracle Cloud into Cloud Manager as a managed instance. The import functionality supports only PeopleSoft environments that were installed using DPKs. Import utilizes the components and configurations set up during the DPK installation, including Relocatable Puppet and environment variables.

You can import all of the nodes in an environment at one time, or create an imported environment and later import other compute nodes separately. If you plan to import nodes to an imported environment as a separate process, the original imported environment must include at least the Database System tier.

You can import these types of nodes:

- ADB-Dedicated database running in OCI.
- ADB-Shared database running in OCI.
- Database running on DB System (Oracle Base Database Service).

  DB Systems on Exadata and VM are supported.
- Full Tier environment running on a compute instance (Infrastructure as a Service).
- Middle Tier on a Linux compute instance (Infrastructure as a Service).
- PeopleSoft Client on a compute instance (Infrastructure as a Service).
- Search Stack on a compute instance (Infrastructure as a Service).
- Windows middle tier on a compute instance (Infrastructure as a Service).

  Windows middle tier is supported only for Process Scheduler used to run nVision.

If you import one of the database tiers you also need at least a Linux middle tier node with App Server in order to create a running environment, as well as to import other nodes without App Server (for example, PeopleSoft Client node, Windows middle tier, and so on). If you import a Full Tier node, you do not need any other nodes, but you may import a PeopleSoft Client node along with a Full Tier node.

Note: Databases created manually on Compute (VM or Bare Metal) are not supported for database tiers. Databases on Compute are supported only as part of Full Tier environments.

Prerequisites

You must fulfill the following prerequisites:

1. Get the following OCI Credentials for the components you want to import from the OCI Console:
   - Database Credentials for DB System
   - Database System OCID
   - Database OCID
   - Database Private IP

     The node for the database contains the private IP.
   - Autonomous Database OCID
   - Full Tier OCID
   - Middle Tier OCID

- PeopleSoft Client OCID
- Search server OCID (OpenSearch or Elasticsearch)
- Windows middle tier OCID

2. Verify requirements for users and groups.

   - DPK installations that were performed by non-root users are not supported for import.
   - DPK installations that were performed by a single user (other than root) are not supported for import.

   The Add Node form is populated with default values for the Windows administrative user, Linux user profiles and groups. After you supply the OCID for the node and click Discover, Cloud Manager crawls the instance and compares the information on users and groups with the default values. If a value is missing or incorrect, you see an error message. You must correct the mistake in the form before you can continue.

   Here are the default values:

   - Default OPC user
   - Default DPK user profiles for Linux -— psadm1, psadm2, psadm3, oracle2.
   - Default DPK Linux groups — oinstall, psft, appinst, dba
   - Default user profiles for Search Stack — esuser

   It is possible to import an environment having custom users and groups that replace the default DPK Linux users and groups. The number of custom users and groups must be the same as the default. The users and groups must be distinct; you cannot reuse a single user or single group for all the default values. For example, if the custom Application Runtime User is psadm222, then psadm222 cannot be used as the PeopleSoft Install User or Application Install User. Similarly, any custom groups must be distinct from other groups in the environment.

3. Copy Cloud Manager SSH public key to all the nodes that will be imported.

   a. Log in to the Cloud Manager instance, for example with Putty, as psadm2.

      ```
      sudo su - psadm2
      ```

      Open and copy the pub key from /home/psadm2/psft/data/cloud/ocihome/keys/cm_adm_pvt_key.pub.

   b. Log on to the Database system as opc, access .ssh/authorized_keys and paste the key from step a into it. This is required to authenticate Cloud Manager to access the Database environment to import the database system.

   c. Log on to each additional node to be imported (Middle Tier, PeopleSoft Client, Search server, Windows middle tier), access .ssh/authorized_keys and paste the key from step a into it.

4. Make sure that the TNS entry is present in tnsnames.ora on the database system you are importing.

5. Drive D: is mandatory for a Windows client.
6. Oratab entry for db home is required in this format <DB_UNIQUE_NAME>:<DB_HOME_PATH>.
7. Add an entry in the /etc/fstab file for the secondary volume (block volume) for each node.

   This is required in order for the secondary volume to be automatically mounted for custom created Linux virtual machines.

You must ensure the following prerequisites are met for importing Exadata into Cloud Manager:

1. If the number of nodes is more than 1, use a private IP of the last node. This applies for the RAC DB system as well.
2. Update the TNS entry for the imported database in $ORACLE_HOME/network/admin/tnsnames.ora.
3. Cloud Manager uses "opc" and "oracle" users for all operations in Exadata.

Importing Nodes for an Entire Environment

To import all of the nodes in an environment at the same time:

1. From the Cloud Manager homepage, click the Environments tile.
2. Click the **Import Environment** button at the top of the Environments page.

   This example illustrates the fields and controls on the Import environment page. You can find definitions for the fields and controls later on this page.

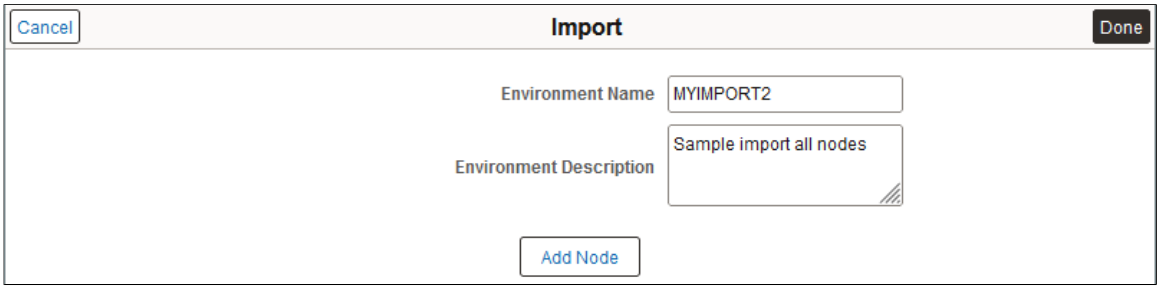

Field or Control	*Description*
Environment Name	Enter a name for the new environment.
Environment Description	Enter a description for the new environment. Ensure that the description does not include any double quotation marks (").
Add Node	Select to add a node to the environment.

3. Click the **Add Node** button.
4. Select the instance type from the drop down list. These types of tiers can be imported:

- ADB-Dedicated
- ADB-Shared
- Database System
- Full Tier

  If you import a Full Tier node, no other nodes are needed for a running environment.
- Middle Tier
- PeopleSoft Client
- Search Stack
- Windows MT (Middle Tier)

5. The Add Node page for the instance type is displayed.

   See the following sections for information on the instance types.
6. Enter the values for the instance type and click **OK**.
7. To add another instance type, click the + icon and select the instance type.

   Note: Add each instance type separately.
8. Once you have included all the nodes for the environment, click **Done**.

   This example illustrates the fields and controls on the Import page with multiple nodes.

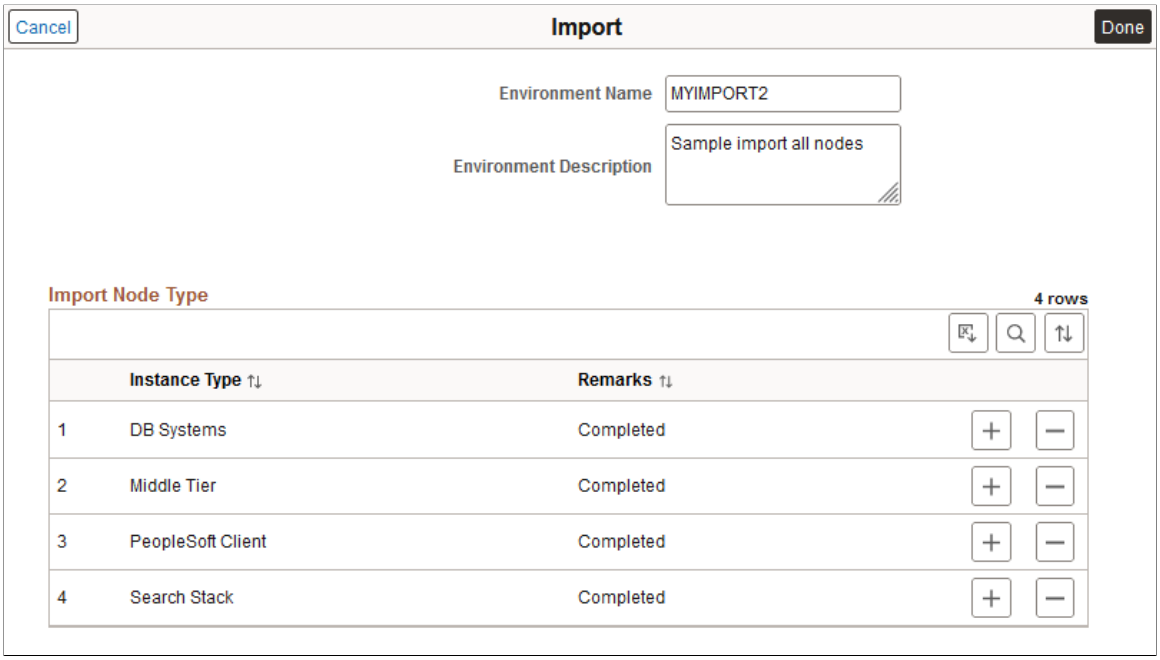

9. The respective card is added to the Environments page with the status of *ImportingMultipleTier*.

10. To view the status of the import process, select Details from the actions menu for the environment, then select Import from the left-side menu.

    See Environment Details - Import Status in this section.

Importing Nodes to an Existing Imported Environment

To continue importing nodes after creating an imported environment:

1. Complete the process to import an environment with at least the Database System node.
2. Select the actions menu for the running imported environment and select **Import Node**.
3. On the Import page, click **Add Node**.

   Note: The Environment Name and Environment Description fields are not editable.

   The nodes (instance types) that you add must be for components that were originally associated with the same database in the selected environment.
4. Select the instance type from the drop-down list.
   - Middle Tier
   - PeopleSoft Client
   - Search Stack
   - Windows MT (Middle Tier)
5. The Add Node page for the instance type is displayed.

   See the following sections for information on the instance types. The fields on the pages change as appropriate for the instance type. Illustrations are given only for a couple of instance types.
6. Enter the values for the instance type and click **OK**.
7. Click **Done** to begin the import process.
8. On the Environments page, the status on the environment card indicates that it is importing the instance, such as *ImportingMiddleTier*, or *ImportingWinClient*.
9. To view the status of the import process, select Details from the actions menu for the environment, then select Import from the left-side menu.

   See Environment Details - Import Status in this section.
10. On the Import page, the steps associated with the original import, such as for the database node, will display as complete. The steps associated with importing the subsequent nodes will display as in progress.
11. To add another instance type, repeat steps 2 through 10.

Note: Complete the import process for each instance type separately.

Database System Instance Type

For the database system node, all the values are mandatory.

This example illustrates the fields and controls on the Add Node page for DB Systems. You can find definitions for the fields and controls later on this page.

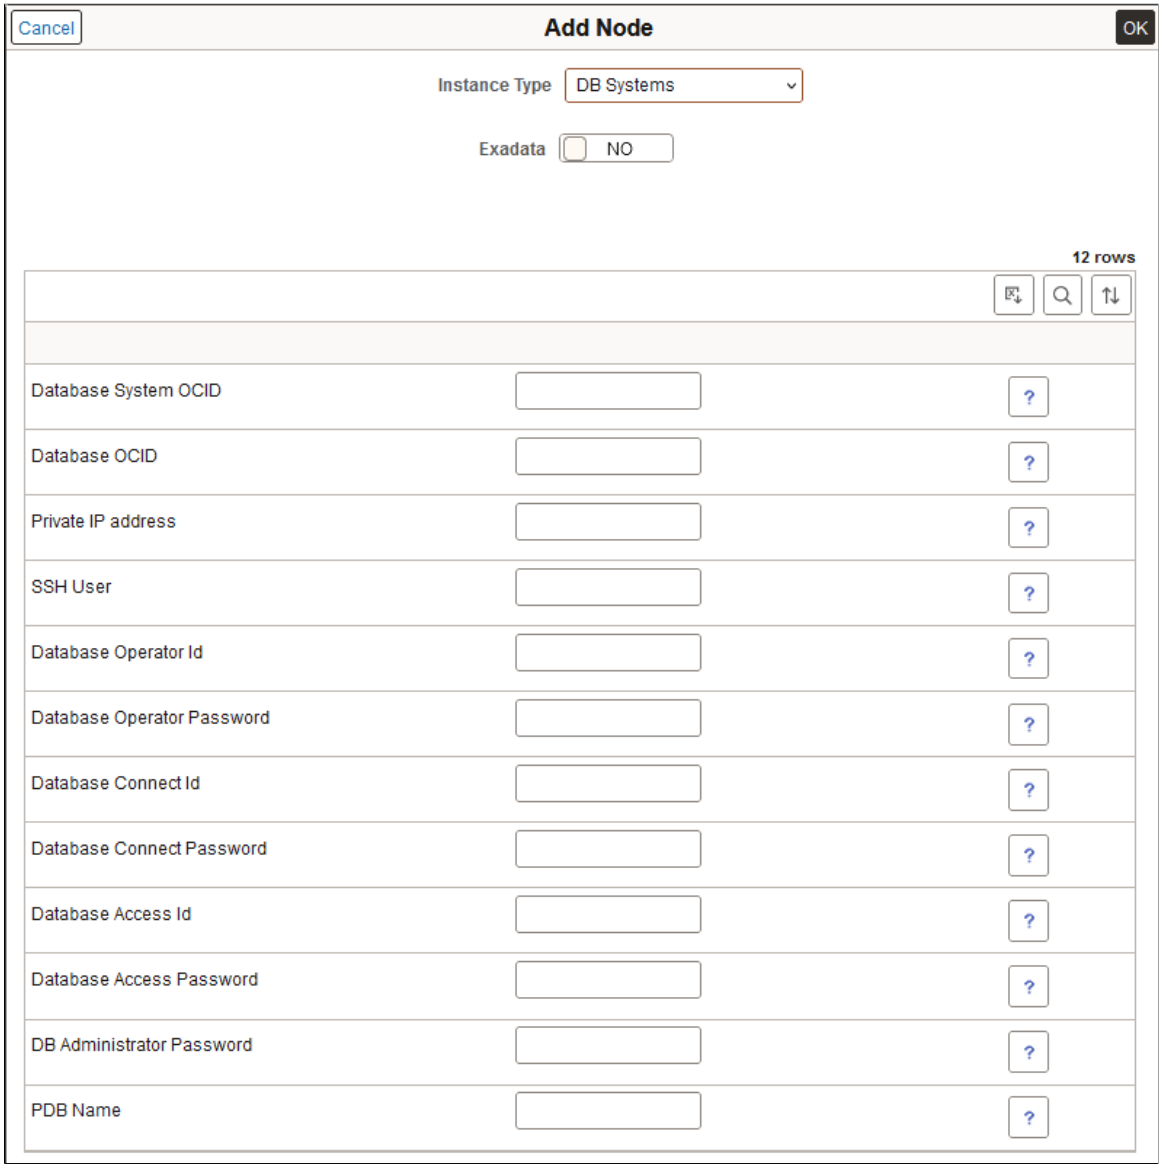

Field or Control	*Description*
Exadata	Select Yes if the DB System is Exadata. An additional field for Container Database Name will be added.

Field or Control	*Description*
Database System OCID	Database System OCID for the target database. For Exadata, if the new resource model consisting of separate infrastructure and VM cluster resources is used, enter the VM Cluster ID.
Database OCID	Database OCID for the target database.
Private IP Address	Private IP address for the target the Database environment.
SSH User	SSH user on the database system being imported
Database Operator ID	Database Operator ID
Database Operator Password	Database Operator Password
Database Connect ID	Database Connect ID
Database Connect Password	Database Connect Password
Database Access ID	Database Access ID
Database Access Password	Database Access Password
DB Administrator Password	Database Administrator Password
PDB Name	Pluggable Database Name This is the database name in the tnsnames.ora file.
Container Database Name	For Exadata DB systems enter the container database name.

ADB Instance Types

The same fields are necessary for ADB-Dedicated or ADB-Shared node, and all the values are mandatory.

This example illustrates the fields and controls on the Add Node page for ADB-Dedicated node. You can find definitions for the fields and controls later on this page.

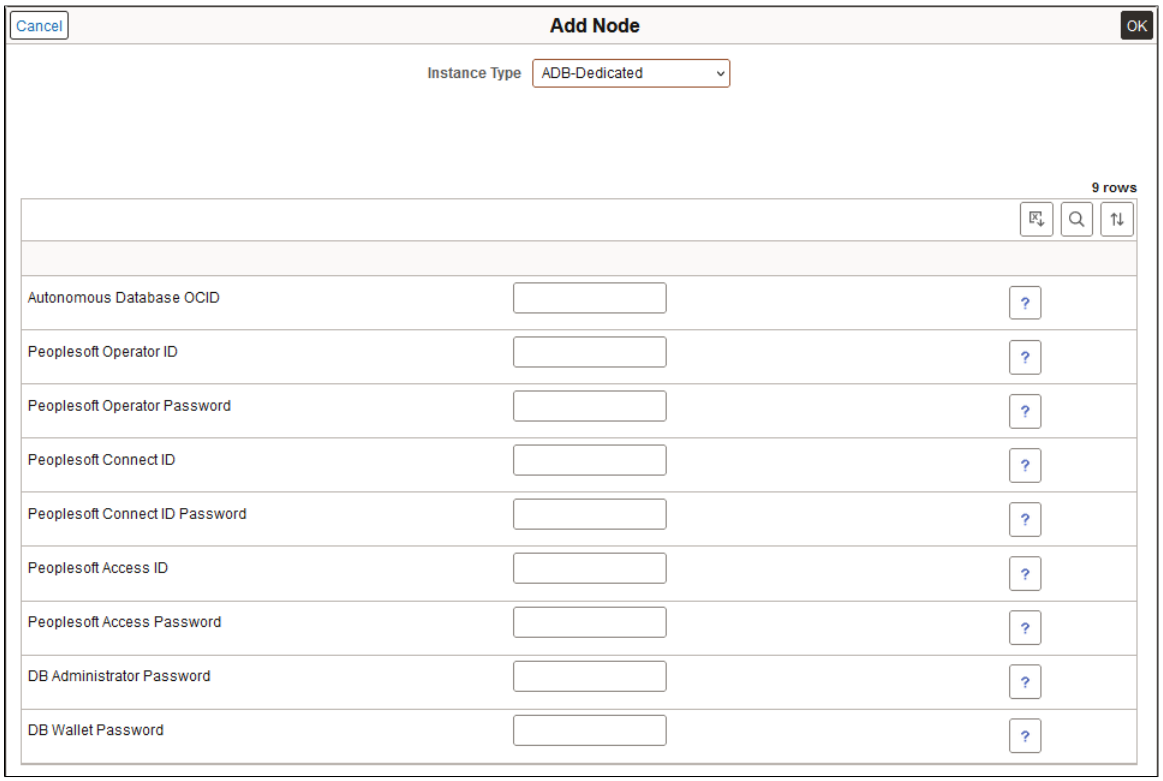

Field or Control	*Description*
Autonomous Database OCID	Autonomous Database OCID for the target database.
PeopleSoft Operator ID	Database Operator ID
PeopleSoft Operator Password	Database Operator Password
PeopleSoft Connect ID	Database Connect ID
PeopleSoft Connect Password	Database Connect Password
PeopleSoft Access ID	Database Access ID
PeopleSoft Access Password	Database Access Password
DB Administrator Password	Database Administrator Password
DB Wallet Password	DB Wallet Password The DB Wallet Password can be any password that adheres to the password policy for wallets for an ADB.

For imported ADB-Dedicated and ADB-Shared nodes, the ADB Wallet Refresh Policy is added to the environment.

This example illustrates the Policies page for an environment with imported ADB-Dedicated or ADB-Shared database node.

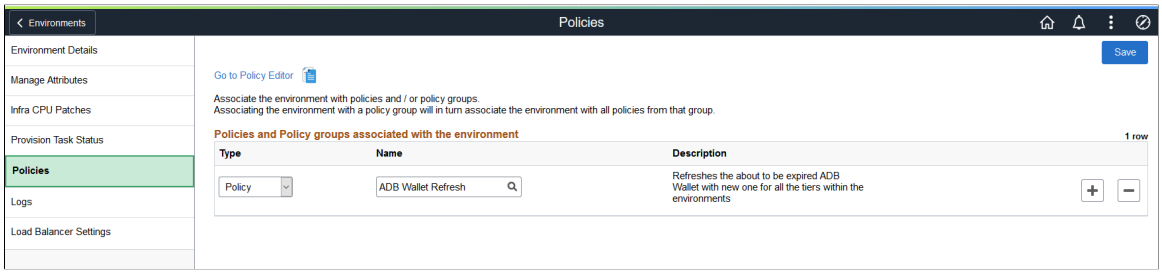

By default the Expiry Day Count is set to 540 (meaning the wallet expires after 540 days). Use the Go to Policy Editor link to change the Expiry Day Count. Select Parameters under Policy Actions.

Full Tier Instance Type

To import a full-tier node.

Note: For information on users and groups, see the product documentation *PeopleSoft PeopleTools <Release> Deployment Packages Installation* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html. Select the PeopleTools release for your environment at the top.

---

1. Click the + (Add) icon on the Import page.
2. Select Full Tier Instance Type.

   The fields are populated with default values for the instance users and groups.
3. Enter the Full Tier OCID.
4. Click **Discover**.

This example illustrates the fields and controls on the Add Node page for Full Tier instance type, Instance Users and Groups [Step 1 of 3].

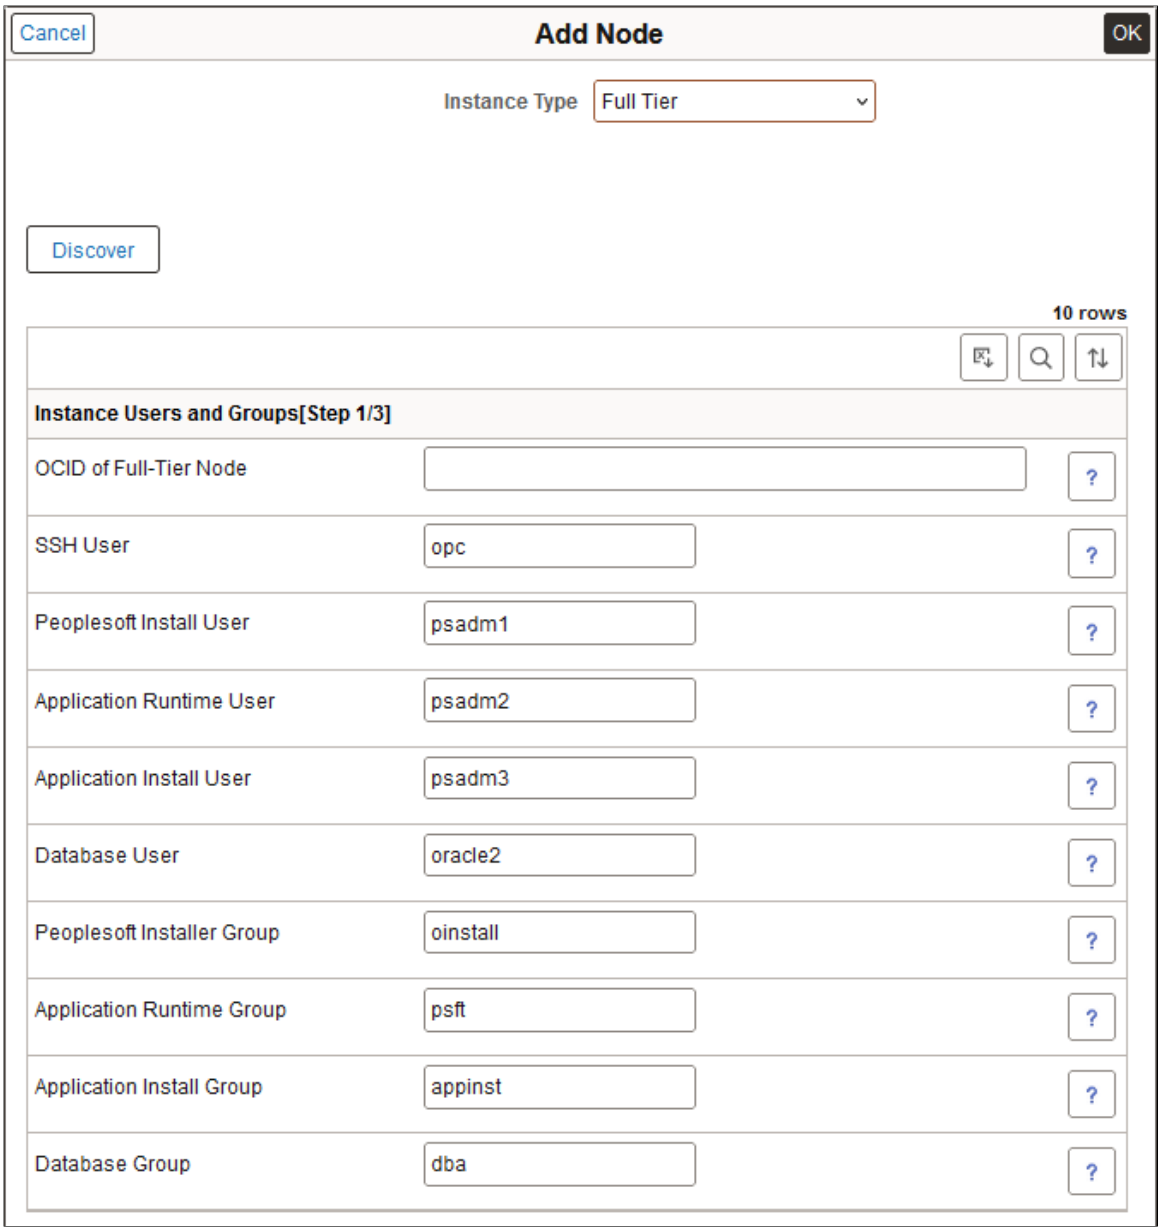

Field or Control	*Description*
OCID of Full-Tier Node	The OCID for the instance in Oracle Cloud Infrastructure console.
SSH User	User that accesses the instance with SSH. The default is opc. If the SSH entry is incorrect, you see an error message and the Discover process will be unable to proceed.

Field or Control	*Description*
Peoplesoft Install User	Linux user with install permission. The default is psadm1. Custom users must be distinct from other PeopleSoft users.
Application Runtime User	Linux user with permissions for application directories. The default is psadm2. Custom users must be distinct from other PeopleSoft users.
Application Install User	Linux user with permissions to install application. The default is psadm3. Custom users must be distinct from other PeopleSoft users.
Database User	The default is oracle2. Custom users must be distinct from other PeopleSoft users.
Peoplesoft Installer Group	The default is oinstall. Custom groups must be distinct from other PeopleSoft groups.
Application Runtime Group	The default is psft. Custom groups must be distinct from other PeopleSoft groups.
Application Install Group	The default is appinst. Custom groups must be distinct from other PeopleSoft groups.
Database Group	The default is dba. Custom groups must be distinct from other PeopleSoft groups.

5. Review the entries and make any necessary corrections.

   The Discover process will connect with SSH to the full tier instance and compare the information on users and groups with the default values. If a value is missing or incorrect, you see an error message. Correct the mistake and click **Discover** again to verify the change.

6. Based on the servers that are present, a pop-up message will appear listing the servers that were discovered. Click OK.

7. Cloud Manager lists the discovered environment paths.

Change the values to match the paths on your environment, if necessary. All of the paths are mandatory except Peoplesoft Cust Path. A blank field indicates that nothing was found. You can enter a valid location for a blank field.

This example illustrates the fields and controls on the Add Node page for Full Tier instance type, Peoplesoft Paths [Step 2 of 3].

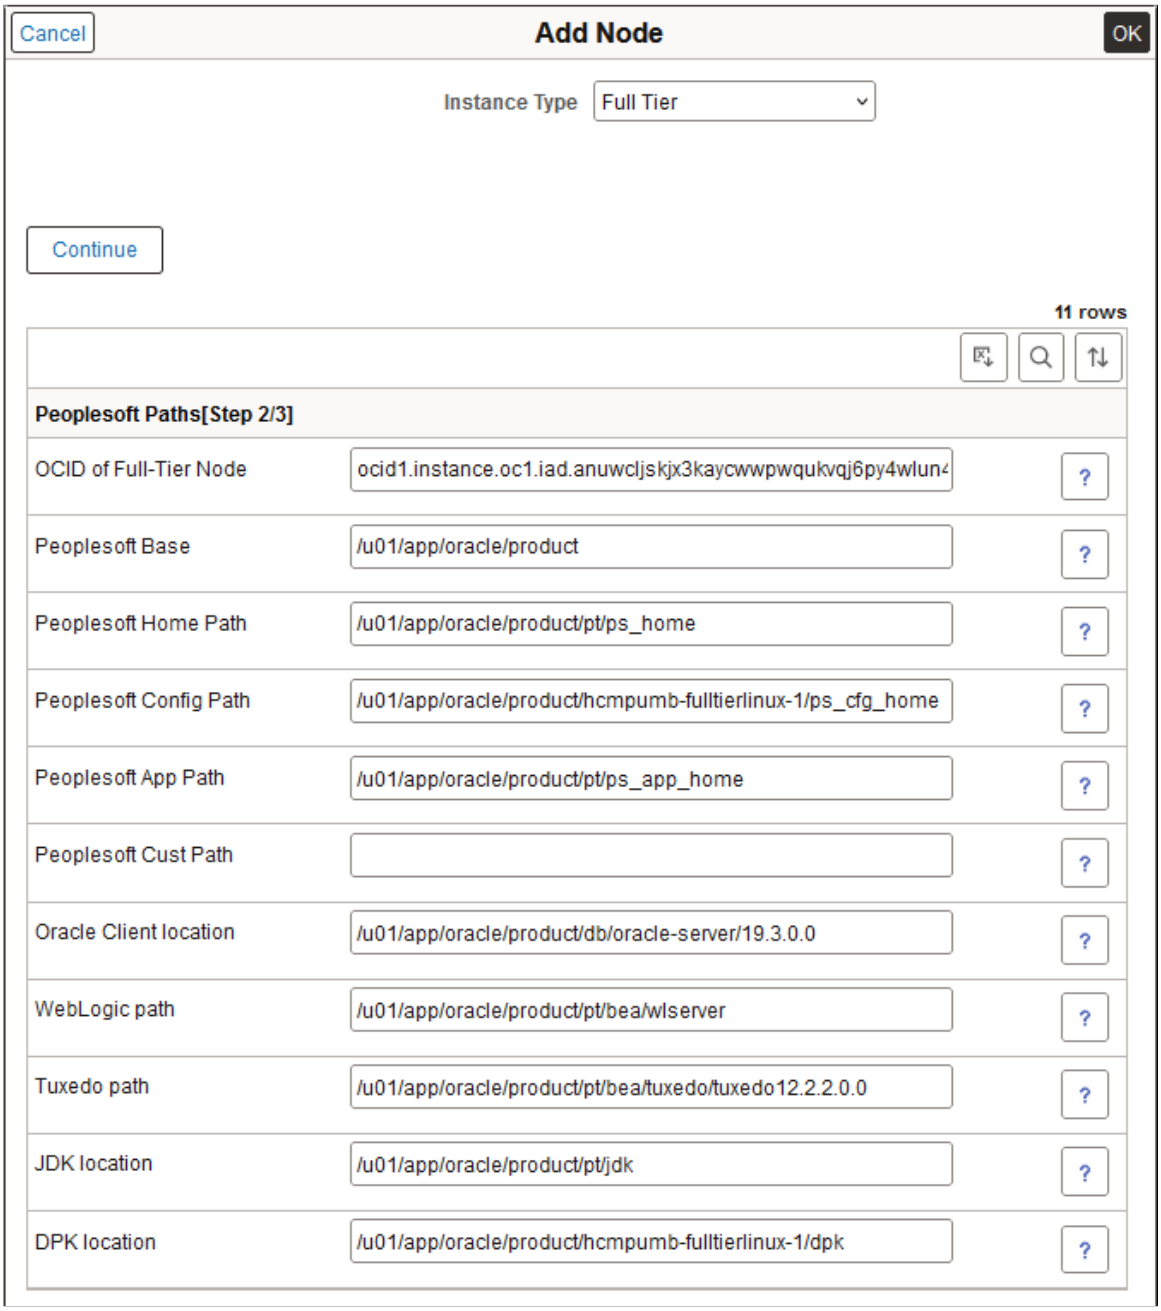

Field or Control	*Description*
Peoplesoft Base	The top-level installation directory; for example, /u01/app/oracle/product.

Field or Control	*Description*
Peoplesoft Home Path	PS_HOME location; for example, /u01/app/oracle/product/pt/ps_home.
Peoplesoft Config Path	PS_CFG_HOME location; for example, /u01/app/oracle/product/<*instance_name*>/ps_cfg_home.
Peoplesoft App Path	PS_APP_HOME location; for example, /u01/app/oracle/product/pt/ps_app_home.
Peoplesoft Cust Path	PS_CUST_HOME location.
Oracle Client location	Oracle Database Client location; for example, /u01/app/oracle/product/db/oracle-server/19.3.0.0.
WebLogic path	Oracle WebLogic location; for example, /u01/app/oracle/product/bea/wlserver.
Tuxedo path	Oracle Tuxedo location; for example, /u01/app/oracle/product/bea/tuxedo/tuxedo12.2.2.0.0.
JDK location	Java JDK location; for example, /u01/app/oracle/product/pt/jdk.
DPK location	DPK location; for example, /u01/app/oracle/product/<*instance_name*>/dpk.

8. Click **Continue**.
9. Enter the credentials for the Full Tier.

This example illustrates the fields and controls on the Add Node page for Full Tier instance type, Peoplesoft Details [Step 3 of 3].

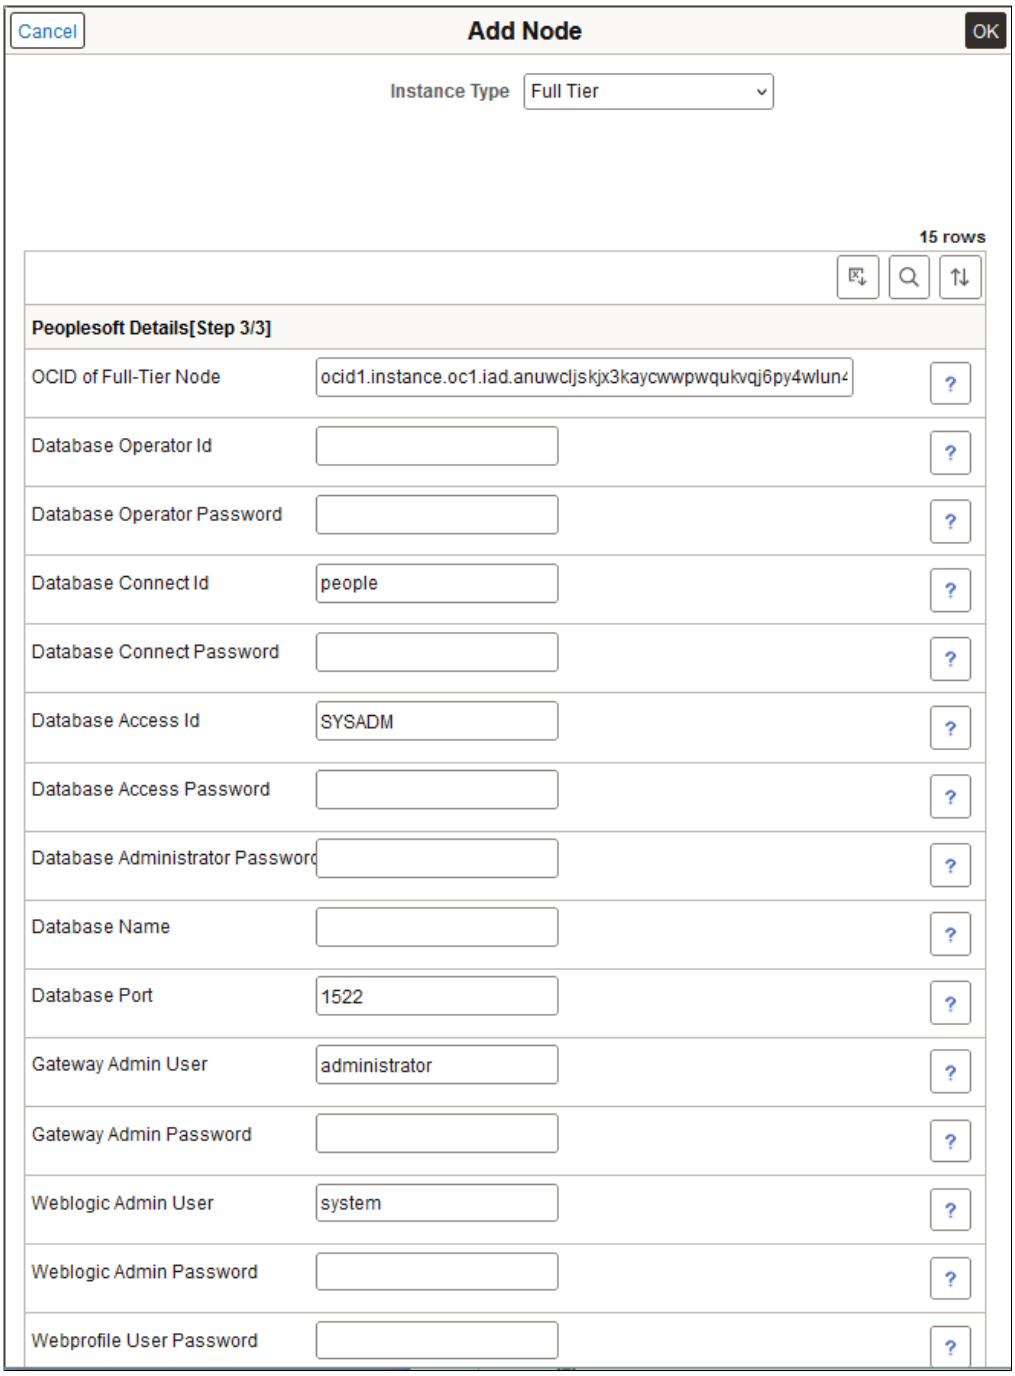

Field or Control	*Description*
Database Operator Id	The user ID that accesses the environment.
Database Operator Password	Enter the password associated with the environment.

Field or Control	*Description*
Database Connect Id	The default is people.
Database Connect Password	Enter the password associated with the environment.
Database Access Id	The default is SYSADM.
Database Access Password	Enter the password associated with the environment.
Database Administrator Password	Enter the password associated with the environment.
Database Name	Enter the name for the database.
Database Port	The default is 1522.
Gateway Admin User	The default is administrator.
Gateway Admin Password	Enter the password associated with the environment.
Weblogic Admin User	The default is system.
Weblogic Admin Password	Enter the password associated with the environment.
Webprofile User Password	Enter the password associated with the environment.

Middle Tier Instance Type

To import a Middle Tier node.

Note: For information on middle tier requirements, see the product documentation *PeopleSoft PeopleTools <Release> Deployment Packages Installation* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html. Select the PeopleTools release for your environment at the top.

---

1. Click the + (Add) icon on the Import page.
2. Select Middle Tier Instance Type.

   The fields are populated with default values for the instance users and groups.
3. Enter the Middle Tier OCID.
4. Click **Discover**.

Field or Control	*Description*
OCID of Mid-Tier Node	The OCID for the instance in Oracle Cloud Infrastructure console.
SSH User	User that accesses the instance with SSH. The default is opc. If the SSH entry is incorrect, you see an error message and the Discover process will be unable to proceed.
Peoplesoft Install User	Linux user with install permission. The default is psadm1. Custom users must be distinct from other PeopleSoft users.
Application Runtime User	Linux user with permissions for application directories. The default is psadm2. Custom users must be distinct from other PeopleSoft users.
Application Install User	Linux user with permissions to install application. The default is psadm3. Custom users must be distinct from other PeopleSoft users.
Database User	The default is oracle2. Custom users must be distinct from other PeopleSoft users.
Peoplesoft Installer Group	The default is oinstall. Custom groups must be distinct from other PeopleSoft groups.
Application Runtime Group	The default is psft. Custom groups must be distinct from other PeopleSoft groups.
Application Install Group	The default is appinst. Custom groups must be distinct from other PeopleSoft groups.
Database Group	The default is dba. Custom groups must be distinct from other PeopleSoft groups.

5. Review the entries and make any necessary corrections.

The Discover process will connect with SSH to the middle tier instance and compare the information on users and groups with the default values. If a value is missing or incorrect, you see an error message. Correct the mistake and click **Discover** again to verify the change.

The Discover process will connect with SSH to the middle tier instance and find the servers (application server, web server, or Process Scheduler) that are present, the PeopleSoft deployment path, and whether COBOL is enabled or not.

6. Based on the servers that are present, a pop-up message will appear listing the servers that were discovered. Click OK.

7. Cloud Manager lists the discovered environment paths.

   Change the values to match the paths on your environment, if necessary. All of the paths are mandatory except Peoplesoft Cust Path. A blank field indicates that nothing was found. You can enter a valid location for a blank field.

Field or Control	***Description***
Peoplesoft Base	The top-level installation directory; for example, /u01/app/oracle/product
Peoplesoft Home Path	PS_HOME location; for example, /u01/app/oracle/product/pt/ps_home
Peoplesoft Config Path	PS_CFG_HOME location; for example, /u01/app/oracle/product/<*instance_name*>/ps_cfg_home
Peoplesoft App Path	PS_APP_HOME location; for example, /u01/app/oracle/product/pt/ps_app_home
Peoplesoft Cust Path	None
Oracle Client location	Oracle Database Client location; for example, /u01/app/oracle/product/db/oracle-server/19.3.0.0
WebLogic path	Oracle WebLogic location; for example, /u01/app/oracle/product/bea/wlserver
Tuxedo path	Oracle Tuxedo location; for example, /u01/app/oracle/product/bea/tuxedo/tuxedo12.2.2.0.0
JDK location	Java JDK location; for example,/u01/app/oracle/product/pt/jdk
DPK location	Location for the DPKs; for example, /u01/app/oracle/product/<*instance_name*>/dpk

8. Click **Continue**.

9. Enter the credentials for the Middle Tier, then click OK.

Field or Control	*Description*
Database Operator Id	The user ID that accesses the environment.
Gateway Admin User	The default is administrator.
Gateway Admin Password	Enter the password associated with the environment.
Weblogic Admin User	The default is system.
Weblogic Admin Password	Enter the password associated with the environment.
Webprofile User Password	Enter the password associated with the environment.

PeopleSoft Client Instance Type

To import a PeopleSoft Client node:

1. Click the + (Add) icon on the Import page.
2. Select PeopleSoft Client Instance Type.
3. Enter the required information.

   This example illustrates the fields and controls on the Add Node page for PeopleSoft Client, Step 1/2.

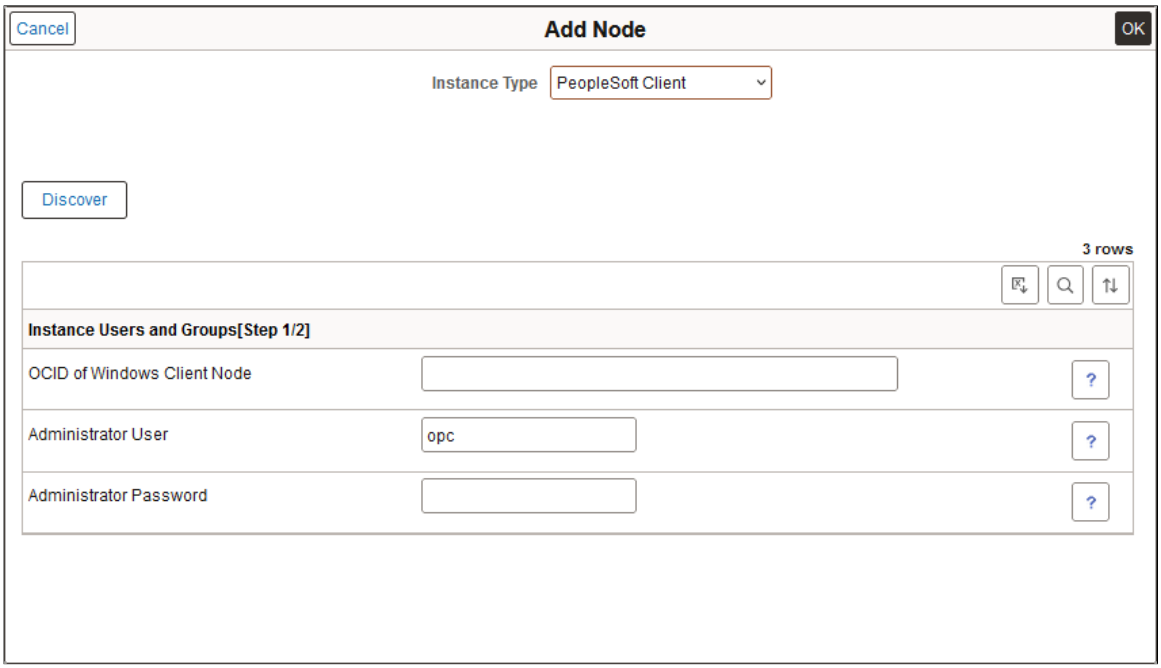

Field or Control	*Description*
OCID of Windows Client Node	The OCID for the instance in Oracle Cloud Infrastructure console.
Administrator User	The default user is opc. If you created an instance with a different user, enter that here. Cloud Manager does not validate the Administrator User, but if the value is incorrect, the import process will fail.
Administrator Password	The password for the Windows administrator.

4. Click **Discover.**

   Cloud Manager lists the discovered environment path.

   This example illustrates the fields and controls on the Add Node page for PeopleSoft Client, Step 2/2.

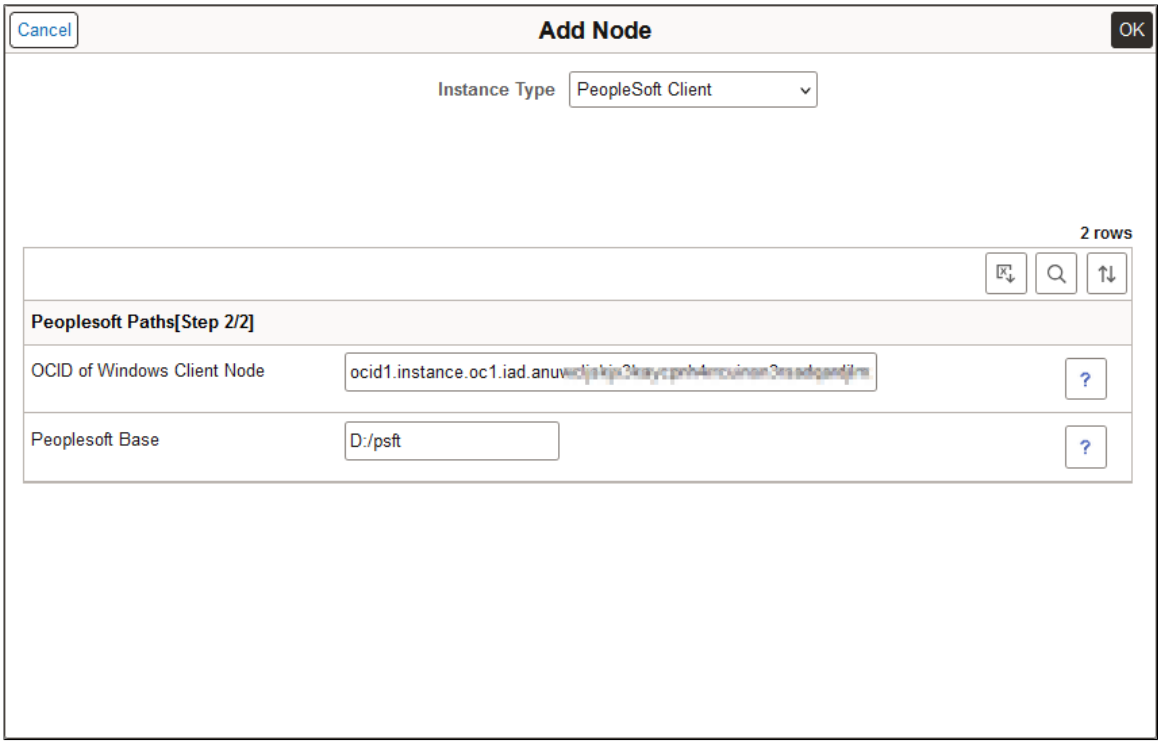

Field or Control	*Description*
Peoplesoft Base	The full path to the base directory that was used in deploying the environment. The default is D:/psft. Drive D: is mandatory for a PeopleSoft (Windows) Client.

5. Review the entries and make any necessary corrections.

The Discover process will connect with SSH to the PeopleSoft Client instance and compare the information on installation paths with the default values. If a value is missing or incorrect, you see an error message. Correct the mistake and click **Discover** again to verify the change.

6. Click **OK**.

Windows MT Instance Type

To import a Windows middle tier (MT) node:

1. Click the + (Add) icon on the Import page.
2. Select Windows MT Instance Type.
3. Enter the required information.

Field or Control	*Description*
OCID of Windows MT	The OCID for the instance in Oracle Cloud Infrastructure console.
Administrator User	The default user is opc. If you created an instance with a different user, enter that here. Cloud Manager does not validate the Administrator User, but if the value is incorrect, the import process will fail.
Administrator Password	The password for the Windows administrator.

4. Click **Discover.**

   Cloud Manager lists the discovered environment paths.

Field or Control	*Description*
Peoplesoft Base	The full path to the base directory that was used in deploying the environment. The default is D:/psft. Drive D: is mandatory for a Windows MT node.
Peoplesoft Home Path	The full path to PS_HOME, which holds the PeopleTools files. For example, D:/psft/pt/ps_home8.61.
Peoplesoft Config Path	The full path to PS_CFG_HOME, which holds the configuration files for the application server, Process Scheduler, and web server (PIA) domains.
Tuxedo path	The full path to the installation directory for Oracle Tuxedo. For example, D:/psft/pt/bea/tuxedo.

Field or Control	*Description*
WebLogic path	The full path to the installation directory for Oracle WebLogic. For example, D:/psft/pt/bea.
JDK location	The full path to the installation directory for Java JDK. For example, D:/psft/pt/jdk.
DPK location	The full path to the directory holding the DPKs used in deploying the environment.

5. Review the entries and make any necessary corrections.

   The Discover process will connect with SSH to the Windows MT instance and compare the information on installation paths with the default values. If a value is missing or incorrect, you see an error message. Correct the mistake and click **Discover** again to verify the change.

6. Click **OK**.

Search Stack Instance Type

To import a Search Stack node:

Note: For information on search stack requirements, see the product documentation *PeopleSoft Deployment Packages Installation for Search Components* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html. Select the PeopleTools release for your environment at the top.

---

1. Click the + (Add) icon on the Import page.
2. Select Search Stack Instance Type.

   The fields are populated with default values for the instance users and groups.

3. Enter the Search Stack node OCID.
4. Click **Discover**.

Field or Control	*Description*
OCID of Search Stack Node	The OCID for the instance in Oracle Cloud Infrastructure console.
SSH User	User that accesses the instance with SSH. The default is opc. If the SSH entry is incorrect, you see an error message and the Discover process will be unable to proceed.
Open Search User	The default is esuser.

Field or Control	*Description*
Open Search Group	The default is esuser.
Wheel Group	This is a native Linux group that is required for the search stack. If there is no wheel group on your environment, you need to add it before importing. The default is wheel.

5. Review the entries and make any necessary corrections.

   The Discover process will connect with SSH to the search stack instance and compare the information on users and groups with the default values. If a value is missing or incorrect, you see an error message. Correct the mistake and click **Discover** again to verify the change.

6. Click **Continue**.

   Based on the servers that are present, a pop-up message will appear listing the servers that were discovered. Click OK.

7. Cloud Manager lists the discovered environment paths.

   Change the values to match the paths on your environment, if necessary. All of the paths are mandatory. A blank field indicates that nothing was found. You can enter a valid location for a blank field.

Field or Control	*Description*
Installation Directory	The top-level installation directory; for example, /u01/app/oracle/product/.
Search Location	The installation location for Elasticsearch or OpenSearch; for example, /u01/app/oracle/product/es/pt/elasticsearch7.10.0
Search Dashboard Location	The installation location for Kibana or OpenSearch Dashboards; for example, /u01/app/oracle/product/es/pt/Kibana7.10.0.

8. Click **Continue**.

9. Enter the credentials for the search stack.

Field or Control	*Description*
Administrator User	The default is esadmin.
Administrator Password	Enter the password for the Elasticsearch or OpenSearch Administrator.

Field or Control	*Description*
Proxy User	The default is people.
Proxy Password	Enter the password for the proxy user.
Cluster Name	The default is ESCL.
Discovery Host Name	The default is 127.0.0.1.
Port	The default is 9200.

Viewing the Import Status

After you initialize the import process, you can check the status on the Import page. In case the process fails, you can view the errors, make corrections, and retry.

To access the Import page:

1. The respective card is added to the Environments page with the status of Importing.
2. To view the status of the import, select Details from the actions menu for the environment, or click the card.
3. Select Import from the left-side menu.

   The list of steps on the status page include discovery and crawling steps for the imported nodes, as well as environment creation and post configuration.

If the steps succeed, the target Database and other imported nodes become a managed instance under Cloud Manager.

This example illustrates the fields and controls on the Import Database Instance page with completed steps.

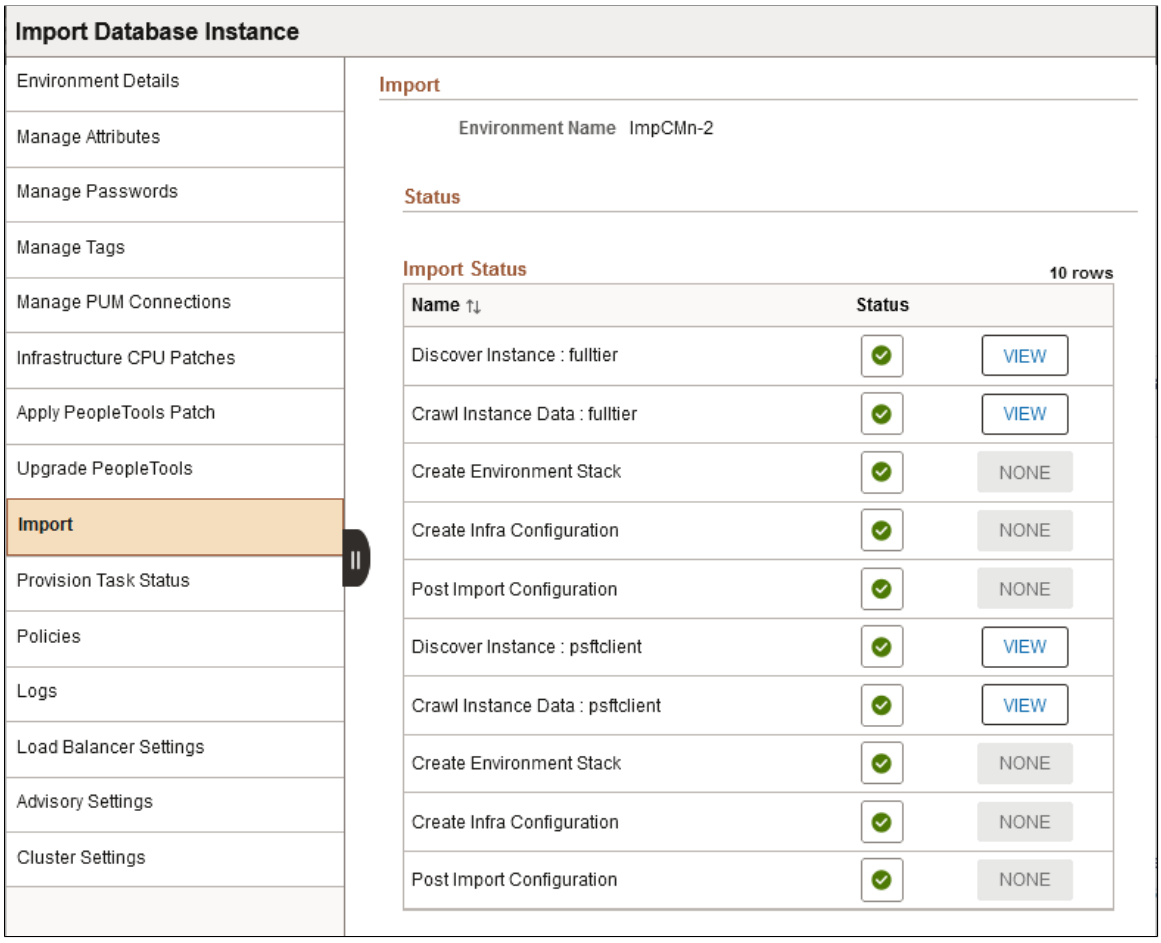

Field or Control	*Description*
View	Click to display the Import Task Data page with the input and derived data for that particular step. See Viewing the Import Task Data page. The View button is available for both successful and failed steps.
Edit	The **Edit** button appears beside the View button when a step fails. Click to edit the input data and retry the step. See Retrying the Import Process for more information.
None	Indicates that the step does not contain any output.
Continue	The **Continue** button appears if a step fails. See Retrying the Import Process for more information.

Field or Control	*Description*
	Pending
	Success
	In Progress
	Failed
	Continue
	Abort
	Step details
	Pause

Viewing the Import Task Data Page

Click the View button beside an import step to review the values. The Input Data section at the top of the page lists data such as users, paths, database name, and so on, which were supplied for the import process. See the previous sections on the various instance types for details.

This example illustrates the top portion of the Import Task Data page, showing the Input Data section.

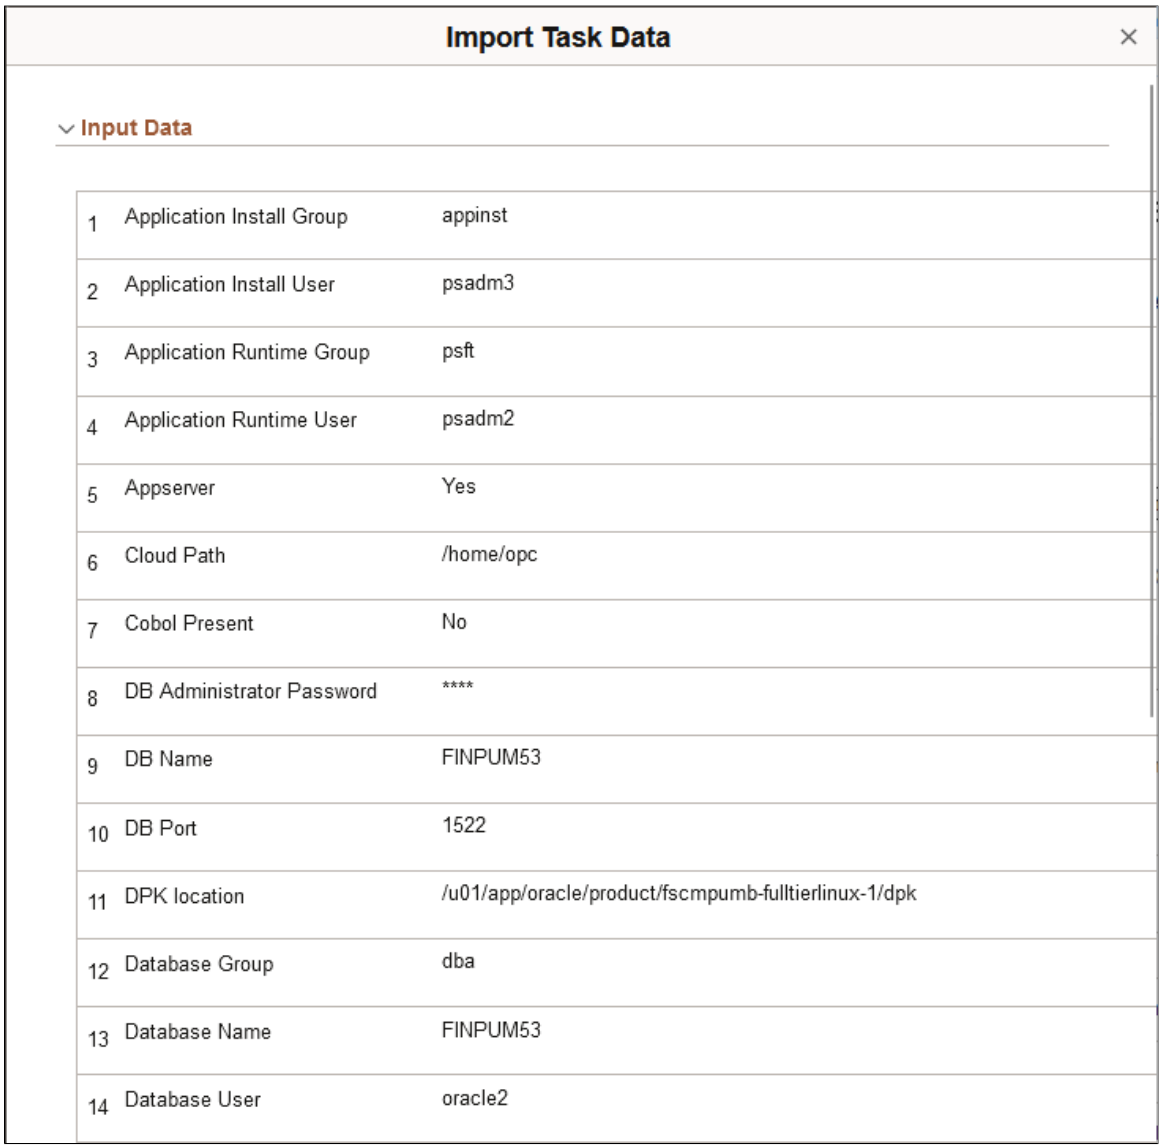

Expand the Derived Data section to view the values that Cloud Manager generates from the input data, such as IB Settings and domain parameters.

This example shows the Import Task Data page, showing the Derived Data section. (The image is truncated on the right side.)

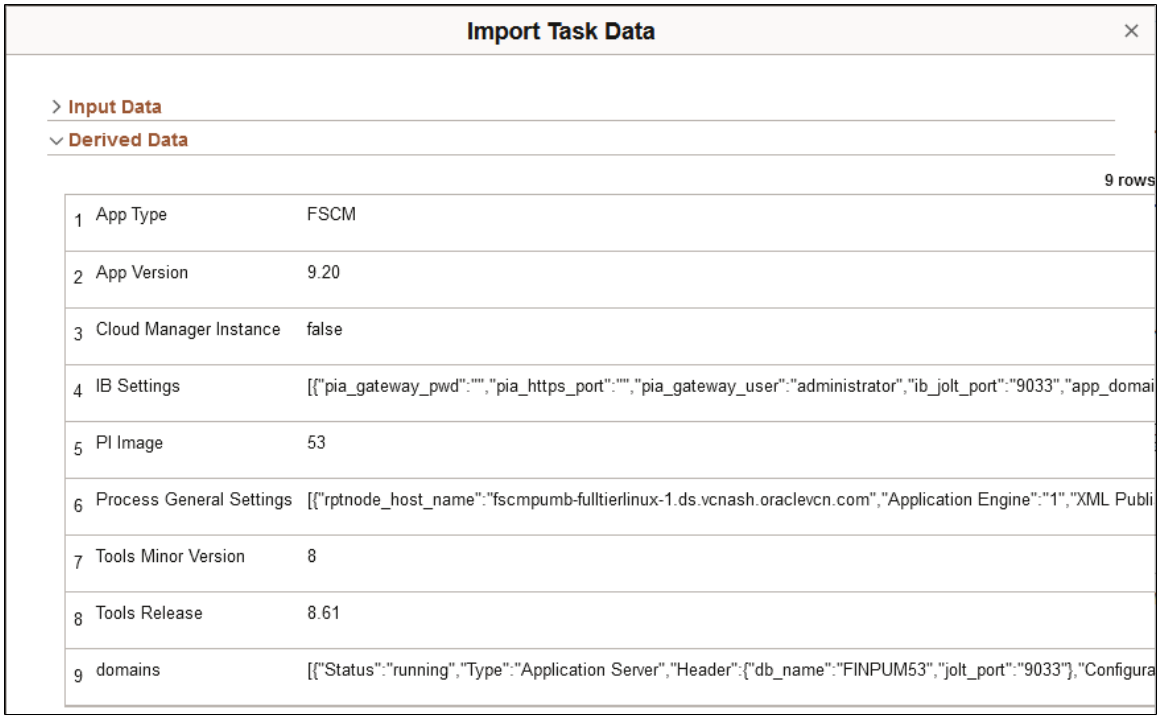

Retrying the Import Process

In case the import process fails, use these steps to retry:

1. Locate the failed step with the failure icon, an exclamation mark in a red circle.

   In this example, the second step, Crawl Instance Data: fulltier, failed.

   This example illustrates the Import Database Instance page showing a process with an error.

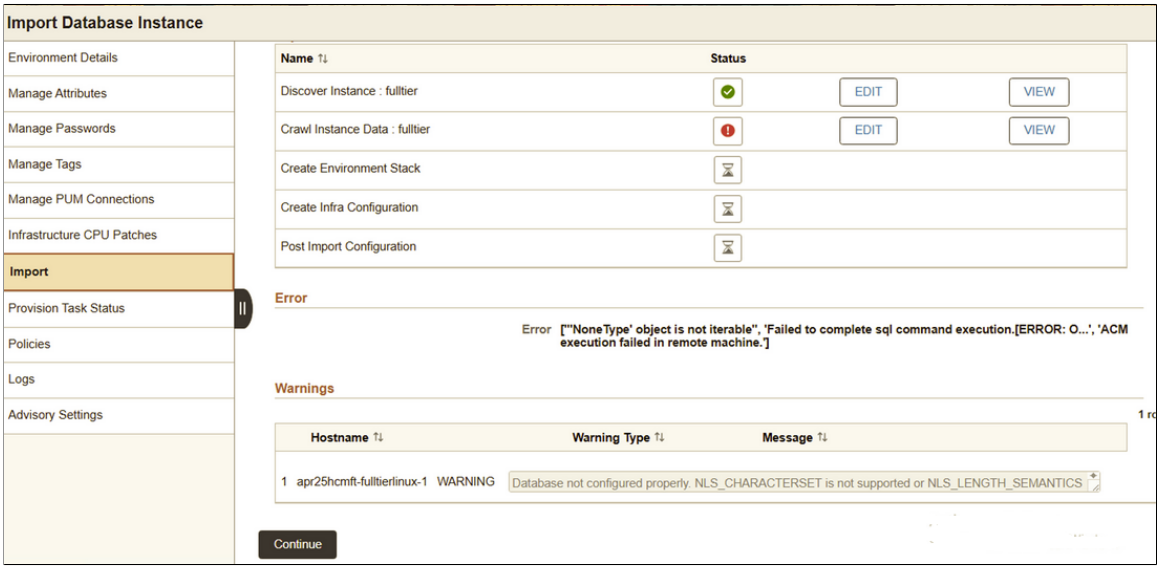

2. Review the Error section below the steps.

3. If necessary, use the log files to determine the failure.

   Select Logs from the left panel to view the logs generated by the import process.

   You can also find Database Import log files in these locations in the environment:

   - Import Database instance — Terraform logs are located in */home/psadm2/psft/data/cloud/cmlogs/envs/<Import envName>*. Log on as psadm2 user to view the logs.
   - Import Database — psp.log is located in */home/psadm2/psft/pt/<pt release number>/appserv/prcs/PRCSDOM/LOGS*. Log on as psadm2 user to view the log.
   - Discover UI logs — psp.log is located in */home/psadm2/psft/pt/<pt release number>/appserv/APPDOM/LOGS*. Log on as psadm2 user to view the log.

4. Click **Edit** for a step to correct the respective input parameters.

   This brings up the Import Task Data page with editable fields.

   Note: You may need to change an input parameter in a step other than the one showing the failure icon.

5. After making the correction, click **Continue** to retry the corrected and subsequent steps.

In case Cloud Manager finds an issue that is not serious enough to halt the import process, you see a Warning section, listing the host name for the affected node, the warning type, and a message with the details. You can make any necessary changes in the environment to resolve the warning after completing the import process.

Post Import Actions

An imported environment supports the following functions:

- Review environment details and manage attributes.

  See Managing Environment Attributes.

- On Demand Scale Up and Scale down.

  See Managing Nodes.

- Start
- Stop
- Delete

  The database node regardless of whether it is running on Database Systems or compute is not deleted.

- Upgrade
- Update

Lifecycle activities like “Apply PeopleTools Patch” and “Upgrade PeopleTools” can be done on the imported environment just like any other Cloud Manager provisioned environment.

Deleting Instances from Imported Environments

If you import an environment and then scale-up the node in Cloud Manager, the following occurs when you delete the scaled-up node:

- Deletes the scaled up node.
- Cleans up the metadata corresponding to the deleted node, which is stored in Cloud Manager.

When you delete a Middle Tier node, Search Stack node, or PeopleSoft Client node, which is not scaled up, from an imported environment, Cloud Manager cleans up the metadata and the instance infrastructure is not deleted from OCI.

Note: Database Systems node is not deleted.

See Managing Nodes

Managing Nodes

Cloud Manager supports on-demand scaling in OCI, which is the ability to scale up or down (horizontal scaling) by adding or removing nodes to an active running PeopleSoft environment as necessary. Using Manage Nodes, you can:

- Add additional middle tiers to a running database or middle tier. (Scale up)
  - Middle tier nodes are added one at a time.
  - Multiple middle tiers (Application Server, Web Server, Process Scheduler Server and Windows) are supported.

  Note: Add or remove node is not supported for full tier environment.
- Remove middle tier node from an environment. (Scale down)
- Add PeopleSoft Client.

  Note: Multiple PeopleSoft Clients are supported.
- Add a Search Stack node or share a search cluster.

  Adding Search Stack node to an environment requires an IB domain that was configured in the environment by Cloud Manager. If not, the Search Stack option will not be available when adding a node through Manage Node option. In such scenario, add a new Middle Tier node with IB enabled and then add the Search Stack node.

  For information on using the Manage Node page for search clusters, see Sharing a Search Cluster Across Multiple Environments and Managing Search Clusters.

  Search Stack supports two nodes:

  - If OpenSearch (or Elasticsearch) is already running in a node, OpenSearch Dashboards (or Kibana) cannot be installed in that node.

    - To install OpenSearch Dashboards (or Kibana), Search Stack is required. So the user has to provision Search Stack and OpenSearch Dashboards (or Kibana) in a new node.

- Stop a node.
- Start a node.

Adding Nodes

To add a node to a running environment:

1. Click the Related Actions button corresponding to the environment.
2. Select Manage Node.
3. Expand the Select Actions section if necessary, and select Add action.
4. Select Type.
5. Select the Operating System.
6. If the type is Middle Tier then there is an option to select an existing MT node from which configuration/custom configuration can be copied for the node being added.
7. Enter the required credentials and settings.
8. Click Submit and confirm.
9. Scale Up process status are:
    - InitiatingScaleup
    - ScaleUpInProgress
    - Running
    - ScaleupFailed

Navigation:

Click the Related Actions button corresponding to the environment. Select Manage Node. The Manage Node page is displayed.

This example illustrates the fields and controls on the Manage Node page for adding a Windows middle tier node.

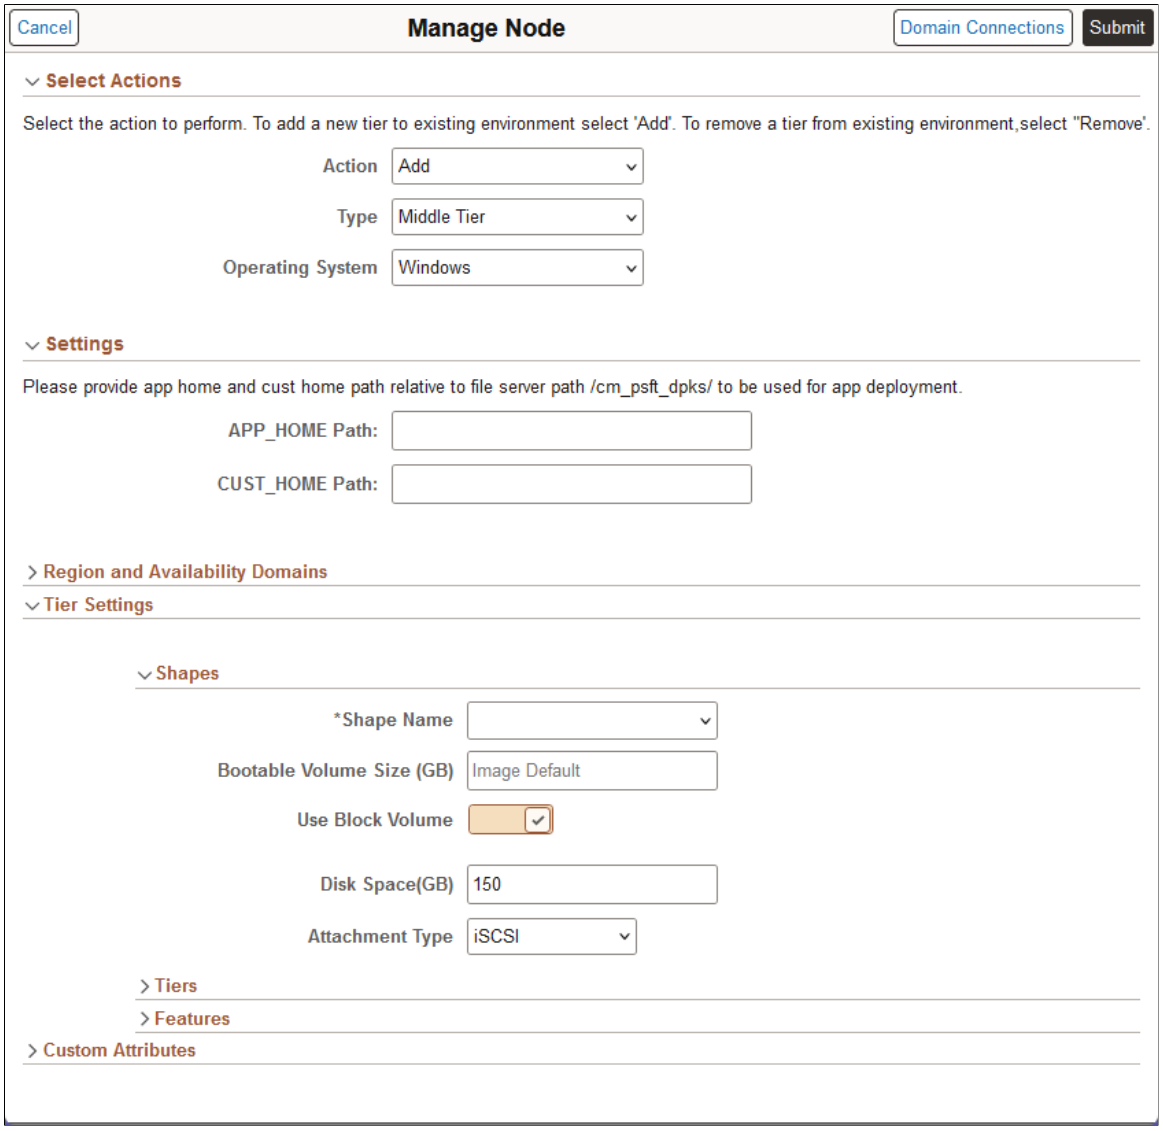

Note: Windows Middle Tier node is only supported for PeopleTools 8.57 and above. It is not supported on PeopleTools 8.55 or 8.56 environments.

The APP_HOME (Application home, or app home) and CUST_HOME (custom home, or cust home) paths are not mandatory when adding a Windows middle tier. However if you want to copy contents from the app home and/or cust home to the Windows middle tier perform the following:

1. Create a directory inside file server mount in the Cloud Manager. For example:
    a. Create a directory ps_app_home_win inside /cm_psft_dpks.
    b. Place the app_home contents as a zip file inside this directory.
    c. Create a directory ps_cust_home_win inside /cm_psft_dpks.
    d. Place the cust_home contents as a zip file inside this directory

2. Provide the path to app home (APP_HOME) and/or cust home (CUST_HOME). For example: ps_app_home_win and ps_cust_home_win.
3. When the Windows middle tier node is provisioned, the contents in the zip file will be extracted to the VM instance.

See Adding Windows Middle Tier Nodes.

When you add a node without copying configuration from an existing node, Cloud Manager fetches passwords from the password group that is mapped to the environment. Such auto-filled passwords cannot be edited from Cloud Manager.

Adding a Search Stack Node

This is an example of adding a search stack node to an environment.

This example illustrates the fields and controls on the Manage Node page for adding a Search Stack node.

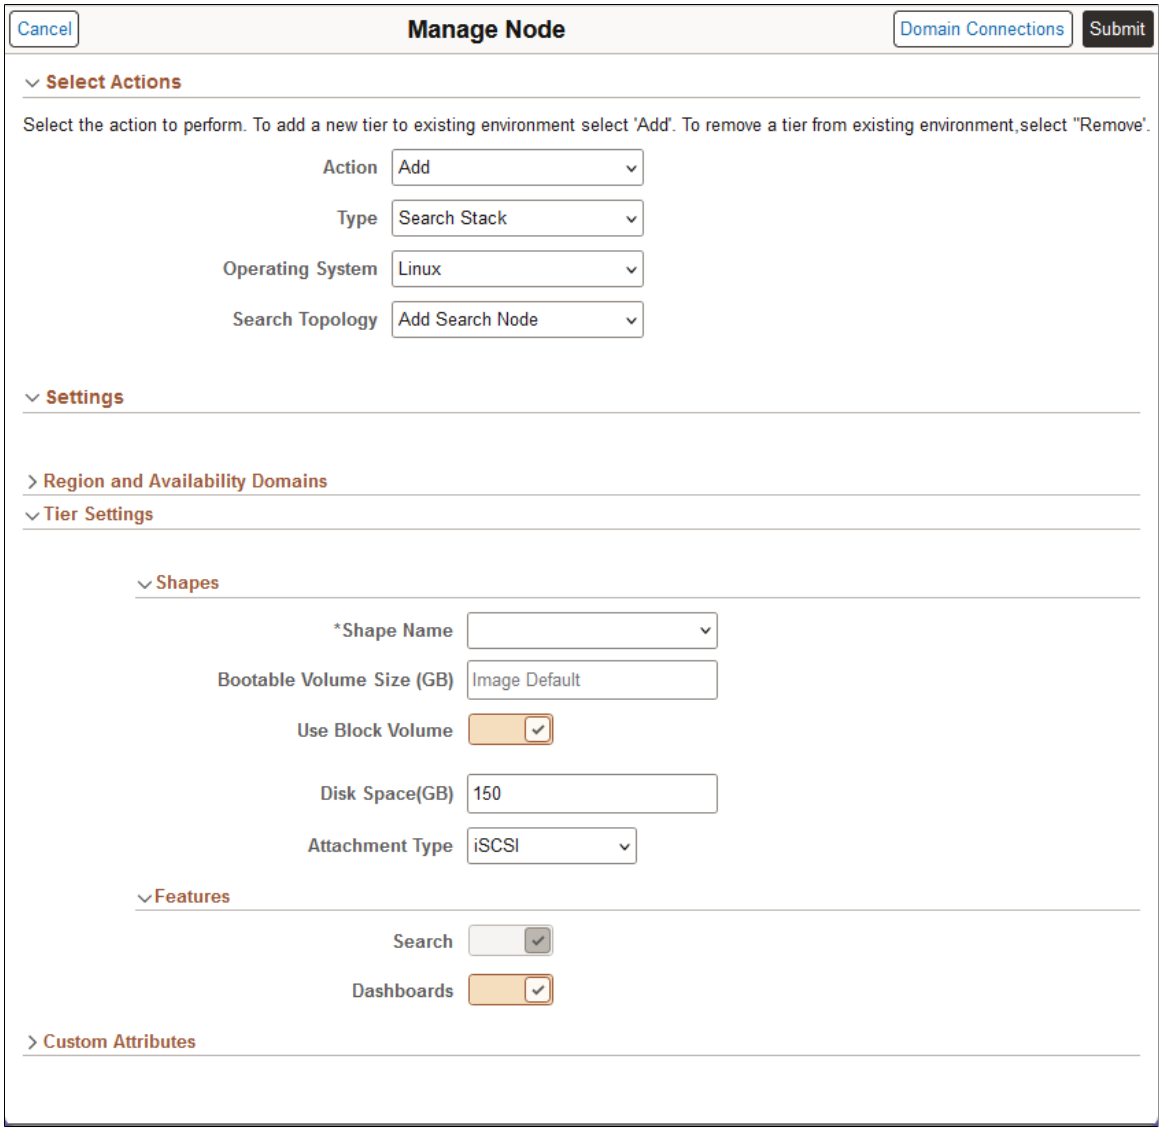

You must select at least one feature. OpenSearch and OpenSearch Dashboards (or Elasticsearch and Kibana) can be added in the same node.

For information on adding a node for an environment with a search cluster, see Managing Search Clusters.

Refer to the section on Creating Topology for more information. See Adding Search Stack Nodes.

Adding a Middle Tier Node

This is an example of adding a middle tier node to an environment.

This example illustrates the fields and controls on the Manage Node page to add a Linux middle tier node.

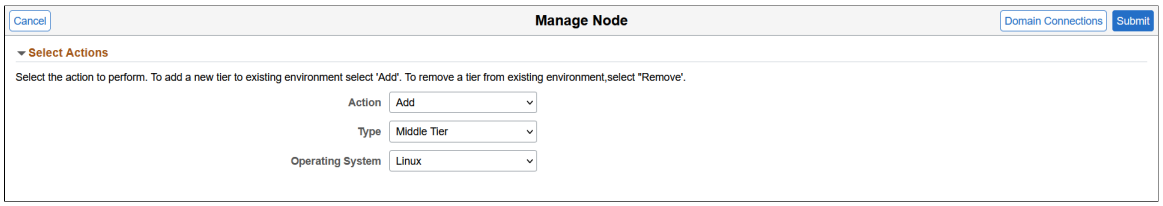

In the Settings section, select the location of the Application DPK.

- Object Store

  When you select Object Store, the DPK Object Name drop down list will display all the application DPK files in Object Store.

  This example illustrates the fields and controls on the Manage Node page where the Application DPK location is Object Store.

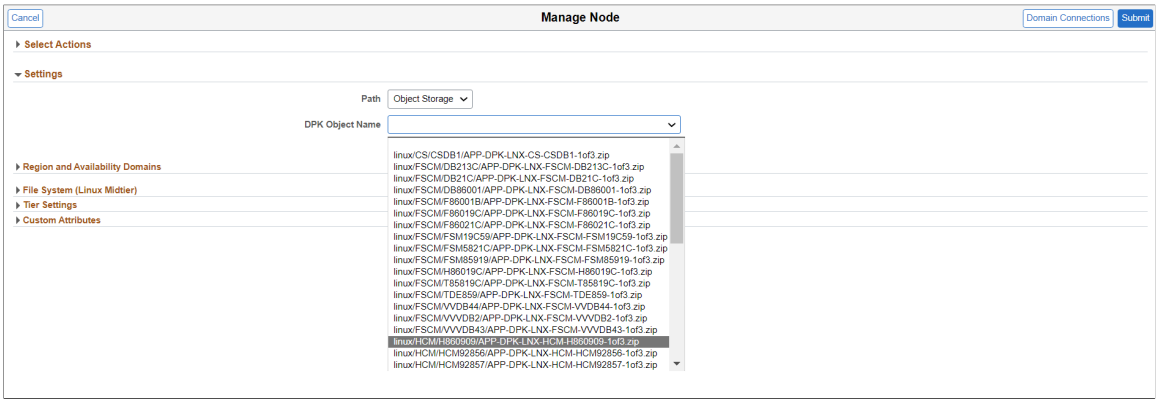

- File Server

  If you manually copied app home and cust home to the file server provide the relative path to the file server.

This example illustrates the fields and controls on the Manage Node page where the Application DPK location is File Server.

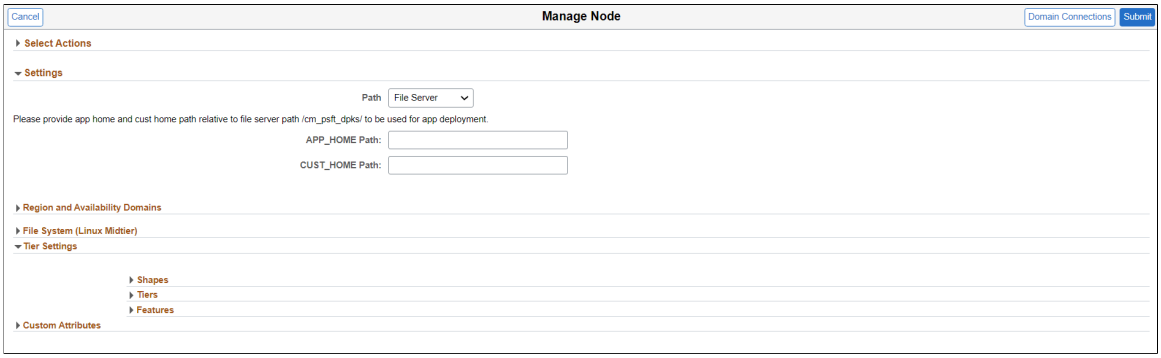

Adding Additional Middle Tier Nodes

This example illustrates the fields and controls on the Manage Node page for adding an additional Linux middle tier node.

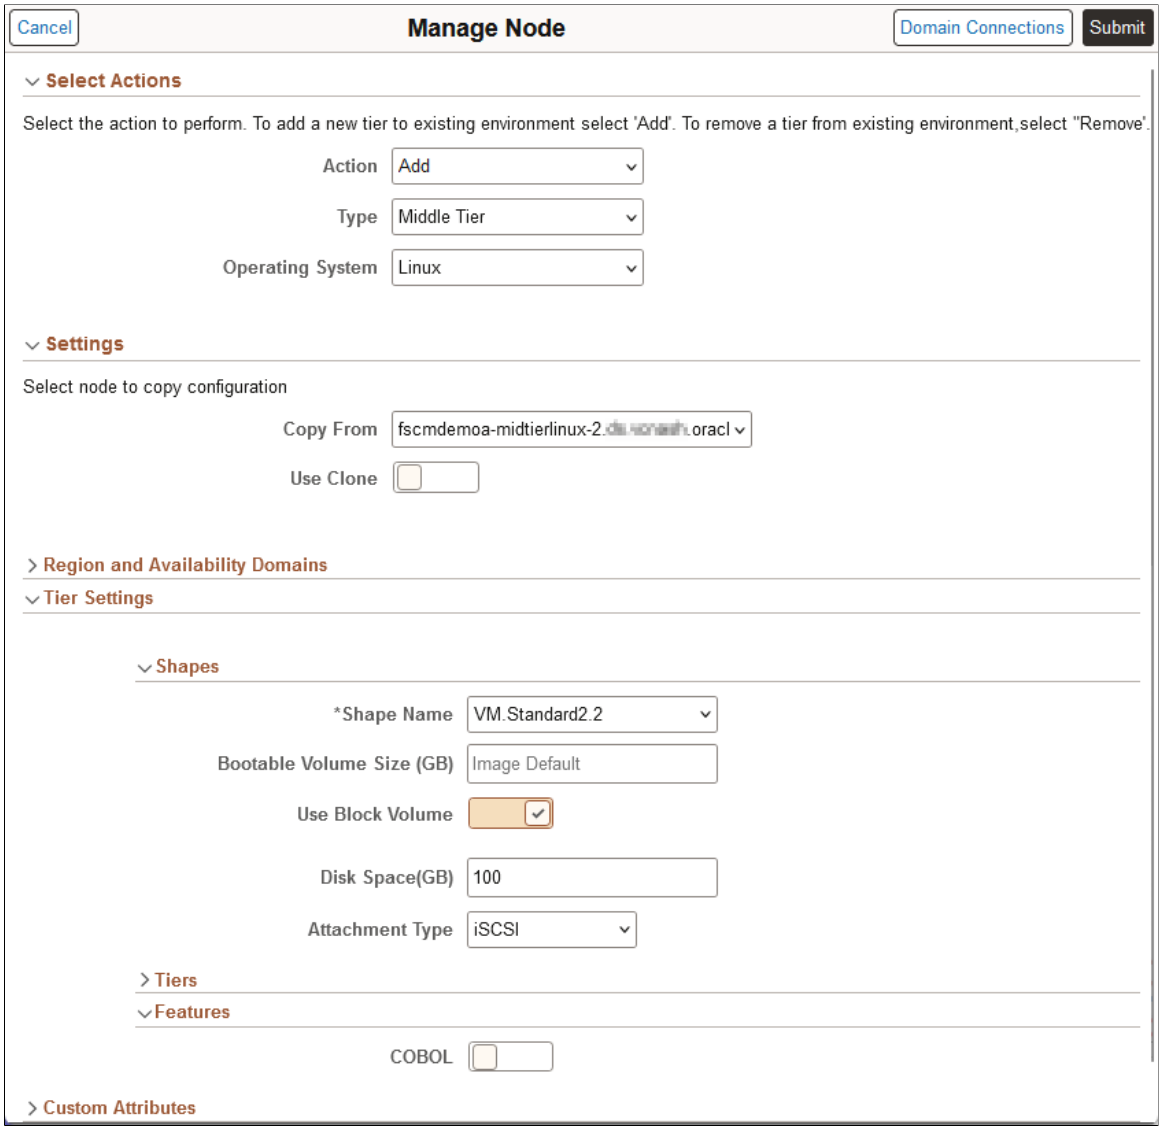

Expand the sections.

Field or Control	*Description*
Action	Available actions are: • Add • Remove • Restart • Start • Stop
Type	Available Types are: • Middle Tier • PeopleSoft Client • Search Stack **Note:** The options are enabled based on which node is already available in the environment.
Operating System	Select either Linux or Windows.
Domain Connections	Configure connections between application server domains and web server domains. See Configuring Domain Connections section in Creating An Environment.
Copy From	If a middle tier exists for the environment, select the middle tier node from the drop down list. The configuration or custom configuration will be retrieved.
Use Clone	Enable this option to reuse the values for Regions and Availability Domains, Tier Settings, and Custom Attributes from the source node that is cloned. If you do not enable this option, you must configure the settings.
Regions and Availability Domains	Defaults to the Regions and Availability settings of the environment to which new node is being added. These fields are read only.

Field or Control	*Description*
Tier Settings	Enter the required Shapes and Tiers. See Environment Template – Select Topology Page
Custom Attributes	Enter Credentials, General Settings, Network Settings, Network Security Group Settings, Fault Domain Settings and Domain Settings. Network Settings for compartment and virtual cloud network are read only. Select the subnet for the primary instance. See Configuring Custom Attributes

Removing Nodes

To remove a node:

1. Click the Related Actions button corresponding to the environment.
2. Select Manage Node.
3. Select Remove action.
4. Select Type.
5. Available Nodes will be displayed.

   Note: Database node can not be deleted.

6. Select the node to remove.
7. Click Submit and confirm.

This example illustrates the fields and controls on the Manage Node page for removing a node.

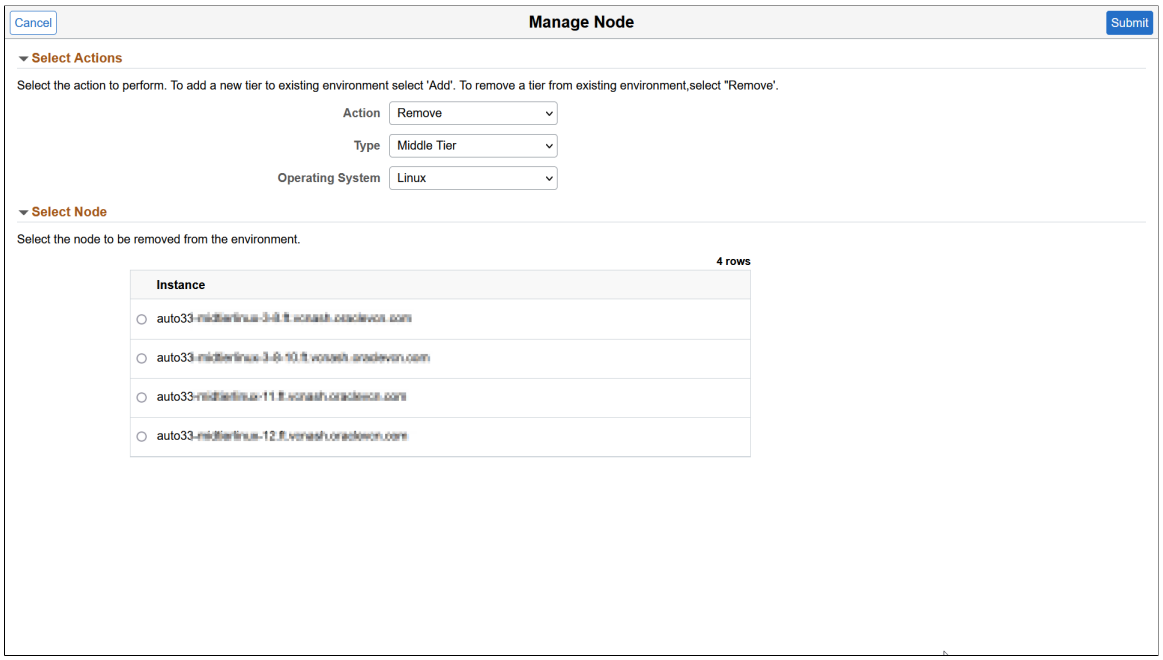

Stopping Nodes

To stop a node:

1. Click the Related Actions button corresponding to the environment.
2. Select Manage Node.
3. Select Stop action.
4. Available Nodes will be displayed.

   Note: Database node can not be stopped.

5. Select the node or nodes to stop.
6. Click Submit.

This example illustrates the Manage Node page to stop a node.

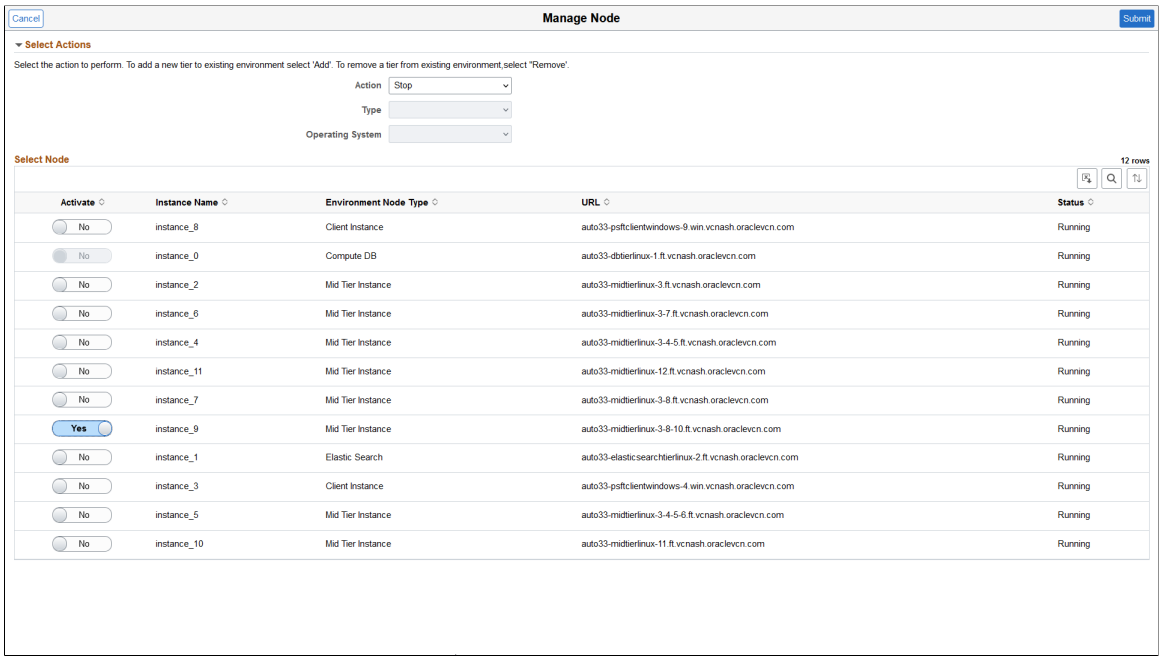

Starting Nodes

To start a node:

1. Click the Related Actions button corresponding to the environment.
2. Select Manage Node.
3. Select Start action.
4. The Activate toggle is available for any nodes in a stopped status.
5. Select the node or nodes to start.
6. Click Submit.

This example illustrates the Manage Node page to start a node.

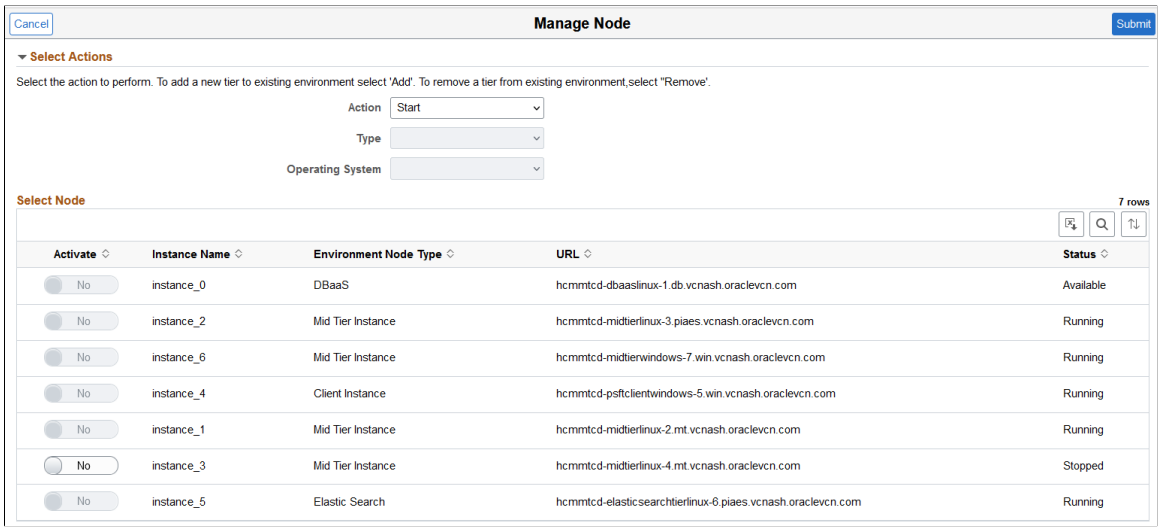

Backing Up and Restoring Environment

Use the Backup and Restore action to take a backup or restore an environment from the backup. Backup action will backup all nodes in the environment.

The time to backup and restore depends on the size of the database. The system will be online and reads are not affected. Import, Clone, Update and Instance restarts might affect the backup process.

Note: Backup and Restore actions are supported for auto-backup enabled Database Service environments with single or multiple PDBs in a CDB. See Configuring Database Backup Settings.

Backing Up the Environment

It is recommended to take a backup:

- Before applying a PeopleTools Patch.
- Before applying a PeopleTools upgrade.
- Before adding or removing a node.

To create a backup:

1. On the Environments page, select related actions for the environment and select Backup/Restore.
2. Any existing backup for the environment will be displayed.
3. Click the Create Backup button.
4. Enter the backup name and click Backup. The backup name has to be unique.

This example illustrates the fields and controls on the Environment Backup page. You can find definitions for the fields and controls later on this page.

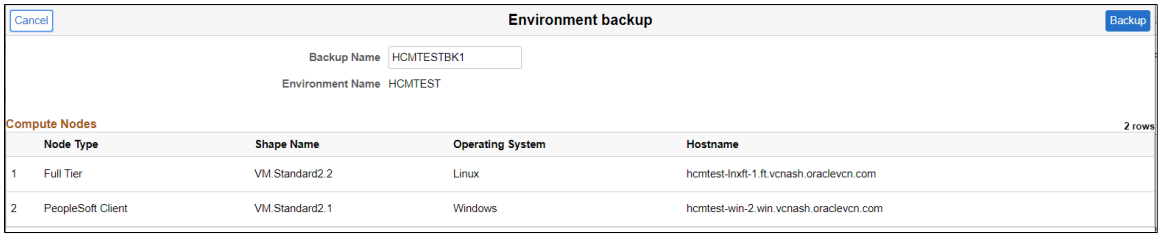

The nodes for the environment are displayed on the Environment Backup page.

The backups for an environment are displayed with the status.

This example illustrates the fields and controls on the Environment Backup page showing the current status. You can find definitions for the fields and controls later on this page.

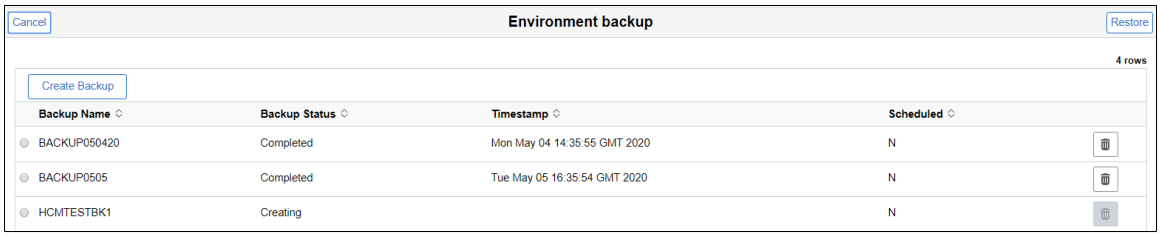

Field or Control	***Description***
Backup Name	Name of the backup
Backup Status	Displays the current status, Creating, In-Progress, Completed or Error.
Timestamp	Time the backup was created.
Scheduled	Currently backups can not be scheduled, so this column will have N.
Delete icon	Click the delete icon to delete an existing backup. Only backups in the Completed status can be deleted. The delete will clear all the block volumes and other dependency resources in OCI.

Restoring an Environment

When you select to restore an environment from a backup:

- The Restore overwrites all the data on the target instance.

- The Target instance will be unavailable during the restore process.

  This example illustrates the fields and controls on the Environments page when an environment is restoring.

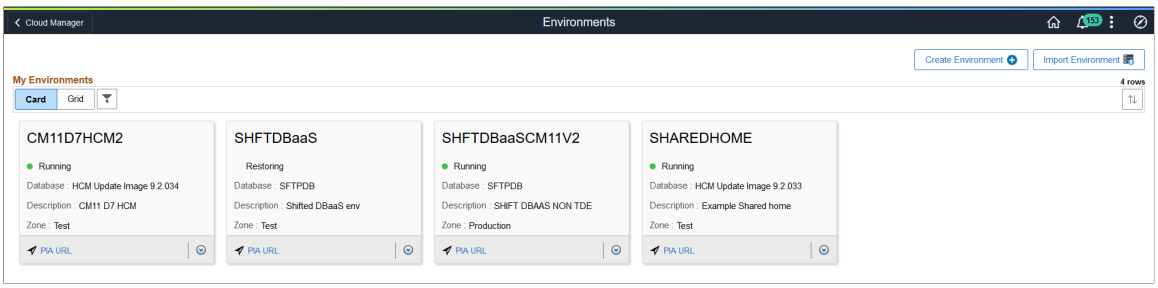

- Only one instance at a time can be restored on a target instance.

Note: When an environment with Load Balancer is restored, you must update the Load Balancer page for the environment. See Configuring Load Balancer Settings.

To restore an environment:

1. On the Environments page, select related actions for the environment and select Backup/Restore.
2. Any existing backup for the environment will be displayed.
3. Select the radio button corresponding to the backup you want to restore.
4. Click the Restore button.
5. All the nodes for the selected environment will be restored to the same instance.

Use Provision Task Status on the environment details page to check the progress of the restore operation.

This example illustrates the fields and controls on the Provision Task Status for restore operation. You can review the deployment tasks.

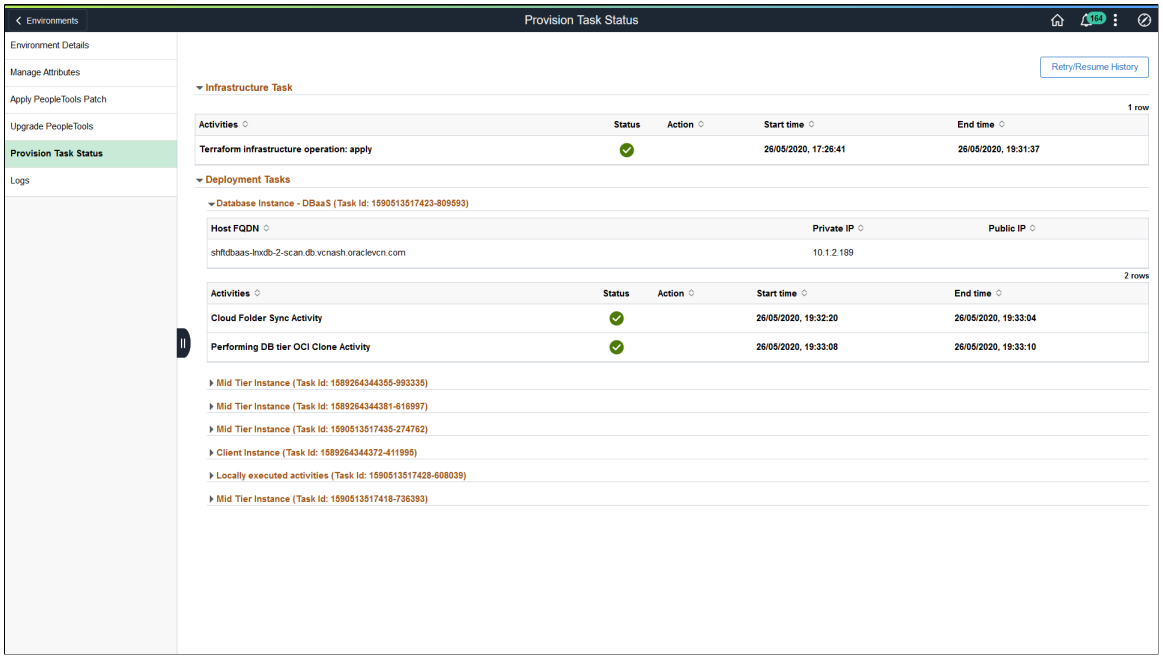

Enabling Disaster Recovery

Use the Disaster Recovery action to replicate an application instance in another region to ensure higher availability of the instance through switchover.

Cloud Manager enables you to replicate the resources of a region where the application instance runs (called "primary region") such as middle-tier and DBS to a standby region, where a switchover happens. This ensures business continuity in the event of an unplanned outage in the application instance.

Disaster recovery service can be enabled only for a DBS environment. You cannot create a standby environment from another standby environment or from a DBS environment that already has a standby environment.

Updating Infrastructure Settings

Navigation: **Cloud Manager Settings** > **Infrastructure Settings** > **OCI Metadata Refresh Across Region**.

To enable disaster recovery service for an environment:

1. Select Yes on the Enable Disaster Recovery field.

This example illustrates the fields and controls in the OCI Metadata Refresh Across Region section on the Infrastructure Settings page.

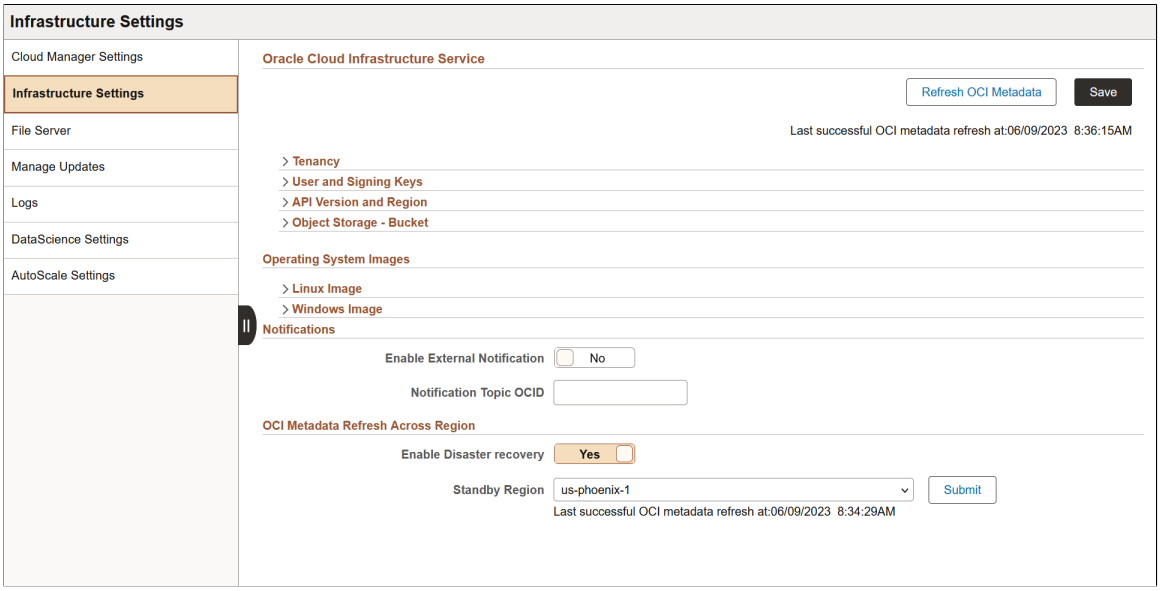

2. Select the Standby Region on the Infrastructure Settings page. Ensure that there is VCN peering between the source environment region and the selected standby region.

3. Click Submit and then save the changes on the Infrastructure Settings page.

Creating a Standby Environment

To create a standby environment, perform the following steps:

1. On the Environments page, select related actions for the environment and select Disaster Recovery. Create Standby Environment page appears.

   This example illustrates the fields and controls on the Create Standby Environment page.

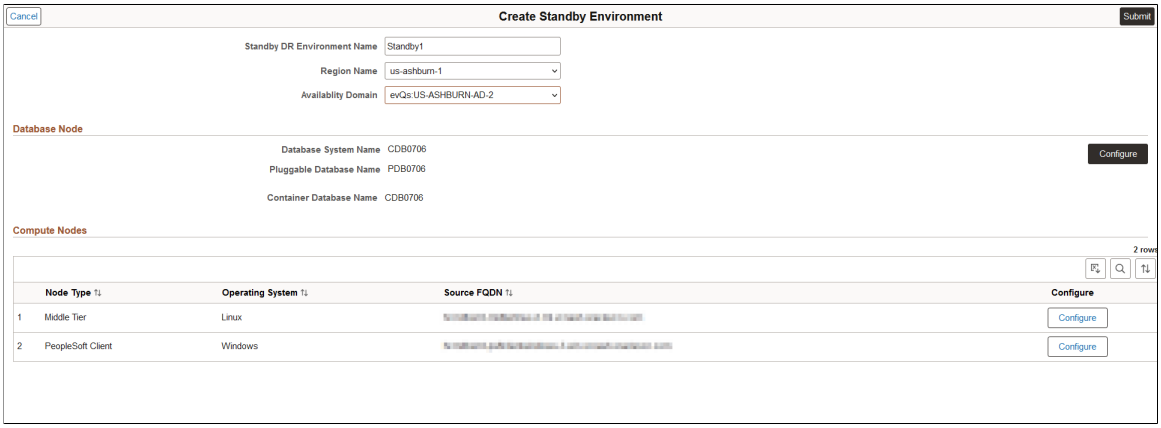

2. Enter the Standby DR Environment Name.

3. Select the standby region from the Region Name drop down.

4. Select the Availability Domain.

5. Click the Configure button in the Database Node section. The Disaster Recovery Database page appears.

   This example illustrates the fields and controls on the Disaster Recovery Database page.

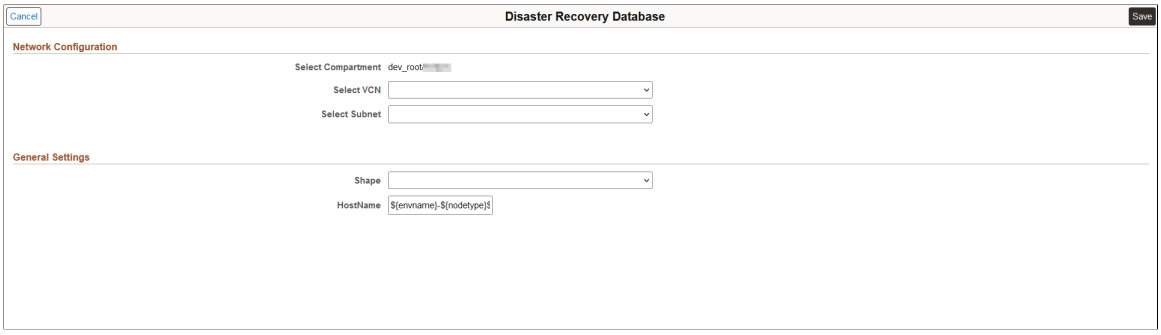

   Select the VCN, Subnet and Shape for the database node of the standby database you want to create. Click Save.

6. Click the Configure button in the Compute Nodes section. The Disaster Recovery Compute Node page appears.

   This example illustrates the fields and controls on the Disaster Recovery Compute Node page.

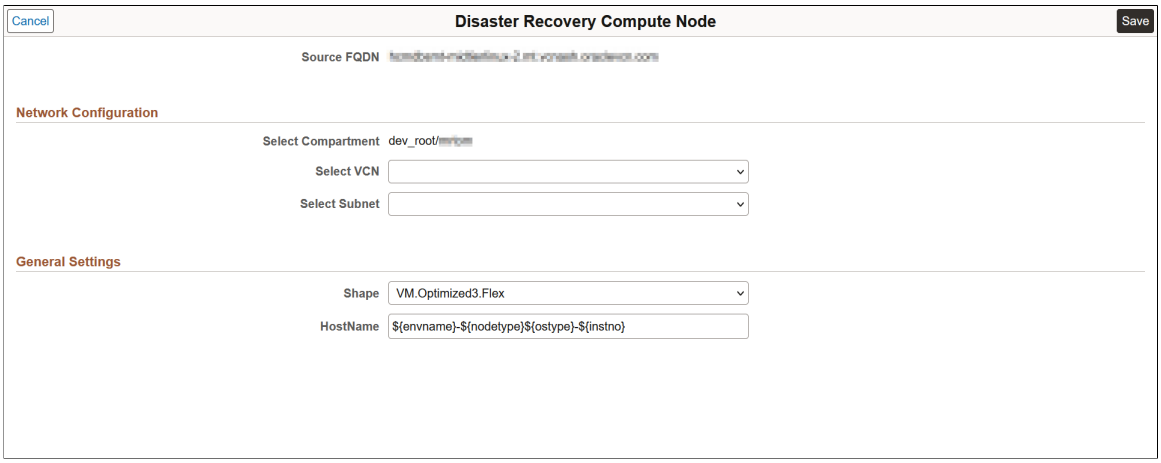

   Select the VCN, Subnet and Shape for the compute node of the standby database you want to create. Click Save.

7. Click Submit. The standby environment creation is initiated.

This example illustrates the fields and controls on the Environments page showing the status of standby environment processing.

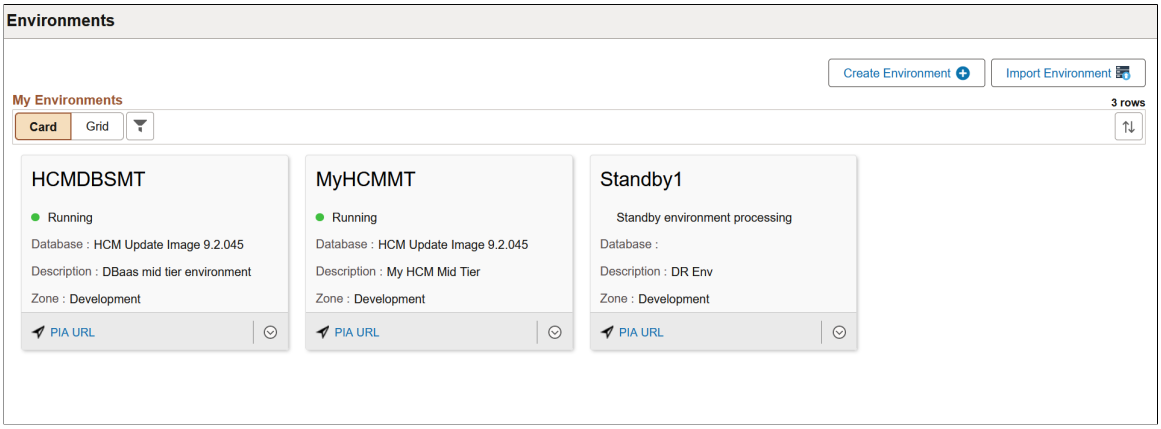

The Environments page displays the status of standby environment. The status of a provisioned standby environment is displayed as *In Standby*.

The Related Actions menu in a standby environment displays the following options:

- Details
- Delete
- Delegate Access

Note: The Environment Details page for a standby environment does not have Refresh Metadata option.

---

8. Connect to the standby instance using SSH. Insert your SSH public key to the compute Linux machines of the standby node. To use the vi editor to open and modify the file to which the SSH public key must be added, use this command:

   ```
   vi .ssh/authorized_keys
   ```

These steps must be performed before a disaster event occurs.

Initiating Manual Failover to the Standby Environment

If a disaster occurs in any primary region, perform the following steps in the OCI console to switch over to the standby environment.

1. Select the standby node in OCI console.
2. Select Data Guard Associations under Resources.
3. Click the Actions icon and select Failover. This makes the standby node the primary node in the database, thus providing disaster recovery. Additionally, the previously selected primary node that had failed becomes the standby node.

Note: You must repeat this process to manually switch over back to the original node, once the disaster event is resolved.

---

4. If the switchover is done in compute systems, start the Cloud Manager applications using the following command:

   ```
   psadmin start -d *all;
   ```

Note: You can use the Oracle Full Stack Disaster Recovery (FSDR) to create and automate disaster recovery plans based on the standby environment created through Cloud Manager. After the standby environment is created, you can use Oracle FSDR to create a disaster recovery protection group and plan to complete the disaster recovery setup. See https://docs.oracle.com/en-us/iaas/disaster-recovery/index.html.

Refreshing DB Systems Environment

You can refresh DB Systems environments managed by Cloud Manager from a backup in Object Store or from another DB system managed by Cloud Manager.

- The RMAN database backup for an on premise environment lifted to Cloud Manager resides in the Object Store.

  See Running Lift Using Hot Backup (RMAN)

- Backups for other DB Systems environments in Cloud Manager are stored in OCI block storage.

  See Backing Up and Restoring Environment

Refresh supports:

- Refresh database only.
- Refresh environment which includes database, PS App Home and PS Cust Home.

Important! The source and target database must be on the same version. Source environment backup must contain spfile.

To refresh an environment:

1. Perform a backup of the environment prior to the refresh. Before starting the Refresh, ensure that the backup has completed.
2. Navigate to the environment card on the Environments page.
3. Select Refresh from the actions for the DB Systems environment you want to refresh.
4. On the Refresh page, select the Source Type.
5. Enter the values for the backup.
6. Click **Refresh**.

Refreshing from Object Store

This example illustrates the fields and controls on the Add Refresh page where Source Type is Object Store. You can find definitions for the fields and controls later on this page.

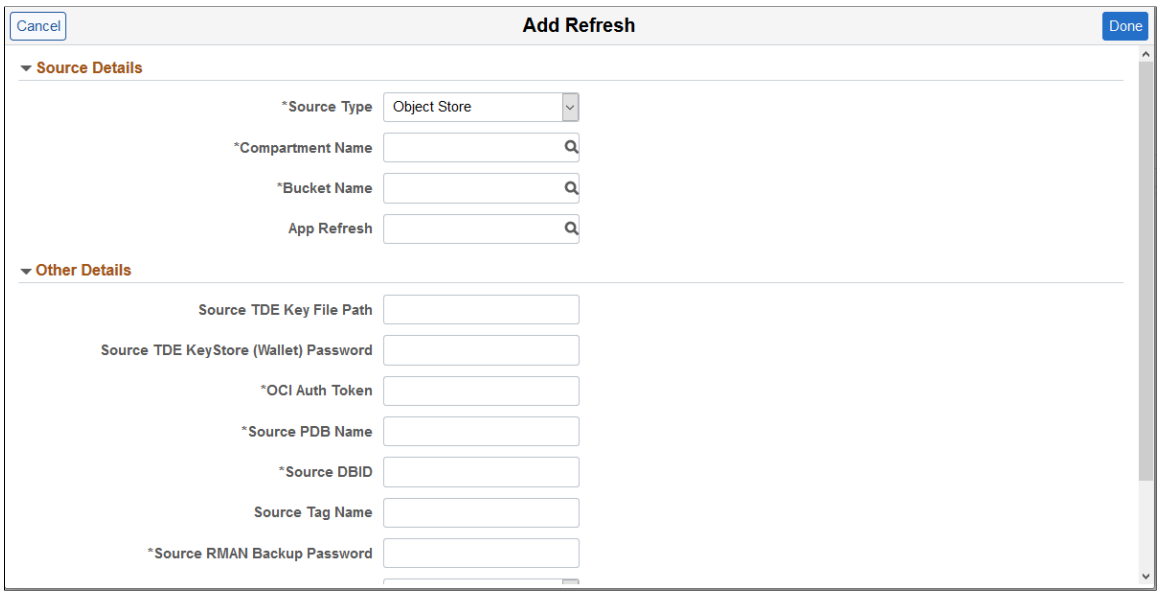

Field or Control	*Description*
Source Type	Select the Object Store from the drop down list.
Compartment Name	Select the compartment used when the backup was taken using DBCLI (database command line interface) or Oracle Database Cloud Backup Module (ODCBM).
Bucket Name	Select the bucket name used when the backup was taken using DBCLI (database command line interface) or Oracle Database Cloud Backup Module (ODCBM).
App Refresh	Select the name of the application backup to refresh PS App Home and PS Cust home. Leave this field blank if you want to refresh the database only.
Source TDE Key File Path	If TDE is enabled on the source database enter the TDE key file path.
Source TDE KeyStore (Wallet) Password	If TDE is enabled on the source database enter the TDE KeyStore (Wallet) Password.
OCI Auth Token	Enter the OCI Auth Token for the source database.

Field or Control	*Description*
Source PDB Name	Select the name used when the backup was taken using DBCLI (database command line interface) or Oracle Database Cloud Backup Module (ODCBM).
Source DBID	The source DBID is available during the RMAN backup stage, user must have the source DBID available.
Source Tag Name	The tag name that was specified during the RMAN backup.
Source RMAN Backup Password	Enter the RMAN backup password. This field only applies to non-TDE databases.
Pre Refresh Custom script	Select an uploaded script to run prior to refreshing the environment. This script is run from a middle tier belonging to the target environment.
Post Refresh Custom script	Select an uploaded script to run post refreshing the environment. This script is run from a middle tier belonging to the target environment.
Target MT DMS Path	Path to Data Mover scripts on the target middle tier node.

Note: Application DPK contains the absolute path in Cloud Manager for PS_APP_HOME and PS_CUST_HOME. If application DPK only includes PS_APP_HOME, only PS_APP_HOME of target will be refreshed, otherwise both PS_APP_HOME and PS_CUST_HOME will be refreshed.

Refreshing from OCI Backup

This example illustrates the fields and controls on the Add Refresh page where Source Type is OCI Backup. You can find definitions for the fields and controls later on this page.

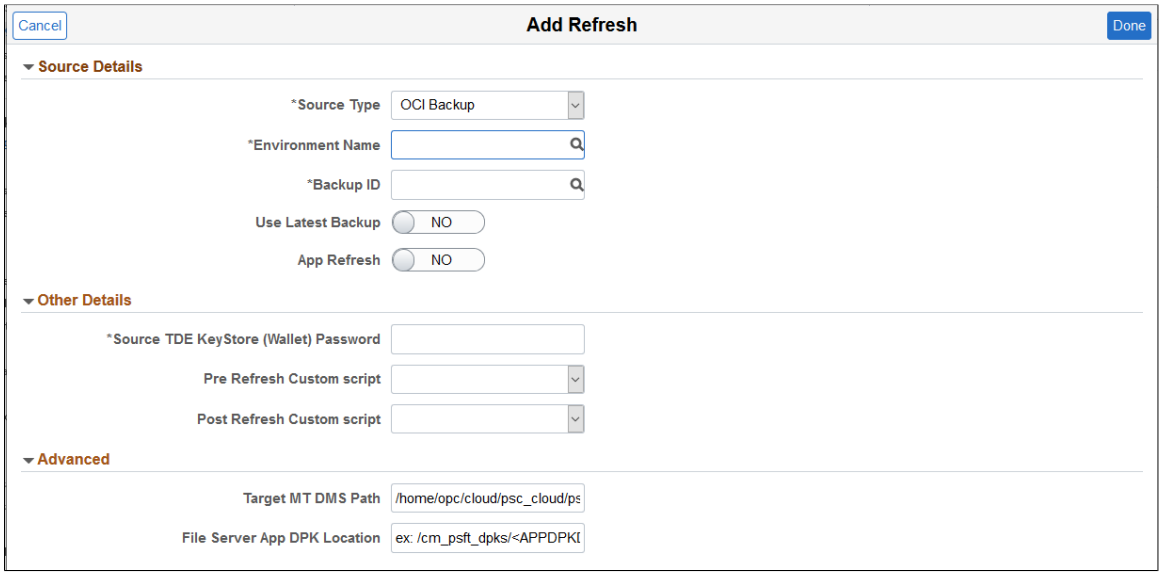

Field or Control	*Description*
Source Type	Select OCI Backup from the drop down list.
Environment Name	Select the environment to refresh from the drop down list.
Backup ID	Select the Backup ID from the drop down list.
Use Latest Backup	Select Yes to use the latest backup.
App Refresh	Select Yes to refresh PS App Home and PS Cust Home as well as the database. Select No for a database refresh only.
Source TDE KeyStore (Wallet) Password	The Source Wallet password is the same as the Database Administrator password.
Pre Refresh Custom script	Select an uploaded script to run prior to refreshing the environment. This script is run from a middle tier belonging to the target environment.
Post Refresh Custom script	Select an uploaded script to run post refreshing the environment. This script is run from a middle tier belonging to the target environment.

Field or Control	*Description*
Target MT DMS Path	Path to Data Mover scripts on the target middle tier node.

Retrieving Failed Refresh

Important! Before doing refresh make sure that an environment backup is done.

To retrieve failed refresh:

1. Select Details > Environment Details and click Refresh.

   Environment may come to Running state.

2. If status updates to Not Started, check the following:
   - PDB, CDB name after refresh may differ from the actual. Correct the names.
   - DB Admin, access usernames and passwords may differ from actual. Correct DB credentials.

3. If status still results in failed status, check Details > Provision Task Status.
   - If failure is in INFRA task, retry from Provision Task Status page.
   - If that does not work, use the backup and restore feature to restore the environment.
   - Make sure spfile is present.

Refreshing ADB Environment

You can refresh an ADB environment from another ADB environment of the same instance type.

To refresh an ADB environment, select Refresh from the actions for the ADB-Dedicated or ADB-Shared environment you want to refresh. The source and target must be the same ADB instance type, ADB-Shared to ADB-Shared or ADB-Dedicated to ADB-Dedicated.

The Source Type field is display-only because the target is created from source directly without taking a backup of the source.

This example illustrates the fields and controls on the Add Refresh page for ADB environment. You can find definitions for the fields and controls later on this page.

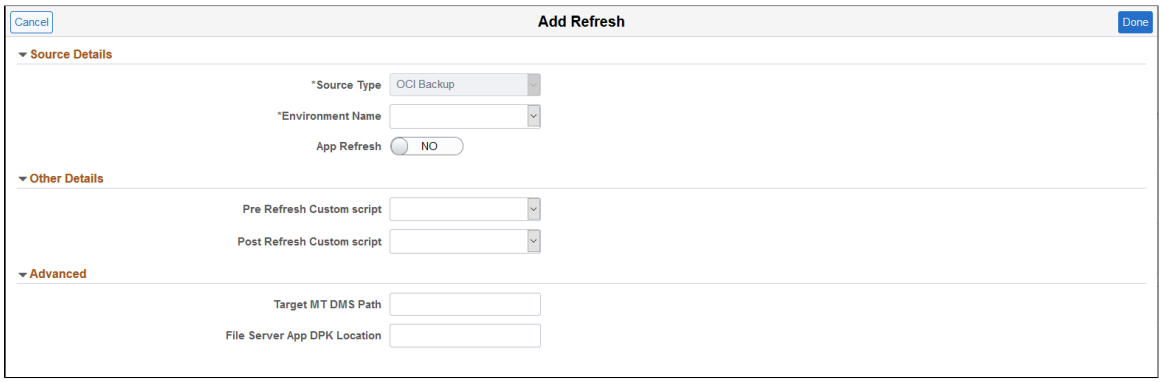

Field or Control	*Description*
Environment Name	Select the environment to refresh from the drop down list.
App Refresh	Select Yes to refresh PS App Home and PS Cust Home as well as the database. Select No for a database refresh only.
Pre Refresh Custom script	Select an uploaded script to run prior to refreshing the environment. This script is run from a middle tier belonging to the target environment.
Post Refresh Custom script	Select an uploaded script to run post refreshing the environment. This script is run from a middle tier belonging to the target environment.
Target MT DMS Path	Path to Data Mover scripts on the target middle tier node.
File Server App DPK Location	Location of Application DPK.

Configuring AutoScale Settings

Use the AutoScale Settings page to control notifications for the following event types:

- Monitoring Event: This notification event is triggered when an erroneous condition occurs during prediction flow.
- Scaling Event: This event is triggered when you need to take an action or an action is already taken by the prediction flow.

The initial notifications about events are sent immediately through push notifications. Bookkeeping is done for all the subsequent occurrences of the same event. A new notification is sent once the selected time interval period has lapsed.

AutoScale Settings can be configured on **Environments** > **Details** > **AutoScale Settings**.

This example illustrates the fields and controls on the AutoScale Settings page. You can find definitions for the fields and controls later on this page.

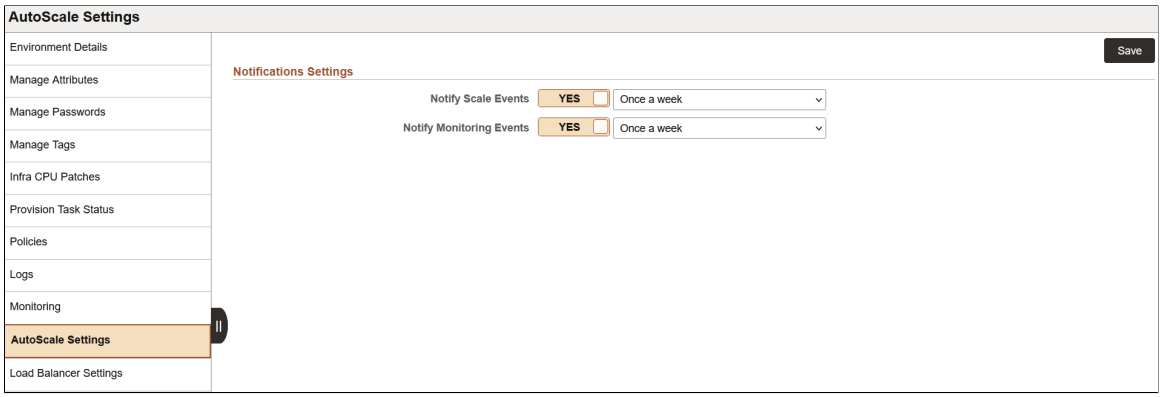

Field or Control	***Description***
Notify Scale Events	Select YES to enable notification for anomaly events and action taken based on anomaly events. You can select the following intervals: • Every Six Hours • Every Three Hours • Every Time • Once Daily • Once a week
Notify Monitoring Events	Select YES to enable notification for issues with the Monitoring flow or model predictions. You can select the following intervals: • Every Six Hours • Every Three Hours • Every Time • Once Daily • Once a week

Cloud Manager sends you an email notification about all the bookkeeping events at the end of the day at around 11:50 p.m. When you select the notification interval as Once a week, you will receive email notification only on Friday. You need to configure the Notification Topic OCID to receive the email notification. See Configuring Cloud Manager Settings for OCI.

Configuring and Reviewing Advisories

Use the Advisory Settings page (ECL_ENVRECM_SET_FL) to configure or update advisory settings for an environment.

Using Advisories includes the following high-level steps:

1. The Cloud Manager administrator enables advisories globally on the Cloud Manager Settings > Advisory Settings page.

   The advisories feature works only for environments with tags.

   See Advisory Settings Page.

2. The environment owner enables advisories for a provisioned environment.

   See Enabling Advisory Settings in this section.

3. When an advisory notification appears on the environment card, the environment owner accesses the list of advisories and takes appropriate action.

   See Receiving Advisory Notifications in this section.

Advisory Type	***Description***
Infrastructure CPU	Apply Infrastructure CPU patches immediately or schedule a policy. There is no expiration date for the Infrastructure CPU advisor.
Midtier scaling	Weekly advisory concerning mid-tier usage. Midtier scaling advisories include an action to create a policy. These advisories expire every week.
Application Level	Review reports of application errors and application crashes. These advisories expire after a day.
Server Level	The Server Level advisory is triggered when the available boot volume percent for the environment drops below a specified level. These advisories expire after a day.

There is a recurring job (ECL_GEN_RECM) and a cron job for generating advisories, which serve the following purposes:

- Expire the generated advisories on the expiry date.
- Generate new advisories.

Enabling Advisory Settings

By default, advisories will be enabled for all advisory types for the environment. If you want to update advisory settings for the environment, click the Related Actions icon for the environment and select Details. Select Advisory Settings from the left panel.

This example illustrates the fields and controls on the Environments > Advisory Settings page. Definitions for the fields and controls appear following the example.

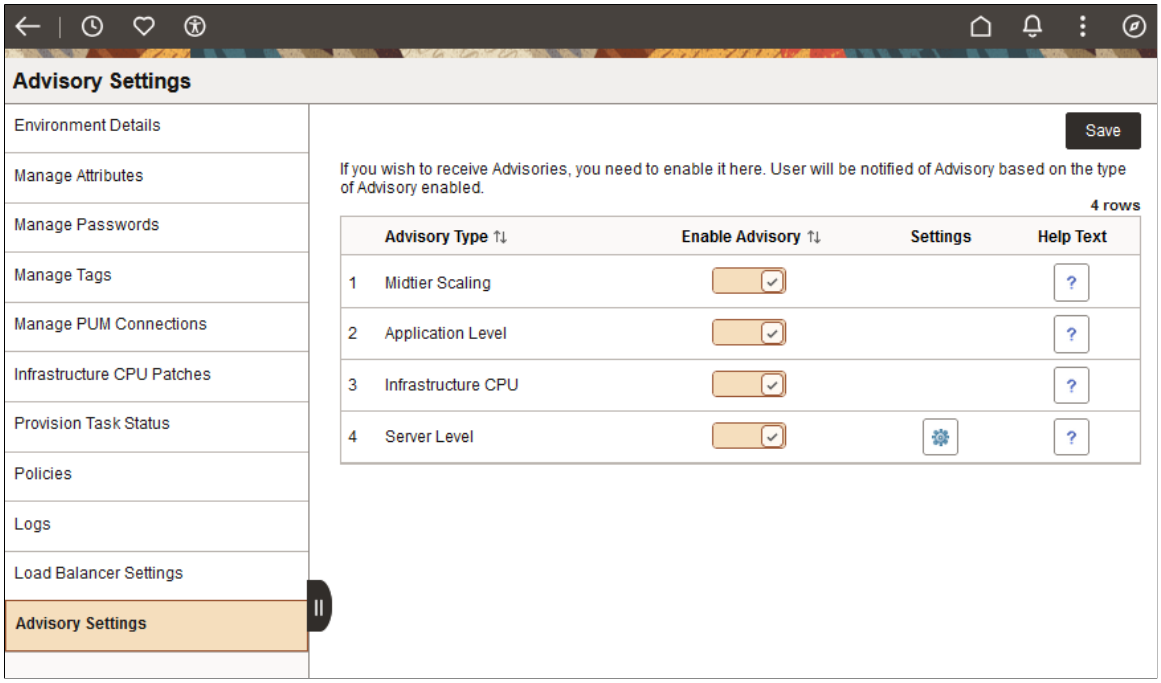

Field or Control	*Description*
Advisory Type	Displays the type of advisory to be enabled or disabled. • Midtier Scaling • Application Level • Infrastructure CPU • Server Level
Enable Advisory	Select this option if you wish to receive notifications with advisories for the respective advisory type.
Settings	Click to display a Settings dialog box to enter the boot volume threshold.
Help Text	Hover over the icon to view the help text that describes the action that happens on enabling advisory.

Click the Settings icon for the Server Level advisory and enter a value for the boot volume threshold percentage. When the available boot volume space drops below this percentage, an advisory is generated.

This example illustrates the Settings dialog box for Server Level advisories.

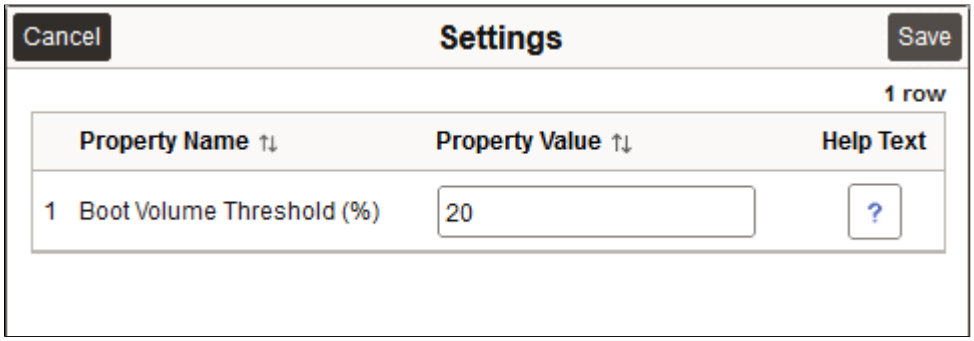

Receiving Advisory Notifications

The advisory is indicated as a badge counter on each environment card. The badge is not displayed if there are no advisories associated with an environment.

This example illustrates the badge counter, in the top right corner of each card, for advisories on the Environments page.

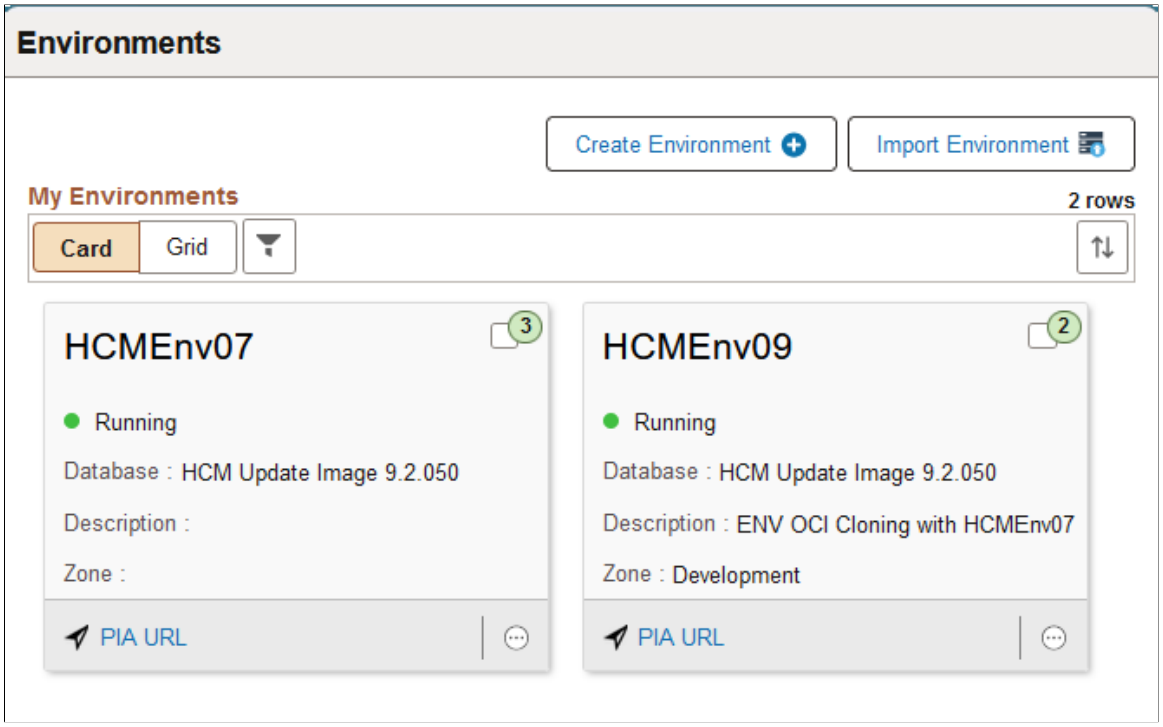

Clicking the badge displays all advisories available for the environment. You can then take appropriate action based on the type of advisory generated. More information is included in the following sections.

This example illustrates the Advisories page with three advisories.

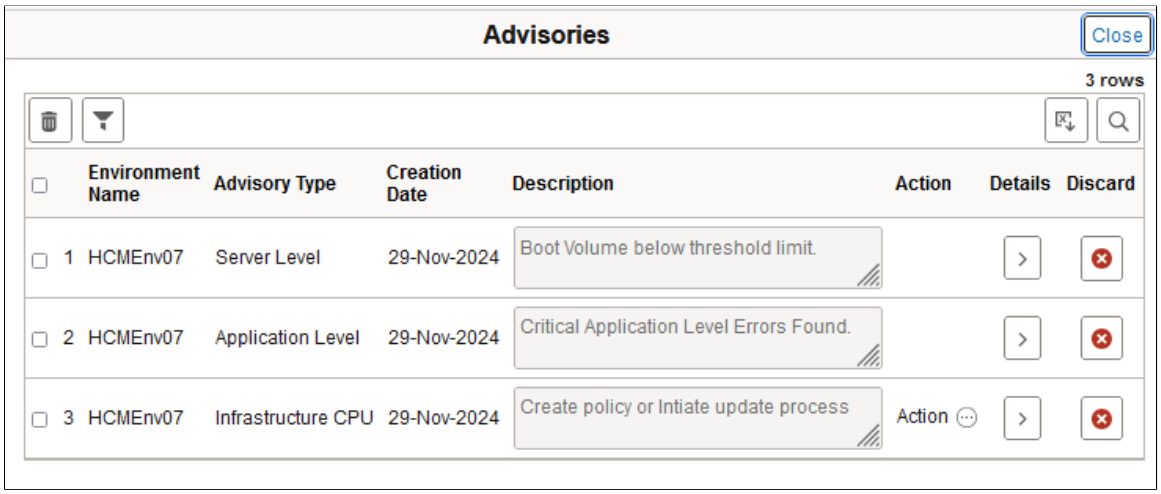

Field or Control	*Description*
Action	Click for a menu with available actions. • Infrastructure CPU — Select the action Apply Update or Create Policy. • Midtier Scaling — Select the action Create Policy. • Server Level and Application Level — There are no actions available from this Advisories page.
Details	Click to review a detailed report generated by the advisory and the advisory history.
Discard	Click to delete the advisory.

Expand the History section on the Advisory Details page to see the date through which the advisory is valid. The section lists any advisories that are not available for review or which have been acted upon. Here are the Advisory States that are included:

- Advisory Action Initiated
- Advisory Expired
- Policy Created Through Advisory
- Advisory Discarded

Using Application Level Advisories

Application Level Advisories are generated for issues such as application crashes or errors. To see a description of the issue, click the Details icon. You can also drill down for more details.

This example illustrates the Advisory Details page for Application Level.

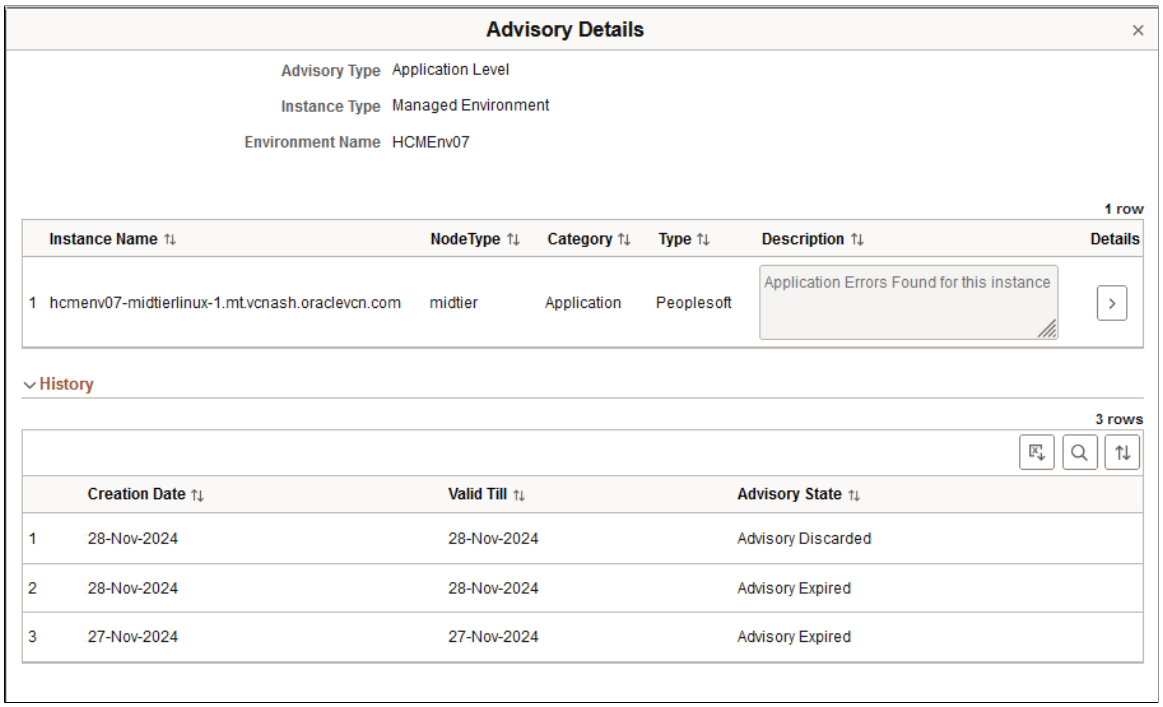

Click the Details icon in the grid at the top to see the More Details page. The page includes:

- The domain where the issue occurred, such as application server, Process Scheduler, or web server (PIA).
- The name of the process affected by the error.
- A more complete error description.
- The complete path to a file, such as a log file for application errors, or a core file for application crashes.

This example illustrations the More Details page for an Application Level Advisory.

Using Midtier Scaling Advisories

This example illustrates the action available for Midtier Scaling Advisory.

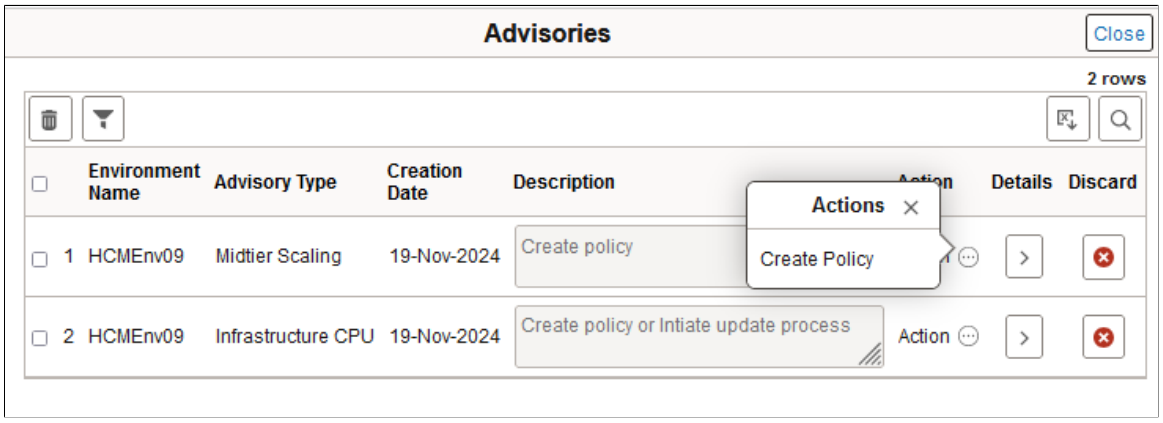

The mid-tier scaling advisory applies to environments that have monitoring enabled and forecasting data present. Cloud Manager provides a weekly advisory for middle tier resources. A mid-tier scaling policy is created for a specific environment.

If you choose to accept the advice to create a new policy, selecting the corresponding action on the Actions field redirects you to the Policy Editor page. See Using Policy Editor.

If a policy already exists, Cloud Manager auto updates the existing forecast data with new forecast data. This policy needs to be recurring, so you must select the recurrence (Environment Midtier Scaling) or set up a custom recurrence under “Schedule”.

On clicking the Details icon, Cloud Manager displays advisory details for the respective environment. The x-axis denotes weekly node forecasting and the y-axis denotes forecasted nodes. The values for each day in the week match the values for parameters Day 01 (Monday) to Day 07 (Sunday) in Policy Action Parameters. See Using Policy Editor.

This example illustrates the Advisory Details page for Midtier Scaling for the selected environment.

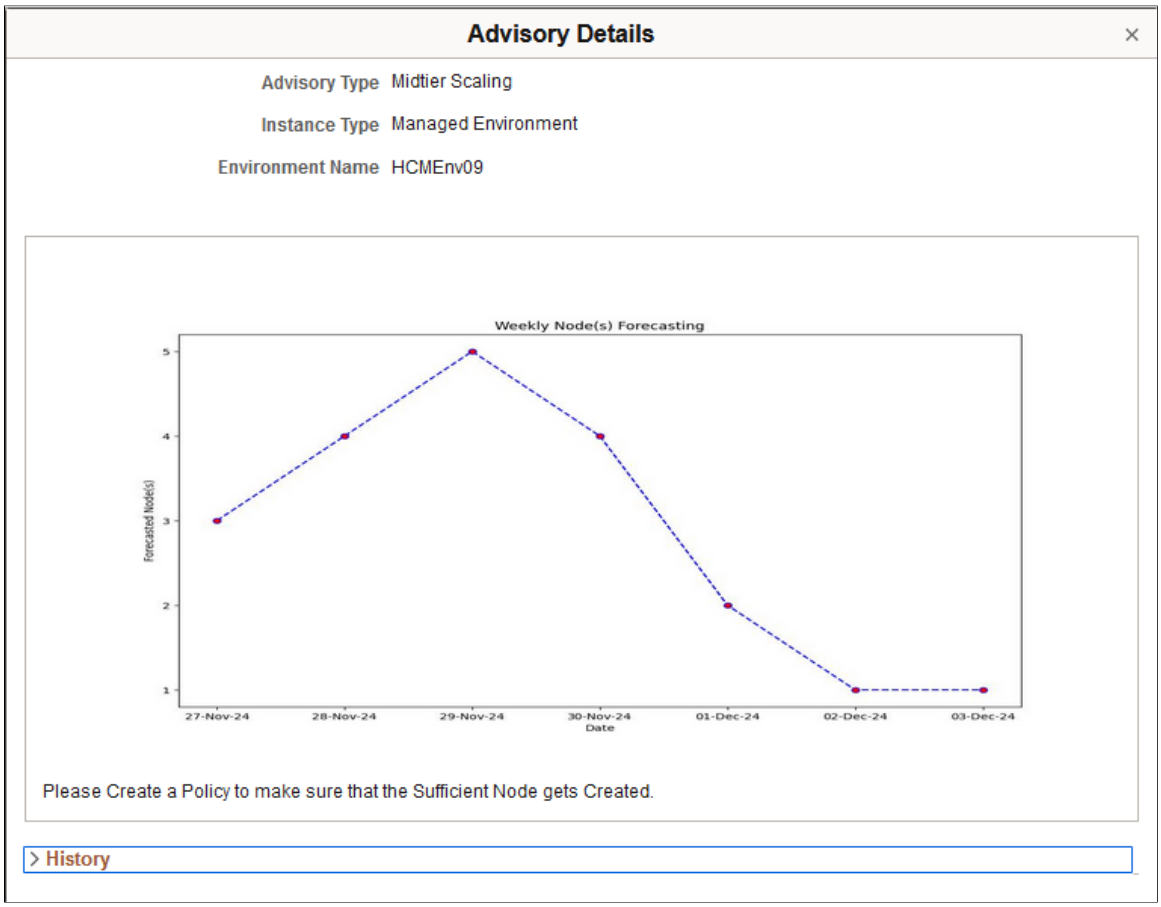

Using Infrastructure CPU Advisories

An Infrastructure CPU advisory is generated when an Infrastructure CPU patch is available in the repository. There are two actions available — Create Policy and Apply Update.

This example illustrates the actions available for an Infrastructure CPU advisory.

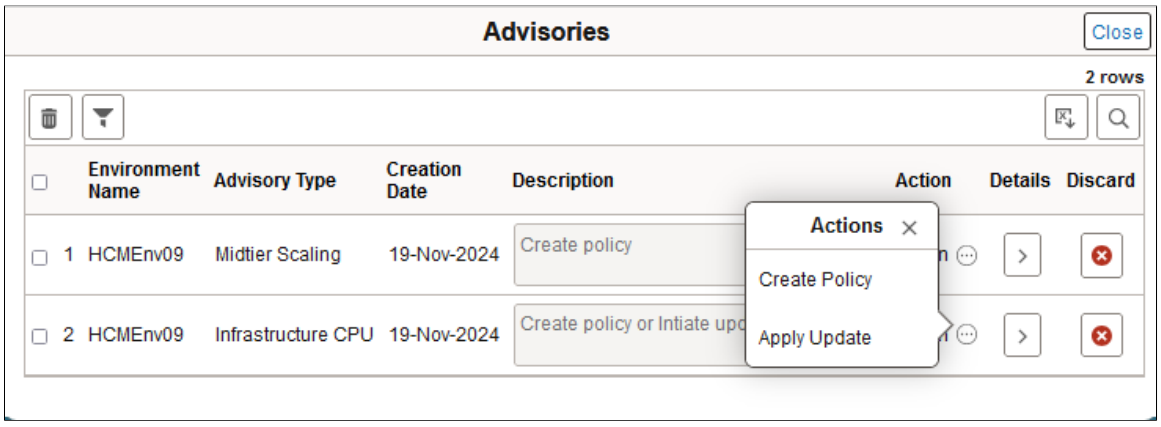

When you receive an Infrastructure CPU advisory, you can apply the Infrastructure CPU update immediately, or create a policy to apply it later. If you choose to apply the Infrastructure CPU update

immediately, select Apply Updates. It will take you to the Infrastructure CPU page in Environment Details. See Applying Infrastructure CPU Patches.

Select Create Policy if you choose to apply the Infrastructure CPU update later. If there is no existing policy, selecting the Create Policy option displays the Policy Editor page. You can review the prefilled policy and click Save to create the policy. You can create a policy with a policy action schedule, which will apply the Infrastructure CPU updates at a specific date and time. If you do not set up a policy action schedule, the policy applies the Infrastructure CPU updates immediately upon download. An Infrastructure CPU policy is created for a specific PeopleTools release such as PeopleTools 8.59 or 8.61. If the new policy is created for Infrastructure CPU, Cloud Manager prompts you to apply the Infrastructure CPU, because this newly created policy will be applicable for the next Infrastructure CPU Download.

If the policy already exists, the Environment Name is added to the existing policy and a message is displayed about this addition.

Click the Details icon for the advisory. The Advisory Details page includes a list of the software components, such as Java, Tuxedo, WebLogic, and Oracle Database client, to be applied with the Infrastructure CPU.

This example illustrates the Advisory Details page for Infrastructure CPU for the selected environment.

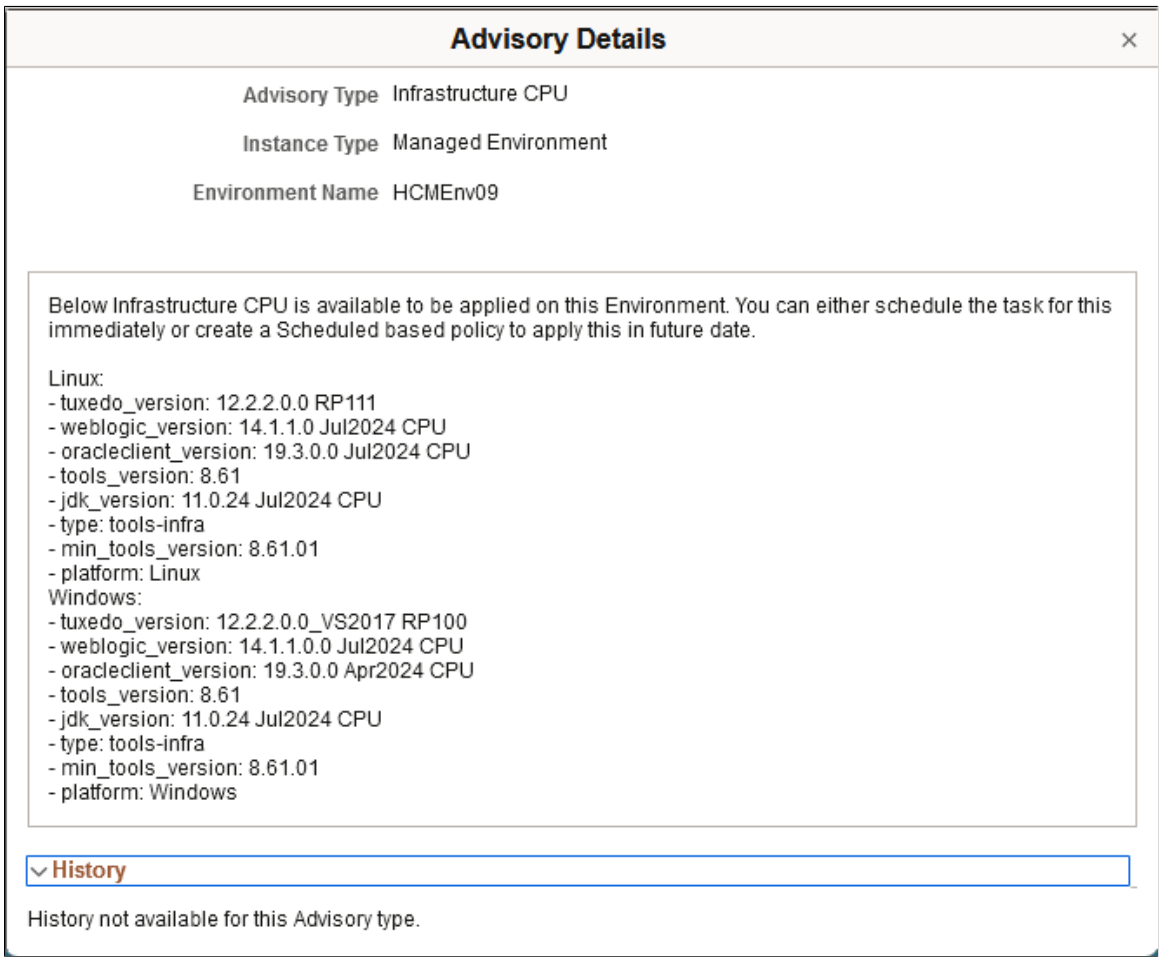

Using Server Level Advisories

Server Level advisories notify you when the available boot volume threshold percentage drops below the threshold set on the Advisory Settings page. The environment owner can consider increasing the boot volume space in OCI. See Enabling Advisory Settings.

This example illustrates the Advisory Details page for Server Level for the selected environment.

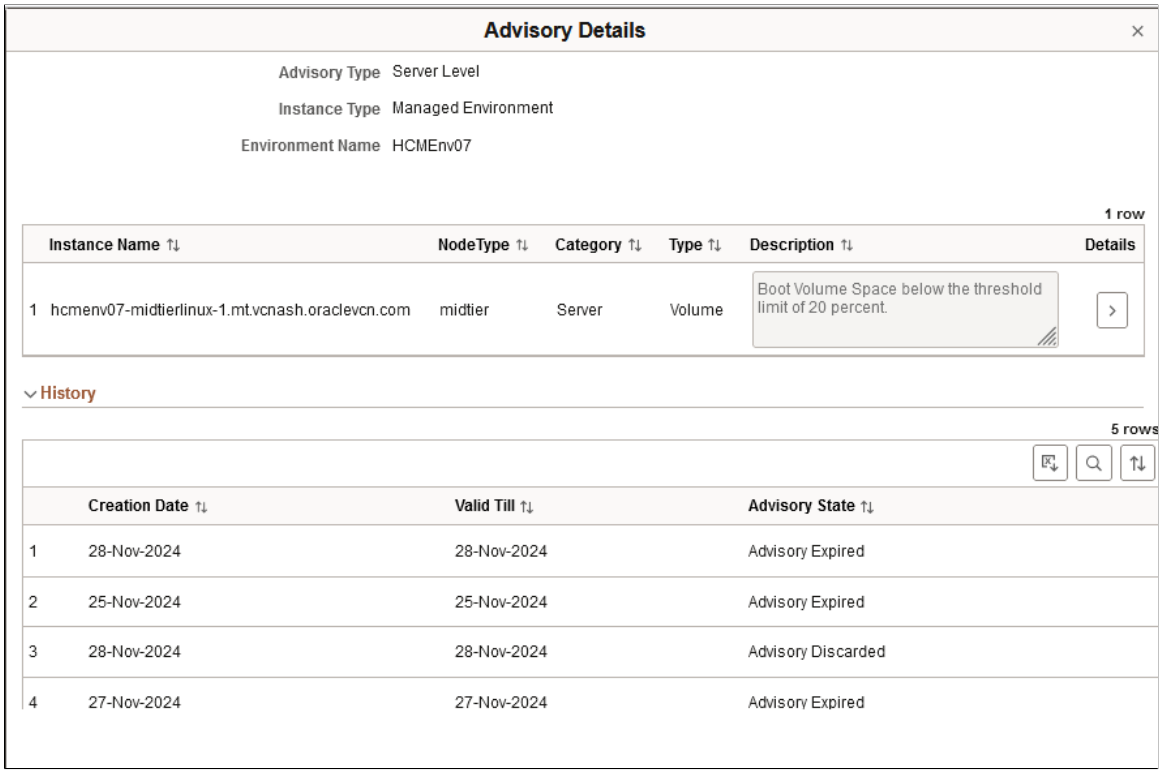

Click the Details icon in the grid at the top of the page. The modal window displays the percentage used for each partition of the boot volume. Use the information to pinpoint which one is over the threshold.

This example illustrates the More Details page for a Server Level Advisory.

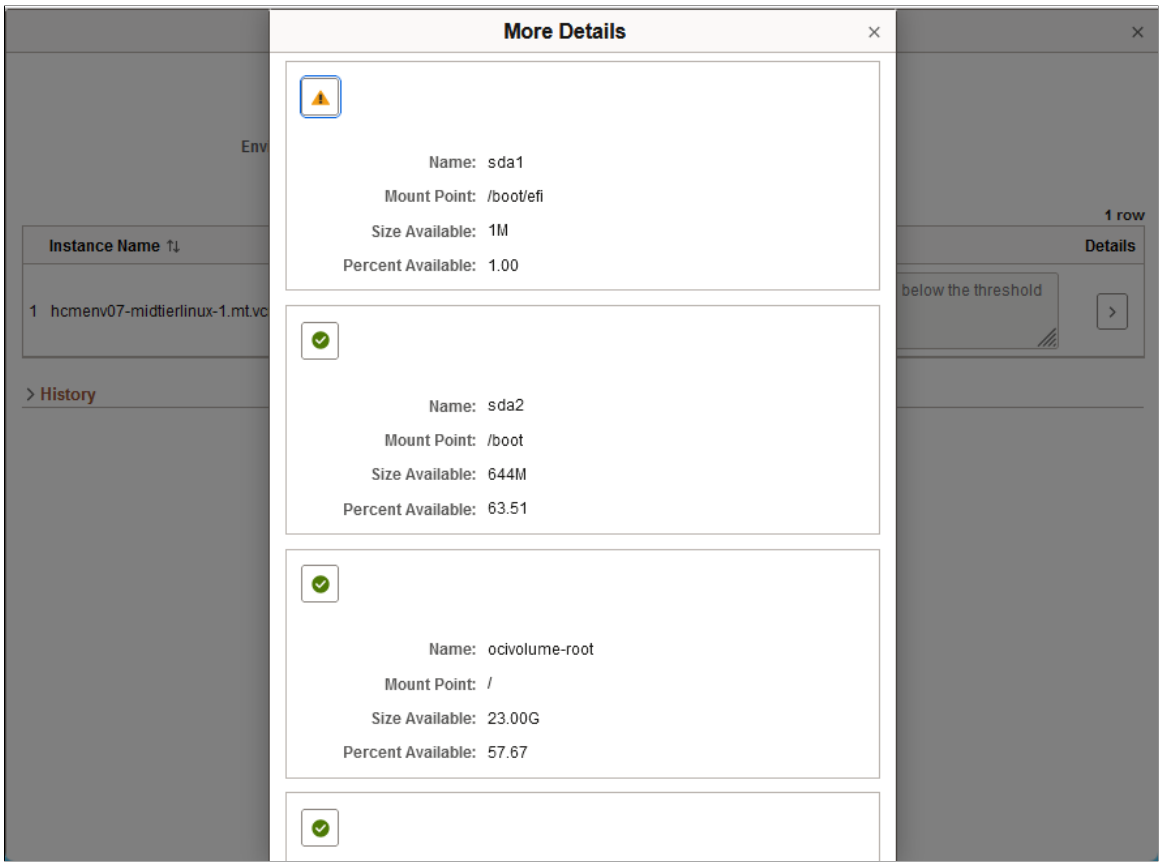

Provisioning and Sharing Search Clusters

A search cluster in Cloud Manager refers to several search nodes that are configured to be used by a provisioned environment.

When you create an environment with a search node, Cloud Manager deploys the node in a single OCI instance (that is, a single Linux VM). Multiple search nodes can be grouped in a cluster. A cluster is identified by a unique name. All of the search nodes must reference the same unique cluster name.

The OpenSearch documentation recommends at least three search nodes for fail-over and high availability. If one search node fails, the search framework will remain available.

This section includes the following topics for using search clusters:

- Provisioning Environments with Search Clusters
- Sharing a Search Cluster Across Multiple Environments
- Managing Search Clusters

For more information, see the product documentation *PeopleTools: Search Technology*, Understanding Clusters in OpenSearch, on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

Provisioning Environments with Search Clusters

Search clusters are supported only for provisioned environments that use OpenSearch. OpenSearch and OpenSearch Dashboards are supported for PeopleTools releases 8.59.21 and later, 8.60.07 and later, and 8.61.

Here are the high-level steps, which are discussed in this section, for provisioning an environment with a search cluster:

1. Add a security rule to open the search cluster transport port on the subnet where search instances are deployed.
2. Create topology with multiple search nodes.
3. Create an environment template that specifies the same search cluster name and port for all search nodes
4. Provision the environment.
5. In the provisioned environment, adjust the replica number in the search instance settings.

Opening the Cluster Transport Port

The node in a cluster communicates through a cluster transport port, which must be opened for the provisioned environment.

The default cluster transport port is 9300. When you provision an environment with a search cluster or add a node to a cluster, the port must be accessible. You enter the port number when creating an environment template.

To enable communication through the cluster transport port, access the details for the VCN used for the environment in the OCI console. Review the security rules for the VCN subnet where search instances are deployed. Verify that the search cluster transport port is accessible.

See the Cloud Manager tutorial Plan the Virtual Cloud Network for PeopleSoft Cloud Manager, on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

Creating Topology for Search Clusters

Create topology with a database node (Database on Compute or DB Systems), one or more mid-tier nodes, and as many search nodes as you want for the cluster. A minimum of three search nodes is recommended for fail-over and high availability. An odd number of search nodes is also recommended.

This example shows a sample topology with database on compute, mid-tier node, and three search nodes.

Topology Information page with three search nodes

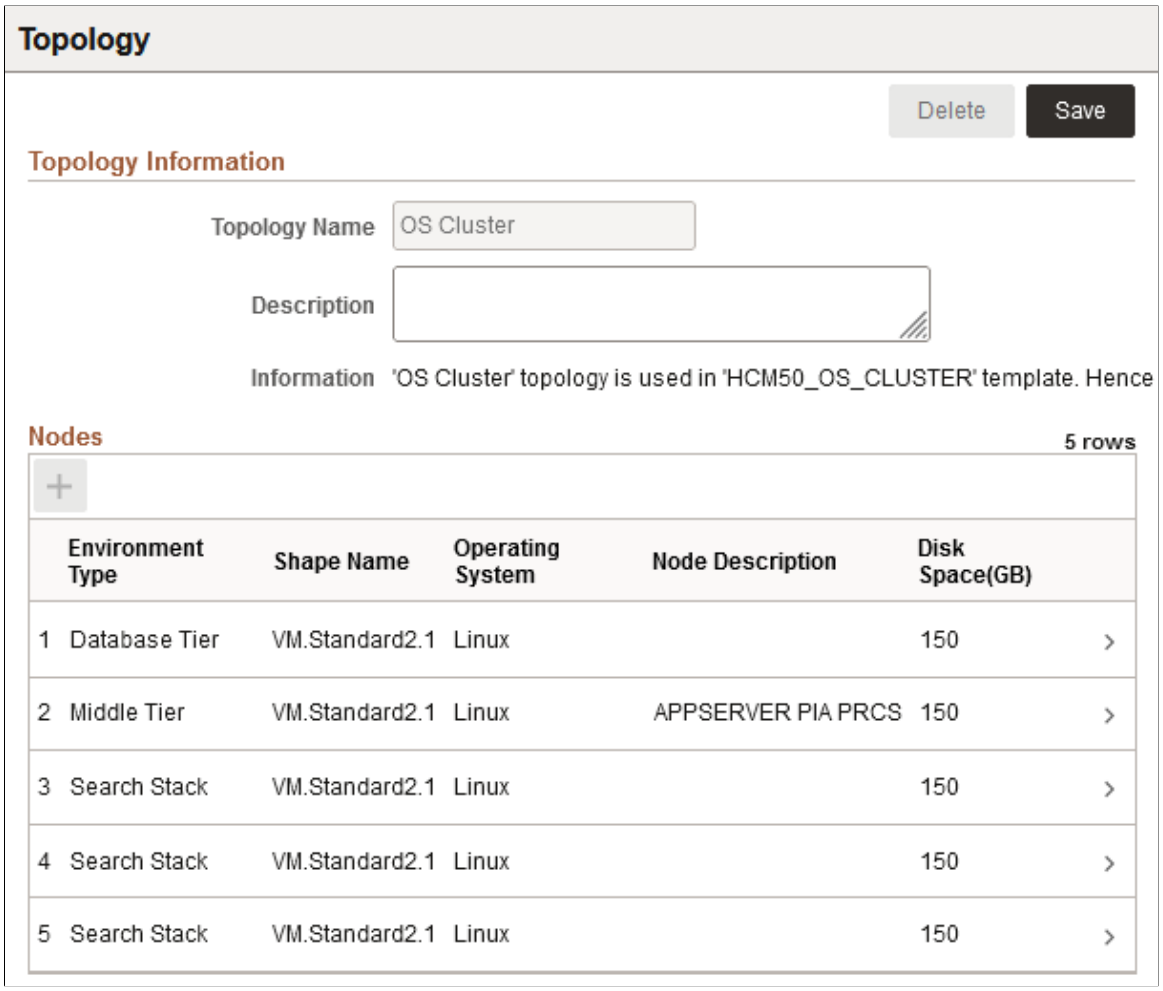

Full-tier environments are not supported for search clusters.

See Creating a New Topology

Creating an Environment Template for Search Clusters

Create an environment template based on the topology with multiple search nodes.

On the Select Topology page, in the Search Stack, General Settings sections, ensure that all of the nodes have the same values for **Cluster Name** (the default is ESCL) and **Cluster Transport Port** (the default is 9300).

This example illustrates the Environment Template - Select Topology page showing the Cluster Name field.

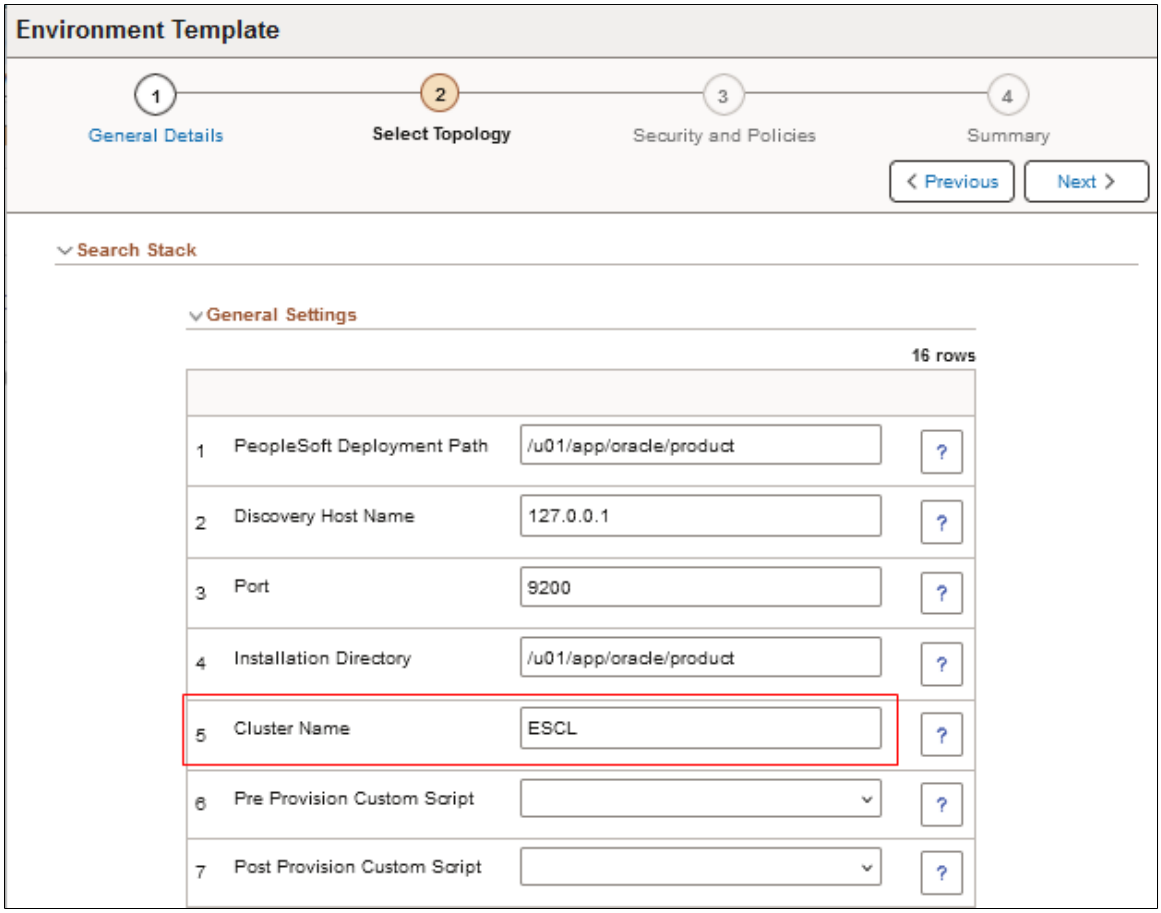

This example illustrates the Environment Template - Select Topology page showing the Cluster Transport Port field.

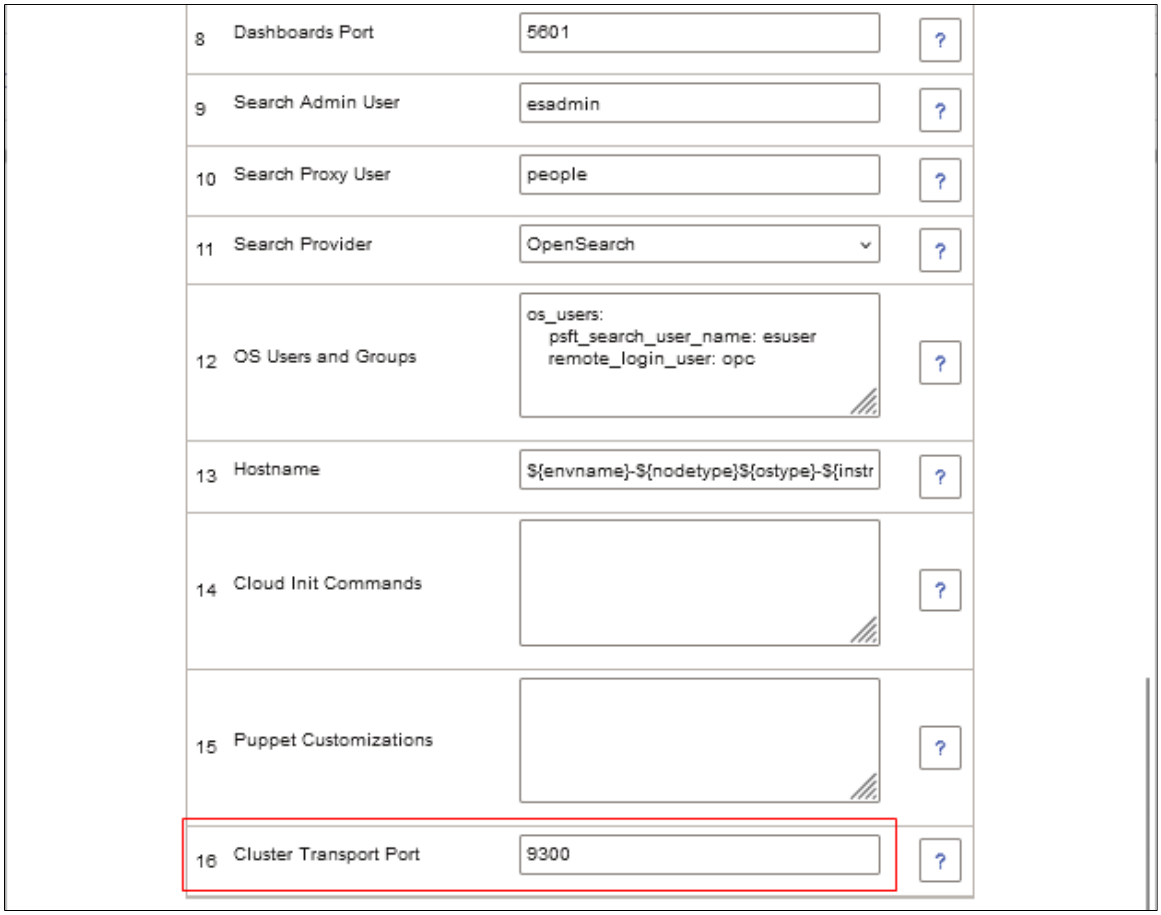

Complete the environment template creation.

See Creating a Template.

Provisioning an Environment with a Search Cluster

Create an environment and complete the configuration.

When you create an environment, the fields are populated with the entries from the selected environment template. Ensure that all of the search nodes include the same cluster name and cluster template port.

After the environment is provisioned:

1. Click the Related Action button for the environment and select Details.
2. Expand the Search section on the Environment Details page to see the search nodes in the cluster.

This example illustrates the Environment Details page for an environment with a cluster comprised of three search nodes.

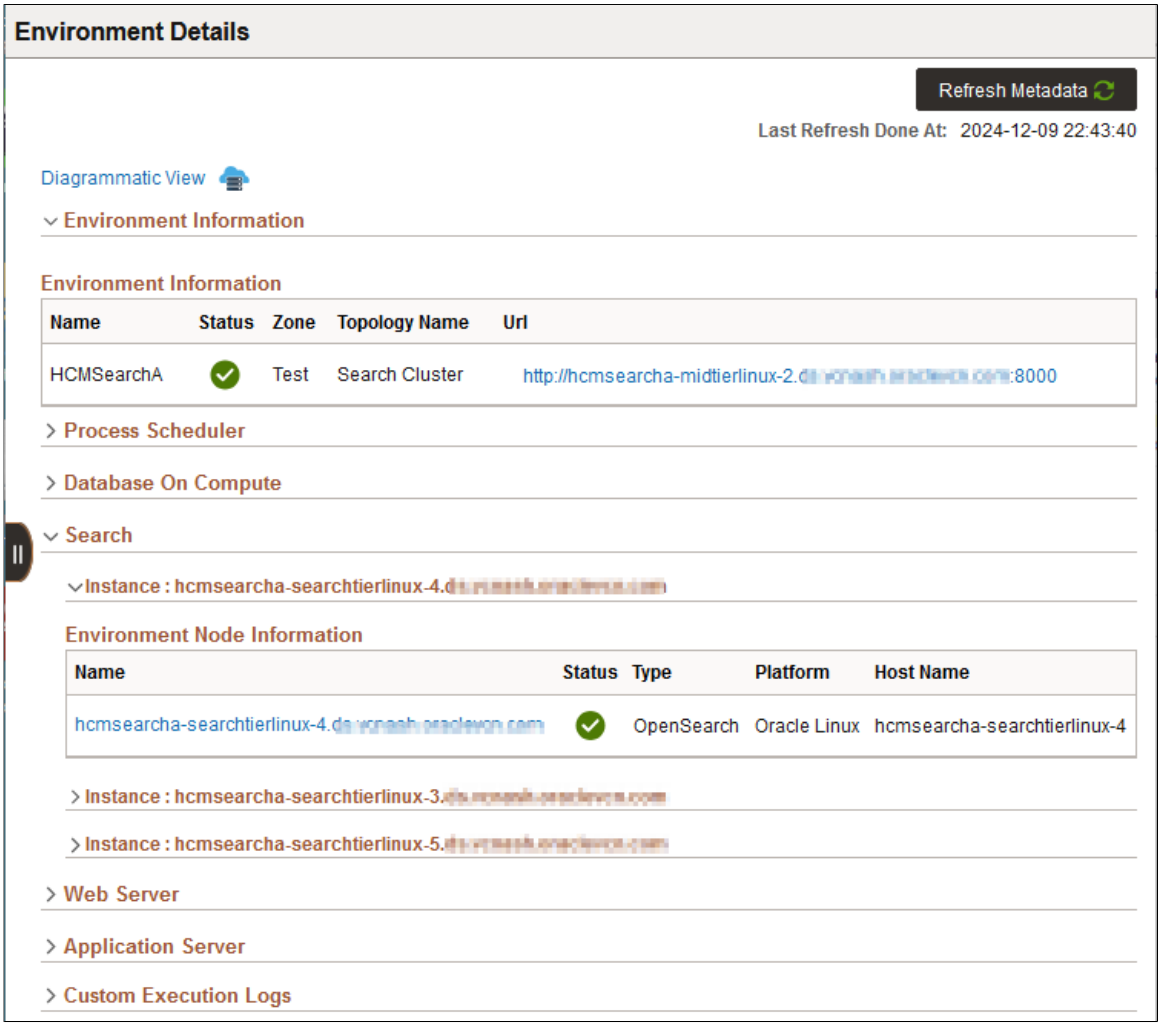

3. Click each instance name for details.

   See Accessing Environment Details

4. Click the URL under Environment Information to sign in to PIA for the environment.

5. Select **PeopleTools** > **Search Framework** > **Search Admin Activity Guide** > **Search Instance Configuration**.

   The PeopleSoft search instance configuration will be updated with metadata for all the nodes in the cluster. This ensures that the search instance configuration works even if a node is down. The Search Instance Properties grid lists the nodes in the cluster.

Note: The PeopleSoft search instance in this context refers to a single instance of the search engine in the PeopleSoft Search Framework.

This example illustrates the Search Instance Properties page with three search instances (1 of 2).

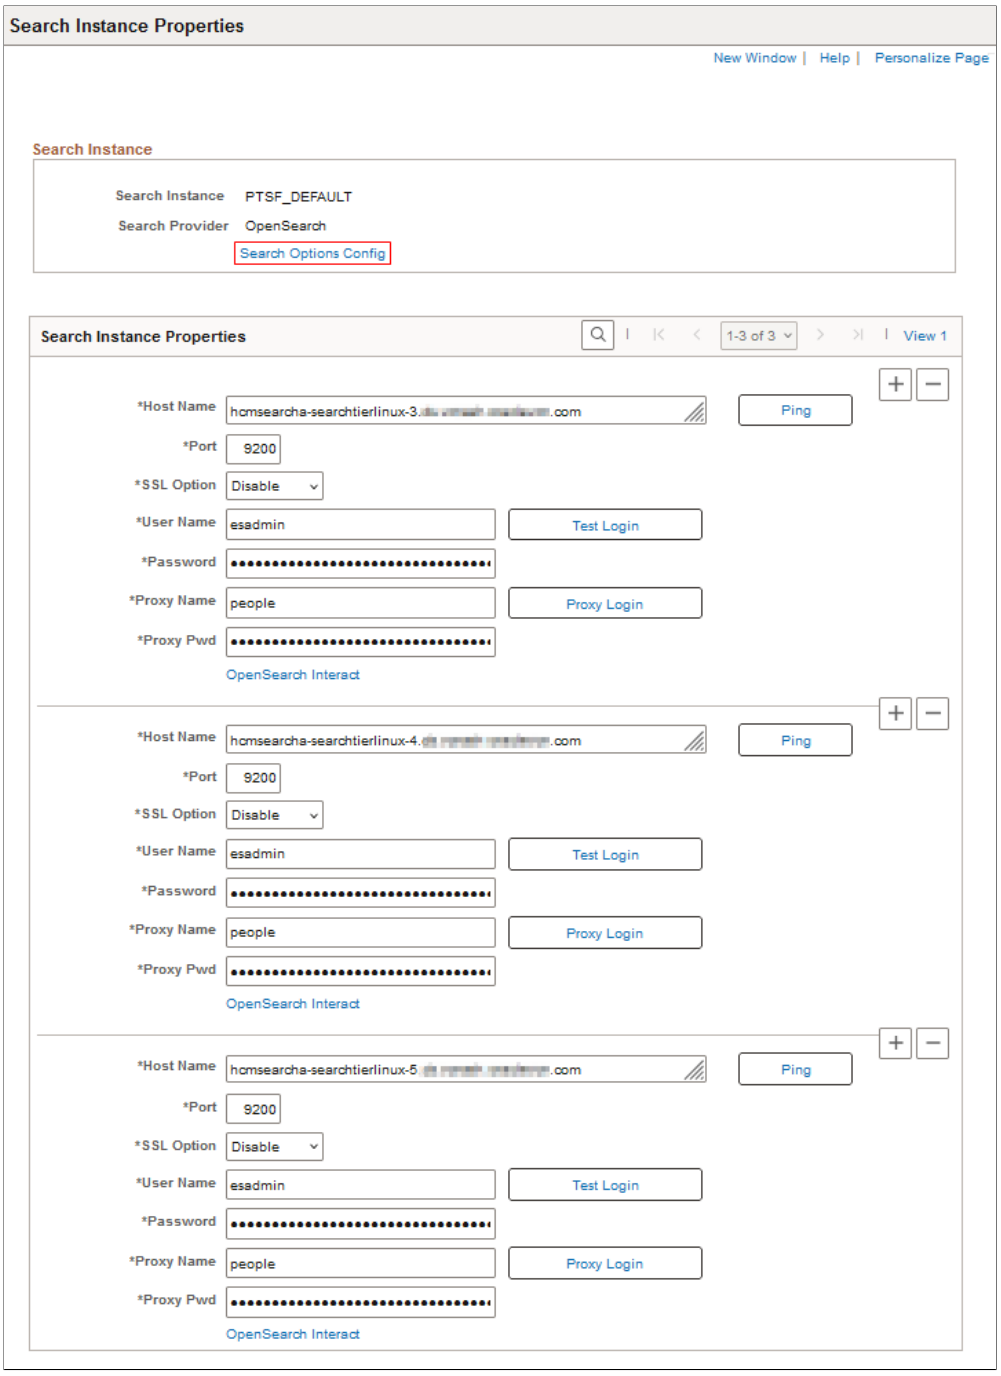

This example illustrates the Search Instance Properties page with three search instances (2 of 2)

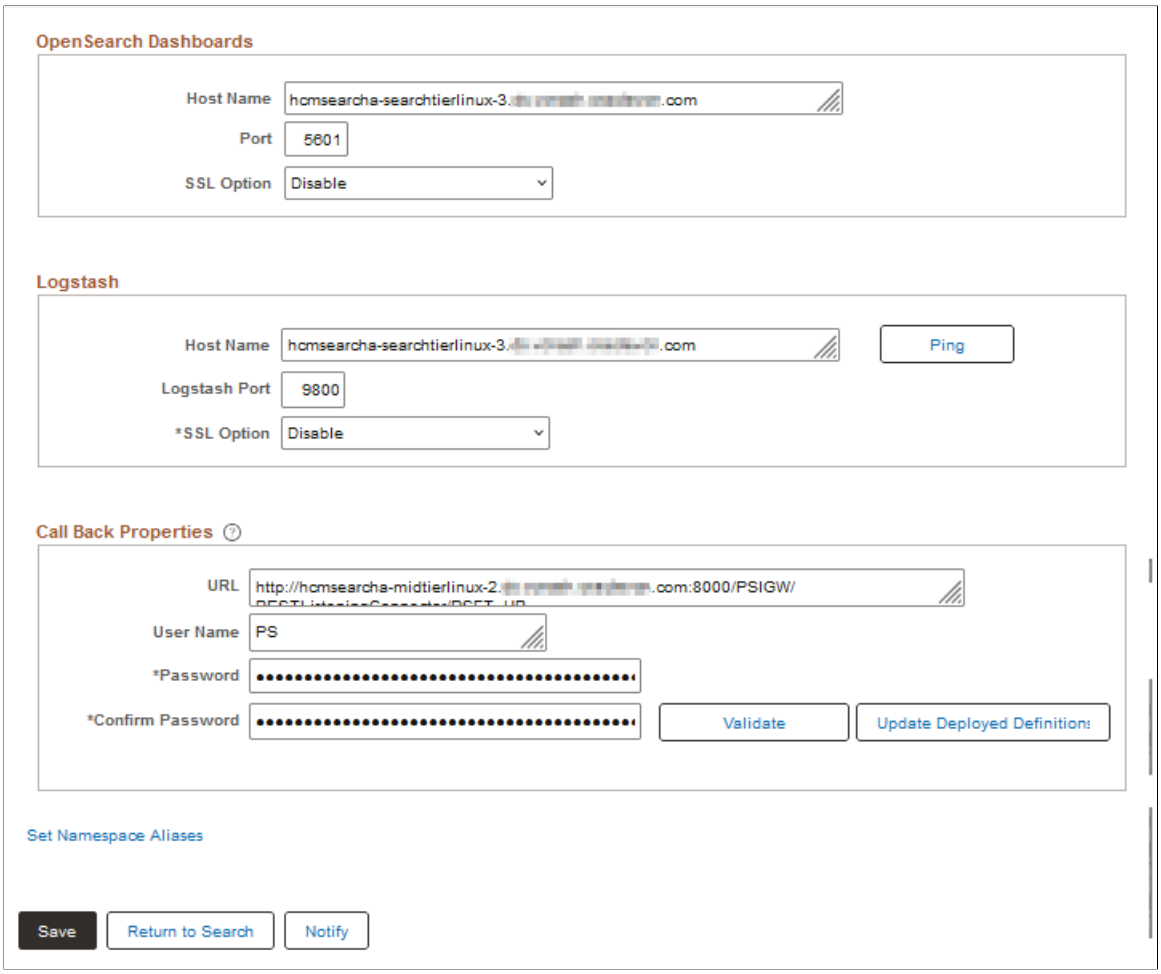

6. Select the link Search Options Config.

   On the Search Options page, ensure that **Number of Replicas** is set to 1 (one) or greater.

This example illustrates the Search Options page with Number of Replicas set to 2.

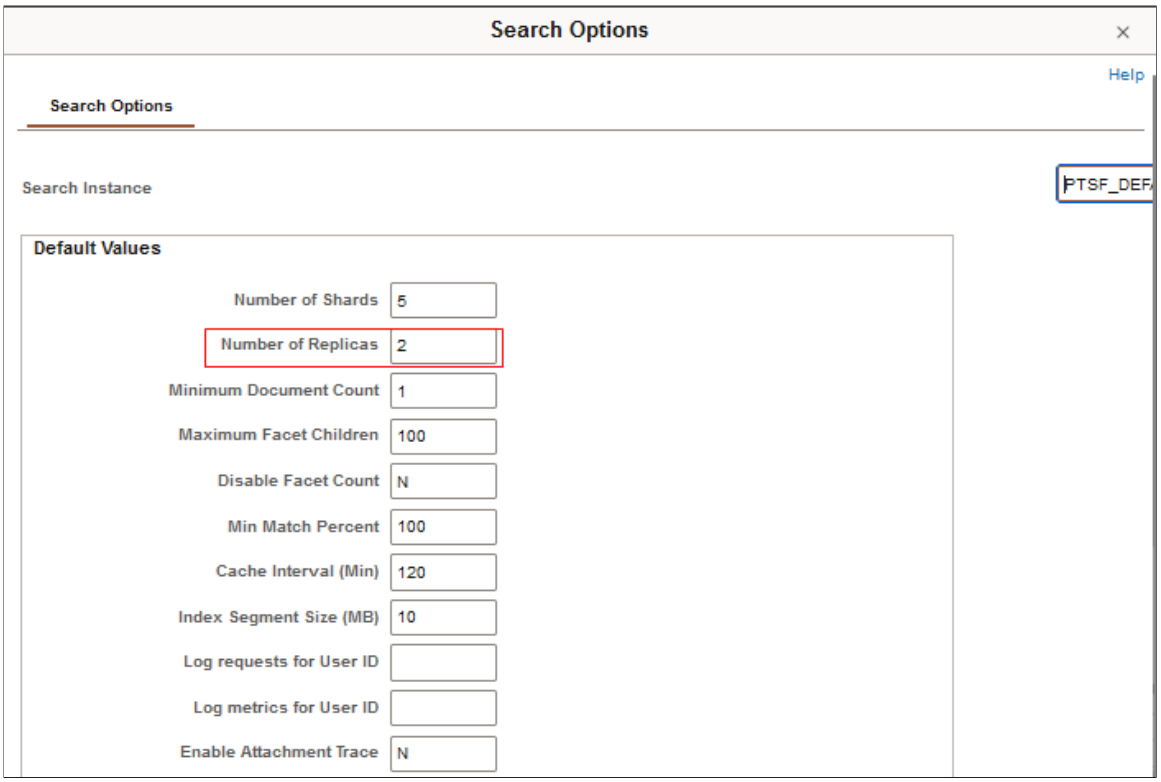

This ensures that search indexes are replicated within the cluster, which protects against loss of data in case a search node fails.

See the product documentation *PeopleTools: Search Technology*, Managing General Search Options and Setting the Number of Replicas, on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html

7. Select **PeopleTools** > **Search Framework** > **Search Admin Activity Guide** > **Administration** > **Deploy/Delete Object**.
8. Select the deployed search definitions, and click **Update**.

   See the product documentation *PeopleTools: Search Technology*, Administering Search Definitions and Search Categories, on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

Sharing a Search Cluster Across Multiple Environments

Set up a search cluster so that it can be used by multiple provisioned environments, referred to here as shared search.

In a shared search setup, search operations for several PeopleSoft environments are performed by a single search cluster. For example, a logging-test cluster could be shared by all test environments and a logging-staging cluster by all staging environments. Shared search clusters can be shared by different PeopleSoft applications, such as FSCM and HCM. Because fewer resources are required for the search component, shared search has the advantage of reducing costs.

The requirements for using shared search are:

- The PeopleTools release must support OpenSearch.

  OpenSearch and OpenSearch Dashboards are supported for PeopleTools releases 8.59.21 and later, 8.60.07 and later, and 8.61.

- The database name for the environment that owns the search cluster and the database names for any environments that share the search cluster must be different.

  Each search index uses the database name as a suffix. This means that multiple data sources from different PeopleSoft applications can be deployed on the same cluster.

- The environment topology includes database tier and middle tier. Full-tier topology is not supported for search clusters.

- The environment must be running for the search cluster sharing to be successful.

- An odd number of search nodes is recommended for a search cluster.

Note: To reset a password on a search node that is part of a shared search cluster, you must manually update it in all environments that share the cluster. See Managing Passwords.

Setting Up Search Cluster Sharing

To share a search cluster, set up the search cluster for the first environment, and then share it from the second.

Here are the high-level steps for sharing a search cluster between two environments:

1. Create an environment with a search cluster.

   See the section Provisioning an Environment with a Search Cluster.

2. Create a second environment whose topology does not include a search cluster or search instances.

3. Use the Manage Node action from the second environment to configure it to share the search cluster on the first environment.

Configuring an Environment to Share a Search Cluster

1. Create topology for an environment with database tier and middle tier. Do not include a search node.

2. Create an environment template and provision the environment.

3. Click the Related Actions button for the environment and select Manage Node.

4. Select **Add** from the Action drop-down list.

5. Select **Search Stack** from the Type drop-down list.

6. Select **Linux** from the Operating System drop-down list.

7. Select **Shared Search** from the Search Topology drop-down list.

Note: This option is available only when there is at least one environment in Cloud Manager that has a search cluster.

8. From the Use Search Cluster in Env drop-down list, select the environment that owns the search cluster.

   This example illustrates the Add Node page for an environment with a search cluster.

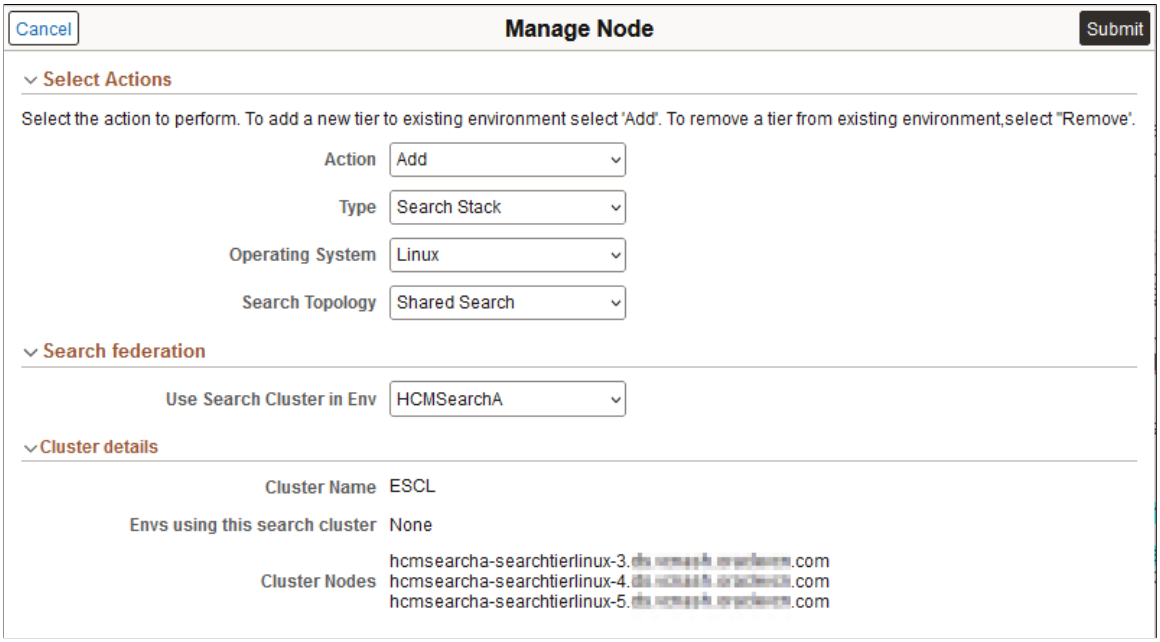

   The Cluster details area displays:

   - Cluster name
   - Environments using the search cluster — Lists the environments that share the search cluster.

     If no other environments share the search cluster, this displays None.
   - Cluster Nodes — Lists the FQDN for all search nodes in the cluster.

9. Click **Submit**.
10. Follow the progress on the Provision Task Status page.

    After the sharing process is complete, the sharing status is displayed on the Environment Details page.

This example illustrates the Environment Details page after sharing a search cluster.

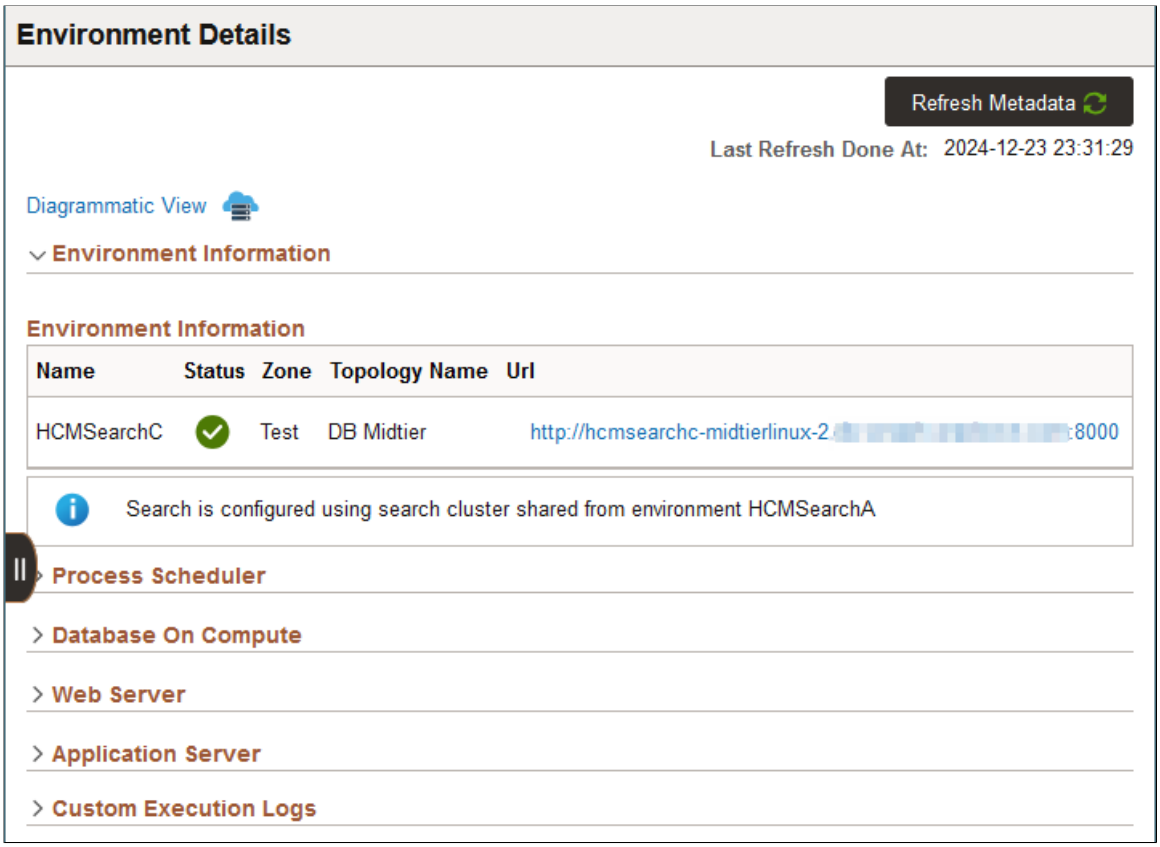

11. Sign in to PIA for the environment and set the number of replicas.

    To set the number of replicas, see the section Provisioning an Environment with a Search Cluster.

Resolving VP1 User Error

For certain PeopleSoft environments such as FSCM, which use VP1 as the default user, the sharing search cluster operation may fail. In case of failure:

1. Click the Related Actions button for the failed environment and select Details.
2. Click the URL under Environment Information and sign in to PIA for the environment.
3. Select **PeopleTools** > **Automated Configuration Manager** > **ACM Configuration Monitor**.
4. If there is a Failed status for the search operation, click the Details icon.
5. If you see the following text in the message, continue with the next step.

   "User Role check VP1. User does not have the role: Report Distribution Administrator"

6. Select **PeopleTools** > **Security** > **User Profiles** > **User Profiles**.
7. Select the VP1 user and add the role ReportDistAdmin, then save the page.

8. In Cloud Manager, go to the details for the PeopleSoft environment, and retry the failed task from the Provision Task Status page.

Removing an Environment from a Shared Search Cluster

Use these steps to unshare—that is, remove an environment from a shared search cluster.

All nodes in a cluster must be up and running in order to remove an environment from a shared search cluster.

1. Click the Related Actions button for the environment and select Manage Node.
2. Select **Remove** from the Action drop-down list.
3. Select **Search Stack** from the Type drop-down list.
4. Select **Linux** from the Operating System drop-down list.
5. Select **Shared Search** from the Search Topology drop-down list.

   The Use Search Cluster in Env field displays the name of the environment that owns the search cluster, the cluster name, and the search nodes, in read-only text.
6. Click **Submit**, which removes the association of the environment with the shared search cluster.

If the process to remove the environment fails, go to Process Monitor. Check for any search-related processes, such as PTSF_GENFEED and related processes, that are queued or in progress. If this is the case, it means that the environment is in the process of deploying search indexes. The system does not allow the unsharing to proceed in this case, to avoid data loss or corruption.

After the search index processes are complete, retry the unsharing from the Provision Task Status page.

Managing Search Clusters

This section describes the following actions for environments with search clusters.

- Add a node to the search cluster.
- Remove a node from the search cluster.
- Clone an environment with a search cluster.
- Import an environment with a search cluster.

Adding a Node to an Existing Search Cluster

Use the Manage Node page to add a search node to an environment. You can add one or more search nodes to an environment that does not have any search nodes to form a cluster. You can also add a search node to an environment with an existing search cluster. All search nodes must be up and running for the operation to work.

Adding a new search node to a managed environment that already has a search node will not form a search cluster, if the managed environment was provisioned in a Cloud Manager image prior to Image19 (Cloud Manager upgrade case). To form a search cluster in such a scenario, a customer needs to remove the existing search node from the managed environment and add new nodes.

Note: On environments that own a search cluster, adding a node only affects the search instance configuration on the environment that owns the cluster. In the environments that share the cluster, you must modify the search instance configuration manually to indicate that the nodes were added. See the product documentation *PeopleTools: Search Technology*.

1. Click the Related Actions button for the environment and select Manage Node.
2. Select **Add** from the Action drop-down list.
3. Select **Search Stack** from the Type drop-down list.
4. Select **Linux** from the Operating System drop-down list.
5. Select **Add Search Node** from the Search Topology drop-down list.
6. Use the Copy From drop-down list to select an existing search node from which to copy the configuration.
7. Supply the necessary attributes and settings.
8. Click **Submit**, which carries out these actions:
   - The node is added to the existing search cluster.
   - All existing nodes in the cluster are updated with the metadata for the added node.
   - The PeopleSoft search instance configuration is updated with the information about the added node.
9. Sign in to PIA for the environment and set the number of replicas.

   If you are adding search nodes to an environment with an existing search cluster, it is recommended to set the number of replicas after all of the search nodes have been added.

   To set the number of replicas, see the section Provisioning an Environment with a Search Cluster.

Removing a Node From an Existing Search Cluster

Use the Manage Node page to remove a search node from an existing search cluster. All search nodes must be up and running for the operation to work.

Note: On environments that own a search cluster, removing a node only affects the search instance configuration on the environment that owns the cluster. In the environments that share the cluster, you must modify the search instance configuration manually to indicate that the nodes were removed.

1. Click the Related Actions button for the environment and select Manage Node.
2. Select **Remove** from the Action drop-down list.
3. Select **Search Stack** from the Type drop-down list.
4. Select **Linux** from the Operating System drop-down list.
5. Select **Remove Search Node** from the Search Topology drop-down list.

6. Select an existing search node to remove.
7. Click **Submit**, which carries out these actions:
   - The node is removed from the existing search cluster.
   - All existing nodes in the cluster are updated to remove the node metadata.
   - The PeopleSoft search instance configuration is updated with the information about the removed node.
8. Sign in to PIA for the environment and update the number of replicas.

   To set the number of replicas, see the section Provisioning an Environment with a Search Cluster.

Cloning an Environment with a Search Cluster

When you clone an environment with a search cluster, a new search cluster is formed.

When you clone an environment with an associated search cluster, you must clone all of the search nodes in the cluster or none of them. You cannot select a subset of search nodes. On the Clone page, ensure that Select to Clone is selected for all the search nodes.

The data in the cloned search cluster will be an exact replica of that in the original search cluster. The name of the search cluster on the cloned environment will be the same as that of the source environment.

When you clone an environment that shares a search cluster (owned by another environment), the cloned environment does not inherit the sharing.

This example illustrates the Clone page with a search cluster.

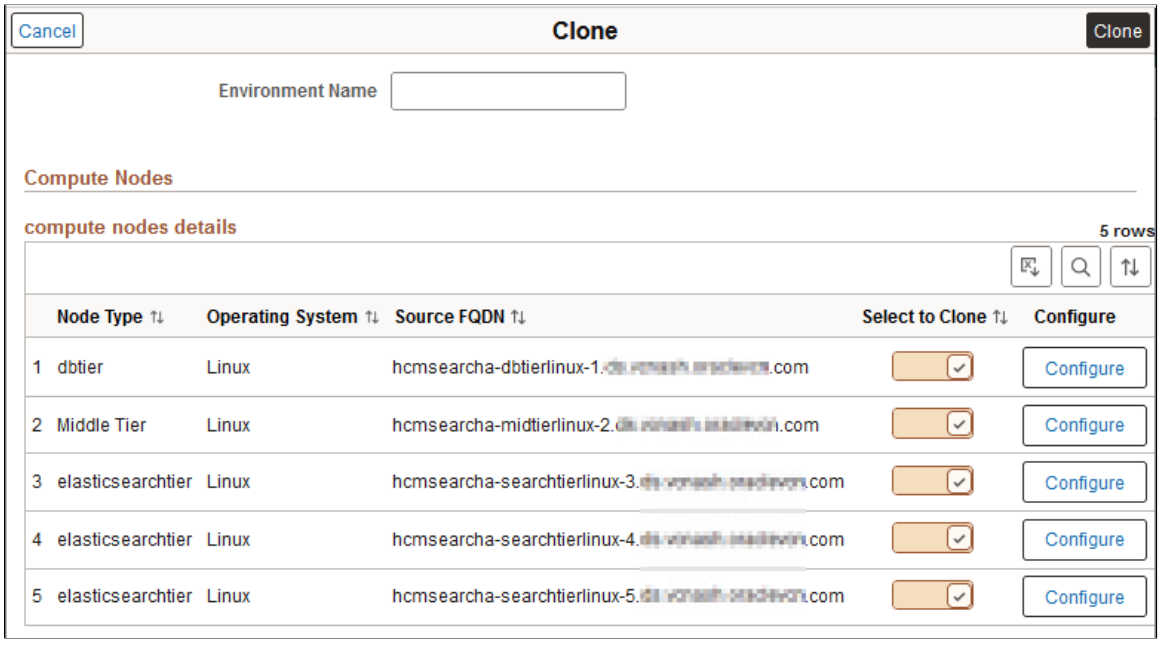

See Cloning Environment.

Importing an Environment with a Search Cluster

When you import an environment with a search cluster, you must import all of the search nodes in the cluster or none of them. You cannot select a subset of search nodes.

See Search Stack Instance Type.

Setting Up Unified Navigation Clusters

A unified navigation cluster setup enables seamless single signon (SSO) and navigation between multiple PeopleSoft applications (HCM, FSCM, ELM, CRM) from a single homepage. Use the Environment Cluster Setup page (ECL_CLUST_SETNG_FL) to create and manage unified navigation clusters.

After defining the portal system and content nodes on the Cluster Settings page, Cloud Manager opens Policy Editor. You verify the required information on the delivered Configure IB PORTAL Cluster policy and run it to complete the setup

A unified navigation cluster in Cloud Manager is comprised of a portal system node and multiple content nodes. You can designate any PeopleSoft application as the portal system node that federates all participating content nodes into a single cluster. A content node is a PeopleSoft application (HCM, FSCM, ELM, CRM) that provides content to the portal system.

Note that the PeopleTools documentation uses the term "content provider" for the content nodes. For more information, see the product documentation *PeopleTools: Portal Technology*, Understanding Unified Navigation, on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

Prerequisites

A unified navigation cluster must fulfill these requirements:

- Common gateway: Each cluster must have a common, or shared, Integration Broker (IB) gateway.

  The IB configuration for each environment specifies the shared gateway URL.Select the PeopleSoft environment with the latest PeopleTools release to serve as the common gateway. All other nodes in the cluster will point to the common gateway for communication.

- Unique content nodes: Each content node (environment) added to a cluster must be unique in terms of its application type.

  For example, you cannot add a second FSCM environment to a cluster with FSCM and HCM environments.

- Integration Broker: At least one node (a host that contains an application server) within each environment in a cluster must have Integration Broker configured to facilitate seamless data exchange.

  For example, if an environment includes multiple mid-tier nodes, ensure that at least one has Integration Broker configured.

- Process Scheduler domain: The environments used for the cluster must include at least one Process Scheduler domain.

The environments can be full-tier or combinations of mid-tier and database tiers. Process Scheduler is required to run the Automated Configuration Manager plugins for the cluster setup.

- Common authentication domain: All environments must use the same authentication domain.

  The authentication domain is part of the template definition, under web server settings.

- Node names: When you configure a unified navigation cluster or shared IB gateway through Cloud Manager, do not use the names of delivered nodes, such as PSFT_EP for FSCM or PSFT_CS for Campus Solutions.

  Enter a name of 5-15 characters. Use only uppercase letters, numbers, and underscores (_) for node names.

  To review the delivered nodes, select **PeopleTools** > **Integration Broker** > **Integration Setup** > **Node Definitions.**. See the product documentation concerning nodes in *PeopleTools: Integration Broker Administration* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

- Common Portal User: The same default portal user (for example, VP1 or PS) must exist on the portal system node and content nodes.

  If the user account is missing from either system, it must be manually created. This step ensures that the user has the necessary access to both the portal system and content node environments, which is essential for proper functionality within the cluster. After creating a new user on a portal system node or content node, restart the node.

- Common Portal User Passwords: The password for the default portal user in the portal system node must match the password for the same user in the content nodes.

  The password for the default content user should also be same in both the systems.

  This synchronisation ensures that both systems can authenticate the same user seamlessly across the cluster. After changing passwords, restart the node.

- Common Gateway keystore password: The password must be the same on the portal system node and content nodes in order for SSO to work.

Creating a New Unified Navigation Cluster

Set up a new cluster with environments created in Cloud Manager. Access the environment that you have chosen as the portal system. Define the portal system and content nodes, and then run a policy to create the cluster.

1. Click the Environments tile on the Cloud Manager home page.
2. Click the environment card to access the details for the environment that you have chosen as the portal system.

   Alternatively, click the Related Action button for the environment and select **Details**.
3. Select Cluster Settings on the left panel of the Environment Details page.
4. Click **Create Cluster**.

On the Create Cluster dialog box, the Portal System field is automatically populated with the environment name.

This example illustrates the Create Cluster dialog box.

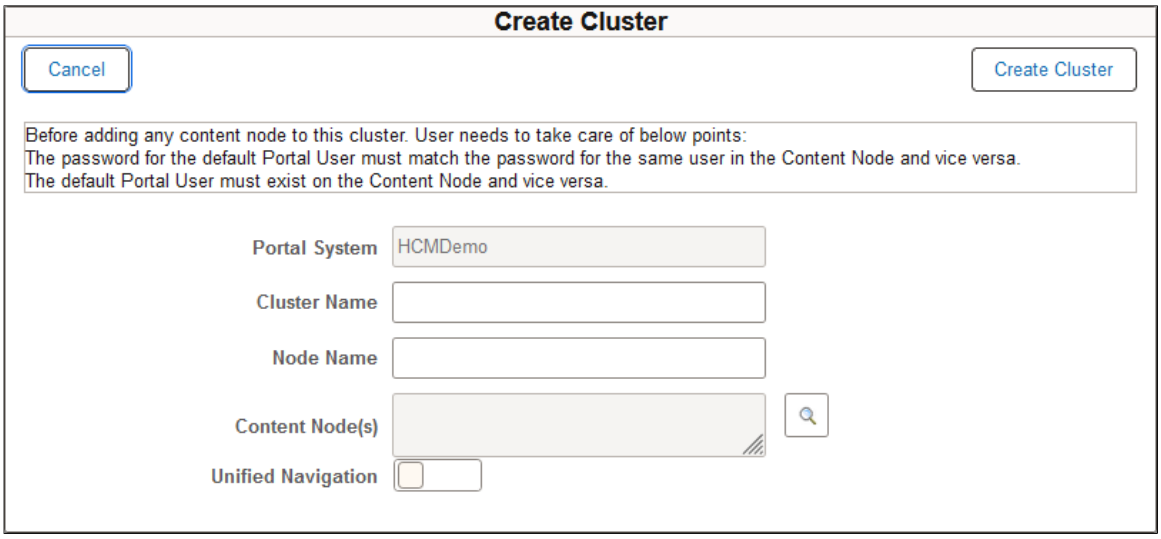

5. Enter a Cluster Name.

   The cluster name is a logical identifier for recognizing the nodes within the cluster. Enter a name between 5-15 characters. Use only letters, numbers, and underscores (_).

6. Enter a Node Name for the Portal System node.

   Enter a name of 5-15 characters. Use only uppercase letters, numbers, and underscores (_) for node names.

7. Click the browse icon and select one or more environments to serve as content nodes.

   You can only use environments that are not already part of another cluster.

   This example illustrates the Environment Name dialog box.

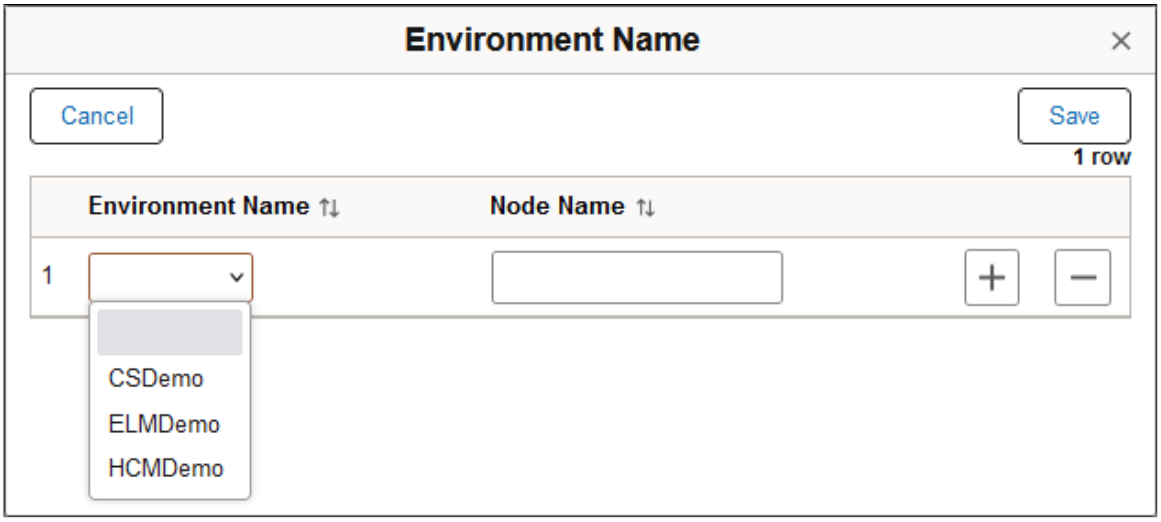

8. Enter a node name.
9. Click the plus icon to add more content nodes, or click **Save** to close the dialog box.
10. To register the menu from the content node into the portal system menu, enable the **Unified Navigation** option.

    When you enable this option, a top-level menu item for the content node will be created in the portal system navigation menu.

    If you do not enable this option, the remote folder from the content node will not be registered in the portal system, but communication between the nodes will still happen through the shared gateway.
11. Click **Create Cluster**.

    Cloud Manager opens the Policy Editor page for the policy Configure IB PORTAL Cluster<internal>. This is a delivered policy that cannot be deleted. It is an Adhoc policy that you run immediately after completing the configuration.

    The page opens with the following populated attributes, which cannot be modified:

    - Name — Configure IB PORTAL Cluster<internal>
    - Policy Object — Environment
    - Policy Type — Adhoc
    - Policy Conditions — Selected Portal System and Content Nodes environments in Running status.

This example illustrates the Configure IB PORTAL Cluster<internal> policy.

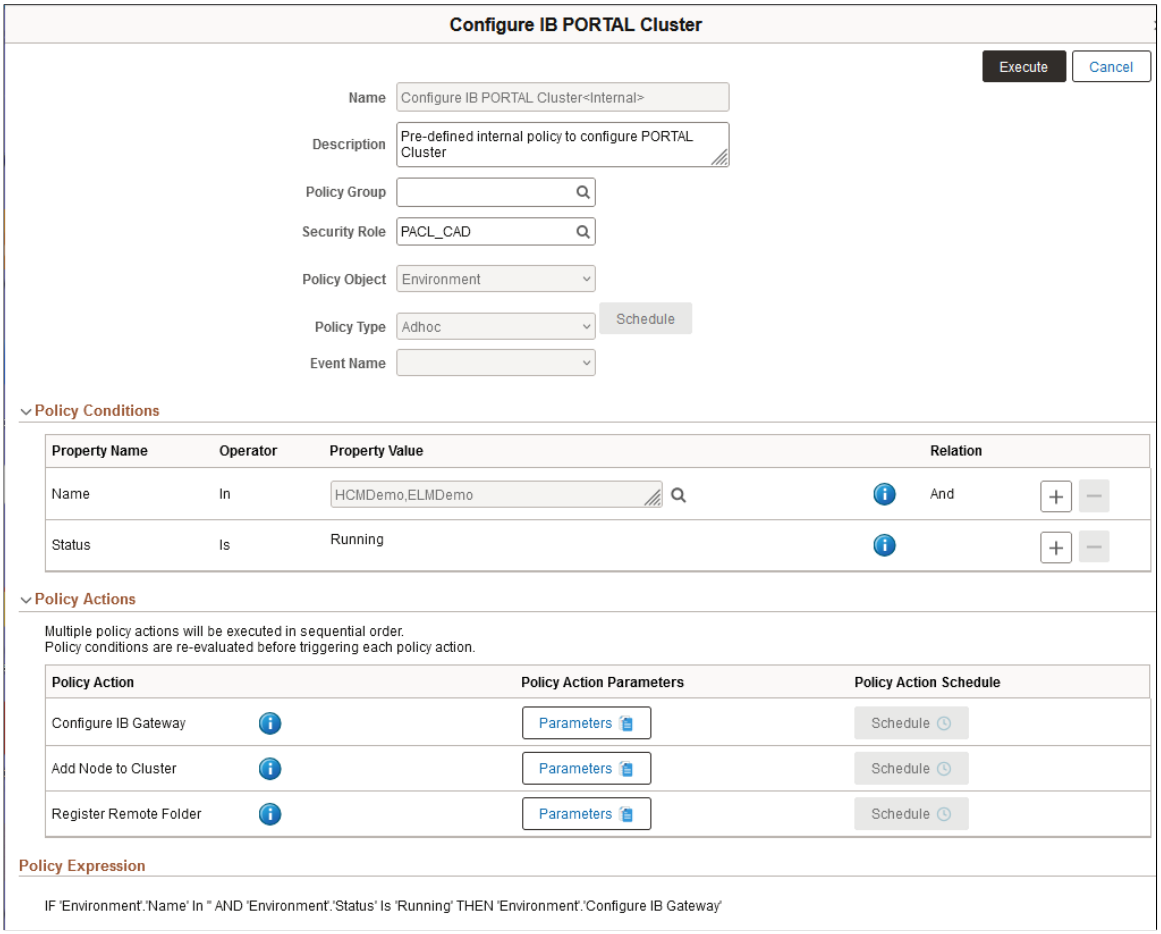

12. Select the Policy Group if desired.
13. Under Policy Actions, click the **Parameters** button for the Configure IB Gateway action.

    Verify or supply the required data and then click **Apply** to return to the policy.

This example illustrates the Policy Action Parameters dialog box.

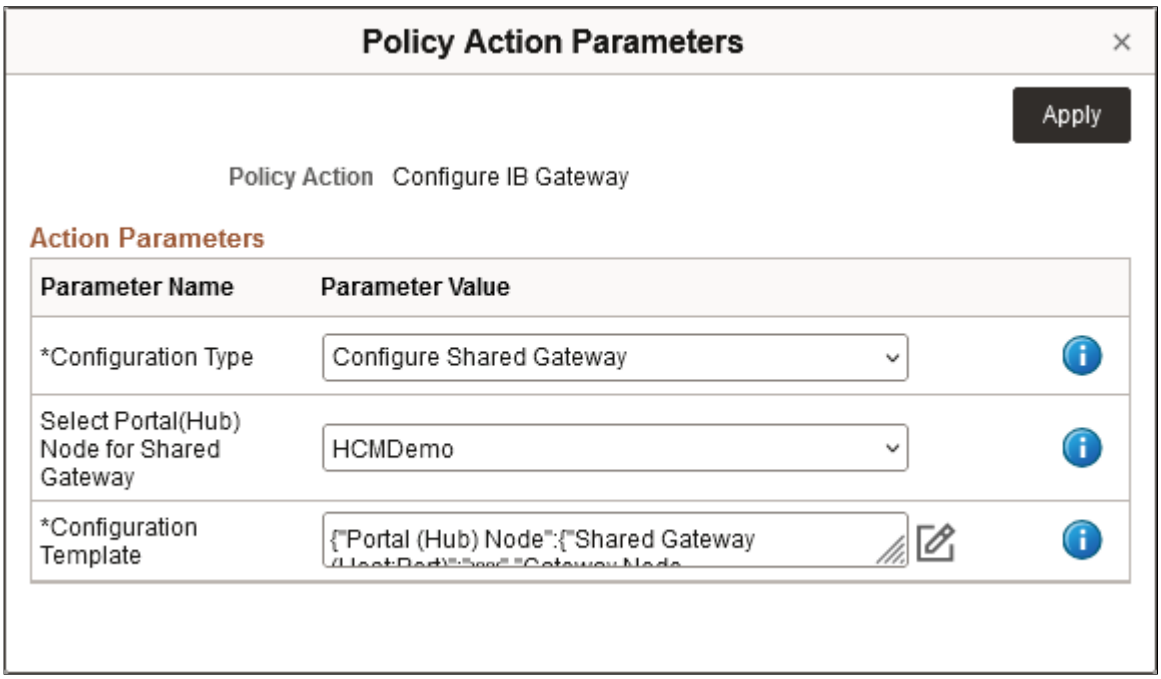

Field or Control	*Description*
Configuration Type	Accept the value Configure Shared Gateway. This is the only acceptable configuration type when creating a unified navigation cluster.
Select Portal (Hub) Node for Shared Gateway	The field is populated with the name for the environment from which you launched the policy.
Configuration Template	Click the edit icon to open a form to define parameters required for IB gateway configuration. The context data section at the top includes data for the portal system and content nodes from the cluster associated with the policy. In the Configuration Template (YAML Format) field verify the values and make any required changes. Copy and paste attribute values from the context data section. Update the values for Shared Gateway (HostPort) and Gateway Node App Server URL. The password values include default masking characters ('***'). If you want to use the same passwords used in provisioning the environment, do not change the values. If you want to change the passwords, do not enter passwords in clear text. You must supply the vault OCID values. The OCIDs do not require single quotes. After supplying the values, click **Save** to return to the Policy Action Parameters dialog box.

This example illustrates the Configuration Template for the Configure IB Gateway policy action.

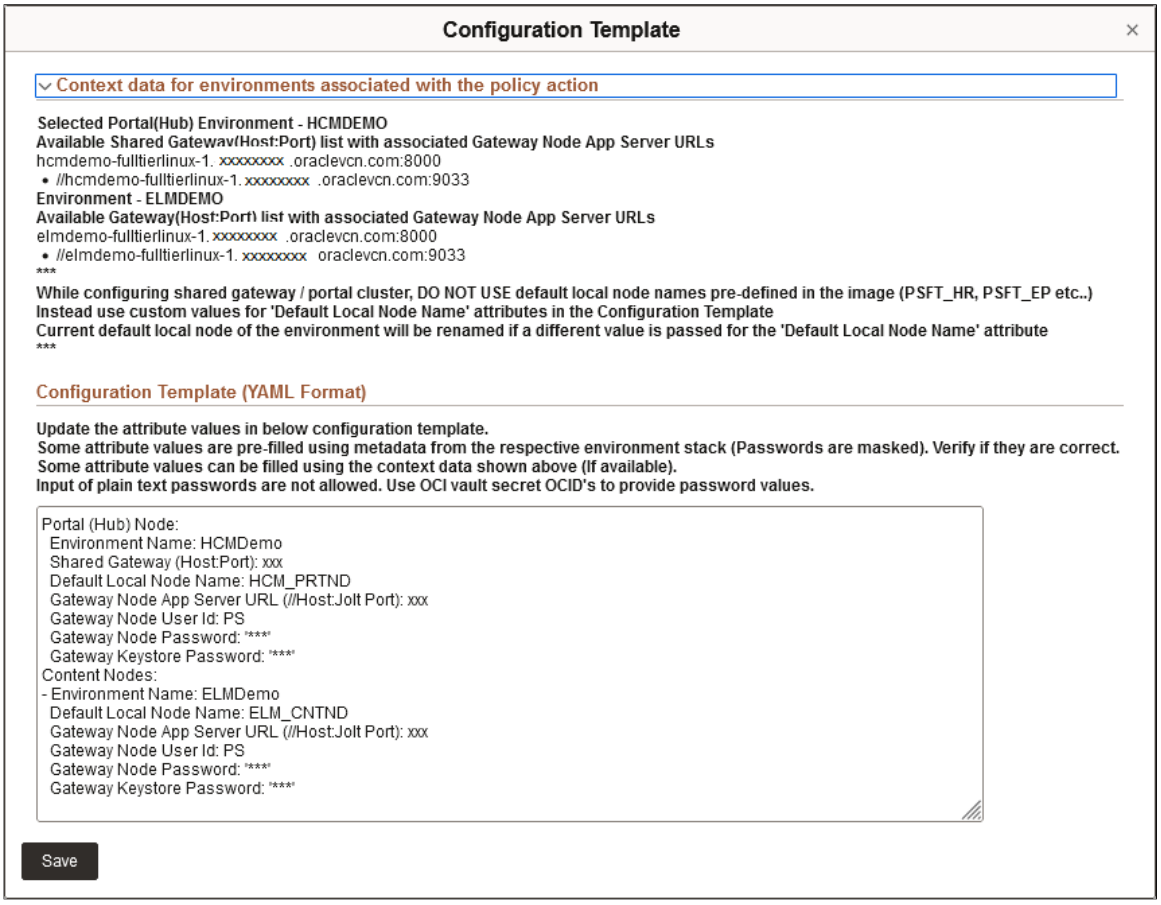

14. Click the **Parameters** button for the Add Node to Cluster action.

    This action registers each content node to the IB network, and configures the portal node in the portal cluster.

    The only parameter is a Configuration Template. Click the edit icon to open the template. The attribute values are populated with the IB node names and environment names for the portal system and content node. You can accept the values, or enter new node names. Review the information and then click **Save** to close.

This example illustrates the Configuration Template for the Add Node to Cluster policy action.

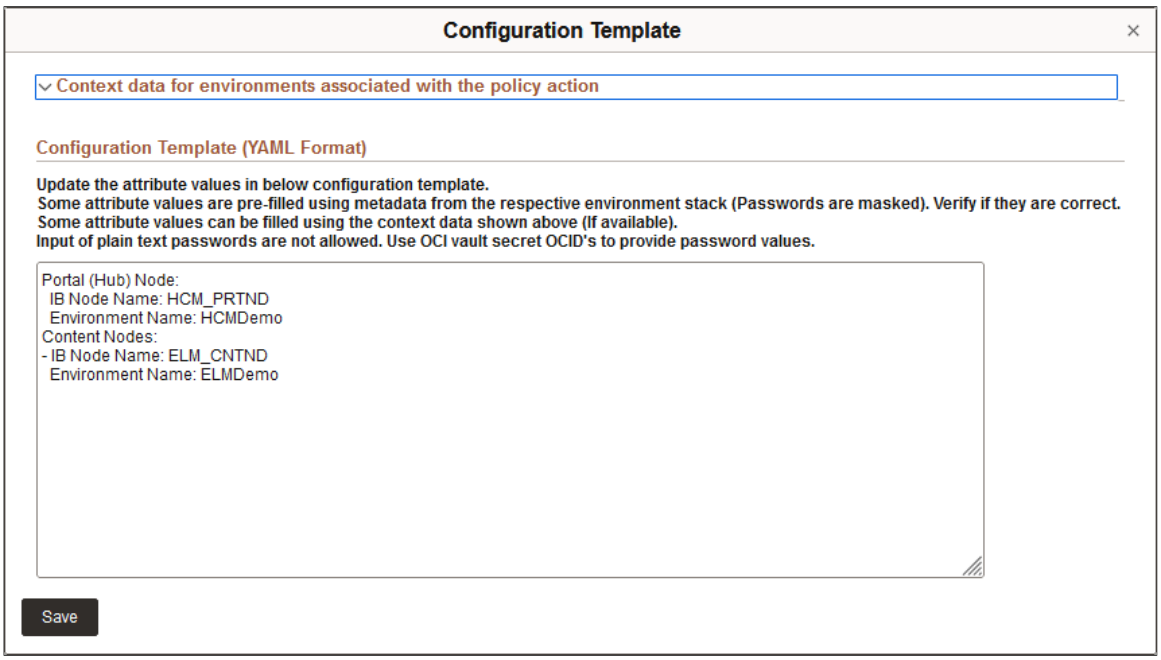

15. Click the **Parameters** button for the Register Remote Folder action.

    The parameters are Unified Navigation and Configuration Template. The Unified Navigation field displays the choice you made on the Cluster Settings page, and cannot be changed here.

    Click the edit icon to open the Configuration Template. The attribute values are populated with the IB node names and environment names for the portal system and content nodes.

    This action sets up the navigation folder for the content node within the portal system navigation. If the cluster was defined with the Unified Navigation option enabled, the Menu Label will be used to group the menu from the content node environment under the top-level of the portal system menu. You can modify the menu label or accept the default. The default menu labels correspond to the PeopleSoft applications—Campus Solutions, Customer Relationship Management, Enterprise Learning Management, Financials and Supply Chain Management, Human Capital Management, and Interaction Hub.

    If you did not select the Unified Navigation option when defining the cluster, the parameter is set to "N", and this policy action runs without registering the navigation folder for the content node.

    On the portal system home page, the home pages for the content nodes will be added to the homepage drop-down selector at the top left of the home page.

    Click **Save** after verifying the information on the Configuration Template, and then click **Apply** to close the Policy Action Parameters dialog box and return to the policy.

This example illustrates the Configuration Template for the Register Remote Folder policy action.

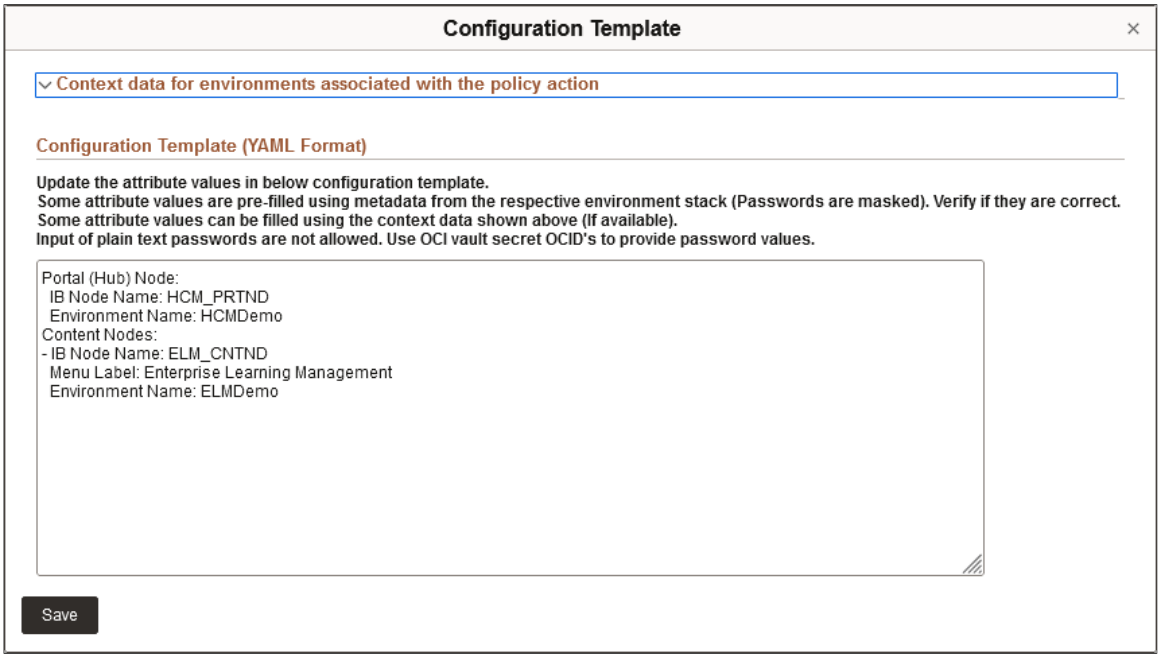

16. Click **Execute** to run the policy.

17. The Policy Monitor page opens. Click the status arrow for the Configure IB PORTAL Cluster policy to follow the progress.

    See Using Policy Monitor.

    You can also follow the progress on the Environments landing page. The environment cards for the portal system and content nodes display these status messages:

    - Configuring Shared Gateway
    - Register remote folder
    - Configure Portal
    - Adding Node to Cluster

    The Provision Task Status page also shows the progress.

18. After the policy runs, stop and start the content node environments.

    To stop an environment, select Stop from the Related Action menu on the environment card. After it is stopped, the Start action becomes available.

    Alternatively, you can use policies to stop and start the environments. You can create a single Adhoc policy with stop and start policy actions, associate the content node environments with the policy, and then run the policy. See Setting Policy Conditions and Actions for Environment Policy Object.

After setting up and running the policy, it is advisable to trigger the policy from the Cluster Settings page of the portal system node. Expand the IB Gateway Configuration section, and click **Configure** to open the page where you can run the policy.

Reviewing Cluster Details

View details after creating a new cluster or adding a node to an existing cluster.

Environment Cards

The card for the portal system environment displays a cluster symbol and the name of the cluster. The card for the content node environment displays the name of the cluster. If the portal system is stopped, the cards for the content node environments display warning symbols ( ).

This example illustrates the environment cards for a portal system environment and two content node environments.

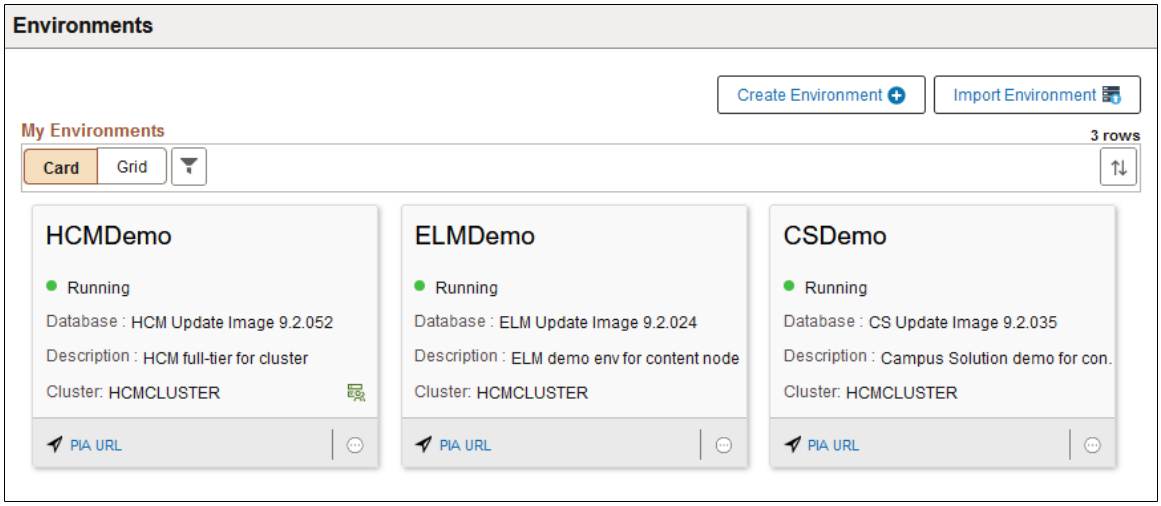

Shared Gateway

After you restart the content node environments, click **Refresh Metadata**. Expand the IB Gateway Configuration section. The node definition points to the shared gateway.

The shared gateway is also specified in each environment. Sign in to the environment and select **PeopleTools** > **Integration Broker** > **Configuration** > **Integration Gateways**.

Cluster Settings Page

When the creation process is complete, the portal system and content nodes, with their environment names and the gateway type, are displayed in separate sections on the Cluster Settings page for each environment in the cluster.

After a cluster is discovered or created, the data in the Cluster Settings page for the content node is read-only and the user cannot take any action.

This example illustrates the Portal System and Content Node(s) sections on the Cluster Settings page.

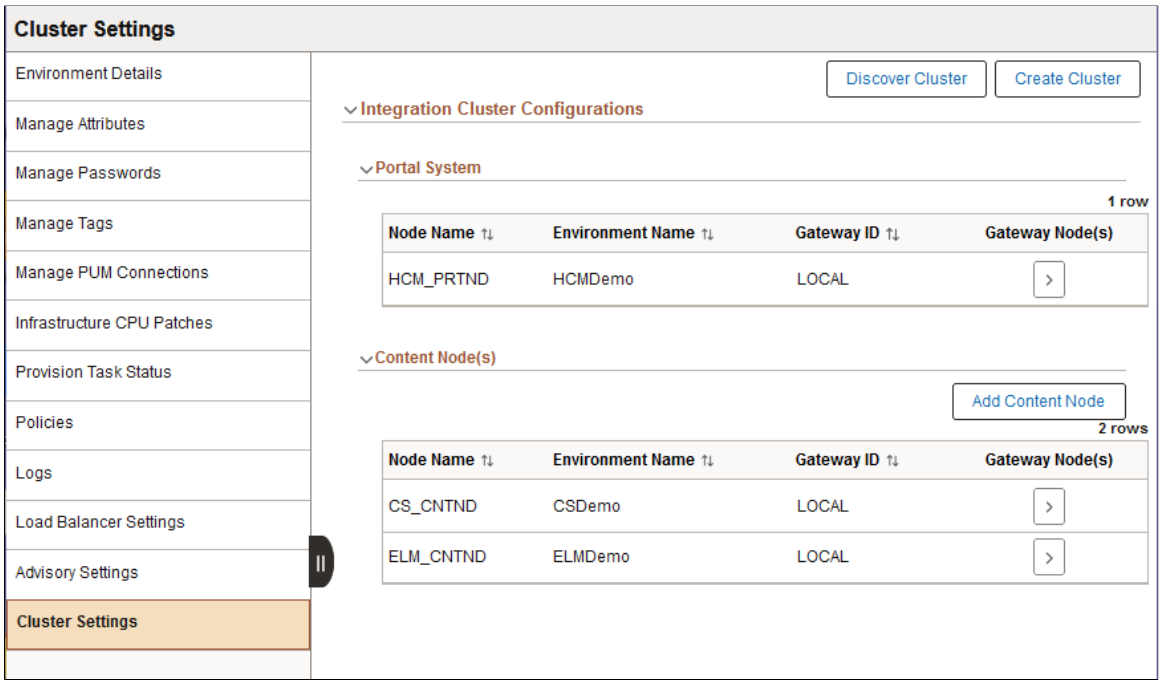

Gateway Node Details

In the row for the portal system or a content node, click the arrow under Gateway Node(s) to view the details of the Integration Broker Gateway node. The details include the Server URL (application server name and Jolt port number), User ID, and PeopleTools version. The Gateway Node(s) page for a content node is similar, but it includes only the specific content node and the portal system.

This example illustrates the Gateway Node(s) page for a portal system node.

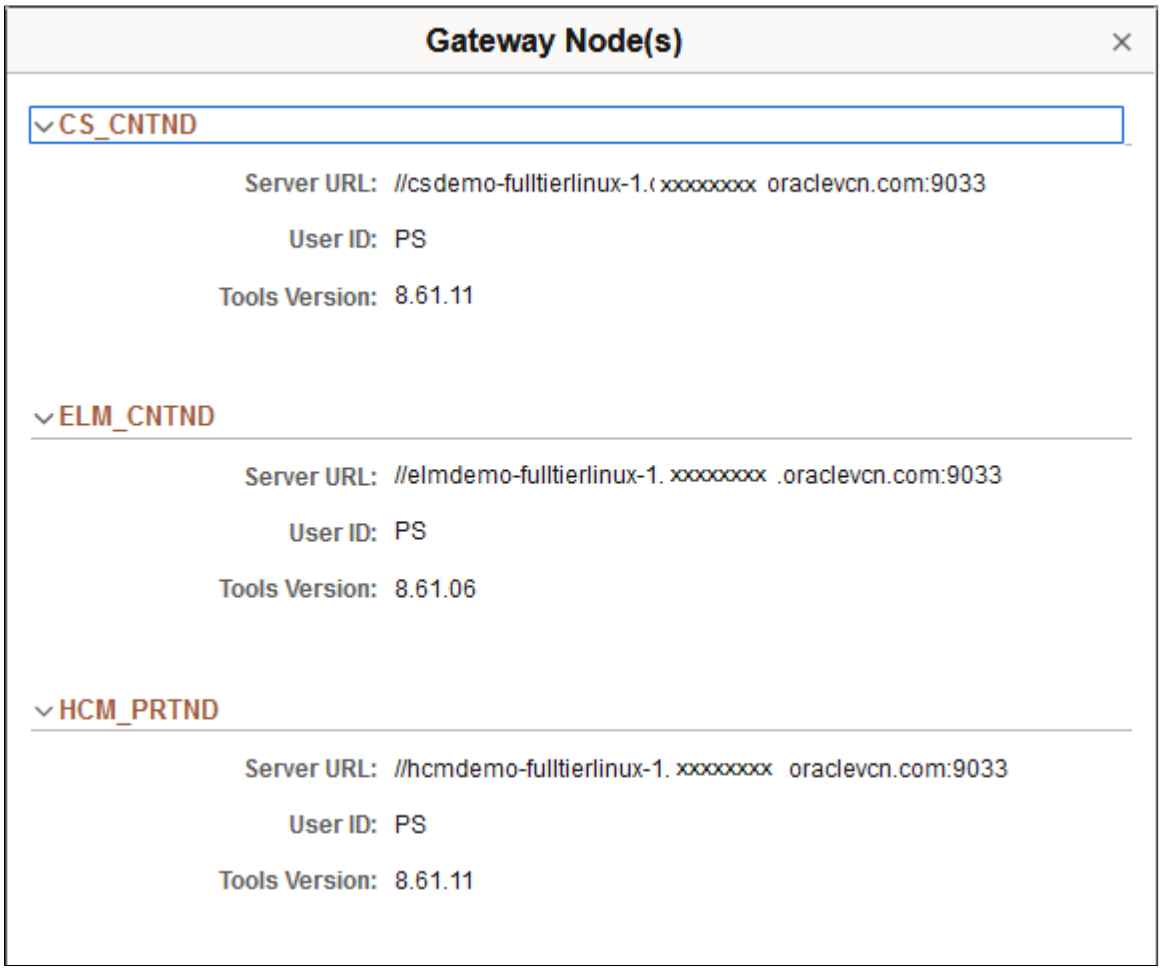

Discovering a Cluster from an Imported Environment

Import PeopleSoft environments and run a discovery process. If the imported environment is part of an existing unified navigation cluster, the cluster discovery process investigates the Integration Broker configuration, gateway URLs, and other integration properties for the imported system.

If the imported environment is not part of a cluster that was created outside of Cloud Manager, you can import it and create the cluster in Cloud Manager. The discovery process also identifies the content nodes in the cluster that are managed by Cloud Manager.

1. Import the portal system and content node environments that make up the cluster.

   See Importing Environment.

2. Click the environment card for the environment that you have designated as the portal system to access the details.

3. Select Cluster Settings on the left panel of the Environment Details page.

4. Click **Discover Cluster**.

5. The Portal System field is automatically populated with the environment name.

This example illustrates the Discover Cluster dialog box.

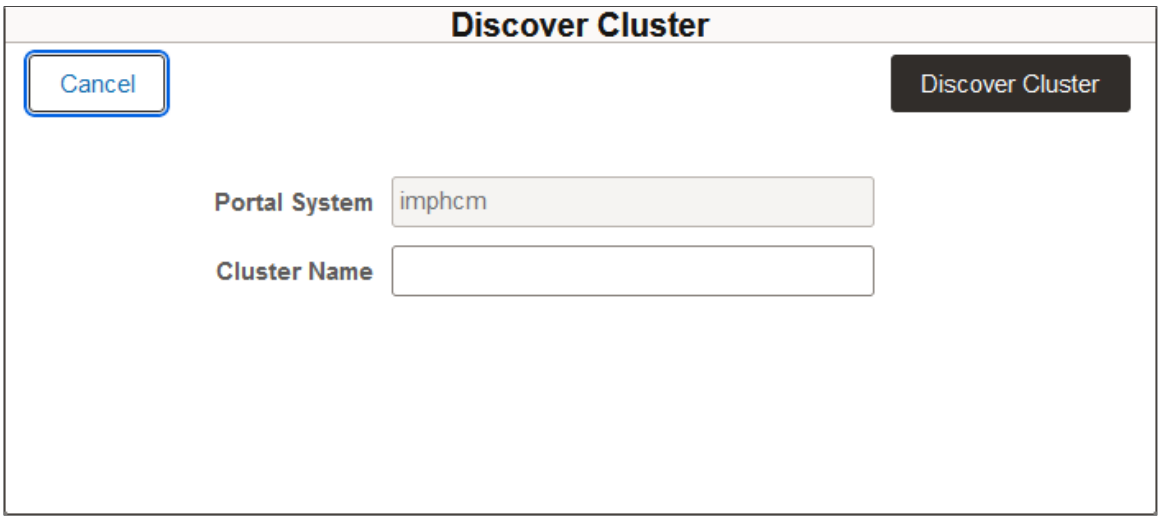

6. Enter a cluster name.
7. Click **Discover**.

When the discovery process is complete, the portal system node and content nodes are displayed in separate sections on the Cluster Settings page, as shown in the example in Reviewing Cluster Details.

Adding a Content Node to an Existing Cluster

You can add a content node to an existing cluster that was imported or created in Cloud Manager.

1. Click the environment card for the environment that you have set up as the portal system node to access the Environment Details page.
2. Select Cluster Settings on the left panel of the Environment Details page.
3. Click **Add Content Node**.

Note: The Add Content Node button appears for an environment that is already in a cluster.

The Portal System, Cluster Name, and Node Name fields are automatically populated.

This example illustrates the Add Content Node(s) dialog box.

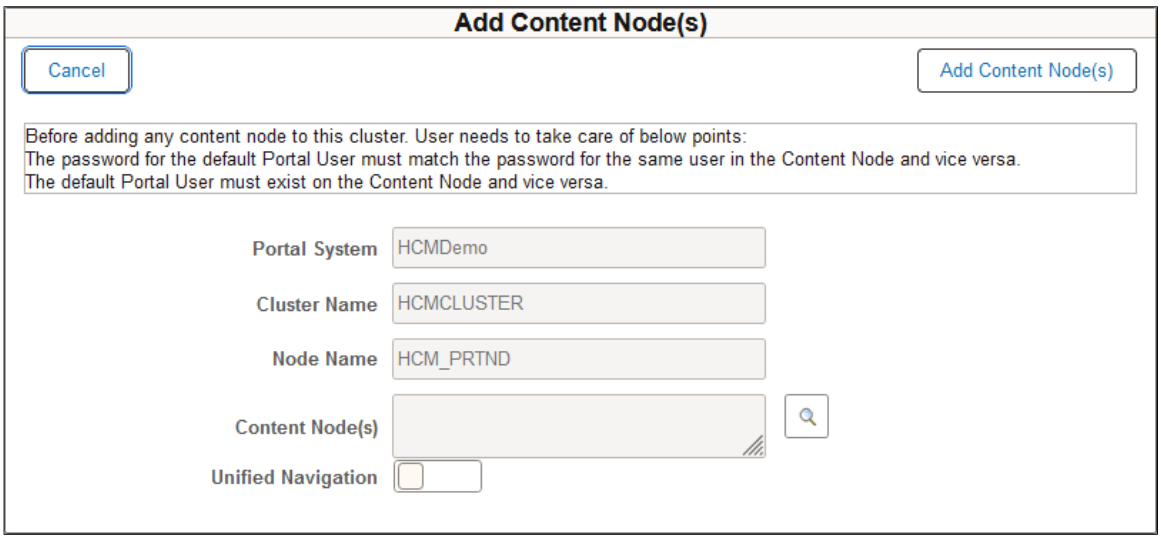

4. Click the Content Nodes(s) browse button and select one or more environments to serve as content nodes.

   You must select only environments that are not already part of another cluster. Do not select an environment if the cluster already contains the same PeopleSoft application.

5. To register the menu from the content node into the portal system menu, enable the **Unified Navigation** option.

   When you enable this option, a top-level menu item for the content node is created in the portal system menu.

6. Click **Add Content Node(s)**.

   Cloud Manager opens the Policy Editor page for the delivered policy Configure IB PORTAL Cluster<internal>. This is a predefined policy that cannot be deleted. The name and Policy Object, Environment, cannot be modified. Because it is an Adhoc policy you can run it immediately. See the policy example in Creating a New Unified Navigation Cluster.

Rediscovering an Imported Cluster

If you have already discovered a cluster but later realize that some participating nodes were not imported, you need to rediscover the cluster to include the missing nodes. To add those nodes, use the following steps.

1. Import the missing node.
2. Click the environment card for the environment that you have set up as the portal system node to access the Environment Details page.
3. Select Cluster Settings.
4. Click **Discover Cluster**.

On the Discover Cluster dialog box the Portal System and Cluster Name fields are populated and cannot be edited. You see a message that Cloud Manager will remove all the data associated with the imported nodes and recreate the configuration.

This example illustrates the Discover Cluster dialog box for an existing cluster.

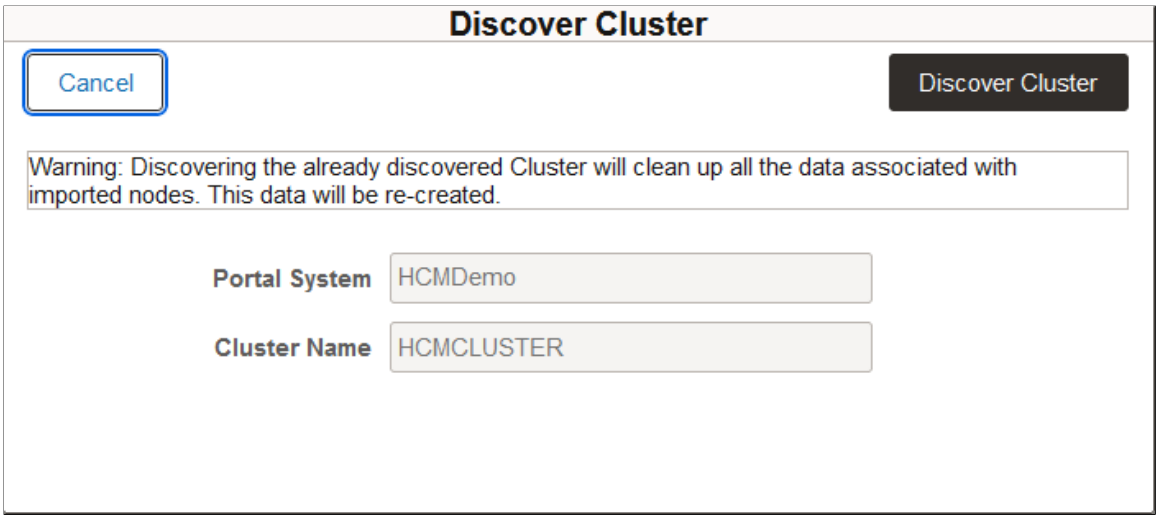

5. Click **Discover**.
6. Select the Gateway Node arrow for the Portal System to verify that the configuration was updated with the added node data.

Deleting a Portal System Environment

To delete an environment that is designated as the portal system in a cluster:

1. Click the Related Action button for the environment and select Delete.
2. A dialog box appears with this message:

   Warning: Current Environment is the Portal System in <Cluster_name> Cluster with Multiple Content environment associated with this IB cluster. Click on delete to ignore the warning and proceed with the Environment Deletion.

3. If you want to delete the environment from Cloud Manager but retain the environment resources in OCI, select Retain OCI Resources.
4. Click Delete to proceed or Cancel to return to the Environments landing page.

When you delete an environment that is designated as the portal system in a cluster, the Gateway URL configured in the content node environments, which is the portal system URL, will be stale. You should reconfigure the content nodes to point to a valid Gateway URL. Do this for all the content nodes that were part of the cluster.

To reconfigure the Gateway URL for content nodes after deleting the portal system node:

1. Click the environment card for the content node environment.

2. On the Environment Details page, expand the IB Gateway Configuration section and click **Configure**.

   This opens the Policy Editor page for the Configure IB Gateway<internal> policy. This is an adhoc policy with one policy action.

3. Click **Parameters** for the Configure IB Gateway policy action.

4. On the Policy Action Parameters dialog box, select configuration type Configure Local Gateway.

5. Click the edit icon for the Configuration Template, and supply the required values.

   The template includes these parameters:

   - Environment Name: <selected environment name>
   - Gateway (Host:Port): xxx
   - Default Local Node Name: xxx
   - Gateway Node App Server URL (//Host:Jolt Port): xxx
   - Gateway Node User Id: xx
   - Gateway Node Password: '***'
   - Gateway Keystore Password: '***'

6. Click **Execute** to run the policy.

Deleting a Content Node Environment

To delete an environment that is designated as a content node in a cluster:

1. Click the Related Action button for the environment and select Delete.

2. A dialog box appears with this message:

   Warning: Current Environment is the Content Node in <Cluster_name> Cluster. Click on delete to ignore the warning and proceed with the Environment Deletion.

3. If you want to delete the environment from Cloud Manager but retain the environment resources in OCI, select Retain OCI Resources.

4. Click Delete to proceed or Cancel to return to the Environments landing page.

Re-creating a Cluster

After performing a PeopleTools upgrade or applying a PeopleTools patch you must re-create the cluster.

If you performed a PeopleTools upgrade or applied a PeopleTools patch on either the content node or portal system node, you must re-create the cluster. You can use the same IB node names as in the original cluster. Follow the same procedure as when creating a new cluster, in Creating a New Unified Cluster.

In addition, after you create a cluster, you have the option to rename the nodes that are part of the cluster, and then re-create the cluster, as follows:

1. On the Environment Details page for the portal system environment, expand the IB Gateway Configuration section and click **Configure**.

   This opens the Policy Editor page for the Configure IB Gateway <internal> policy.

2. Click **Parameters** for the Configure IB Gateway policy action.

3. On the Policy Action Parameters dialog box, select Rename Node as the configuration type.

4. Click the edit icon for the Configuration Template, and supply the required values.

   Rename all the non-local IB nodes in the portal system, which corresponds to the Local Default node of each content node. The Configuration Template includes these parameters:

   - Environment Name: — Enter the environment name of the portal system.
   - Current Node Name: — Enter the current non-local IB node name for the portal system node.
   - New Node Name: — Enter the new non-local IB node name for the portal system node.
   - Is Default Local Node (Y / N): — Enter N.

5. On the Environment Details page for each content node environment, expand the IB Gateway Configuration section and click **Configure**.

   This opens the Policy Editor page for the Configure IB Gateway <internal> policy.

6. Click the edit icon for the Configuration Template, and supply the required values.

   Rename the non-local IB node in each content system, which corresponds to the Local Default node of the portal system. The Configuration Template includes these parameters:

   - Environment Name: — Enter the environment name of the content node.
   - Current Node Name: — Enter the current non-local node name for the content node.
   - New Node Name: — Enter the new non-local IB node name for the content node.
   - Is Default Local Node (Y / N): — Enter N.

7. On the Environment Details page for the portal system environment, select the Cluster Settings page and create a new cluster with the new node names.

Related Links

Applying PeopleTools Patch
Upgrading PeopleTools
Configuring IB Gateway

Chapter 7

Using Orchestration Manager

Understanding Orchestration Manager

Define and manage policies using Orchestration Manager features. Administrators can automate frequently-used activities such as starting and stopping environments, scaling up or down, and applying maintenance. Policies can be scheduled or implemented in real time by connecting through events.

Click the Orchestration Manager tile on the Cloud Manager home page to access the pages to manage policies.

This example illustrates the Orchestration Manager tile.

Orchestration Manager features are described in these topics:

- Using Policy Editor
- Setting Up Auto Scaling
- Adding a Policy with Multiple Actions
- Adding a Policy with Custom Actions
- Using Policy Monitor
- Creating Policy Groups

Policies are based on conditions and actions. The Policy Object that you choose, either Repository Artifact or Environment, controls the conditions, actions, and action parameters you can select.

For example, to set up a backup for a specific time, select:

- Policy Object — Environment
- Policy Type — Schedule

- Condition — Environment + Running
- Action — Backup
- Action Parameter — Backup prefix

To set up a policy that runs when an Infrastructure CPU is downloaded to the repository, select:

- Policy Object — Repository Artifact
- Policy Type — Event
- Event Name — Infra DPK Download
- Condition — Tools + 8.61 + Linux
- Action — Apply CPU Patches
- Action Parameter — Environment Name

After defining policies, the policies can be associated with environments using one of the following:

- Specify the environments for the policy when creating the policy using the Policy Editor.

  See Setting Policy Conditions and Actions for Environment Policy Object
- Specify policies for an environment from the Environment Policies page.

  See Associating Policies with Environment
- Specify OCI tags as criteria for policy.

  See Managing Tags
- Associate environments or templates with Policy Groups.

  See Creating Policy Groups
- Specify the policies in the Environment Template used to create the environment.

  See Environment Template – Security and Policies Page

Using Policy Editor

To access the Policy Editor page (ECL_POLICY_EDITOR), select the Orchestration Manager tile.

Use the Policy Editor page to add policies, as well as manage existing policies.

See Managing Policies.

Policies will be displayed by Policy Group when available. Policy groups are optional and can be used to group policies together for display, and to facilitate associating related policies. Any policies that are not associated with a specific policy group will be shown under Default Policies.

See Creating Policy Groups.

Adding a Policy

When you add a policy, the owner is the user who defined the policy. The owner can be an administrator user or a self-service user.

Policies defined by an administrator can be applied on any policy object artifact. For example, if an administrator adds a policy for stopping environments, then all the environments available in the Cloud Manager instance (irrespective of which user created the environment) can be associated with the policies.

Policies created by a self-service user can be applied to the environments created by that user.

To add a policy:

1. Select the Orchestration Manager tile.
2. Click the **Add Policy** button.
3. Define and save the policy on the Policy Editor page.

This example illustrates the fields and controls on the Add Policy page.

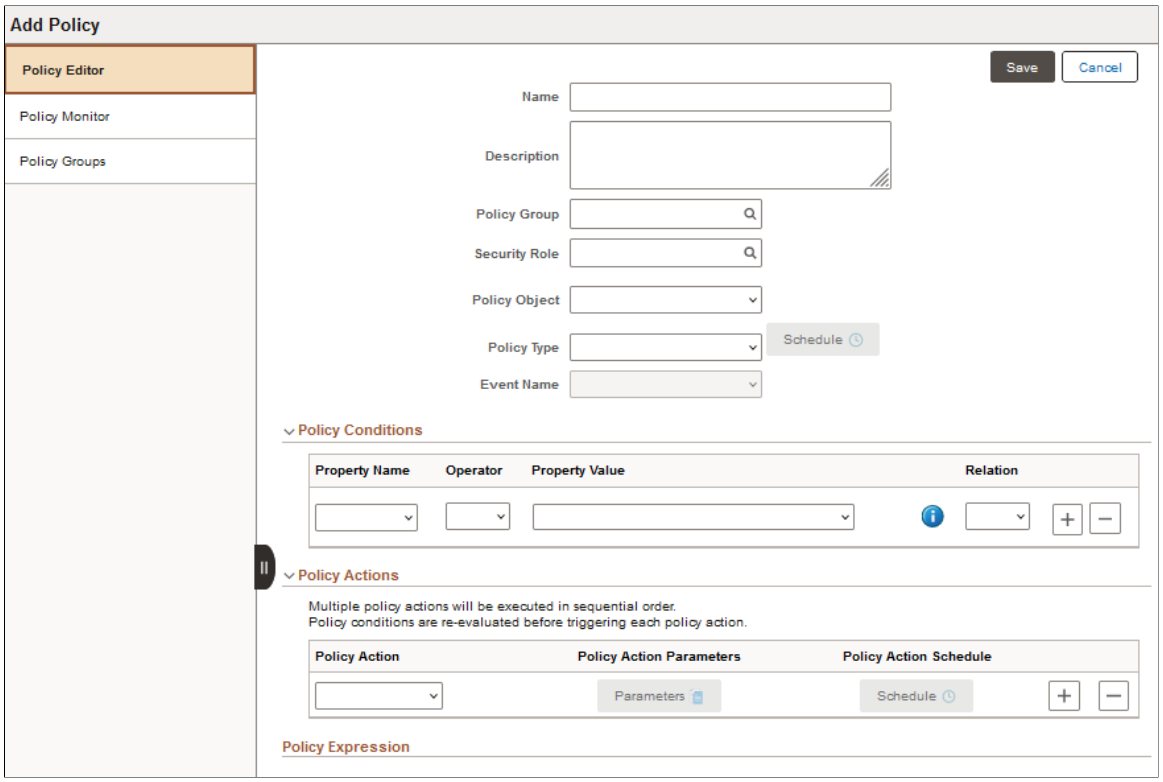

Field or Control	***Description***
Name	Enter a name for the policy.
Description	Enter a description for the policy.

Field or Control	*Description*
Policy Group	Select an existing policy group. Policy groups are used to group policies together. The policy does not need to belong to a group. See Creating Policy Groups
Security Role	Select the security role that can edit the policy.
Policy Type	Select the policy type. • Schedule Policy can be scheduled for a start date and time or recurrence. • Event Based on an event, for example PRPs are downloaded. • Adhoc The user can run the policy at any time.
Schedule	This button is available when the policy type is Schedule. Select to set the policy schedule. See Setting Policy Schedule
Policy Object	A Policy object exposes properties and actions, which are used by the Policy editor to set policy conditions and action for the policy. Built-in policy objects are: • Environment Used for creating policies for life cycle activities on environments, for example start and stop. See Setting Policy Conditions and Actions for Environment Policy Object. • Repository Artifact Used for creating policies based on artifacts in the repository, such as downloaded HCM images. See Setting Policy Conditions and Actions for Repository Artifact Policy Object.

Field or Control	*Description*
Event Name	If Policy Type is Event and Policy Object is Environment, select one of these Event Names. • AddNode • BackUp • Clone • Delete • InfrastructureCPU • PTP (PeopleTools Patch) • PTU (PeopleTools Upgrade) • Provisioning • RemoveNode • Restore • ScaleDown • ScaleUp • Start • Stop If Policy Type is Event and Policy Object is Repository Artifact, select one of these Event Names: • Infra DPK Download • Patch Download • PUM DPK Download
Policy Expression	When you save a policy, this field displays a statement representing your choices. This field also displays error messages in case of incorrect choices.

Setting Policy Schedule

When you select the Schedule button from the Add Policy page, you can select to create an ad hoc schedule or a recurrence.

Important! Set the base time zone in PIA (**PeopleTools** > **Utilities** > **Administration** > **PeopleTools Options**) to match the time zone of the Cloud Manager Database.

Adhoc

Select Adhoc if you want to schedule the policy for a specific data and time.

This example illustrates the fields and controls on the Schedule page for an Adhoc policy schedule.

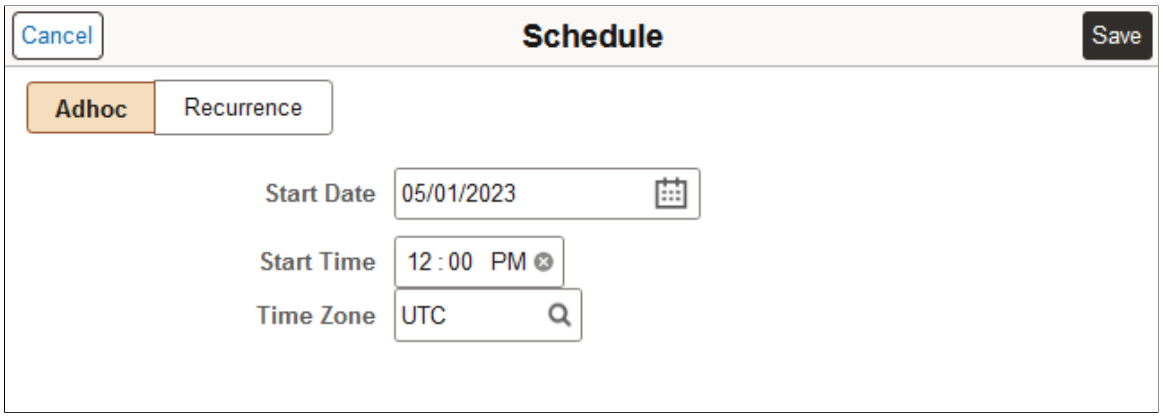

Field or Control	*Description*
Start Date	Select the start date.
Start Time	Enter the start time.
Time Zone	Select the time zone.

Recurrence

Select Recurrence if you want the policy to be run on a recurring schedule.

This example illustrates the fields and controls on the Schedule page for a Recurrence policy schedule.

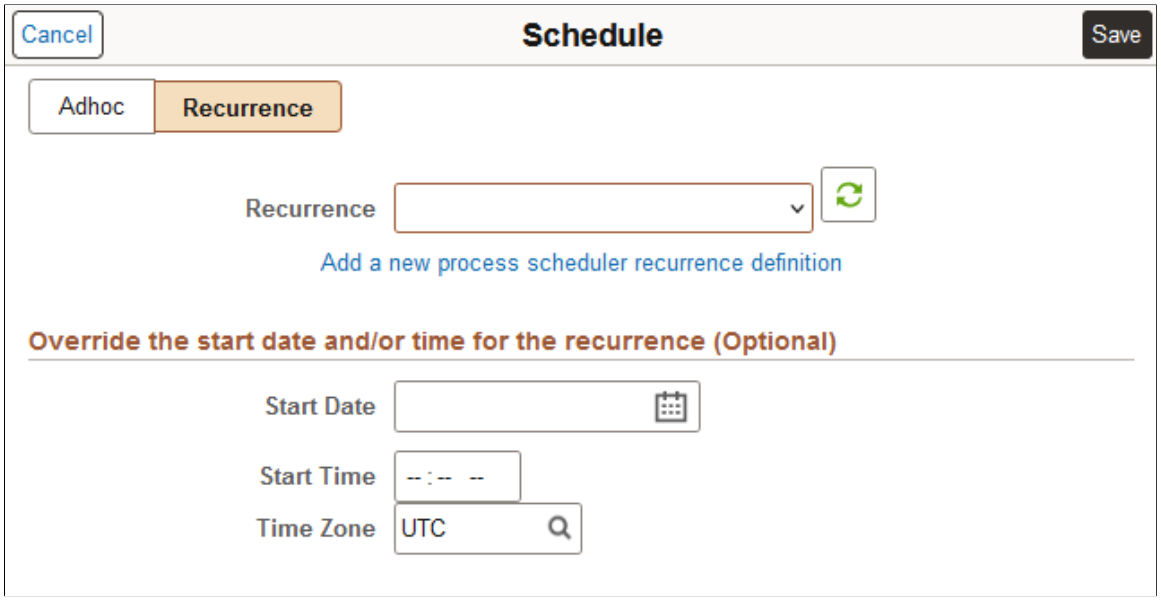

Field or Control	*Description*
Recurrence	Select an existing Process Scheduler recurrence schedule.
Add a new process scheduler recurrence definition	Select this link if you want to add a new recurrence definition. This link will open the Process Scheduler Recurrence Definition page.
Start Date	Optionally, you can enter a start date for this recurrence.
Start Time	Optionally, you can enter a start time for this recurrence.
Time Zone	Optionally, you can enter a time zone for this recurrence.

See the product documentation *PeopleTools: Process Scheduler*, Defining Recurrence Definitions.

Setting Policy Conditions and Actions for Environment Policy Object

Environment policies can be based on a schedule or an event. Policies where the Policy Object is Environment, the Policy Type is Event, and the Event names are ScaleUp and ScaleDown, are used for auto scaling. See Setting Up Auto Scaling.

This example illustrates the fields and controls on the Add Policy page where the Object Type is Environment.

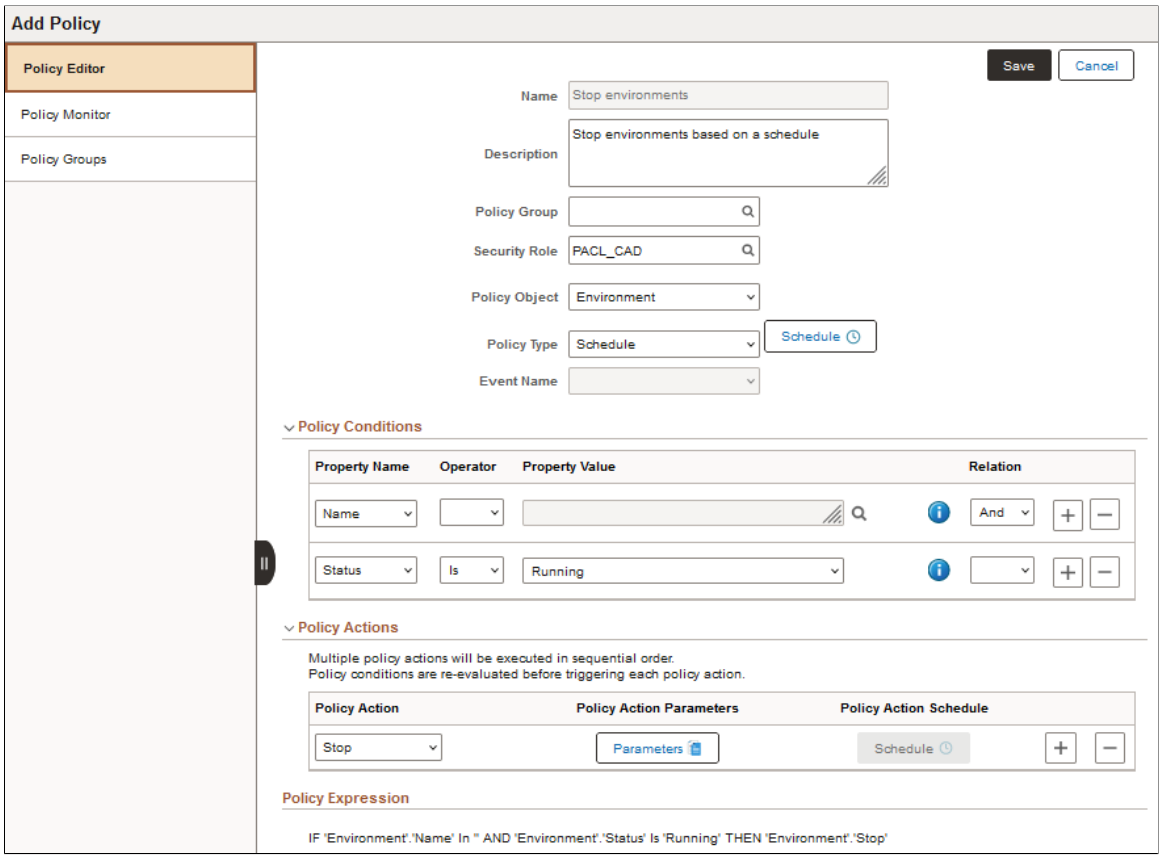

Field or Control	***Description***
Property Name	Available Property Names are: • Name • Status • Tag **Note:** Status is not applicable for Delete events.
Operator	Depends on Property Name, operators include: • In • Is • Matches
Property Value	Select the Property Value(s) from the drop-down list or lookup depending on the Property Name. Enter a single value or a comma separated list of values.

Field or Control	*Description*
	Mouse over this icon for an explanation of the property and example values.
Relation	Specify the relationship between one condition and the next with operators And, Or. You can also use parentheses to group conditions.
	Select to add another condition.
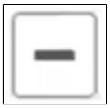	Select to delete a condition.
Policy Action	Select the Policy Actions for Environment Policy Objects. The available Policy Actions and Action Parameters are listed in the next section.
Parameters	Use to select the action parameters for the selected Action. Actions and Action Parameters are listed in the next section.
Policy Action Schedule	Select this button to schedule policy actions. For scheduled environment policies, this button is not available for use. Use the Schedule button next to Policy Type.

Policy Actions for Environment Policy Object

Depending on the policy action selected, the action parameters are displayed. Mandatory actions parameters are prefixed with an asterisk (*) and must be entered. Mouse over the information icon for the action parameter for information on how to enter the parameter value.

Policy Action	*Action Parameters*
Backup	Backup Prefix: Enter a text string to be appended to the database backup.

Policy Action	*Action Parameters*
Custom Action	Define Source (Define any one): Select one of the following as the source for the custom action. • PeopleCode Handler (Application Class). • Repository File • Inline Command(s) Set Execution Content: Enter the content needed to run the custom action. • Operating System • Node Type • Input (JSON) See Adding a Policy with Custom Actions.
Health Check	Environment Names Check All Domains are All (Required) Check if PIA is Accessible (Required) Check Load Balancer Status (Required) Health Check action must be followed by Send Email Notification action, so that the user receives the Health Check Status by email. **Note:** This action is not applicable for Delete, ScaleUp, or ScaleDown events.

Policy Action	***Action Parameters***
Refresh	Environment Name (Required) Use Latest Backup Backup OCID Perform Middle Tier Refresh Source TDE KeyStore (Wallet) Password (Required) Pre Refresh Custom Script Post Refresh Custom Script Target MT DMS Path File Server App DPK Location See Refreshing DB Systems Environment
RefreshADBWallet	Expiry Day Count
ScaleDown	Scale Type (Required) –- Auto Scale Nodes to Scale Scale Limit (Required) Day 01 (Monday) through Day 07 (Sunday) Application Server Domain (Required) Process Scheduler Domain PIA Domain
ScaleUp	Scale Type (Required) –- Auto Scale Nodes to Scale Scale Limit (Required) Day 01 (Monday) through Day 07 (Sunday) Application Server Domain (Required) Process Scheduler Domain PIA Domain

Policy Action	*Action Parameters*
Send Email Notification	Notification Topic OCID To use this action, enable External Notifications on the Infrastructure Settings page. If the Notification Topic OCID is not supplied here, the value will be taken from the Infrastructure Settings page. If there is no value in either place, no email will be sent. See Notifications in Infrastructure Settings Page. If the previous action is Health Check, the user will receive an email with the Environment Health Check Status report. if the previous action is any other action, then the user will receive an email informing them that the action is completed. If this is the first action, then the user will receive an email informing them that the Life Cycle process is complete.
Start	None
Stop	None

Adhoc Policy for Configuring IB Gateways

Cloud Manager enables you to define and run on-demand policies using the Adhoc policy type. You can run Adhoc policies, used for purposes such as restarting a group of environments, at any time without waiting to schedule an event. The pre-defined internal policies can be used to perform IB configurations on managed environments.

This example illustrates the fields and controls in the Configuration Type field on IB Gateway Configuration.

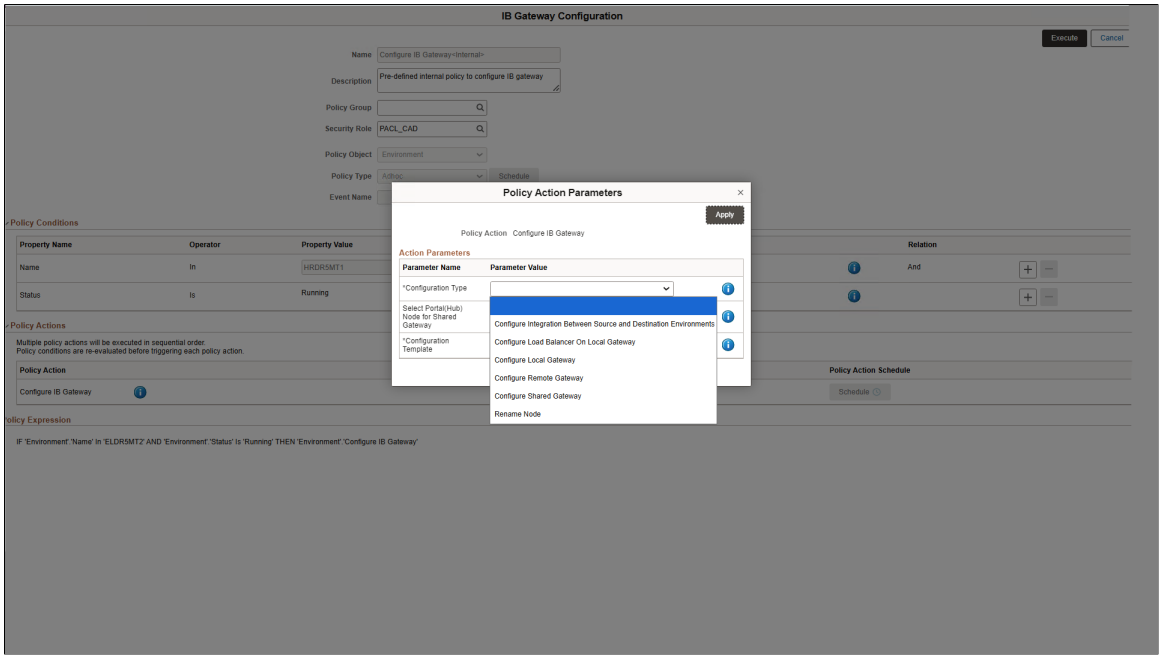

Policy Action Parameter	***Definition***
Configuration Type	Select a configuration type according to the requirement. The available options are: • Configure Integration between Source and Destination environments • Configure Load Balancer on Local Gateway • Configure Local Gateway • Configure Remote Gateway • Configure Shared Gateway • Rename Node
Configuration Template	This value is partly pre-populated from the environment stack and depends on the Configuration Type selected. It is displayed in the form of JSON data. You can modify the content, if needed.
Select Portal (Hub) for Shared Gateway	Select the portal or hub node from the drop-down list with the environment names associated with the policy. This parameter is applicable only for Shared Gateway configuration type.

Starting and Stopping Individual Nodes

To apply the Start/Stop Policy Action on individual nodes from an environment:

1. Select **Name** as the first Property Name on the Policy Conditions section.
2. Select the Operator.
3. Select the environments where you want to start or stop individual nodes from Property Value drop-down.
4. Select **Start** or **Stop** from the Policy Action drop-down. Using the + button you can create multiple policy actions, which are run in sequence.

   Note: To start a node, the environment that contains the node must be in Running state.

5. Click **Parameters** button on the Policy Action Parameters field. Policy Action Parameters dialog box appears. The environments you selected are listed.
6. Select the desired nodes from the listed environments. The default value is **All** on the **Selected Nodes** field. If you keep the default value, it becomes an environment start/stop policy action.

   Click the lookup icon to select other nodes.

   This example illustrates the fields and controls on Policy Action Parameters dialog box when Stop policy action is selected.

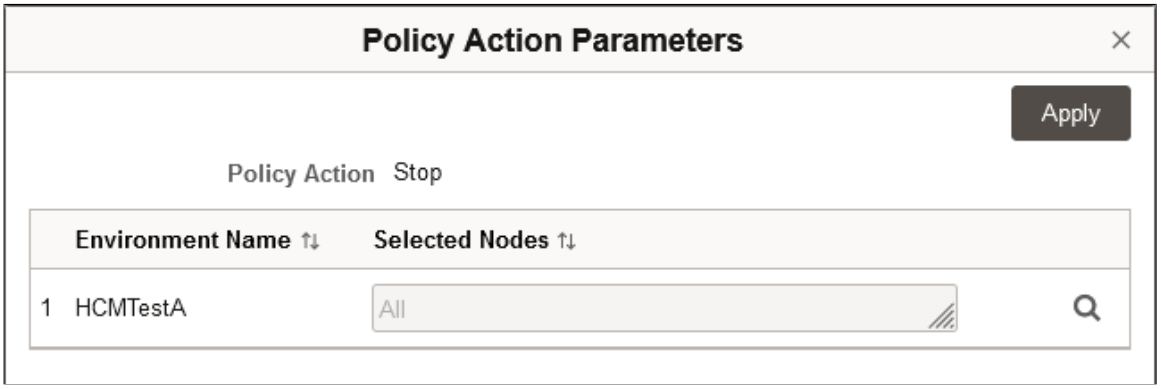

   If the Policy Type is selected as **Event**, select the Policy Action Schedule to schedule the policy action.

7. Click Save. The policy for starting or stopping nodes is successfully added.

Setting Policy Conditions and Actions for Repository Artifact Policy Object

When you select Repository Artifact, the only Policy Type is Event. An event-based policy can be triggered by the event or the policy action can be scheduled for a specific time or recurrence.

Repository Artifact events are:

- Patch Download
- Infra DPK Download

- PUM DPK Download

  See Adding a Policy to Provision PUM Environments.

By default when you create a policy for a repository artifact, the policy is triggered when the artifact is downloaded. Use the scheduling option to delay running the event.

Important! You must not apply Infra DPK Download event-based policies directly. It is recommended to apply such policies by setting a schedule for the policy action.

Infrastructure DPK Download example:

- You do not want to disrupt the environments during normal business hours when the download takes place.
- Set a schedule to run the apply on the weekend.
- If Infrastructure DPK gets downloaded in the middle of a week, CPU patches will be applied on the environments only during the weekend.

This example illustrates the fields and controls on the Add Policy page where the Object Type is Repository Artifact.

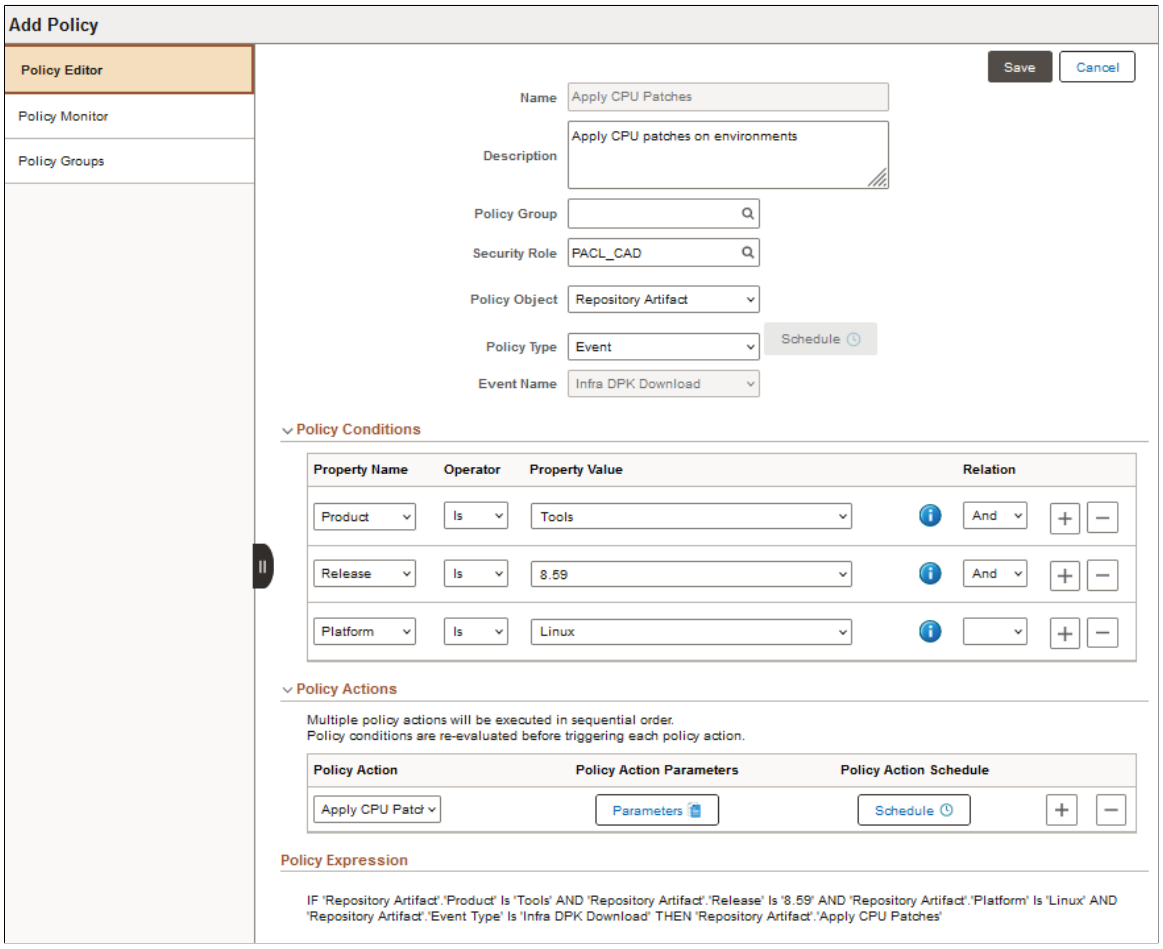

Policy Conditions for Repository Artifact

Field or Control	*Description*
Property Name	For Policy Object Repository Artifact, property names are: • Platform • Product • Release
Operator	Is
Property Value	Select value from the drop-down list.
Relation	Specify the relationship between one condition and the next with operators *And*, *Or*. You can also use parentheses to group conditions.
	Select to add another condition.
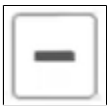	Select to delete a condition.

Policy Actions for Repository Artifact

An event-based policy will run when the event triggers the policy. Select the Schedule button to schedule the policy action for a specific time or recurrence.

Policy Action	*Parameters*
Apply CPU Patches	Environment Names
Apply PRPs	Environment Names
Backup	Environment Names Backup Prefix

Policy Action	***Parameters***
Custom Action	Environment Names (Required) PeopleCode Handler (Required) Input (JSON)
Define and Upload Target	Select PUM Target (Required) Custom IB Local Gateway Node of Target
Delete Old PUM Source	None
Health Check	Environment Names Check All Domains are up (Required) Check if PIA is Accessible (Required) Check Load Balancer Status (Required) Health Check action must be followed by Send Email Notification action, so that the user receives the Health Check Status by email.
Migrate PUM Metadata	Select Old PUM Source (Required) Custom IB Local Gateway Node of Old PUM Source Auto Discover Old PUM SOurce on next policy run (Required) Migrate PUM Metadata from Old PUM Source (Required) Upload Targets from Old PUM Source (Required)
Provision PUM	Environment Name Prefix (Required) Environment Template Name (Required) Environment Database Name Prefix (Max length 3)

Policy Action	*Parameters*
Send Email Notification	Environment Names Notification Topic OCID (Optional) To use this action, enable External Notifications on the Infrastructure Settings page. If the Notification Topic OCID is not supplied here, the value will be taken from the Infrastructure Settings page. If there is no value in either place, no email will be sent. See Notifications in Infrastructure Settings Page. This action must follow another policy action, or no email will be sent. If the previous action is Health Check, the user will receive an email with the Environment Health Check Status report. If the previous action is any other action, then the user will receive an email informing them that the action is completed.
Start	Environment Names
Stop	Environment Names

Setting Parameters When Selecting Environments for Policy Actions

The required parameters vary depending upon the Policy Actions. For actions such as Start, Stop, Apply PRP, or Apply CPU, you must select an environment for the policy to work on. To set the parameters:

1. Click the Parameters button for the Policy Action.
2. Click the Lookup icon.
3. You can select an environment based on Environment Name or Tag.

   For an environment, click the Lookup icon and select the environment.

   For a Tag, select a Tag Namespace from the drop-down list, and then select a tag.

This is an example of the Set Parameter Values page when Tag is selected.

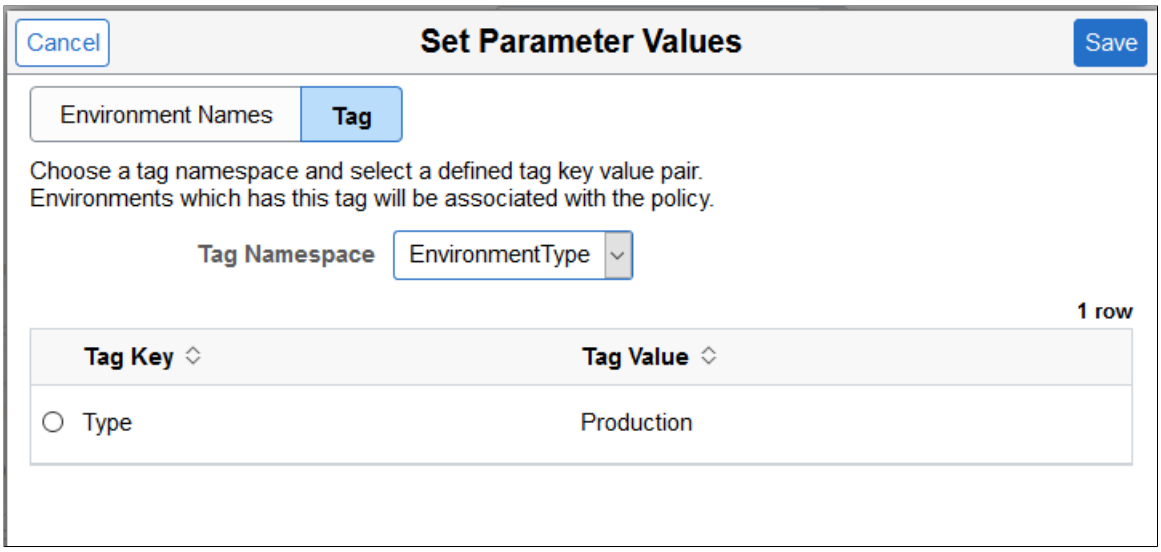

4. Click Save.

   The policy action parameter value is updated.

   This is an example of the Policy Action Parameters page

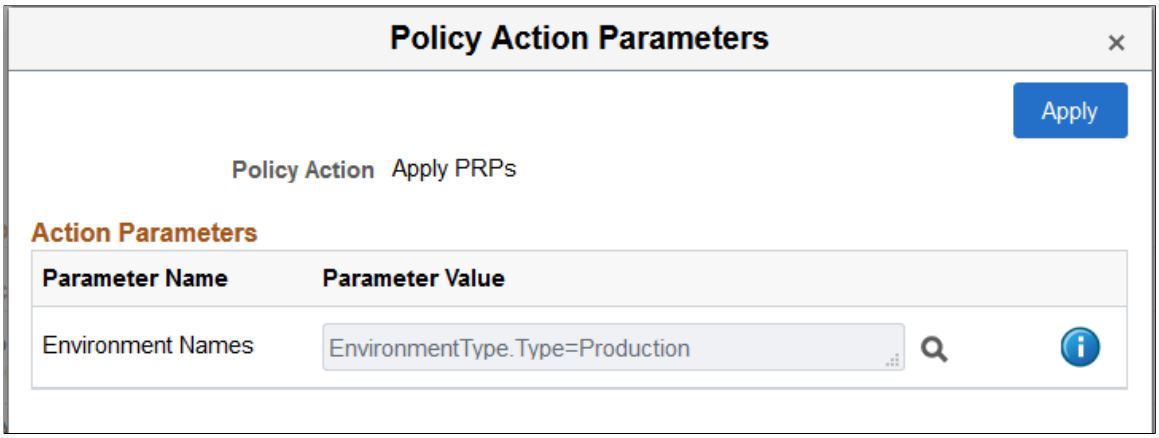

5. Click Apply to save the parameters.

Scheduling Event Based Policy

Select the Schedule button to schedule the policy for a specific time or recurrence.

Setting a schedule for an event based policy is completely optional. By default an event based policy will run in real time as and when the event occurs. Add a schedule for an event based policy only when needed.

This example illustrates the fields and controls on the Schedule page.

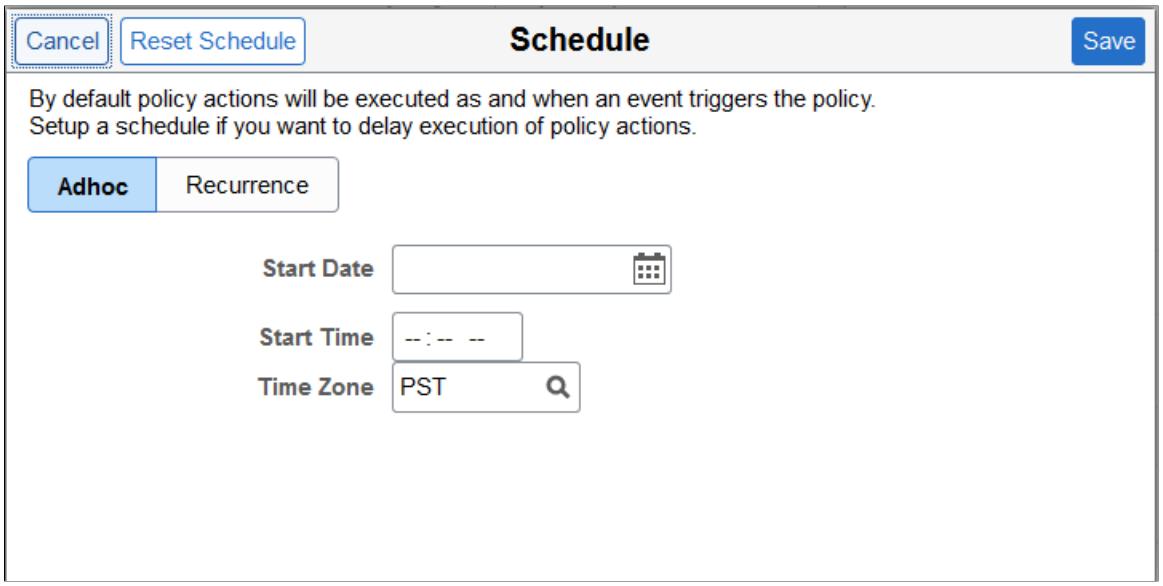

Select Adhoc or Recurrence to schedule the policy. Use the Reset Schedule to reset the schedule.

Managing Policies

All defined policies are available and grouped by policy group.

You can edit the delivered default policies for your requirements and save them to a new name, or use them as examples to create new policies.

The Default Policies section shows policies that are not associated with a policy group.

This example illustrates the fields and controls on the Policy Editor page. You can find definitions for the fields and controls later on this page.

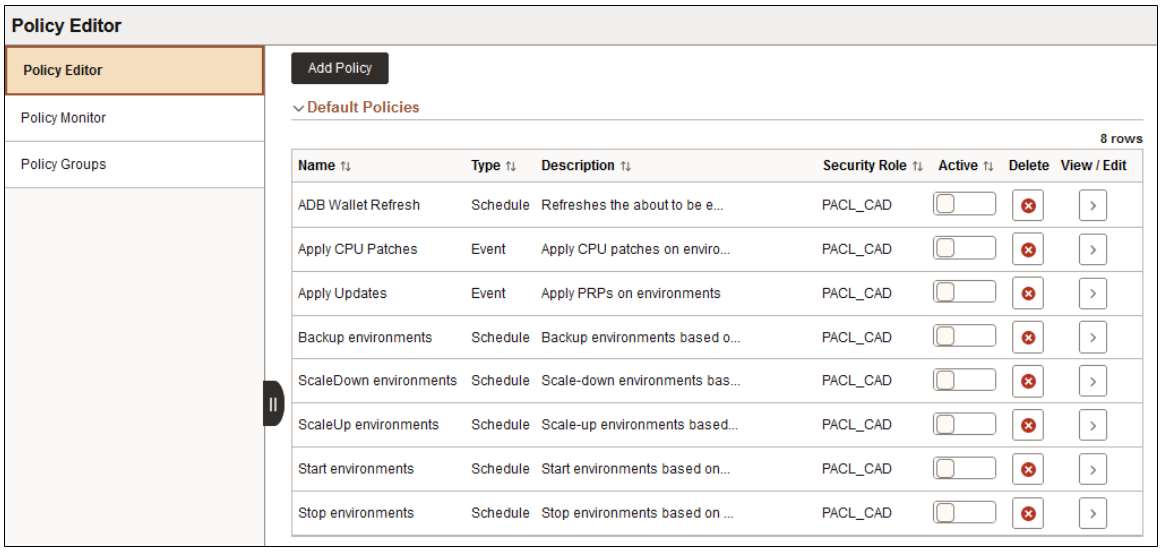

Field or Control	*Description*
Active	Enable the switch to activate or inactivate a policy. When a policy is active the switch displays a check mark. Once a policy is inactivated, it will be suspended and will not be run.
Delete	Use the Delete icon to remove a policy.
View/Edit	Click the icon (>) to view or edit a policy. The user must belong to the security role assigned to the policy in order to edit the policy. If the user does not belong to the security role, the user can view the policy, but can not edit it.

Editing the Policy

When you select to edit or view the policy, the policy is displayed.

This example illustrates the fields and controls on the Add Policy page when you select to edit or view a policy.

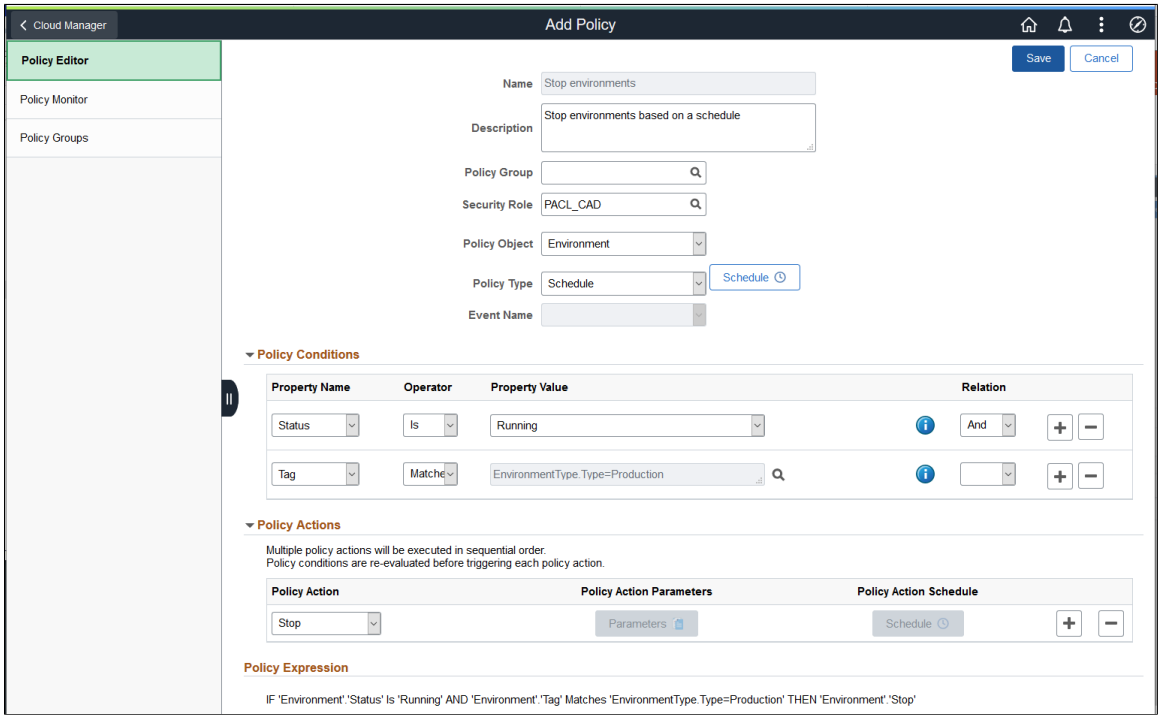

Note: Environments can be added to or removed from a policy directly from the Environment Details page. See Associating Policies with Environment.

Setting Up Auto Scaling

Cloud Manager integrates with Oracle Data Science services to take advantage of machine learning techniques. This enables Cloud Manager to learn from past usage and detect the current state and take corrective action if any, so that your users experience smooth, reliable performance.

Auto Scaling Prerequisites

Prerequisites include:

- Subscription to Oracle Data Science.
- Set up Data Science in OCI.

  See the tutorial *Create Data Science Resources for Auto Scaling in Cloud Manager (Optional)* at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.
- Configure Data Science in Cloud Manager Settings.

  See Data Science Settings Page
- Enable Monitoring Flag in Cloud Manager Settings Page.

  See Cloud Manager Settings Page

Setting Up the Target Managed Instance

The environment that is managed by the auto scaling policy must meet these requirements:

- Target managed instance must include OpenSearch (or Elasticsearch) instance with OpenSearch Dashboards (or Kibana).

  Note: The search server machine (OpenSearch or Elasticsearch) should have network access to connect to OCI telemetry (monitoring) API. The custom image used for OpenSearch or Elasticsearch must include OCI Python SDK. The Marketplace Linux image includes the OCI Python SDK. See Infrastructure Settings Page.
- Target managed instance must be on PeopleTools 8.58 or above.
- Enable Monitoring on the target managed instances.

  There are several ways to enable monitoring for an environment:

  - Enable monitoring on the template used to create the environment.

    See Environment Template – Select Topology Page
  - Enable monitoring after the environment is created.

    See Monitoring Environments

  This starts collecting performance data and indexes that information in an OpenSearch (or Elasticsearch) instance.

- Train the system with a minimum load of typical business transactions for at least 3 or 4 weeks in order to collect enough samples for prediction.

  Note: When the number of database connections exceeds the number of available processes, the connection may fail intermittently, resulting in ORA-12520 error. This error can be resolved by increasing the value of process parameter.
  See Intermittent TNS-12520 or TNS-12519 or TNS-12516 Connecting via Oracle Net Listener (Doc ID 240710.1) and JDBC Connections Fail with ORA-12520 (Doc ID 2660207.1) for more information.

- ECL_DATA_UPD runs weekly to collect performance data for the instance and pushes the performance data to the Object Storage bucket.
- ECL_ML_JOBS runs weekly to pull this performance data, train the model and publish the Machine Learning Model in Data Science.

Adding Auto Scaling Policy

Once the Data Science model is trained, define the auto scaling policy.

For an auto scale policy, use the following:

This example shows the Add Policy page with a sample Auto Scale policy.

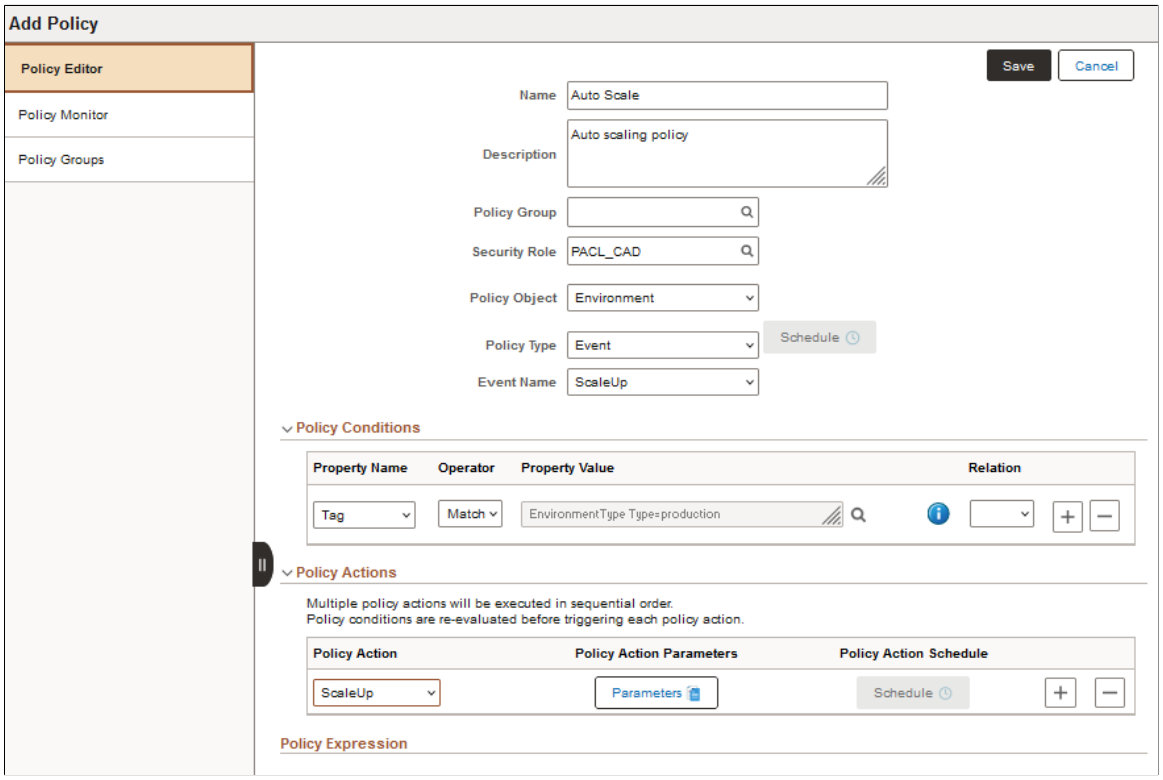

Field	***Value***
Policy Object	Environment

Field	*Value*
Policy Type	Event
Event Name	ScaleUp (positive anomaly) or ScaleDown (negative anomaly)
Policy Conditions	Select the environments by name or tag.

This example illustrates the Policy Action Parameters page for ScaleUp.

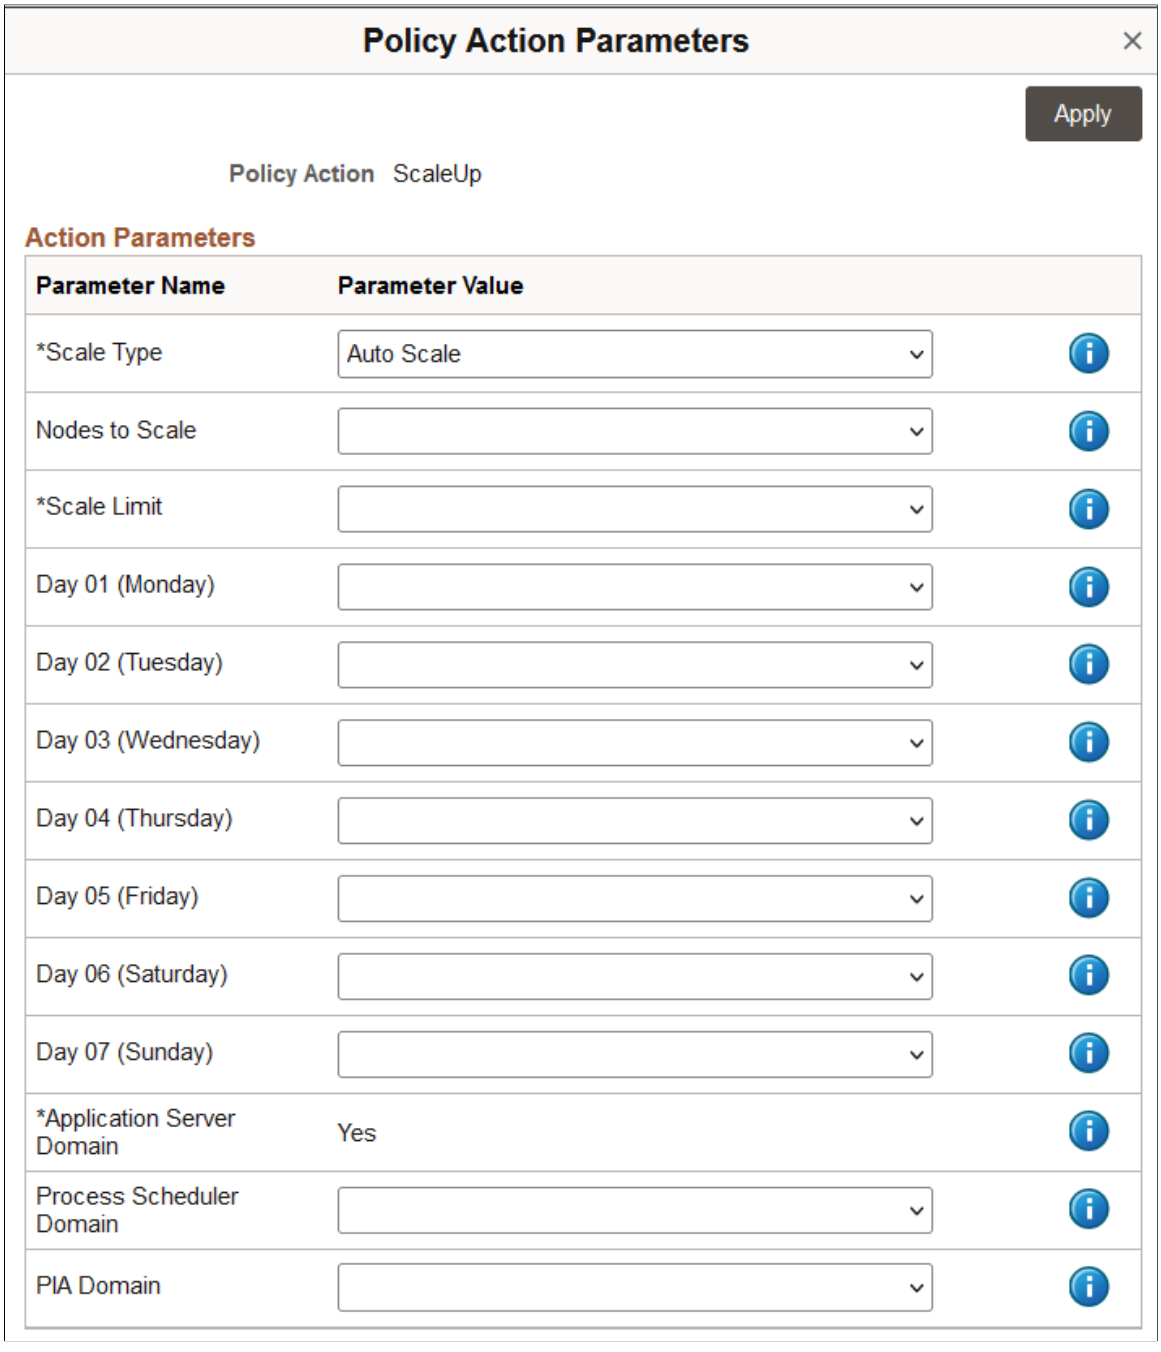

Field or Control	*Description*
Scale Type	Select Scale Type: • Auto Scale System will automatically scale when a positive or negative anomaly is detected. • Notification System will send the user a notification, which can also be viewed in the notifications alert window.The user must manually take corrective action based on the notification.
Nodes to Scale	Number of nodes to be scaled. Set to Auto when deploying an auto scale policy.
Scale Limit	For a ScaleUp policy set this field to the maximum number of nodes that can exist in the environment. For a ScaleDown policy set this field to the minimum number of nodes that should remain on the environment.
Day 01 (Monday) through Day 07 (Sunday)	Select the desired number of nodes for the day.
Application Server Domain	Select Yes to scale Application Server Domain. (Required)
Process Scheduler Domain	Select Yes to scale Process Scheduler Domain.
PIA Domain	Select Yes to scale PIA Domain.

Detecting Anomaly

When environment monitoring is enabled, Cloud Manager will send performance data every 5 minutes to Data Science prediction API.

If Scaling policies are active and Scale Type is set as Auto Scale, when there is a confirmed positive anomaly, additional nodes(s) will be added. Adding a new node typically takes 15 to 20 minutes.

System will add nodes until the Scale Limit set in the Scale Up policy is reached in case of successive scale up events.

After a scale up event is triggered, the system, by default, waits a lock down period of three hours before bringing the system back to the base configuration by removing nodes. The system needs to detect a negative anomaly after the lock down period in order for the system to remove nodes. System will remove nodes until the Scale Limit set in the Scale Down policy is reached.

Note: Nodes can be removed below the base configuration level if the Scale Limit for a Scale Down event is less than the number of nodes in the original environment.

If Scaling policies are active and Scale Type is set as Notification, a notification will be sent to you when there is a confirmed positive or negative anomaly. The notification can also be viewed in the notifications alert window. You must manually take corrective action based on the notification. On clicking the notification, you can view a graph that includes recommendations for increasing or decreasing the number of nodes according to the current capacity and usage.

This example illustrates the chart representing the current capacity requirement and cost comparison for adding nodes to an environment. It appears on clicking the notification specific to an environment on the Alerts section under Notification panel on the home page.

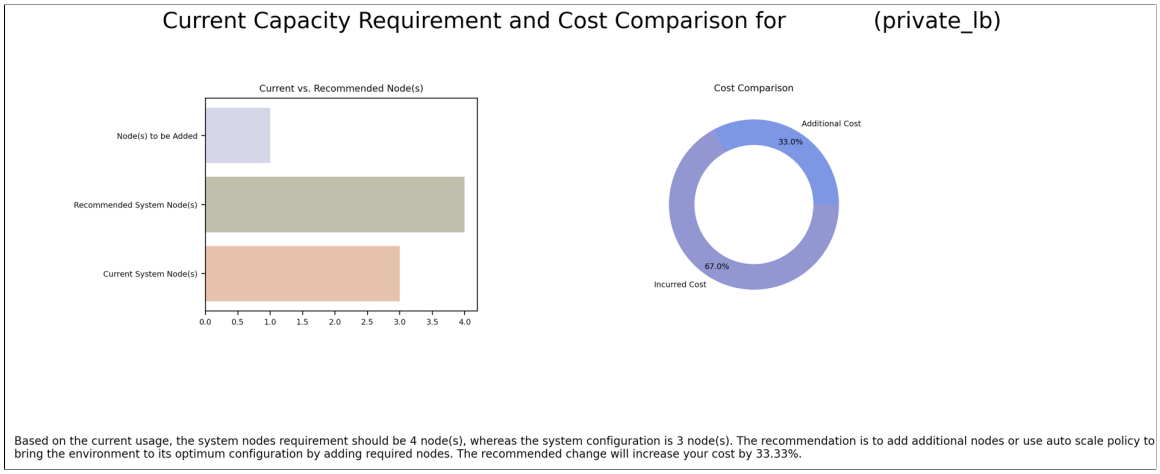

This example illustrates the chart representing the current capacity requirement and cost comparison for removing nodes from an environment. It appears on clicking the notification specific to an environment on the Alerts section under Notification panel on the home page.

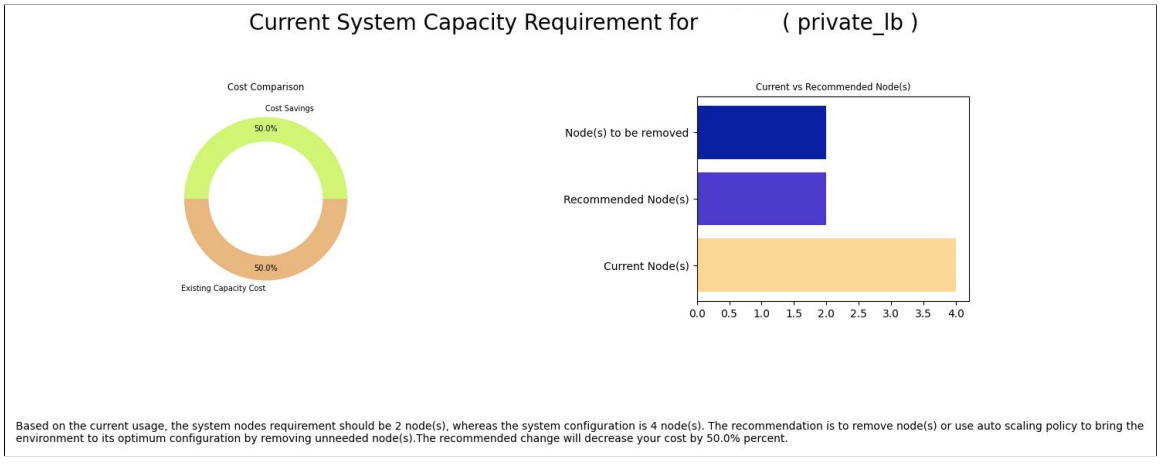

Adding a Policy with Multiple Actions

You can create policies that perform multiple sequential actions. Policy conditions are reevaluated before triggering each policy action.

Here are sample use cases for multiple action policies:

- Create a policy that starts an environment followed by a scale up action to run every Monday morning at 7 a.m.

  This policy would make sure that the environment automatically starts and adds an extra node to handle more load on the system.

- Create a policy to scale down followed by a stop environment to run Friday at 9 p.m.

  This policy would scale down and stop automatically, thereby eliminating resource usage over the weekend and providing a cost savings.

- Create a policy to back up the environment and then refresh.

  This policy will ensure the environment is properly backed up and available if the refresh fails.

- Create a policy that is run when an Infrastructure CPU DPK is downloaded in the Repository, to back up the environment, apply Infrastructure CPU, and send an email notification.

This is an example of a policy that will back up an environment and then refresh the environment.

This example illustrates the fields and controls on the Add Policy page for a policy containing multiple actions when the Policy Object is Environment.

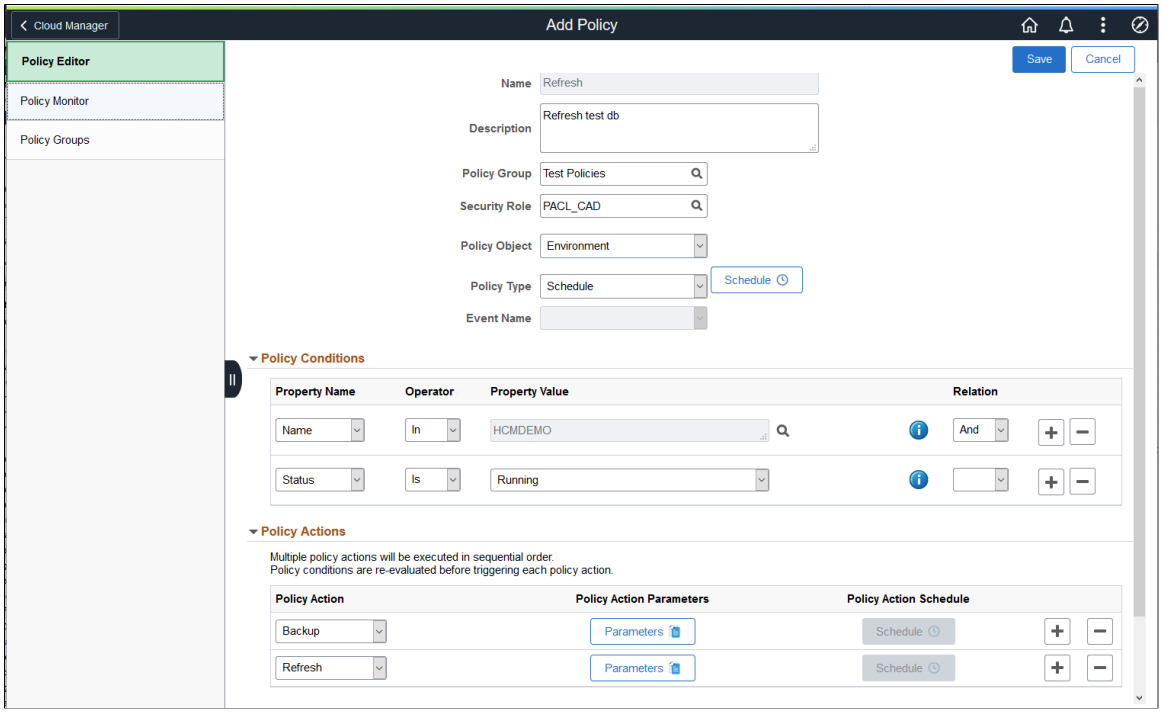

This is an example of a policy with multiple actions for a Repository Artifact. When a PeopleTools 8.60 Infrastructure CPU is downloaded, the policy will stop the specified environment, take a backup, start the environment, apply Infrastructure CPU patches, and then send an email notification.

This example illustrates the fields and controls on the Add Policy page for a policy containing multiple actions when the Policy Object is Repository Artifact.

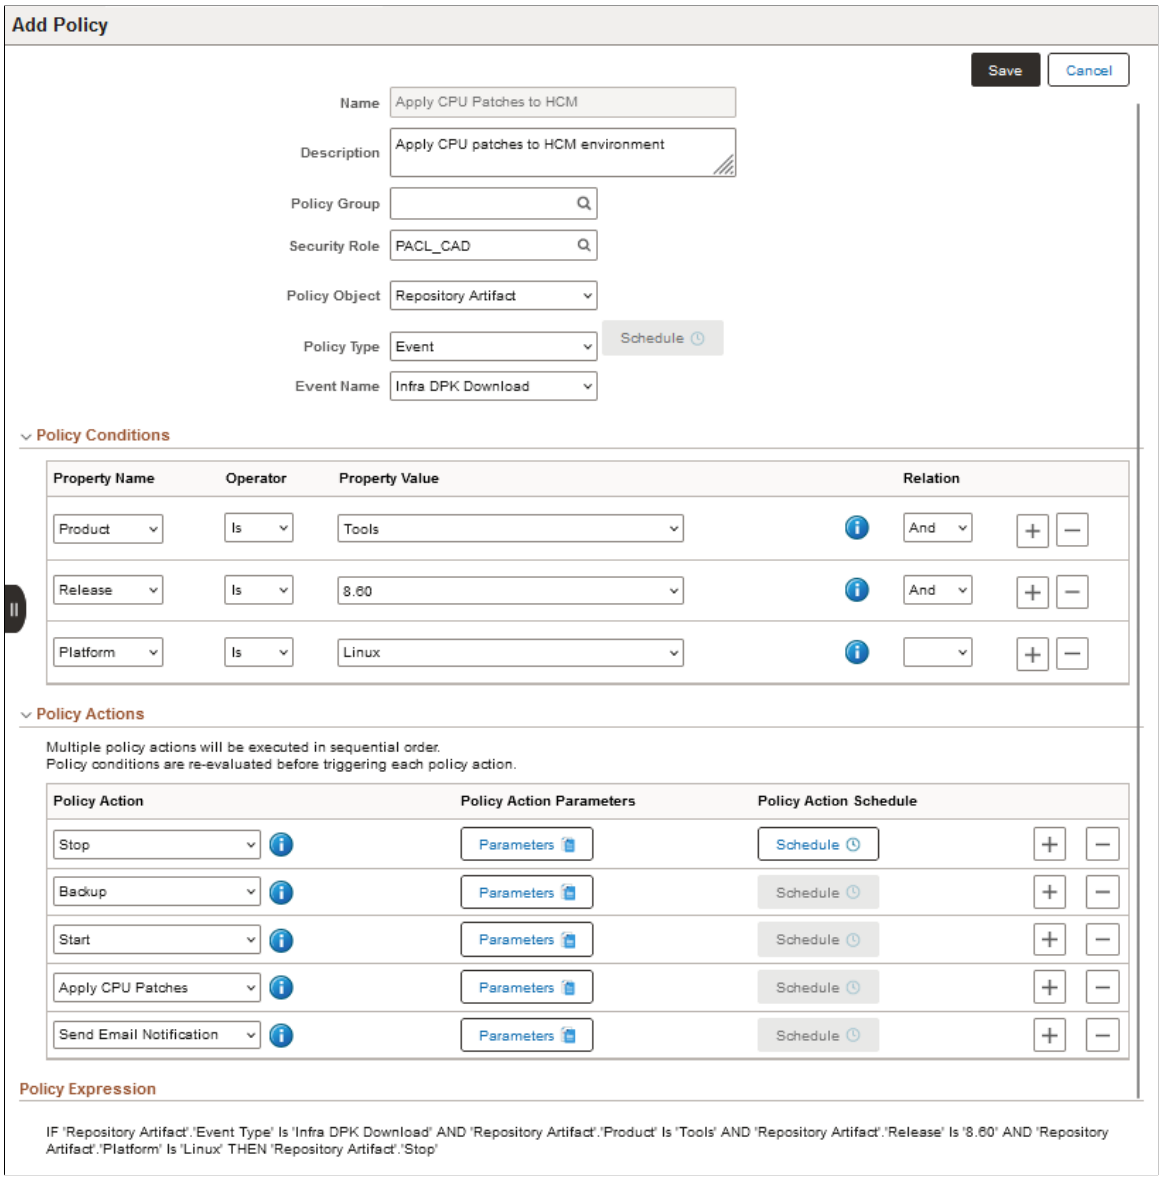

Adding a Policy with Custom Actions

You can define a custom action and include it as part of a policy action. For example, you can write a custom action to perform external health checks on provisioned environments or run OCI REST API calls.

Custom actions are supported for Environment and Repository Artifact policy objects, and scheduled or event-based policy types. The allowed parameters depend upon the policy object.

Parameter Type	*Parameter*	*Description*
Define Source: Select one of these options. If you select more than one you receive an error when you save the page.	PeopleCode Handler (Application Class)	Available for Repository Artifact and Environment policy objects. See the following sections for more details on using PeopleCode Handler.
	Inline Commands	Available for Environment policy objects and life-cycle events. To use inline commands for custom actions set Policy Object to Environment and Policy Type to Event. Enter one or more comma-separated commands. For Linux nodes the commands will be interpreted as shell (.sh) commands. For Windows nodes the commands will be interpreted as batch (.bat) commands. This is not applicable for Delete event.
	Repository File	Available for Environment policy objects and life-cycle events. Upload the file to the Repository before creating the policy action. See Upload Custom Scripts Page. To use an uploaded repository file for custom actions set Policy Object to Environment and Policy Type to Event.
Set Execution Content. Enter the content needed to run the custom action.	Input (JSON)	Available for Repository Artifact and Environment policy objects. If used with the Repository File parameter, you can use predefined variables to set environment variables. See Using Environment Variables with Custom Actions Based on Repository Files or Command Lines.
	Operating System	Available for Environment policy objects and life-cycle events. Select Linux or Windows.

Parameter Type	*Parameter*	*Description*
	Node Type	Available for Environment policy objects and life-cycle events. Select Database, DBSystem, Middle Tier, Full Tier, Search Stack, or PeopleSoft Client. This is optional. If the Operating System is specified and Node Type is not specified, all nodes of the chosen operating system will be used.

The following guidelines apply to custom actions:

- Ensure that custom scripts exit with the correct status so that the proper script processing status is propagated to Cloud Manager. This ensures that the correct status is shown on the environment card.
- Custom actions based on Repository Files or Command Line(s) are not supported for Stop events.
- Custom actions based on Repository Files or Command Line(s) are not supported for schedule-based policies.
- Custom actions based on Repository Files are not applicable when the Policy Object is Repository Artifact.
- If a custom action based on Repository Files or Command Line(s) is associated with a Delete event, the script will run locally on Cloud Manager. For other events, the script is run remotely, on the managed environment node.

Using Environment Variables with Custom Actions Based on Repository Files or Command Lines

As part of policy processing of custom actions based on Repository Files, Cloud Manager sets environment variables that will be available in the custom script being run.

Cloud Manager sets the following environment variables as part of a Custom Action:

Environment Variable	*Description*
PSFT_ADDED_NODE	Identify a newly-added node. This is supported only for Add Node policies.
PSFT_LIFE_CYCLE_EVENT_NAME	Event name for which the action is triggered.
PSFT_POLICY_NAME	Policy name

See Upload Custom Scripts Page for a list of the environment variables set by Cloud Manager during custom action processing and examples of using the environment variables.

You can target the implementation of a custom script or command line(s) on specific environment nodes by setting the following variables in the Input (JSON) field:

Environment Variable	*Allowed Values*	*Description*
MIDTIER_TYPE	• all — Target all domains. • appserver — Target nodes with only the application server domain. • prcs — Target nodes with only the Process Scheduler domain. • pia — Target nodes with only the PIA domain. • appbatch — Target nodes with the application server and Process Scheduler domains.	This variable is only applicable if the NodeType selected on the Policy Action Parameter page is Middle Tier or if Node Type is left blank. Use this variable to target Middle Tier nodes for specific domain combinations. If this variable is not supplied then all Middle Tier nodes are selected for a particular operating system type.
INSTANCE	• SINGLE — Target only one node. • NEWLY_ADDED — Target only the newly-added node.	This variable is applicable to Middle Tier, Search Stack, and PeopleSoft Client. For Middle Tier, this variable will be used along with MIDTIER_TYPE for node selection. This variable will not be applicable for DBsystems, Database, or Full Tier as the Number of Instance will always be one in those cases.

This example illustrates the Policy Action Parameters dialog box with Input (JSON) using the INSTANCE and MIDTIER_TYPE environment variables.

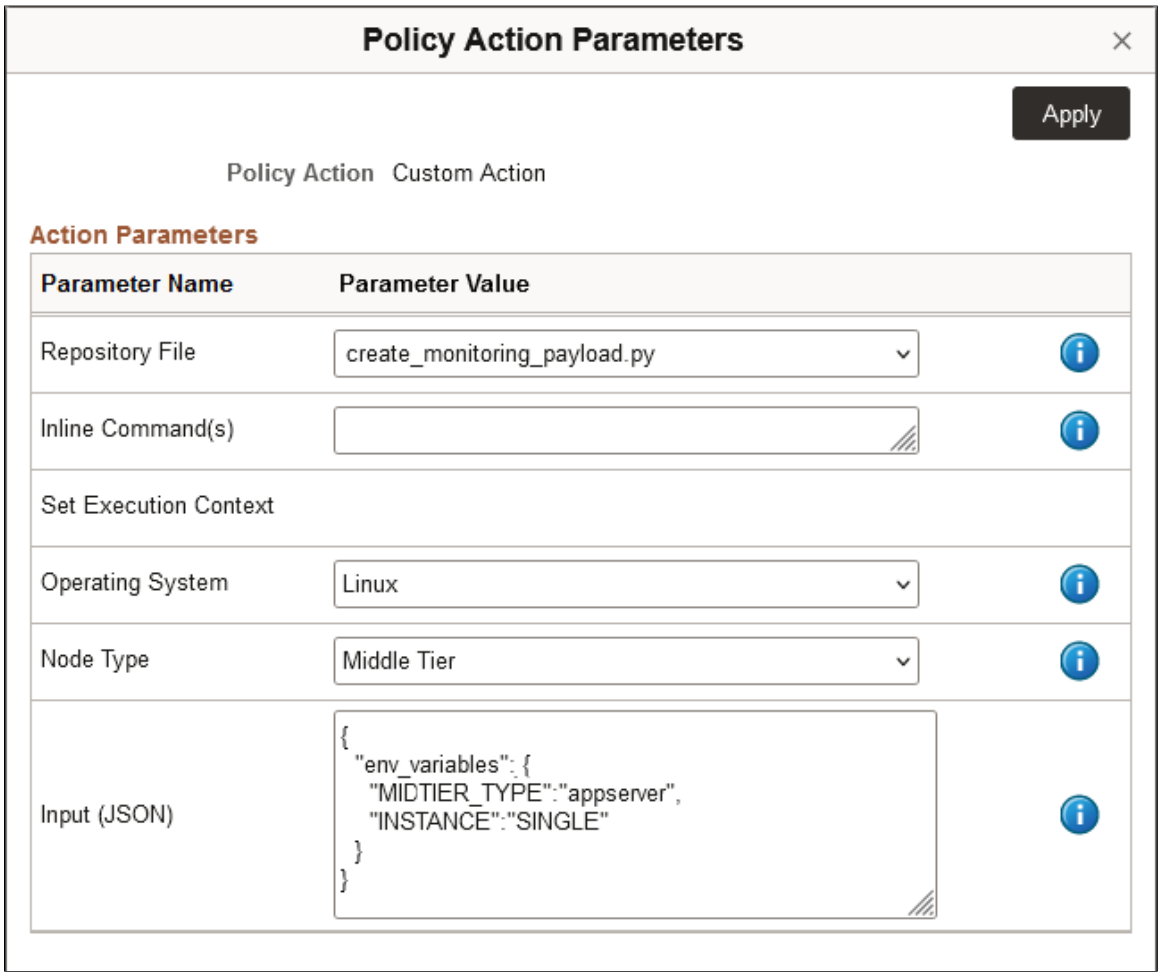

Here is an example of using the delivered environment variables. The Policy Action Parameters page includes these entries:

- Repository File = Python file (create_monitoring_payload.py)
- Node Type = Middle Tier
- Operating System = Linux
- Input (JSON) includes MIDTIER_TYPE = appserver and INSTANCE = SINGLE

The environment variables in the Input (JSON) directs the script to act on a single Linux Middle Tier node that has only an application server.

Defining a Policy Action with PeopleCode Handler

Cloud Manager delivers a PeopleSoft application class, ECL_CM:Governance:CustomActionHandler, which acts as the superclass for any custom policy action handler class. You are responsible for writing the processing logic for the custom action.

To set up a policy for a custom action, define a PeopleCode application class extending from the class ECL_CM:Governance:CustomActionHandler. When the policy is run, a predefined method of the instance of the handler class is invoked.

See the product documentation *PeopleTools: PeopleCode API Reference,* Understanding Application Classes on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

You can create a policy with multiple actions that combines several custom actions, or a mixture of custom actions and predefined actions, in sequence.

The base PeopleSoft application class is structured with input and output methods. When you set up a policy with multiple actions they are run one after another in the order you define. This gives you the option to write a series where the output from one policy action is passed to the next policy action as input.

Using Application Designer to Edit PeopleCode

Use Application Designer, which is included with the PeopleSoft Client, to write the PeopleCode application class.

To install a PeopleSoft Client:

1. Subscribe to the Interaction Hub (IH) download channel in the Repository.
2. Create a template to deploy the PeopleSoft Interaction Hub 9.1 database on a PUM topology.

   Note: The IH image should be based on the same PeopleTools release as the current Cloud Manager update image.

   The delivered PUM topology includes a PeopleSoft Client node on Microsoft Windows.
3. Create an environment from the IH template.

   Make a note of the Windows Client IP on the Environment Details page in the PeopleSoft Client section.
4. Use RDP to connect to the PeopleSoft Client.

   To determine the IP address of the PeopleSoft Client, go to the Environment Details page for the environment. Click Refresh Metadata and then select Diagrammatic View. Hold the mouse pointer over the PSoft Client box and look for the public ip in the table.
5. Follow the instructions in the product documentation to use Application Designer.

   See the product documentation *PeopleTools: PeopleCode Developer's Guide* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

Defining a PeopleCode Application Class for Custom Actions

Here is the definition of ECL_CM:Governance:CustomActionHandler class. Descriptions of the methods are included as comments (marked by /* at the beginning of the comment and */ at the end).

When you define the policy action for custom actions, you are required to provide input in JSON format. The contents in the Input (JSON) field is then passed as input through the "Execute" method (in bold font).

The SavePolicyActionOutput (in bold font) API can be used to save the policy action output, which can be seen from the Policy Monitor page. If you combine multiple custom actions, the method GetPreviousPolicyActionOutput (in bold font) is used to read the output produced by the previous custom policy action.

```
/***
*** Custom policy action handler classes should extend from this class and implemen⇒

t the Execute method.
*** Policy governance framework will invoke the Execute method of the handler class⇒

 while executing the policy
***/
class CustomActionHandler
   method CustomActionHandler();

   /* Property that can be used by subclasses to set a description for the custom h⇒

andler class */
   property string description;

protected
   method GetPolicyObjectName(&policyInstanceObj As string) Returns string;
   method GetPolicyArtifactId(&policyInstanceObj As string, &policyArtifactObj As s⇒

tring) Returns string;
   method SavePolicyActionOutput(&policyInstanceObj As string, &policyArtifactObj A⇒

s string, &output As string);
   method GetPreviousPolicyActionOutput(&policyInstanceObj As string, &policyArtifa⇒

ctObj As string) Returns string;

   method Execute(&policyInstanceObj As string, &policyArtifactObj As string, &hand⇒

lerInputObj As string);

private
   instance object &policyUtils;

end-class;

method CustomActionHandler
   &policyUtils = CreateJavaObject("com.peoplesoft.pa.cl.governance.policyengine.Po⇒

licyUtils");
   %This.description = "";
end-method;


/***
*** Description: Get the policy object name of the executed policy
*** Parameters:
*** &policyInstanceObj : JSON representation of the policy instance object containi⇒

ng policy instance metadata
***/
method GetPolicyObjectName
   /+ &policyInstanceObj as String +/
   /+ Returns String +/
   Local string &policyObjectName = "";
   Local JsonParser &jParser = CreateJsonParser();
   Local boolean &ret = &jParser.Parse(&policyInstanceObj);
```



```
   Local JsonObject &policyInstanceJsonObj = &jParser.GetRootObject();
   If &policyInstanceJsonObj.IsExist("policy_object_name") Then
      &policyObjectName = &policyInstanceJsonObj.GetAsString("policy_object_name");
   End-If;

   Return &policyObjectName;
end-method;


/***
*** Description: Get the policy artifact id associated with the executed policy
*** Parameters:
*** &policyInstanceObj : JSON representation of the policy instance object containi⇒

ng policy instance metadata
*** &policyArtifactObj : JSON representation of the policy artifact object associat⇒

ed with the executed policy
***/
method GetPolicyArtifactId
   /+ &policyInstanceObj as String, +/
   /+ &policyArtifactObj as String +/
   /+ Returns String +/
   Local string &policyArtifactId = "";
   Local string &policyObjectName = %This.GetPolicyObjectName(&policyInstanceObj);
   Local JsonParser &jParser = CreateJsonParser();
   Local boolean &ret = &jParser.Parse(&policyArtifactObj);
   Local JsonObject &policyArtifactJsonObj = &jParser.GetRootObject();

   /* We are handling only environment artifacts for now */
   If &policyObjectName = "Environment" And
         &policyArtifactJsonObj.IsExist("name") Then
      &policyArtifactId = &policyArtifactJsonObj.GetAsString("name");
   End-If;
   If &policyArtifactId = "" And
         &policyArtifactJsonObj.IsExist("Product") And
         &policyArtifactJsonObj.IsExist("Release") Then
      &policyArtifactId = &policyArtifactJsonObj.GetAsString("Product") | "_" | &po⇒

licyArtifactJsonObj.GetAsString("Release");
   End-If;

   Return &policyArtifactId;
end-method;


/***
*** Description: Save the policy action output. Policy action output can be seen fr⇒

om Policy Monitor page.
*** Parameters:
*** &policyInstanceObj : JSON representation of the policy instance object containi⇒

ng policy instance metadata
*** &policyArtifactObj : JSON representation of the policy artifact object associat⇒

ed with the executed policy
*** &output : Output generated for the policy action execution
***/
method SavePolicyActionOutput
   /+ &policyInstanceObj as String, +/
   /+ &policyArtifactObj as String, +/
   /+ &output as String +/

   Local string &policyInstanceId = "";
   Local string &policyObjectName = "";
   Local string &jobId = "";
   Local JsonParser &jParser = CreateJsonParser();
   Local boolean &ret = &jParser.Parse(&policyInstanceObj);
   Local JsonObject &policyInstanceJsonObj = &jParser.GetRootObject();
```

```
   If &policyInstanceJsonObj.IsExist("policy_object_name") Then
      &policyObjectName = &policyInstanceJsonObj.GetAsString("policy_object_name");
   End-If;
   If &policyInstanceJsonObj.IsExist("policy_instance_id") Then
      &policyInstanceId = &policyInstanceJsonObj.GetAsString("policy_instance_id");
   End-If;
   If &policyInstanceJsonObj.IsExist("policy_action_jobid") Then
      &jobId = &policyInstanceJsonObj.GetAsString("policy_action_jobid");
   End-If;

   &policyUtils.updatePolicyActionOutput(&policyInstanceId, &policyObjectName, &pol⇒

icyArtifactObj, &jobId, &output);
end-method;


/***
*** Description: Get output of previous policy action
*** NOTE: At present, we will return output from the previous custom policy action ⇒

(if any). Defined policy actions currently does not produce any output as such.
*** This API is given a generic name, so that when we allow output from defined pol⇒

icy actions (in future), we can re-use the same API
*** Parameters:
*** &policyInstanceObj : JSON representation of the policy instance object containi⇒

ng policy instance metadata
*** &policyArtifactObj : JSON representation of the policy artifact object associat⇒

ed with the executed policy
*** Returns the output of previous policy action as a string. Empty string if previ⇒

ous policy action is not found or if it does not produce any output
***/
method GetPreviousPolicyActionOutput
   /+ &policyInstanceObj as String, +/
   /+ &policyArtifactObj as String +/
   /+ Returns String +/
   Local string &output = "";
   Local string &policyInstanceId = "";
   Local string &policyObjectName = "";
   Local string &jobId = "";
   Local JsonParser &jParser = CreateJsonParser();
   Local boolean &ret = &jParser.Parse(&policyInstanceObj);
   Local JsonObject &policyInstanceJsonObj = &jParser.GetRootObject();
   If &policyInstanceJsonObj.IsExist("policy_object_name") Then
      &policyObjectName = &policyInstanceJsonObj.GetAsString("policy_object_name");
   End-If;
   If &policyInstanceJsonObj.IsExist("policy_instance_id") Then
      &policyInstanceId = &policyInstanceJsonObj.GetAsString("policy_instance_id");
   End-If;
   If &policyInstanceJsonObj.IsExist("policy_action_jobid") Then
      &jobId = &policyInstanceJsonObj.GetAsString("policy_action_jobid");
   End-If;

   &output = &policyUtils.getPreviousPolicyActionOutput(&policyInstanceId, &policyO⇒

bjectName, &policyArtifactObj, &jobId);
   Return &output;
end-method;


/***
*** Description: Custom policy action handler classes should implement this method.⇒

 This method will be invoked while executing the policy action
*** Parameters:
*** &policyInstanceObj : JSON representation of the policy instance object containi⇒
```



```
ng policy instance metadata
*** &policyArtifactObj : JSON representation of the policy artifact object associat⇒

ed with the executed policy
*** &handlerInputObj : JSON representation of the input provided in the policy acti⇒

on parameters, while defining the policy
***/
method Execute
   /+ &policyInstanceObj as String, +/
   /+ &policyArtifactObj as String, +/
   /+ &handlerInputObj as String +/
end-method;
```

Here are definitions for the parameters included in the Execute and SavePolicyActionOutput methods:

Parameter Name	***Parameter Definition***	***Example***

policyInstanceObj	JSON string representation of policy instance (implemented policy) details.
	
	The following attributes are available in the JSON object:
	
	• policy_object_name: Policy object defined when creating the policy (Environment or Repository Artifact).
	
	• policy_action_jobid: Job ID of the orchestration job for the policy action handler.
	
	• policy_instance_id: Policy instance ID
		⇒
		
		{
		"policy_object_name": ⇒
		
		"Environment",
		"policy_action_jobid":⇒
		
		"1695014483158_548813",
		"policy_instance_id": ⇒
		
		"0a64a811-a53a-4931-bc46-4⇒
		
		82049339d71"
		}

policyArtifactObj	JSON string representation of the policy artifact associated with the policy instance (executed policy)
	
	The following attributes are available in the JSON object:
	
	• name: Managed environment name associated with the policy.
	
	• owner (Optional): Owner of the managed environment.
	
	• status (Optional): Status of the managed environment.
		⇒
		
		{
		"owner": "CLADM",
		"name": "HCMProd",
		"status": "Running"
		}

handlerInputObj	JSON string representation of the input defined for the custom policy action. This is the input added in the policy action parameter dialog box.	none

Reviewing a Sample PeopleCode Application Class

This is an example of a PeopleCode application class for a custom policy action, EmailNotificationHandler, which sends email notifications during policy processing.

Before invoking the policy, the user must set up OCI Notification and update the Notification Topic OCID in Cloud Manager Infrastructure Settings page.

See Infrastructure Settings Page.

All custom policy action handler classes should implement the "Execute" method. This method will be invoked when the policy runs.

```
import ECL_CM:Governance:CustomActionHandler;

class EmailNotificationHandler extends ECL_CM:Governance:CustomActionHandler
   method EmailNotificationHandler();
   method Execute(&policyInstanceObj As string, &policyArtifactObj As string, &hand⇒

lerInputObj As string);
end-class;


method EmailNotificationHandler
   %Super = create ECL_CM:Governance:CustomActionHandler();
   %This.description = "Custom policy action handler class for sending email notifi⇒

cations";
end-method;

method Execute
   /+ &policyInstanceObj as String, +/
   /+ &policyArtifactObj as String, +/
   /+ &handlerInputObj as String +/


   Local string &artifactId = %Super.GetPolicyArtifactId(&policyInstanceObj, &polic⇒

yArtifactObj);
   Local JsonParser &jsonParser1 = CreateJsonParser();
   Local boolean &ret = &jsonParser1.Parse(&policyInstanceObj);
   Local JsonObject &policyInstance = &jsonParser1.GetRootObject();
   Local string &policyInstanceId = &policyInstance.GetAsString("policy_instance_id⇒

");
   Local string &curdatetime = DateTimeToLocalizedString(%Datetime, "yyyy-MM-dd HH:⇒

mm:ss");
   Local string &policyName = "";

   /* Populate email title and body */
   Local string &topic = "Topic-PTU";
   Local string &title = "Policy execution notification";
   SQLExec("select ecl_name from ps_ecl_policy_mon where ecl_rp_id = :1", &policyIn⇒

stanceId, &policyName);

   Local string &body = Char(10) | Char(10) | "Policy execution details below" | Ch⇒

ar(10);
   &body = &body | Char(10) | "Policy Name : " | &policyName | Char(10);
   &body = &body | Char(10) | "Environment: " | &artifactId | Char(10);
   &body = &body | Char(10) | "Performed at: " | &curdatetime | Char(10);

   /* Add any custom data passed to the policy in case if needed */
   If &handlerInputObj <> "" Then
      Local JsonParser &jsonParser2 = CreateJsonParser();
```



```
      &ret = &jsonParser2.Parse(&handlerInputObj);
      Local JsonObject &handlerInput = &jsonParser2.GetRootObject();
      If &handlerInput.IsExist("notes") Then
         &body = &body | Char(10) | "Notes: " | &handlerInput.GetAsString("notes") ⇒

| Char(10);
      End-If;
   End-If;

   /* Send email notification */
   /* Refer CM documentation on how to setup a notification topic for email notific⇒

ations */
   Local JavaObject &notificationObj = CreateJavaObject("com.peoplesoft.pa.cl.commo⇒

n.OCINotificationServiceImpl");
   Local JavaObject &messageObj = CreateJavaObject("com.peoplesoft.pa.cl.common.Mes⇒

sage");
   &messageObj.setBody(&body);
   &messageObj.setTitle(&title);
   &notificationObj.send(&topic, &messageObj);

   /* Expose policy action output if needed */
   Local string &actionOutput = "Successfully executed EmailNotificationHandler. Ma⇒

il content: " | &body;
   %Super.SavePolicyActionOutput(&policyInstanceObj, &policyArtifactObj, &actionOut⇒

put);

end-method;
```

Defining a Policy Action with Custom Action

To add a policy with a custom action:

1. In Application Designer, write PeopleCode for the desired action based on the ECL_CM:Governance:CustomActionHandler class.
2. In Cloud Manager, define a policy and select Custom Action as the Policy Action.
3. Click the Parameter button for the Policy Action and supply these parameters to associate the PeopleCode application class with the policy action.

Field or Control	***Description***
PeopleCode Handler class	Enter the name of the PeopleCode application class to be invoked when the policy action is run.
Input (JSON)	(Optional) Enter the JSON string that is to be passed to the handler class when the policy action is run.

Field or Control	*Description*
Environment Name	This parameter is available for policies with Repository Artifact policy object. Click the search icon to display the Set Parameter Values page. Select a tab and use one of these methods to identify the environment: • Environment Names: Click the search icon to select a provisioned environment to associate with the custom action. • Tag: Select a Tag Namespace, and then select a tag that is associated with a provisioned environment.

4. Initiate the policy and follow the status on Policy Monitor.

   See Using Policy Monitor.

   Click the Policy Action Output icon (>) to view the output of the policy in JSON format.

Using Policy Monitor

Use the Policy Monitor page to track the execution status of the policies.

Select the View/Edit icon (>) to view the policy execution status.

This example illustrates the Policy Monitor page with policy execution action.

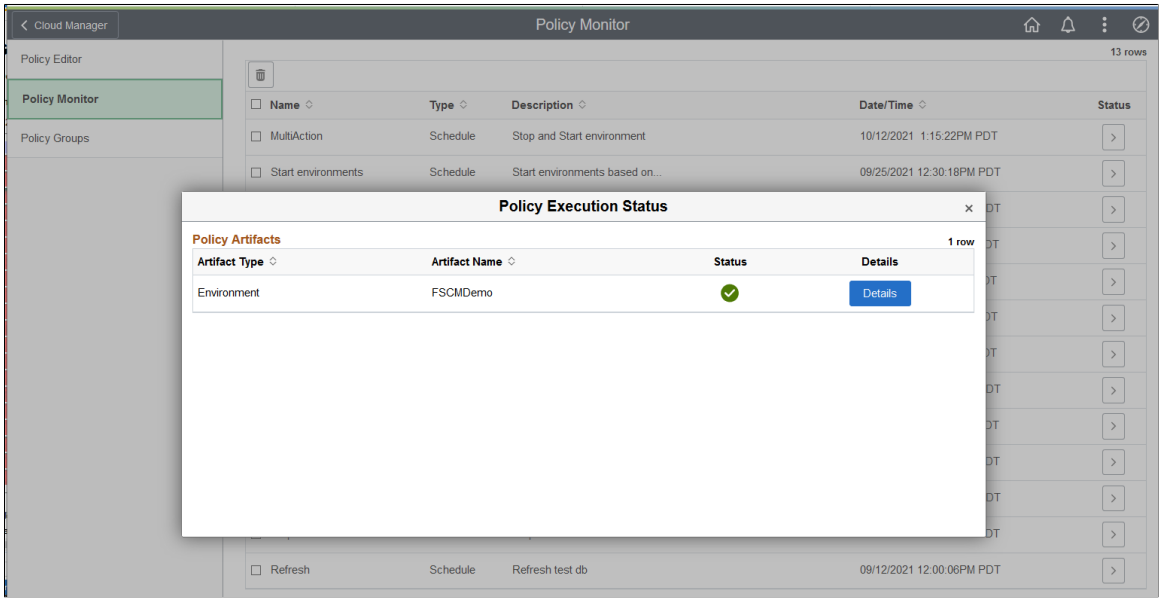

Click the Details button to view the policy action execution status.

This example illustrates the fields and controls on the Policy Action Execution Status page.

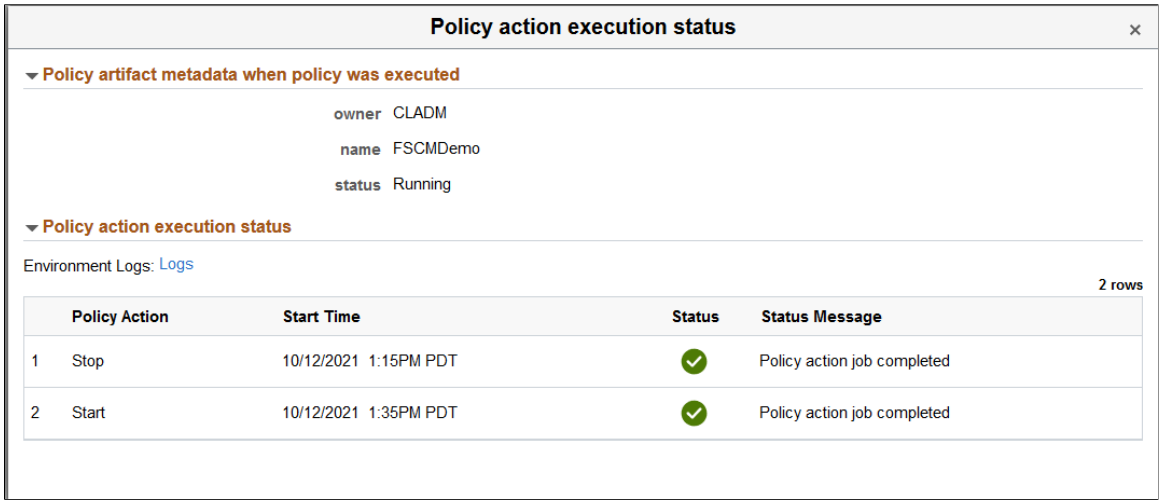

Creating Policy Groups

Policy groups are optional and are used to group policies together for display and to facilitate associating related policies.

To view or add a policy group, select the Orchestration Manager tile, then select Policy Group.

Enter a name and description for a policy groups. Add additional groups if necessary, and click **Save**.

This example illustrates the fields and controls on the Policy Groups page.

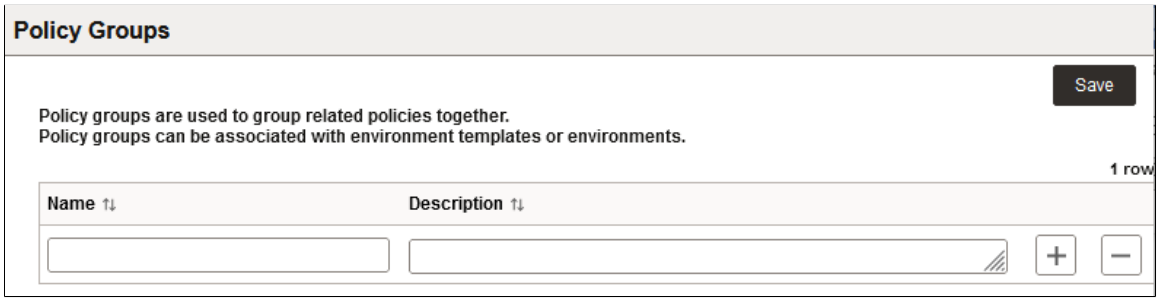

After adding a policy group, you can select it when creating an environment or environment template. Adding a policy group to an environment template adds all of the policies in that group. The environments created from the template will have these policies auto enabled.

See Accessing Environment Details and Environment Template – Security and Policies Page.

Policy groups with policies that are associated with an environment or template will be displayed on the Policy Editor page under Default Policies.

This example illustrates the fields and controls on the Policy Editor page that contains a Policy Group.

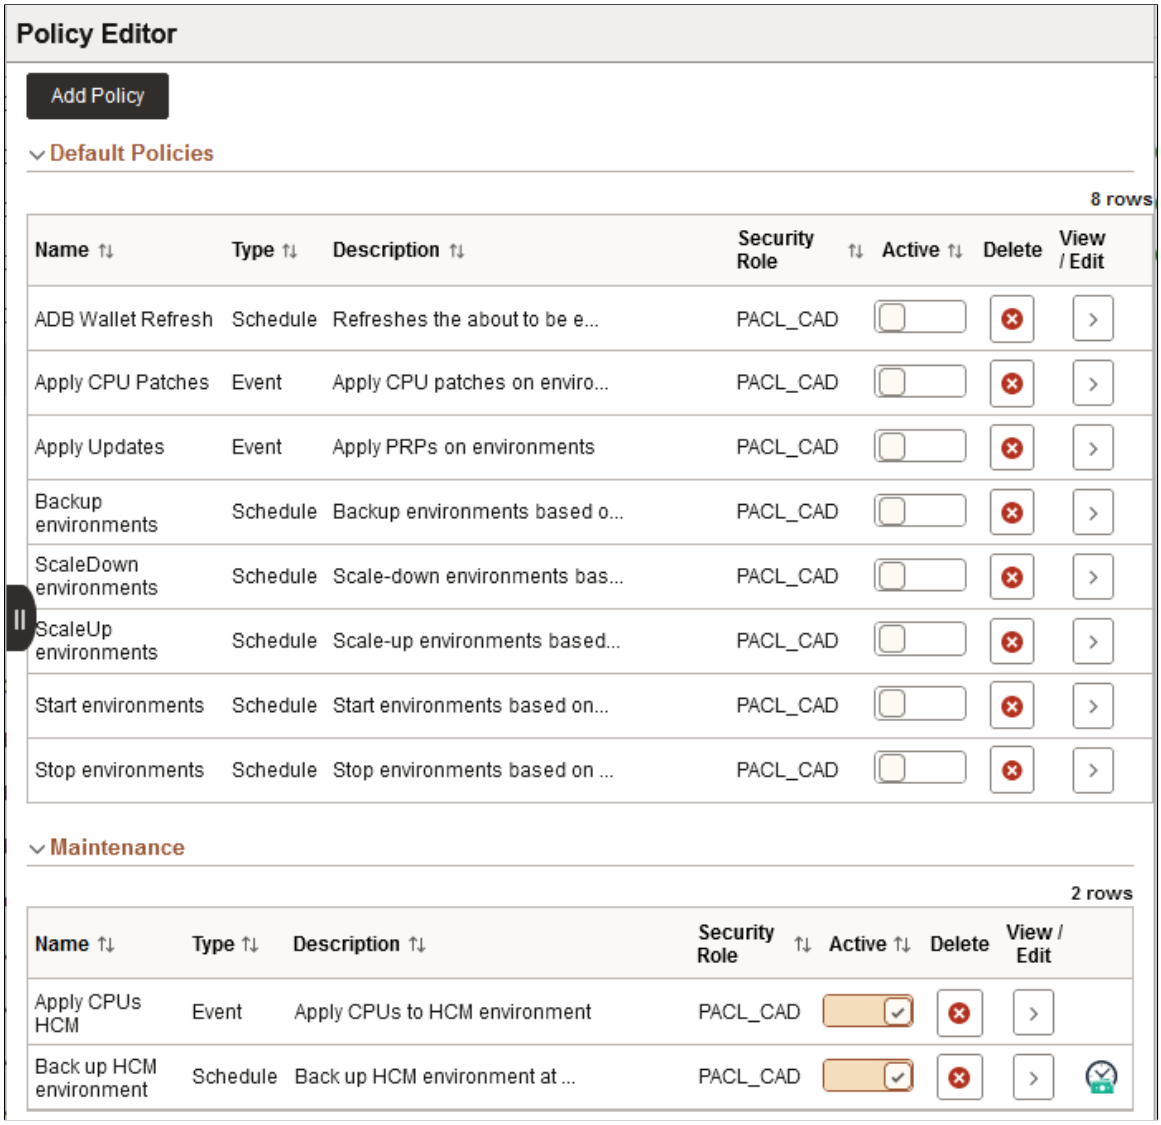

Configuring Self-Managed Update Images

Cloud Manager enables you to automate the process of provisioning PUM environments and keeping them current, which ensures that the latest PUM environment is available for you with all the PUM metadata and that the existing targets are migrated from the old PUM source. You can also automatically upload new targets on to the current PUM source, configure load balancer on the Current PUM source, as well as delete the old PUM source.

After defining the policy, you do not need to modify the policy again when new PeopleSoft Update Images are released, because a new PUM environment is provisioned using the downloaded PeopleSoft Update Image DPK whenever a new PeopleSoft Update Image is downloaded in Cloud Manager for the application pillar. Following this, all PUM metadata is migrated from the old PUM source environment based on the previous PeopleSoft Update Image, and all the target metadata is uploaded.

Note: For policy automation to work, old PUM source must be on PeopleTools 8.60.10 or higher and targets must be on PeopleTools 8.59.14 or higher.
You must ensure that current PUM source, old PUM source and targets are in the same subnet. If they are in different subnets, you must ensure that the current PUM source is able to make PeopleSoft IB REST calls to both the old PUM source and targets. This is required for the policy automation to work.

The following policy actions are added on the Repository Artifact policy object to enable automation of provisioning PUM environments:

1. Provision PUM.

   Cloud Manager automatically applies all the PRPs downloaded along with the PeopleSoft Update Image on the provisioned PUM environment, which is assigned as the current PUM source.

   This is a mandatory policy action. The policy action parameters are:

Policy Action Parameters	***Description***
Environment Name Prefix	Enter a prefix to identity the provisioned environment. When the policy provisions an environment, the environment name will include the prefix in the format <prefix><release number><image number> The name on the tile for the provisioned environment will reflect the downloaded PUM image (the latest image number), and the description on the provisioned environment tile will specify that it is auto-provisioned through a policy.
Environment Template Name	Select an existing PUM full-tier environment template for the application. The template definition may specify an earlier update image. The policy uses the template and substitutes the latest downloaded PUM image in the template definition. The template itself is not modified. The drop-down list includes only templates that have an associated password group. See Creating a Template.
Database Name Prefix	A prefix for the environment database name, with maximum length of three letters. This parameter is optional. If no value is given for this parameter, the environment is provisioned using the database name defined in the environment template. Even though the parameter is optional, adding a value for this parameter is highly recommended.

   See Enabling Selective Adoption in Cloud Manager.

2. Migrate PUM metadata.

   Cloud Manager invokes multiple PUM Automated Updates (PAU) REST APIs on the Current PUM source for each of these listed operations:

a. Define old PUM source

b. Migrate PUM metadata

c. Upload targets

If the target database is on PeopleTools 8.59 for patch 14 or higher, the service operation for uploading metadata works only after you assign service security for that operation. To assign service security:

a. Navigate to **PeopleTools**>**Integration Broker** >**Integration Setup**>**Service Operation Definitions**.

b. Type PTIA_UPLOAD_METADATA_POST in the Service Operation field and click Search.

This example illustrates the fields and controls on the Service Operations - Search page.

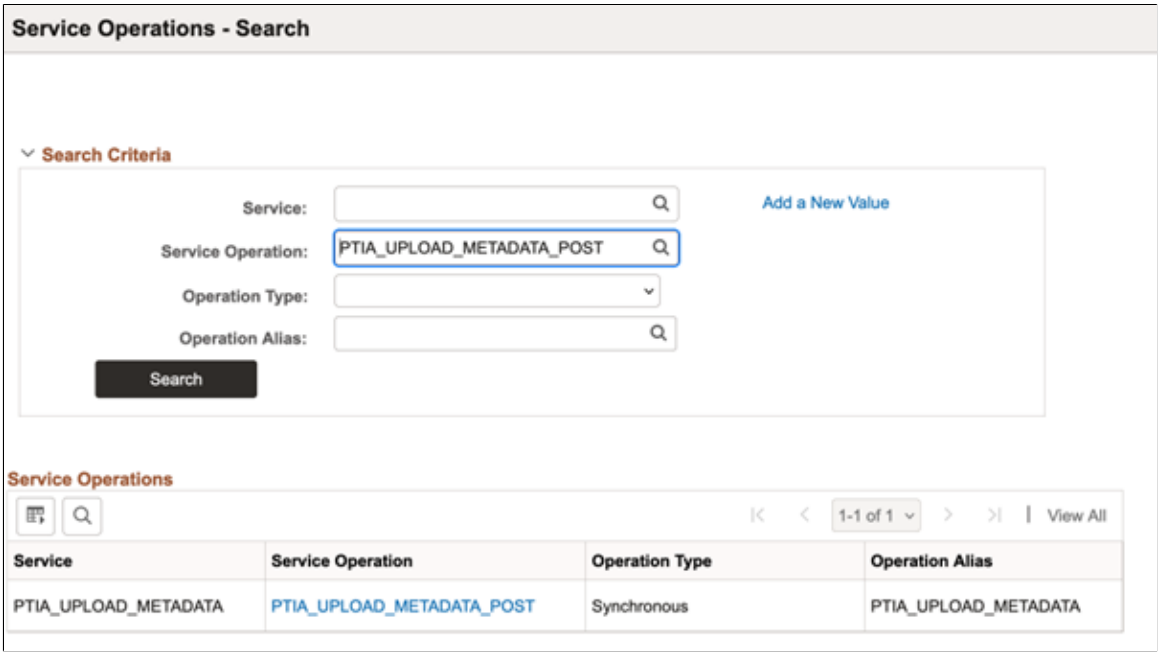

c. Click the Service Operation link. The Service Operations page appears.

This example illustrates the fields and controls on the Service Operations page.

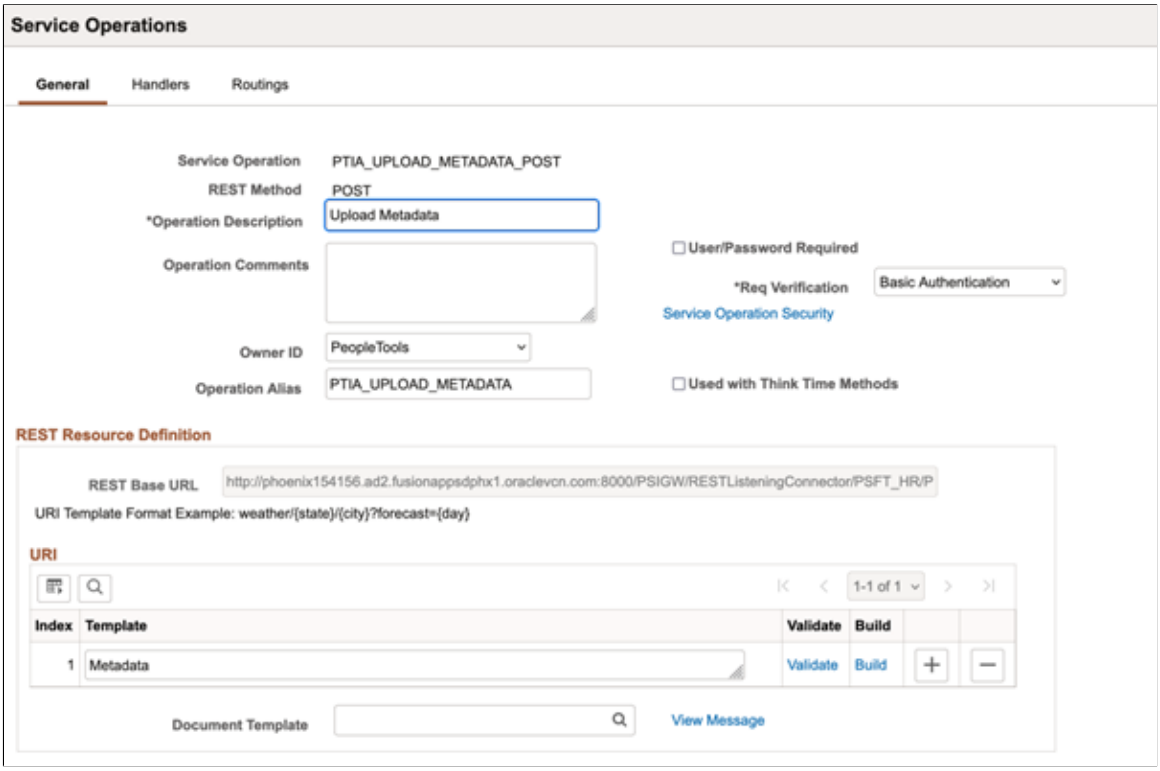

d. On the General tab of Service Operations page, click Service Operations Security link. Web Service Access page appears.

e. On the Web Service Access page, check if PTPT4300 is assigned in the Permission List. If not, select PTPT4300 in the Permission List field and provide it Full Access. Click Save.

This example illustrates the fields and controls on the Web Service Access page.

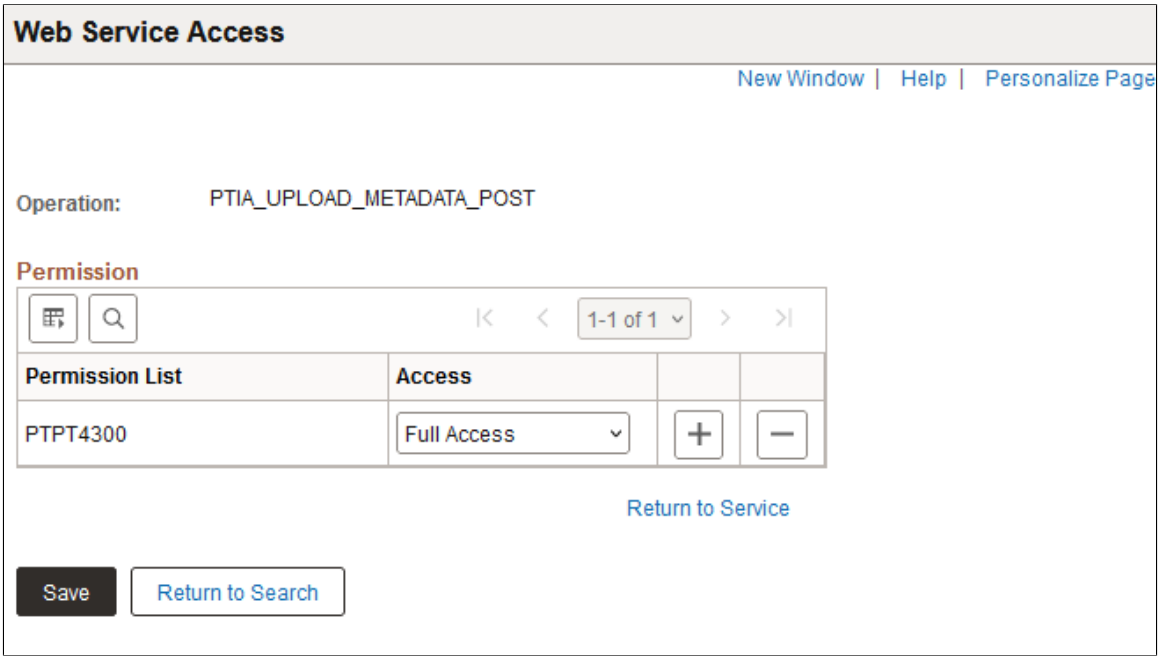

This is a mandatory policy action, which includes the following PAU operations:

Policy Action Parameters	***Description***
Select old PUM source	Select the old PUM source environment.
Custom IB local gateway node of old PUM source	If the old PUM Source environment has a custom IB local gateway node (If the environment is not provisioned through CM and has custom configuration), then specify the custom IB node name.
Auto-discover old PUM source on next policy run	If the expectation is to automatically switch the Current PUM source environment as the old PUM source when the policy is triggered for a new PeopleSoft Update Image, then set "Yes" for this parameter. If user would not want the policy to automatically switch old PUM source environment, then set the value to "No".
Migrate PUM metadata from old PUM source	Flag indicating that this policy action migrates PUM Metadata from old PUM source (This parameter is read-only). Cloud Manager migrates data for all PUM Metadata types such as customisation Repo Data, Package Data, PAU Data, and Test Repo Data by default.
Upload targets from old PUM source	Flag to indicate whether to upload all targets from old PUM source. The default value is "Yes". If you would like to only migrate PUM metadata from the old PUM source but do not want to upload existing targets, you can set the parameter to "No".

This example illustrates a sample policy definition on the Policy Editor page.

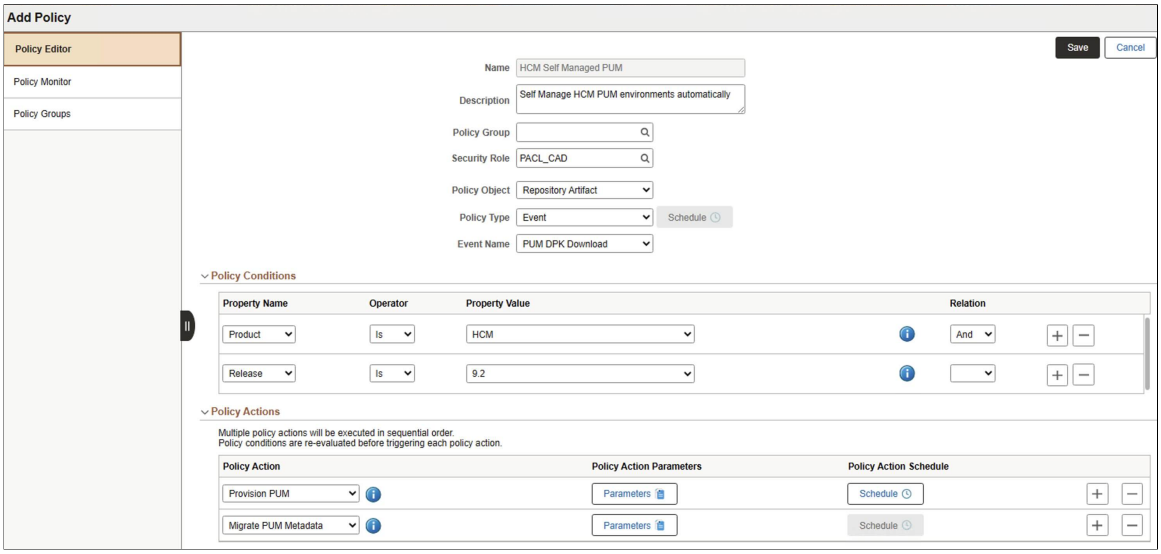

3. Define and upload target.

   Cloud Manager invokes PAU REST APIs to define and upload the new target. This is an optional policy action.

Policy Action Parameters	***Description***
Select PUM Target	Select a new Target environment.
Custom IB Local Gateway Node of Target	If the target environment has a custom IB local gateway node (If the environment is not provisioned through CM and has custom configuration), then specify the Custom IB node name.

   If you need to define more than one Target environments, this policy action can be repeated and you can specify the new target environment in each policy action.

   Note: Adding a new target through this policy action works only for target environments based on PUM full tier. Any number of targets with different topologies can be added directly from PUM source environments through available LCM user interfaces.

4. Delete the old PUM source.

   This is an optional policy action to remove the old PUM environment after the configurations are transferred. There are no parameters associated with it.

5. Custom Action to configure load balancer on the current PUM source.

   This is an optional policy action.

Policy Action Parameters	*Description*
Environment Names	$CURRENT_PUM is a policy variable representing the name of the current PUM source environment. As the current PUM source name is not known while defining the policy, the policy variable serves the purpose of associating the current PUM source name with the Invoke Handler policy action.
PeopleCode Handler (Application Class)	ECL_CM:Governance:CustomAction:ConfigurePUMLB. This custom PeopleCode application class is used to perform load balancer configuration on the current PUM source.
Input (JSON)	This custom action would require the following as input JSON: `{ "oci_loadbalancer_ocid": "xxx", "oci_loadbalancer_compartment_ocid⇒ ": "xxx", "oci_loadbalancer_backendset_name"⇒ : "xxx", "oci_loadbalancer_listener_name": "⇒ xxx", "psft_loadbalancer_fqdn": "xxx", "psft_web_server_domains(domain_na⇒ me:port)": "xxx"}` The attributes of input JSON are as follows: • oci_loadbalancer_ocid : OCID of the load balancer • oci_loadbalancer_compartment_ocid : OCID of the compartment where the load balancer exists • oci_loadbalancer_backendset_name : Name of the load balancer backend set • oci_loadbalancer_listener_name : Name of the load balancer listener • psft_loadbalancer_fqdn : FQDN of the PIA load balanced URL • psft_web_server_domains(domain_name:port) : Comma separated list of web server domain entries to be added into the backend set. Each entry should be of format Domain:Port.

Chapter 8

Managing Alerts and Notifications

Viewing Alerts and Notifications

Alerts and Notifications are used to notify users that new patches have been downloaded and are available for Cloud Manager, as well as the patch details and priority.

Note: You must subscribe to the IH_91_Linux download channel.

Use the Manage Updates page in Cloud Manager Settings to apply the PRPs. See Updating Cloud Manager Overview

The Notifications and Alerts tile is displayed on the Cloud Manager Homepage. If new patches are available the tile will indicate the number of notifications.

This example illustrates the Notifications and Alerts tile on the Cloud Manager homepage.

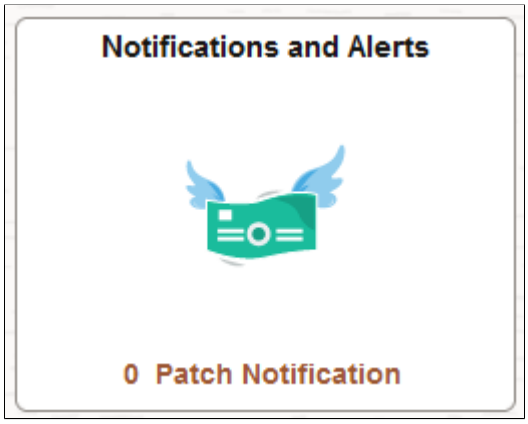

Using the Patch Notifications Page

Select the Notifications and Alerts tile to view the Patch Notifications page.

Note: Only unapplied patch notifications are displayed.

This example illustrates the fields and controls on the Patch Notifications page. You can find definitions for the fields and controls later on this page.

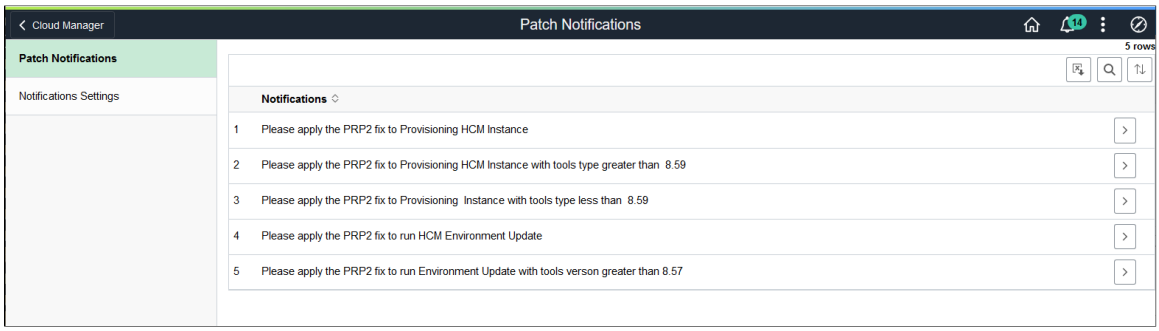

Patch notifications are displayed. Use the > button to view the notification information.

This example illustrates the fields and controls on the Notification Information. You can find definitions for the fields and controls later on this page.

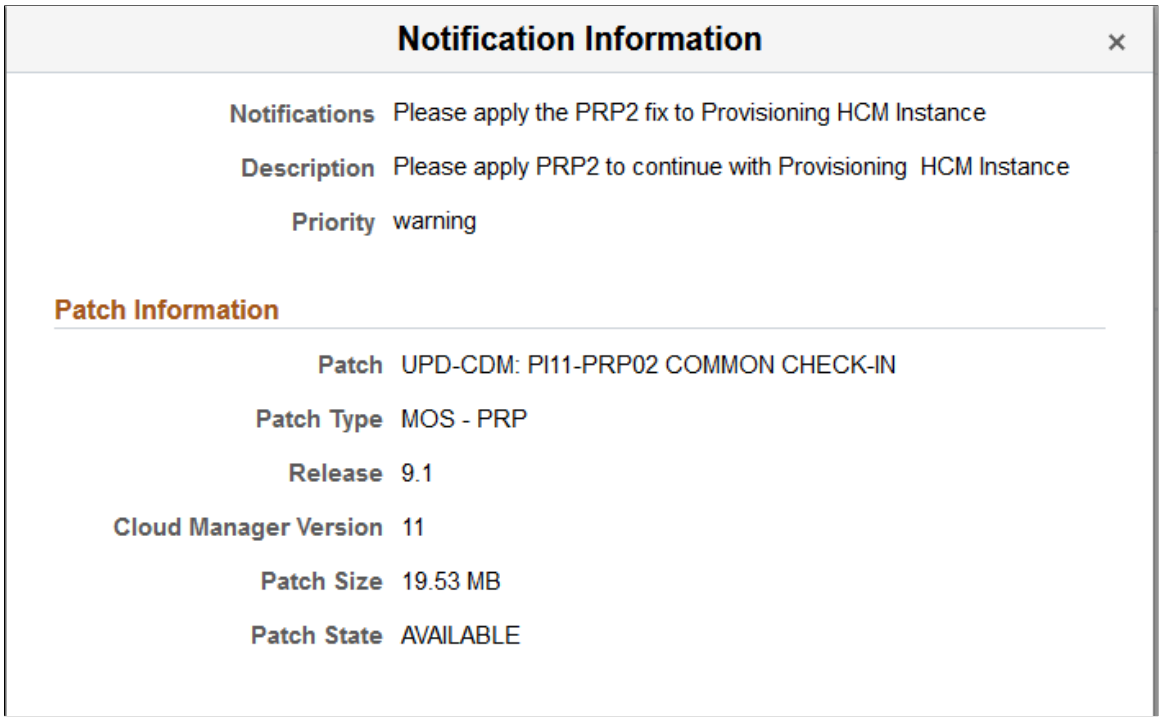

The notification information includes a description and priority, as well as the patch information.

Enabling Notifications

When you enable notifications, you will be alerted if you perform an activity in Cloud Manager that is affected by the patch. For example, you want to create a new environment or apply a PeopleTools patch. To enable notifications, select the Notifications and Alerts tile, then select Notification Settings.

This example illustrates the fields and controls on the Notification Settings page. You can find definitions for the fields and controls later on this page.

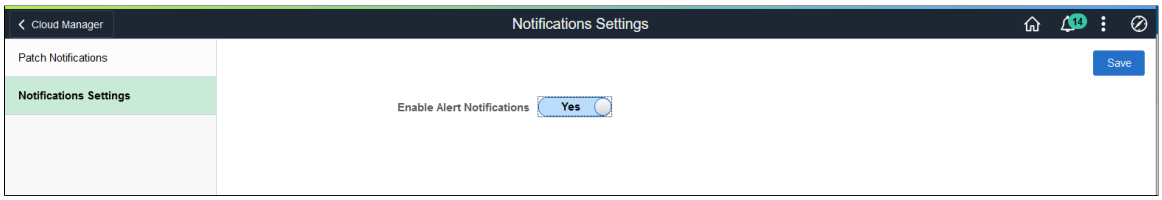

Field or Control	*Description*
Enable Alert Notifications	Select *Yes* to enable notifications. If you select *No*, you will not receive a notification when performing an activity affected by the patch.

Example Creating Environment

If notifications are enabled and a patch affects creating an environment, a message will be displayed.

This example illustrates a notification on the Create Environment page, when a patch is available and notifications are enabled.

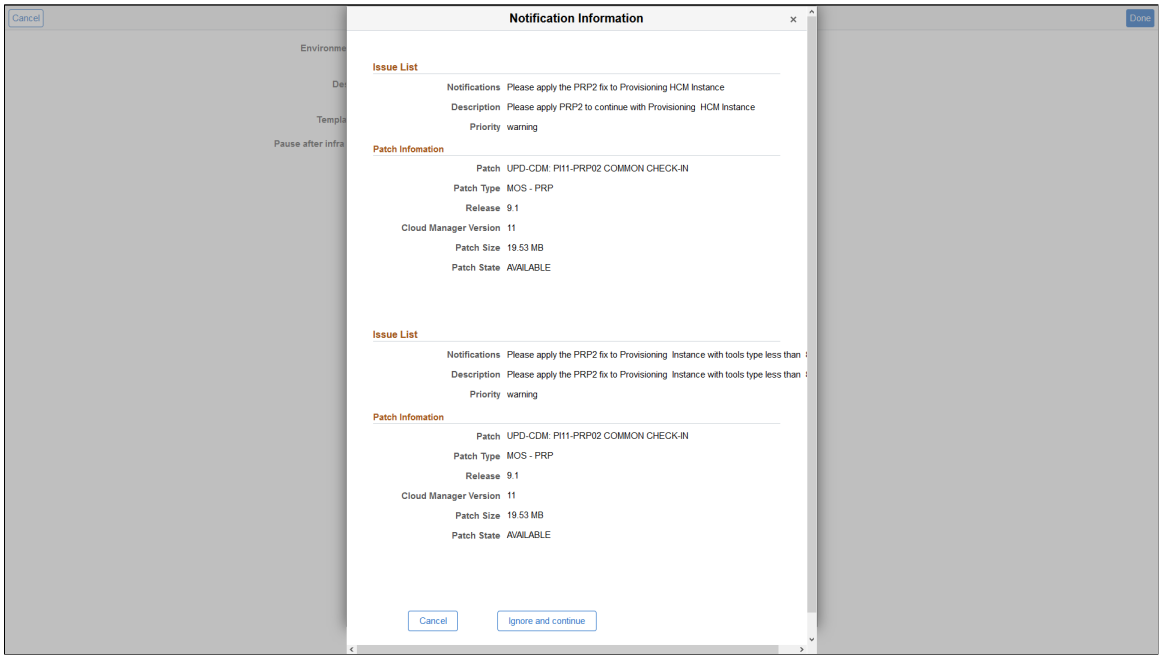

Field or Control	*Description*
Cancel	Select to cancel. Apply the patch using the Manage Updates page.

Field or Control	*Description*
Ignore and continue	Select to ignore the notification and continue.

Example Apply PeopleTools Patch

If notifications are enabled and a patch affects applying a PeopleTools patch, a message will be displayed.

This example illustrates a notification on the Apply PeopleTools Patch page, when a patch is available and notifications are enabled.

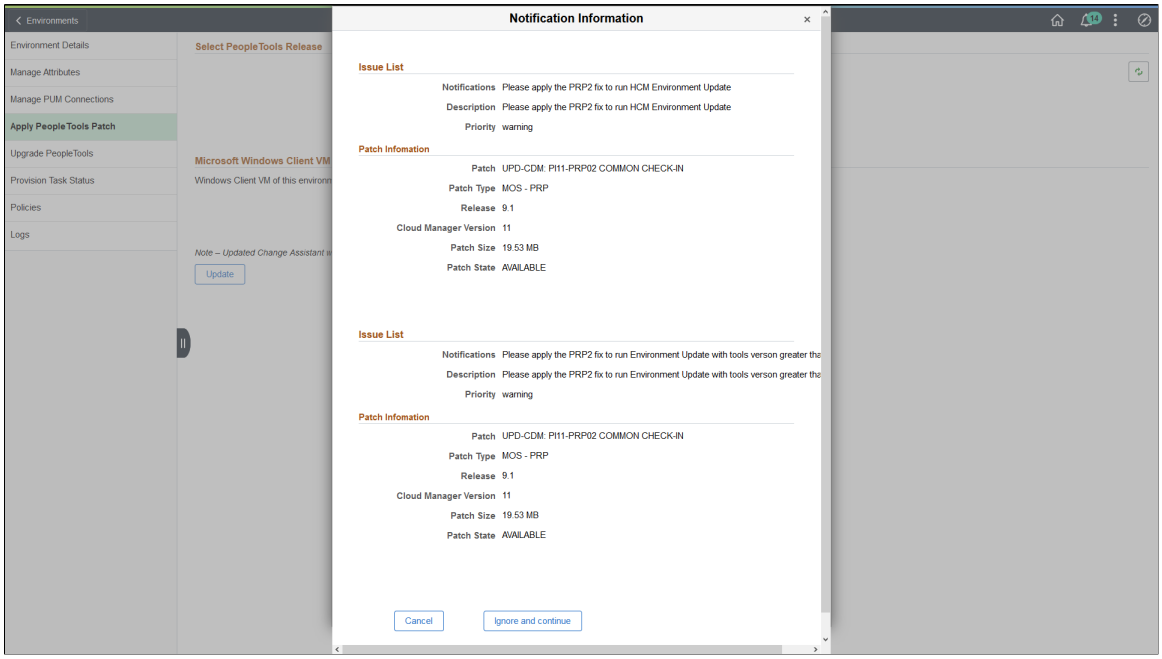

Field or Control	*Description*
Cancel	Select to cancel. Apply the patch using the Manage Updates page.
Ignore and continue	Select to ignore the notification and continue.

Chapter 9

Using the Lift and Shift Process to Migrate On-Premises Environments to Oracle Cloud

Understanding the Lift and Shift Process

The Lift and Shift process in Cloud Manager enables the automated migration of on-premises PeopleSoft environments to Oracle Cloud.

Important! The native Lift and Shift process using Oracle Database Cloud Backup Module (ODCBM) is not supported from Cloud Manager Image 20. It is recommended to use Zero Downtime Migration (ZDM) for real-time database upgrade without incurring any downtime. All major Oracle database types — including Autonomous Database, DBS, Bare Metal, Exadata, and ExaScale — fully support ZDM.

Migration to Cloud is achieved in two steps:

- Lift: Using the lift utility provided in Cloud Manager, PeopleSoft Application environment data (for example, PS_APP_HOME, PS_CUST_HOME) is packaged into DPK format. The PeopleSoft Oracle database is backed up using RMAN, and is uploaded to Oracle Object Storage.

  The Lift utility provided in Cloud Manager lifts the application tier (middle tier) and packages it into a DPK. The database tier is independently packaged into a separate DPK.

  Cloud Manager supports database lift using hot backup. Hot backup is performed with RMAN (Recovery Manager) using ODCBM that is bundled with the lift utility.

  Note: The DPKs that were lifted using older versions of Cloud Manager, may not be successfully shifted in later Cloud Manager versions.

- Shift: Cloud Manager downloads the lifted DPKs and RMAN backup. It then creates a new environment on Oracle Cloud. Once shifted, customers can use Cloud Manager to further manage, scale up or scale down, or clone these environments.

  Note: Before doing a Shift action, Lift and Shift topology must be updated with the right VM shape for each node.

Overview of the Lift and Shift process

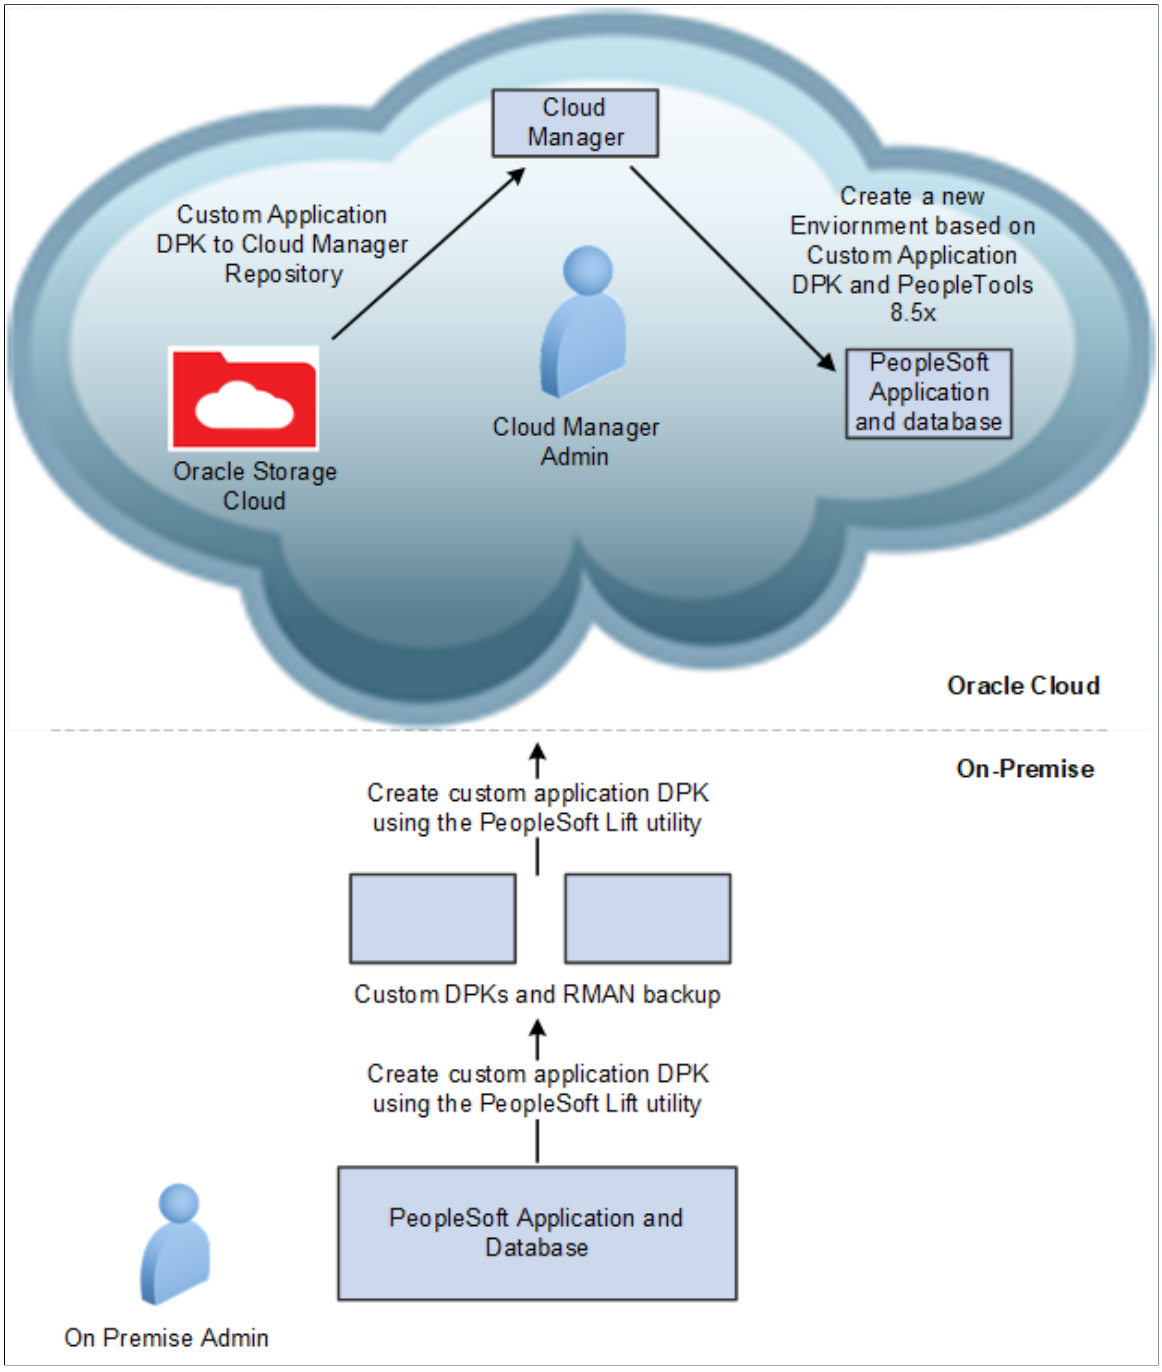

Customers will download the Lift utility from Cloud Manager and run it on an on-premises environment to create and upload customer application DPKs to the Oracle Cloud Service. Then using Cloud Manager, they use the customer application DPK to create a running application environment intact with all the customizations that have been done on-premises. It is a two-step process that simplifies days of laborious tasks. The Lift and Shift process is helpful to migrate many of your different environments such as demo, development, test, and training environments to the Oracle Cloud. Once an environment has been lifted, you can provision as many separate instances as you need.

To migrate a PeopleSoft environment from on-premises to Oracle Cloud using Cloud Manager, it must be running PeopleSoft application version 9.2 or above on a supported Linux operating system. The database

must be on Oracle 19c or later. Follow the support guidelines for PeopleSoft PeopleTools and PeopleSoft applications on My Oracle Support Certifications and on the PeopleSoft Cloud Manager Home Page, My Oracle Support, Doc ID 2231255.2.

Note: To shift an environment, you must log on to Cloud Manager with a User ID that includes the Cloud Manager PeopleSoft Administrator role (PACL_PAD). The default User ID CLADM includes this role.

Using the Lift Process to Migrate an Environment to the Oracle Cloud Infrastructure (OCI)

Pages Used to Migrate the Environment to Oracle Cloud

Page Name	*Definition Name*	*Usage*
Lift and Shift Tile	ECL_LAS_HOME_FL_GBL (CREF for the tile)	To access Lift and Shift landing page.
Lift and Shift Page	ECL_LAS_HOME_FL	The landing page containing the lift utility and the lifted containers.

Lift and Shift Tile

Use the Lift and Shift tile (ECL_LAS_HOME_FL_GBL) to access Lift and Shift landing page. The Lift and Shift tile is delivered as part of the Cloud Manager home page.

This example illustrates the Lift and Shift tile.

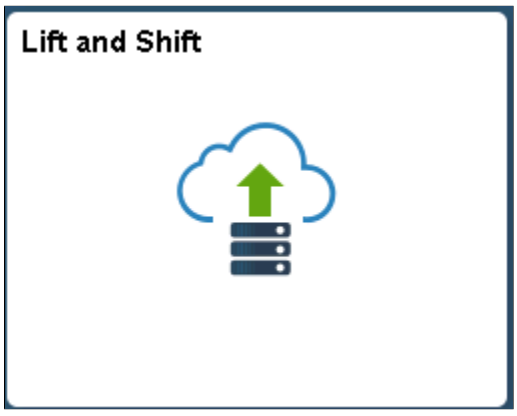

Lift and Shift Page

Use the Lift and Shift landing page (ECL_LAS_HOME_FL_GBL) to view and access the lifted environments (uploaded customer DPKs in Oracle Cloud for Cloud Manager).

Navigation:

Click the Lift and Shift tile on the delivered Cloud Manager Fluid Home page. The Lift and Shift page is displayed.

This example illustrates the fields and controls on the Lift and Shift page.

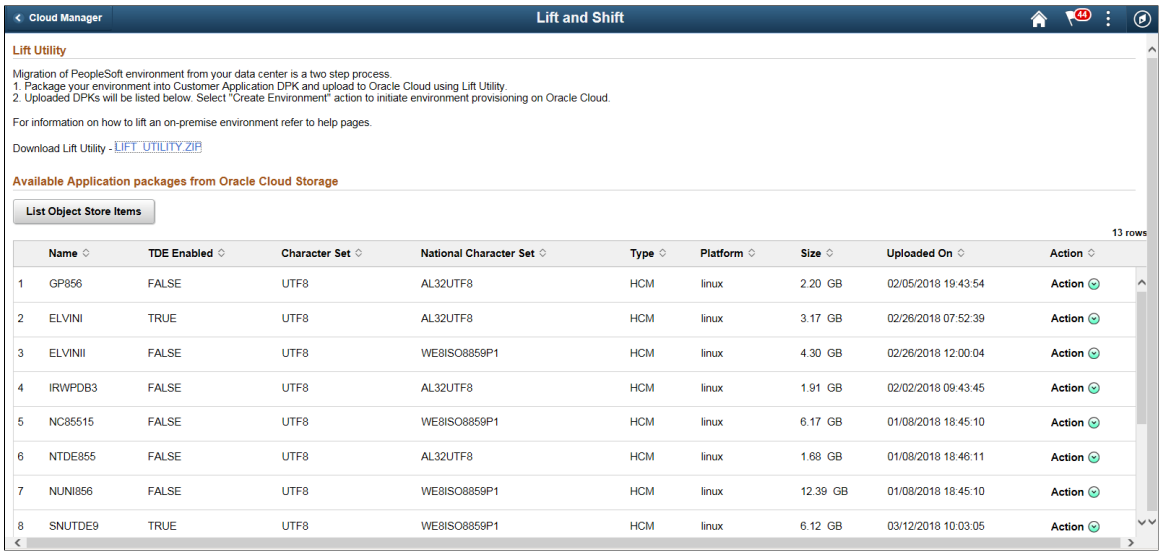

Note: Currently, in PeopleSoft Cloud Manager, an updated version of the Lift utility is available that captures more details from on premise environment. DPKs that were lifted earlier using Lift Utility from Cloud Manager Update Image 8 or older can no longer be deployed in Cloud Manager Update Image 10 and later. Hence, you must delete those old DPKs and do a lift operation again on the on-premises environments.

Field or Control	***Description***
Name	Name of the lifted environment.
TDE Enabled	Whether the database has encrypted tablespaces or not.
Character Set	The database character set used for lift operation.
National Character Set	Whether the database is unicode or non unicode. AL32UTF8 indicates unicode database and the value WE8ISO8859P1 indicates non unicode database.
Type	Shows the PeopleSoft application product pillar.
Platform	Indicates the Operating System platform.
Size	Total size of the lifted DPKs. **Note:** Assume that if the lifted DPK size is K, then the disk size should be 2.5 times K.

Field or Control	*Description*
Uploaded On	The date and time on which the DPKs were uploaded in Oracle Cloud.
Action	Use this button to perform a variety of related actions, such as viewing the details of each of the lifted DPKs, provisioning a new environment, and to delete a lifted DPK.
List Object Store Items	Click this button to refresh the lifted application list and make it current.

Downloading the Lift Utility

Navigate to the Lift and Shift page and click on the LIFT_UTILITY.ZIP. Please make sure you have applied any PRPs or fixes to Cloud Manager before downloading the utility. Copy the utility to the on-premise PeopleSoft system that needs to be migrated.

This example illustrates the link to the Lift utility.

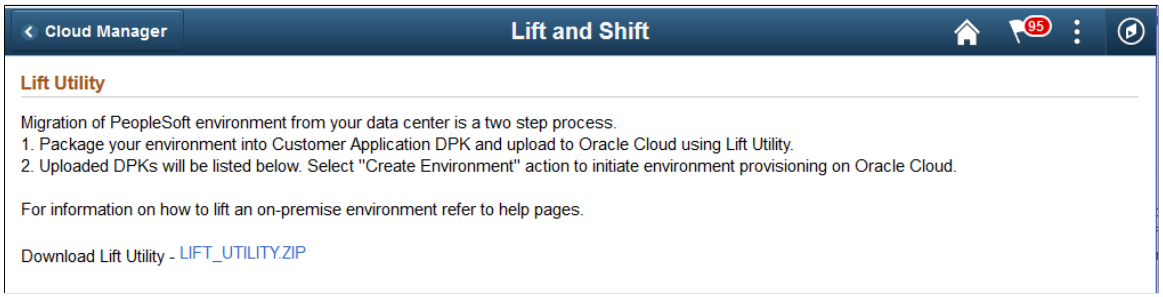

Installing Lift Prerequisites

Required Lift Prerequisites Applications

Lift prerequisite applications include:

- Python 3.6
- YUM Modules (gcc, libffi-devel, openssl-devel, zlib, wget)
- PIP Modules (oci-cli, pyyaml, xmltodict, requests, ensurepip)
- Java 1.8 (For database tier instance for RMAN)

Using the Automatic Lift Prerequisite Utility

The Lift prerequisite install feature will verify and install the applications required to perform Application Lift and Database (RMAN) Lift based on the user input.

This Lift prerequisite validation will be triggered in silent mode when the user triggers the Lift utility (psft-osl.sh) to validate the user environment before invoking the Lift to OCI.

Warning! Automatically installing the pre-requisites may update or overwrite any existing version on the system. Review the packages that will be installed automatically.

Requirements for using this utility:

- Ensure the necessary proxy and ports are set in order to access Internet, download and install the required prerequisite applications.
- Script must be triggered by the root user to install the applications.
- If the customer doesn't permit installation or in an event of failure during installation, the script will exit and the customer would need to install the prerequisites manually.

  See Manually Installing Lift Prerequisites

To automatically install the Lift prerequisite applications:

1. Download and copy the Lift Utility to the on-premises instance.
2. Log in as root to the on-premises instance.
3. Extract the Lift utility zip to a temporary folder and set full permissions to the folder.

   ```
   $ mkdir /tmp/CM9_LIFT
   $ unzip LIFT_UTILITY.zip -d <LIFT_UTILITY_PATH>
   $ chmod -R 777 <LIFT_UTILITY_PATH>
   ```

4. Navigate to the below path:

   ```
   $ cd <LIFT_UTILITY_PATH>/setup
   ```

5. Trigger the Lift Prerequisite install script.

   ```
   $ sh psft-lift-setup.sh
   ```

6. Generate the Oracle Cloud Infrastructure Auth Token. See Generating Oracle Cloud Infrastructure Auth Token.

Script Examples for Automatic Lift Prerequisite Utility

When you run the utility, it will ask if this is a database environment or an application environment.

```
Is <environment> a Database Environment: (yes/no):
```

- Application environment (answer no at above prompt)

  The system will scan the environment and indicate the applications that need to be installed.

This example illustrates applications that are required for an application lift.

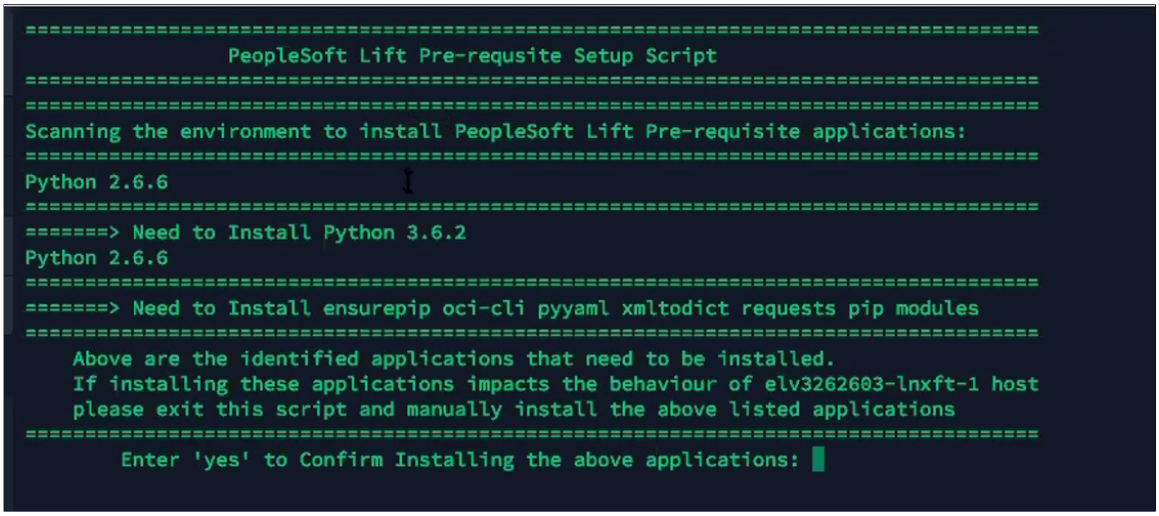

- Database environment (answer yes at above prompt)

  The system will scan the environment and indicate the applications that need to be installed.

  This example illustrates applications that are required for a database lift.

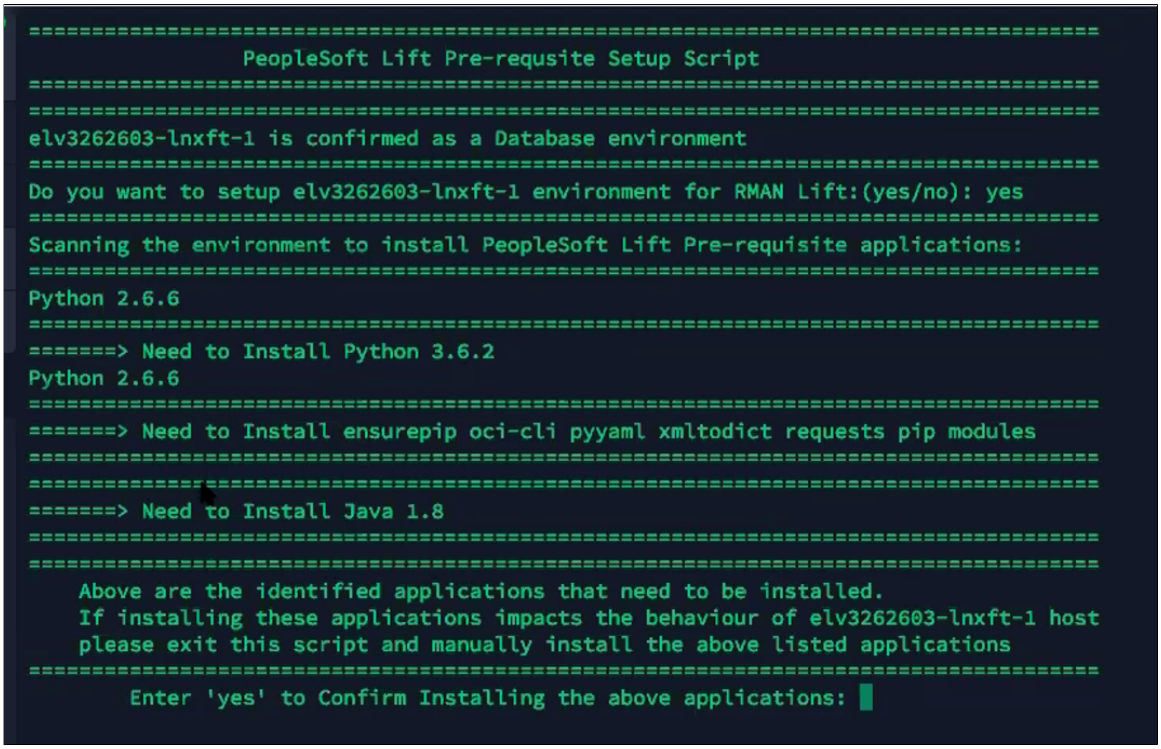

- To confirm installing the applications, enter yes at the prompt.

  ```
  Enter 'yes' to Confirm Installing the above applications:
  ```


Warning! If installing the applications impacts the behavior of your environment, then enter no. You will need to install the applications manually.

During installation of the prerequisite applications the utility will prompt for user confirmation, in order to continue installing certain dependencies. Not installing dependencies may result in Lift failure.

When all the applications are successfully installed, you will get a message “Complete!”.

This example illustrates all desired applications for PeopleSoft Lift are installed.

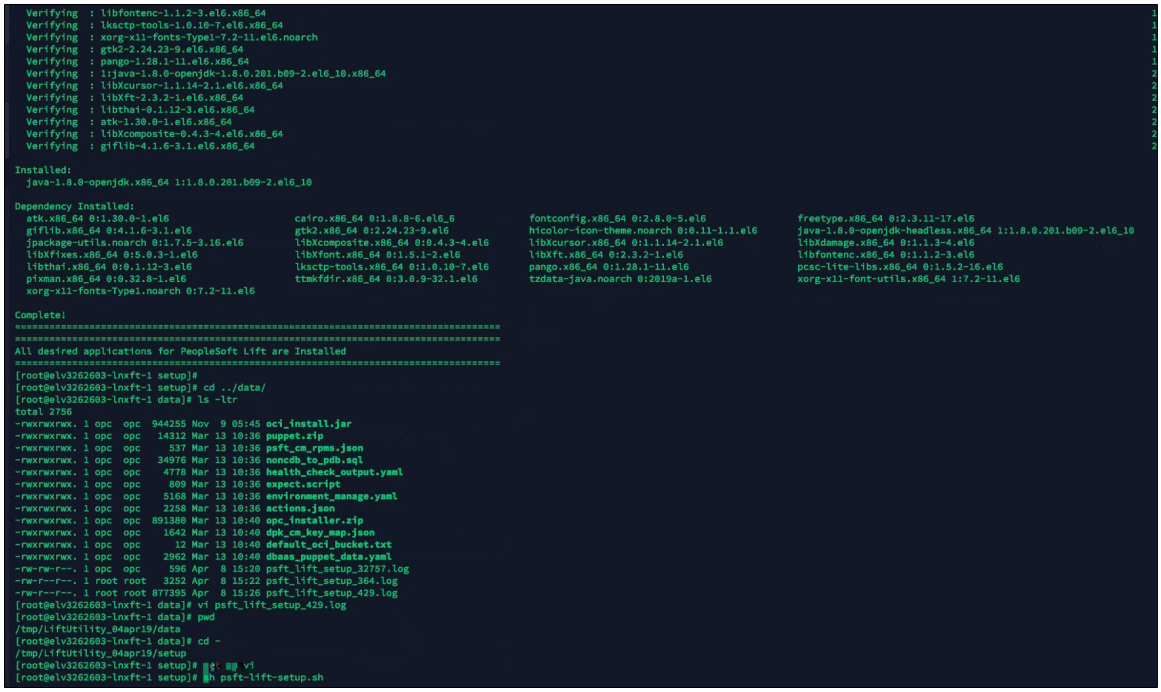

You can list the files and view the logs that were created for the install. The logs located at <Lift_Utility>/data/psft_lift_setup_<PID>.log

When you are ready to run the Lift Utility, it will verify that all the prerequisites have been installed.

Manually Installing Lift Prerequisites

Note: Use the manual method in case the automated method fails to install all prerequisites.

In the on-premises PeopleSoft instance, you must perform the following steps to manually install the lift prerequisites:

1. Extract the Lift Utility in a certain path on the respective application and database instance:

   ```
   mkdir -p <LIFT_UTILITY_PATH>
   unzip LIFT_UTILITY.zip -d <LIFT_UTILITY_PATH>
   ```

2. Install Python 3.6:

   a. Remove any old Python files present within the lift base directory by running the command.

   ```
   rm -rf <LIFT_UTILITY_PATH>/lnx_python
   ```



```
mkdir -p <LIFT_UTILITY_PATH>/lnx_python
```

b. Install the prerequisites by running the following commands:

```
sudo yum install gcc
sudo yum install libffi-devel
sudo yum install openssl-devel
sudo yum install zlib
sudo yum install wget
```

c. Download Python 3.6.2 by running the following commands:

```
cd <LIFT_UTILITY_PATH>
wget https://www.python.org/ftp/python/3.6.2/Python-3.6.2.tgz
tar xzf Python-3.6.2.tgz
cd Python-3.6.2.tgz
```

d. Configure and compile the source by running this command:

```
./configure --prefix=<LIFT_UTILITY_PATH>/lnx_python
make altinstall
```

e. Create a soft link for the Python executable by running the following commands:

```
cd <LIFT_UTILITY_PATH>/lnx_python
ln -s bin/python3.6 python
```

f. Set environment variables. Do the following:

```
export PYTHON_HOME=<LIFT_UTILITY_PATH>/lnx_python
export PYTHONPATH=<LIFT_UTILITY_PATH>/lnx_python
export PATH=<LIFT_UTILITY_PATH>/lnx_python/bin:<LIFT_UTILITY_PATH>/lnx_py⇒

thon/:$PATH
export LANG=en_US.utf-8
export LC_ALL=en_US.utf-8
```

g. Install PIP with this command:

```
<LIFT_UTILITY_PATH>/lnx_python/python -m ensurepip
```

3. Install the below PIP packages.

a. Install the oci-cli package with this command:

```
pip install oci-cli
```

b. Install PYYAML with this command:

```
pip install pyyaml
```

c. Install XMLTODICT with this command:

```
pip install xmltodict
```

4. Install Java version 1.8 (JRE) using Java Official Documentation.

Verify Java is Installed by running the below commands.

```
$ java -version
    java version "1.8.0_144"
    Java(TM) SE Runtime Environment (build 1.8.0_144-b01)
    Java HotSpot(TM) 64-Bit Server VM (build 25.144-b01, mixed mode)
```



```
$ echo $JAVA_HOME
    /usr/lib/jvm/java-1.8.0-openjdk/jre
```

5. Generate the Oracle Cloud Infrastructure Auth Token. See Generating Oracle Cloud Infrastructure Auth Token.

Performing Application Lift

Application Lift means lifting the middle tier which consists of the Application Server, Web Server and Process Scheduler of PeopleSoft Application.

1. Use only the PeopleSoft Admin user (for example, psadm2) to perform Application Lift.
2. Make sure to have sufficient free disk space for Application Lift (based on PS_APP_HOME and PS_CUST_HOME size). A minimum disk space of 10GB is required.
3. Ensure PS_APP_HOME and PS_CUST_HOME directories are available.
4. Ensure that the user running the Lift utility has the permission to create files or directories at the user's home directory, Lift utility directories, and the destination directory where the DPKs are saved, /tmp, PS_APP_HOME, and PS_CUST_HOME directories.

Note: Installing OCI-CLI is a prerequisite for the lift utility. See Installing Lift Prerequisites

---

To perform the one-step Lift automation procedure for the application:

1. Download the Lift utility from the Lift and Shift page. For this, perform the following:
   a. Navigate to the Lift and Shift tile.
   b. Copy the “LIFT_UTILITY.zip” utility to the target machine to perform lift.

   Note: If you have recently updated Cloud Manager with any PRPs that has fixes to the lift utility, then SSH to the Cloud Manager instance and delete the stale zip file from /tmp/LIFT_UTILITY.ZIP.

   ---

2. Navigate to the below folder after extracting the LIFT_UTILITY.zip and set permissions:

```
chmod -R 777 <LIFT_UTILITY_PATH>
cd <LIFT_UTILITY_PATH>/setup
```

3. For Linux, run the **sh psft-osl.sh** command to perform lift.
4. Select 1 at the prompt to select the type of environment to lift.

This example illustrates the prompt to lift the Application Environment.

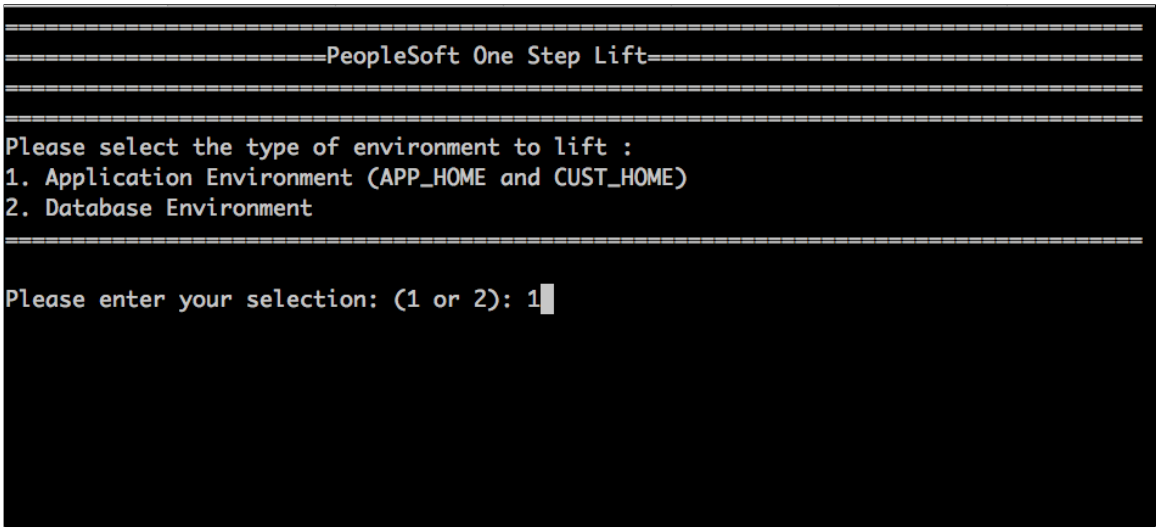

5. To create the PeopleSoft App Server DPK, you need to provide the database name (or PDB name in case of supported Oracle multitenant databases) and destination directory.

   Note: If the utility is unable to fetch the data from the environment (for example, app_type/ oracle_home), it will prompt the user to input the same.

6. Choose any one of the below options:

   1. Create and Save DPK in APP/DB Environment.

   2. Create, Save DPK in APP/DB Environment and Upload the DPK to Oracle Cloud Infrastructure (Object Storage).

7. If you select option 1. Create and Save DPK in APP/DB Environment, you will need to manually upload the DPK to Object Storage. See Uploading the DPK Manually to Oracle Cloud Infrastructure.

8. If option 2 to upload the DPK to Oracle Object Storage is selected, then the script prompts the user to input the Oracle Cloud tenancy credentials as mentioned below in order to upload the DPK once created:

   See Locating OCI Credentials

   – Oracle Cloud Infrastructure Region Name

   – Oracle Cloud Infrastructure Tenancy Name

   – Oracle Cloud Infrastructure Tenancy ID

   – Oracle Cloud Infrastructure User ID

   – Private Key Location, indicates the API signing private key that was created during CM configuration and must be copied to the instance where lift utility will be run. This input refers to the full path to the file.

   – Passphrase, refers to the passphrase that was used to encrypt the API signing keys.

Note: You need to manually copy the key file or copy the key file contents and save locally in the machine where you perform the lift. This is the corresponding Private Key to the Public Key that was set in the API Keys of the user setting.

9. After Application Lift is complete, the following Application DPK will be created based on the PeopleTools version on the application instance.

   APP-DPK-<platform>-<app_type>-<db_name>-1of3.zip.

   Note: The APP-DPK*-3of3.zip will not be created as part of the Lift utility, however the APP-DPK*-3of3.zip DPK will be available from the PeopleTools DPK when the shift is triggered from Cloud Manager.

The Lifted DPKs created are available in the destination directory. If you had chosen to create and upload DPK to Oracle Object Storage, then the uploaded DPKs are available in Oracle Object Storage and listed in the Lift and Shift page of Cloud Manager.

Performing the Database Lift

The Database lift means lifting the database of PeopleSoft instance as a hot backup using RMAN.

Note: Database Lift using hot backup can only be performed on a database instance that has access to the internet.
It is recommended to bring the database patch level of the on-premises environment equivalent to that of the database patch level of the Oracle Database Cloud Service before starting the Lift and Shift process. If the patch levels are different, then Cloud Manager will try to either rollback or update the patch. It is possible that there could be some incompatibilities during lift and shift due to rollbacks or updates. Users will then need to manually verify and rectify it.

Installing OCI-CLI is a prerequisite for the lift utility. See Installing Lift Prerequisites

Considerations Before Running Database Lift

1. Lift can be performed on the DB instance (Local Lift) only.
2. Ensure to use only the Database owner user (for example, oracle) to perform DB Lift.
3. Ensure to have sufficient free disk space for DB Lift based on DB size.
4. Oracle Database 19c and later is supported.
5. Ensure to take the back up of your Database environment and the RMAN configurations before performing DB Lift. Optionally, it is recommended to use a clone of the environment for the Lift operation.

   Note: During the Lift process, the Oracle Database will not be shut down.

6. Ensure to back up the ORACLE_HOME.

7. Ensure that the user running the lift utility has permission to create files/directories at the user's home directory, Lift utility destination directory where the DPKs are saved, /tmp, and ORACLE_HOME directory.

Note: If you want to encrypt the database before lifting using TDE, see Encrypting Tablespaces Using Transparent Data Encryption.

Using RMAN for Hot Backup Database Lift

RMAN Lift and Shift supports the following Oracle Databases:

- Oracle 19c and later
- Container Databases
- Unicode and Non-Unicode Databases
- TDE enabled Database
- Database on ASM

Note: During Lift and Shift with RMAN the Oracle Database Cloud Backup Module (ODCBM) is used in the background to perform full RMAN Backup and Restore operations.

During Database Lift with RMAN the Lift Utility will create a fresh Bucket in the OCI Object storage and then the RMAN Backup of the source (on-premises) Database environment will be compressed and encrypted before being wired to the bucket in the OCI Object Storage. Along with the RMAN backup the "APP-DPK-<platform>-<app_type>-<db_name>-2of<X>.zip" is created to capture the sqlpatches, database parameter file and other metadata information.

The "APP-DPK-<platform>-<app_type>-<db_name>-2of<X>.zip" will be small, since we are not packaging the Database files (*.dbf) within this zip.

RMAN Lift and Shift does not support:

- RAC Database
- RMAN (L1) Incremental Backup.

Note: It is recommended that the DB Tier Database is started with a spfile.
Ensure the Database “Archive Log Mode” is enabled (My Oracle Support Doc ID 371139.1).
Ensure the proxy (if needed) is correctly specified and the proxy authentication does not have any special characters.
To perform an RMAN Lift of a PeopleSoft environment that was already lifted using the older version of Cloud Manager (Lift Utility), it is recommended to delete the existing Lifted DPK from the Cloud Manager “Lift and Shift” UI and then trigger a RMAN lift for the same.

Running Lift Using Hot Backup (RMAN)

To perform the Lift using Hot Backup:

1. Navigate to the Lift and Shift tile.

2. Download and copy the “LIFT_UTILITY.zip” to the target machine to perform lift.

Note: If you have updated Cloud Manager with PRP (find PRP name/number), then SSH to Cloud Manager VM and delete the stale zip file from /tmp/LIFT_UTILITY.zip.

---

3. Navigate to the below folder after extracting the LIFT_UTILITY.zip and set permissions:

```
chmod -R 777 <LIFT_UTILITY_PATH>
cd <LIFT_UTILITY_PATH>/setup
```

4. Export Java home.

```
export JAVA_HOME=/tmp/java18/jre
```

5. For Linux, run the **sh psft-osl.sh** command to perform lift.

6. At the prompt “Do you want to Lift the Application Environment”, enter “N” to Lift the DB instance.

7. At the prompt input the following:

   – Container Database name (applicable only for Multitenant database)

   – Database name

8. Performing Lift using Hot Backup requires the OCI details:

   See Locating OCI Credentials

   – Oracle Cloud Infrastructure Region Name

   – Oracle Cloud Infrastructure Tenancy Name

   – Oracle Cloud Infrastructure Tenancy ID

   – Oracle Cloud Infrastructure User ID

   – Private Key Location, indicates the API signing private key that was created during CM configuration and must be copied to the instance where lift utility will be run. This input refers to the full path to the file.

   – Passphrase, refers to the passphrase that was used to encrypt the keys.

Note: You need to manually copy the key file or copy the key file contents and save locally in machine where you perform a lift. This is the corresponding Private Key to the Public Key that was set in the API Keys of the user setting.

   – OCI User name

   – OCI Auth Token

Note: Ensure not to delete this OCI Auth Token (for the user) from OCI console, because the User name/Token will be used for Lift and Shift process. The OCI Auth Token is also required for a refresh from on-premises to Cloud.

   – OCI Infrastructure Compartment ID (OCID)

9. To create the PeopleSoft Database DPK, you need to provide the following:

   – Container Database name (Applicable only for supported versions of Oracle multitenant database)

   – Database name

   – Number of channels (threads) for RMAN backup (Value between 1–8)

   – TDE KeyStore (Wallet) Password (If TDE is enabled on source Database)

   – RMAN Backup Encryption Password (If TDE is not enabled on source Database)

   – DB environment Proxy Host (* If Any)

   – DB environment Proxy Port (* If Any)

   – Destination Directory

   Note: If the utility is unable to fetch the data from the environment, it will prompt the user to input the same.

10. The script then displays the details captured from the user and prompts for the user's confirmation to proceed. The utility allows the user to modify the above listed inputs, if required.

11. The entire process is logged into psft_lift_session_<session_name>_<session_count>_<PID>.log file.

12. After DB Lift is complete, the following DB DPK will be created.

    APP-DPK-<platform>-<app_type>-<db_name>-2of3.zip

    Note: The APP-DPK*-3of3.zip will not be created as part of the Lift utility, however the APP-DPK*-3of3.zip DPK will be available from the PeopleTools DPK when the shift is triggered from Cloud Manager.

Uploading the DPK Manually to Oracle Cloud Infrastructure

During the process to upload the lifted APP DPKs to OCI Object Storage, if you chose to only create and save the DPK in the APP environment, you will need to manually upload the DPKs.

To manually upload the DPKs to OCI Object Storage:

1. Set the following environment variables:

   - export PYTHON_HOME=<*LIFT_UTILITY_PATH*>/lnx_python.
   - export PYTHONPATH=<*LIFT_UTILITY_PATH*>/lnx_python.
   - export PATH=$PATH:<*LIFT_UTILITY_PATH*>/lnx_python/bin.

2. Create an OCI_Config file with the below contents:

   a. [DEFAULT]

b. user=<*user OCID*>

c. fingerprint=<*Finger print*>

d. key_file=<*private key file location*>

Note: You can use the same API Signing Key pair that was created when setting up Cloud Manager, or you can create a new one. If you create a new pair then, you must add the newly created public API key under the User Settings using OCI UI.

---

e. pass_phrase=<*Passphrase for the private key*>

f. tenancy=<*tenancy OCID*>

g. region=<*region name*>

For example:

```
[DEFAULT]
 user=ocid1.user.oc1..aaaaaexxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
 fingerprint=36:d6:c7:9b:d4:21:d7:ad:10:70:4f:58:b7:70:0f:fb
 key_file=/tmp/key.pem
 pass_phrase=XXXXXXXXXX
 tenancy=ocid1.tenancy.oc1..aaaaaxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
 region=us-ashburn-1
```

3. If you are uploading for the first time, create the container **psft_oci_las** with the following command:

```
<LIFT_UTILITY_PATH>/lnx_python/bin/oci --config-file /tmp/oci_config os bucket⇒
 create -ns <tenancy name> --name psft_oci_las --compartment-id <Compartment I⇒
D>.
```

For example,

```
<LIFT_UTILITY_PATH>/lnx_python/bin/oci --config-file /tmp/oci_config os bucket⇒
 create -ns mycloud --name psft_oci_las --compartment-id ocid1.compartment.oc1⇒
..aaaaaaxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
```

4. Run the following command to upload the APP DPK. Replace the variables in the command with the actual file and path names:

Note: The <Bucket Name > should be *psft_oci_las*. Do not specify any other bucket name.

```
<LIFT_UTILITY_PATH>/lnx_python/python upload_dpk_to_oci.py -d <Tenancy Name> -⇒
c psft_oci_las -s <Source folder containing DPK file> -t <Target Folder Name>-⇒
f <INI file location generated during lift operation> -g <Full path of oci con⇒
fig file>
```

Variable	***Description***
-d	Tenancy name to which the DPKs will be uploaded to.

Variable	*Description*
-c	psft_oci_las —The container to which the DPKs will be uploaded. This value should not be changed.
-s	Source folder where the DPK files are saved during lift.
-t	Target folder Name on OCI. Should be Platform/AppType/DBName where AppType is application type [HCM,FSCM, ELS, ELM, CRM] and DBName is the name of the database. For example: linux/HCM/MYHCMDB. **Important!** Ensure that the target folder name is as shown in the example above. There must be no preceding or trailing '/' in the target folder path.
-f	INI file location that was generated during lift operation.
-g	Path to OCI config file that will be used to connect to OCI to upload DPKs.

Locating OCI Credentials

When performing the Lift process, you will be prompted for OCI credentials.

Locating Oracle Cloud Infrastructure Tenancy and Region Name

From the OCI home page, use the navigation menu in the upper left to navigate to your cloud resources.

To view the tenancy details, open the Profile Menu and click Tenancy:<your_tenancy_name>.

Field or Control	*Description*
	Profile menu

This example illustrates the fields and controls on the Tenancy and Region.

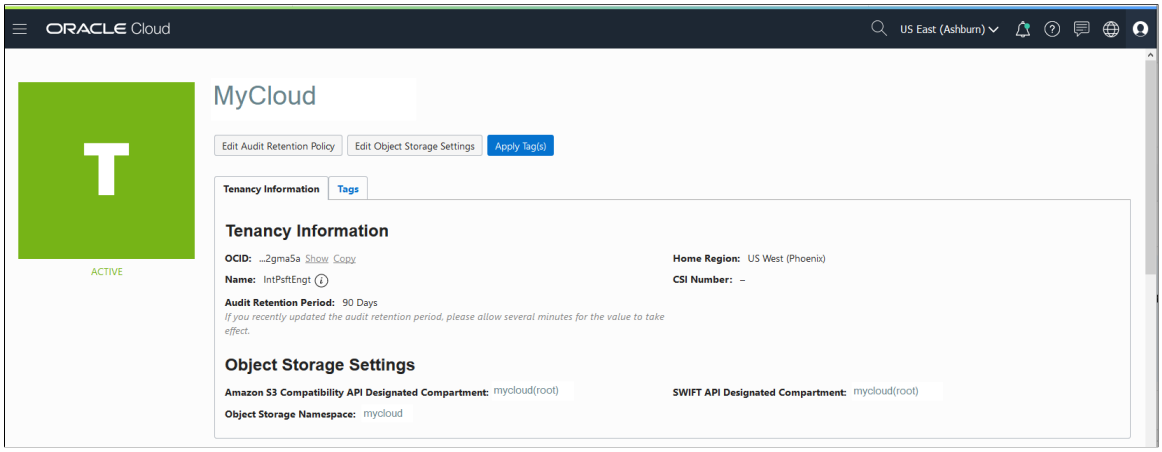

You can use the Copy link to copy the tenancy OCID.

Locating Oracle Cloud Infrastructure User ID

To view the user details open the Profile Menu and click User Settings.

This example illustrates the fields and controls on the User Details page.

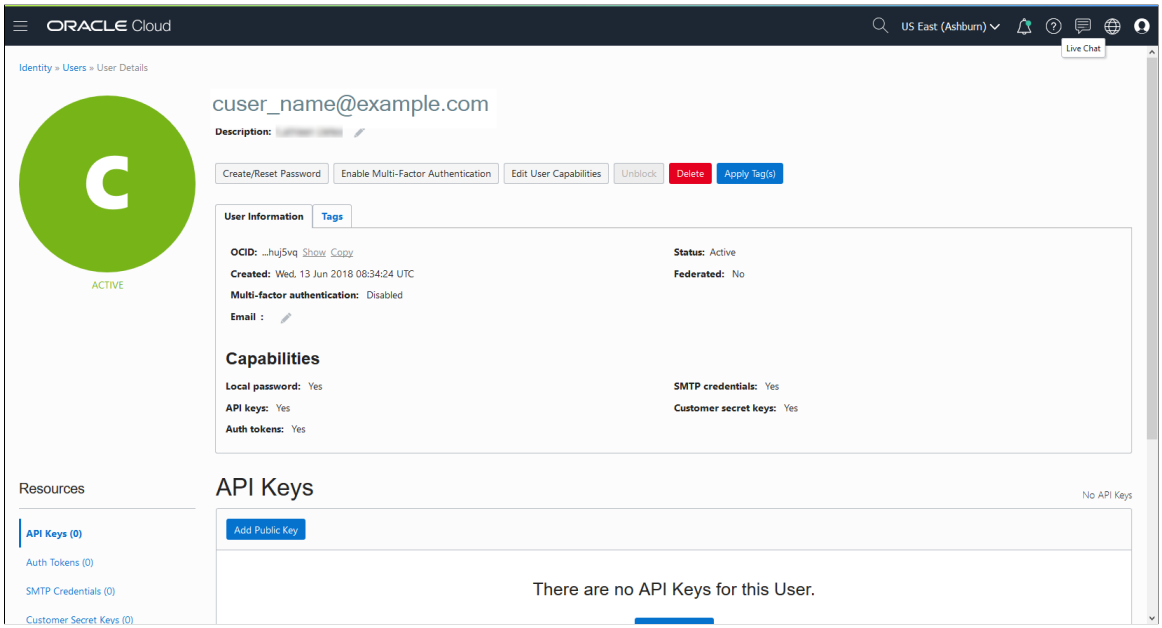

You can use the Copy link to copy the user OCID.

Locating Oracle Cloud Infrastructure Fingerprint

To locate the OCI fingerprint:

1. Open the Profile Menu and click User Settings.

2. Select API keys from the menu displayed on the left.

This example illustrates the API Keys page, which contains the fingerprint.

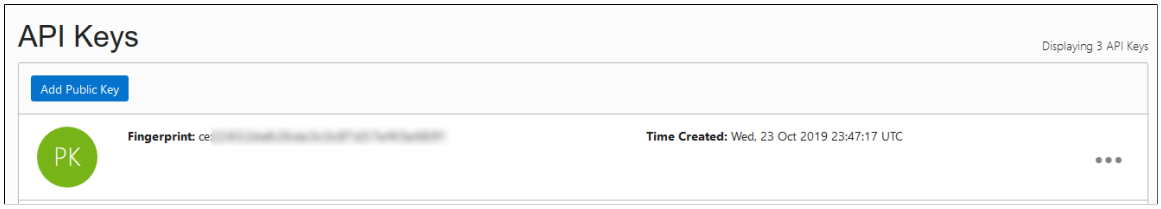

Generating Oracle Cloud Infrastructure Auth Token

To generate the OCI Auth Token:

1. Click Profile Menu in the top-right corner of the Console, and then click User Settings.
2. Select Auth Tokens from the menu displayed on the left.
3. Click on "Generate Token" and save the Auth Token displayed.

Important! Be sure to note or save the Auth Token immediately as you will not be able to retrieve the Auth Token once the page is closed.

This Auth Token will be needed for performing Lift and Shift.

Note: Refrain from deleting the Auth Token after performing the Lift as this new token will be used to perform Shift. If Auth Token is deleted after performing a Lift, you will need to generate a fresh Auth Token and perform a fresh Lift.

This example illustrates Auth Tokens page in OCI console.

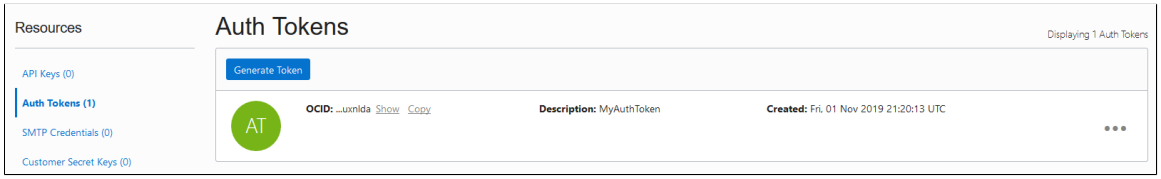

Deleting Oracle Cloud Infrastructure Bucket and Objects

To delete the OCI Bucket Objects and the Bucket for RMAN Lift:

1. Delete all OCI Bucket Objects (This does not prompt for confirmation (--force))

```
oci os object bulk-delete -ns <Tenancy_Name> -bn <Bucket_Name> --config-file <⇒
oci_config> --force
```

2. Delete all OCI Bucket

```
oci os bucket delete -ns <Tenancy_Name> --name <Bucket_Name>  --config-file <o⇒
```



```
ci_config>
```

Example:

```
Tenancy_Name = Intxxxxx

Bucket_Name = PSPDB_1348135976_25092018

oci_config = <oci_config_filepath>

oci_config_filepath:

 [DEFAULT]
 user=ocid1.user.oc1..aaaaaxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
 fingerprint=36:d6:c7:9b:d4:21:d7:ad:10:70:4f:58:b7:70:0f:fb
 key_file=/tmp/key.pem
 pass_phrase=XXXXXXXXXX
 tenancy=ocid1.tenancy.oc1..aaaaaxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
 region=us-ashburn-1
 oci_username=xxxxx.xxxx@oracle.com
 oci_token=YYYYYYYYYYYYY
 oci_tenancy_name=intxxxxx
```

Using the Shift Process to Provision the Migrated Environment on the Oracle Cloud

Use the Shift process to deploy the packaged environment in Oracle Cloud.

Prerequisites

- The Lift and Shift topology must be modified with the required size and disk capacity of the database and middle-tier nodes. If shifting to DBaaS, then modify the Lift and Shift - DBaaS topology.

  Note: The disk space of the database node must be configured based on the size of the lifted database. The recommended disk space on the database node is at least 2.5 times the lifted database size.

- During the Shift process, Cloud Manager can update the PeopleTools patch of the lifted environment. To update the PeopleTools patch during shift, make sure to have the required PeopleTools DPK already downloaded and available in the repository.
- The Shift process makes use of the latest PeopleSoft Update Image for the application type. For example, if your lifted environment is an HCM environment, then make sure you have the latest HCM PeopleSoft Update Image downloaded in the repository.
- Before shifting, the Lift and Shift related topologies must be edited and saved to add shape name and disk capacity where applicable.

  See Editing an Existing Topology

  Lift And Shift topology is supported only when the database version associated with the lifted environment is Oracle 19c and the database is non-TDE. For all database versions other than Oracle 19c, the Lift and Shift - DBaaS topology is always used for shift.

Note: Verify the Lift and Shift topology; be sure to select the right topology based on the choice of database to be created on DBaaS. You also need to verify the sizing and disk space based on the lifted DPK size and desired environment; a minimum allocation should be provided. For database node, you need to provide a size that is equivalent to 2.5 times of the actual lifted DPK size (not zipped).

- For Exadata database:

  You need to add the SSH public key for the Cloud Manager user on all Exadata cluster VMs. See Adding SSH Keys to a VM Cluster in the Oracle Cloud Infrastructure documentation.

  On the Infrastructure Setting page, click the Refresh OCI Metadata button to sync the Exadata DB Systems provisioned in OCI. After refreshing the metadata, the instance will appear in the DB Systems section on the provisioning page.

- The DB Admin password and the Wallet password are the same. If you want to change the Wallet password you must do that manually.

- The database operator IDs used during the Shift operation should have specific permissions to perform various actions. The permissions are listed below:

  - For ACM (Automated Configuration Manager) — ACM administrator
  - For IB (Integration Broker) — Integration administrator
  - For Elasticsearch or OpenSearch — Search Administrator, Search Server, Search Developer
  - For Process Scheduler — PeopleSoft Administrator, ProcessSchedulerAdmin, ReportDistAdmin
  - For Portal — PeopleTools, Portal Administrator

- Supported databases are listed in the Support Matrix for Shift Provisioning on Target Database posted on the PeopleSoft Cloud Manager Homepage.

- You must log on to Cloud Manager with a User ID, such as the default User ID CLADM, which includes the Cloud Manager PeopleSoft Administrator role (PACL_PAD).

- For Government Cloud:

  FIPS (Federal Information Processing Standards) is disabled during Shift. After the Shift process completes you can enable FIPS by following these steps:

  1. Log in to the DBS instance.
  2. Change user to *oracle*.
  3. Open file "$ORACLE_HOME/ldap/admin/fips.ora".
  4. Change SSLFIPS_140=TRUE.

     SSLFIPS_140 value was set to FALSE during Shift.

- Install JDK 1.8 on Cloud Manager.

Some DB Systems include JDK 11 by default, whereas the Shift process requires JDK 1.8. Use these steps to download and install JDK 1.8, before performing Shift operations. This procedure is needed once for each Cloud Manager instance.

1. Download JDK 1.8 from this Oracle site: https://www.oracle.com/java/technologies/javase/javase8-archive-downloads.html

   The file name to download is jdk-8u202-linux-i586.tar.gz.

2. Log in to the Cloud Manager instance and save the JDK file in /tmp.

3. Extract the .tar.gz file using the following command:

   ```
   sudo tar xvzf /tmp/jdk-8u202-linux-i586.tar.gz -C /cm_psft_dpks/
   ```

   The command copies the extracted directory to /cm_psft_dpks, which is the mounted directory for the Cloud Manager File Storage Service (FSS).

4. Rename the extracted directory to jdk1.8.

   ```
   sudo mv jdk1.8.0_202 jdk1.8
   ```

5. To verify the path is correct, run the command:

   ```
   /cm_psft_dpks/jdk1.8/bin/java -version
   ```

   You should see a message such as:

   ```
   java version "1.8.0_202"
   Java(TM) SE Runtime Environment (build 1.8.0_202-b08)
   Java HotSpot(TM) Server VM (build 25.202-b08, mixed mode)
   ```

Pages Used to Provision the Migrated Environment on the Oracle Cloud

Page Name	*Definition Name*	*Usage*
Lift and Shift – Create Environment Wizard	ECL_LAS_GENERAL_FL	Use the Lift and Shift – Create Environment wizard to perform shift operation by means of a guided process.
Lift and Shift – Advanced Options Page	ECL_LAS_ADV_FL	Use Lift and Shift – Advanced Options page for defining target database details.
Lift and Shift – Custom Attributes Page	ECL_LAS_CUSTATR_FL	Use Lift and Shift – Custom Attributes page for defining the custom attributes as per the lifted environment.
Lift and Shift – Review and Submit Page	ECL_LAS_REVIE_FL	Use Lift and Shift – Review and Submit page to review and submit the entered environment details.

Lift and Shift Page

Once an environment is lifted, it will be available on the Lift and Shift page. Click the List Object Store Items button to view all items.

This example illustrates the fields and controls on the Lift and Shift page. You can find definitions for the fields and controls later on this page.

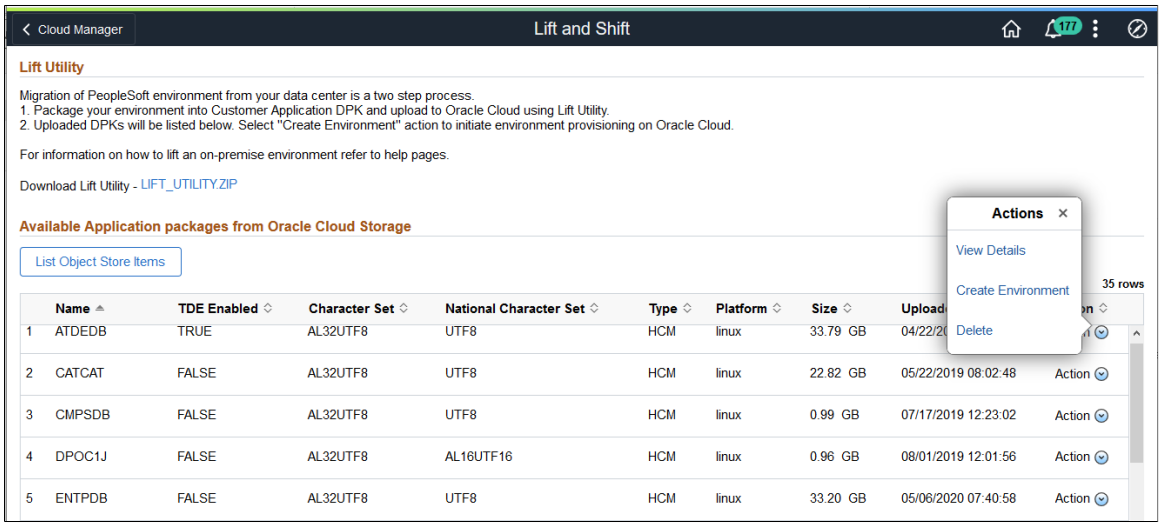

For description of this page see Lift and Shift Page

Lift and Shift actions:

Field or Control	***Description***
View Details	Click to verify the lift details. The type of the DB zip file (2of3.zip or 2of2.zip file) will be *db_rman*.
Create Environment	Click to Shift the environment. See Lift and Shift – Create Environment Wizard.
Delete	Click to delete.

This example illustrates the fields and controls on the View Details page, showing the type as db_rman for the database zip file.

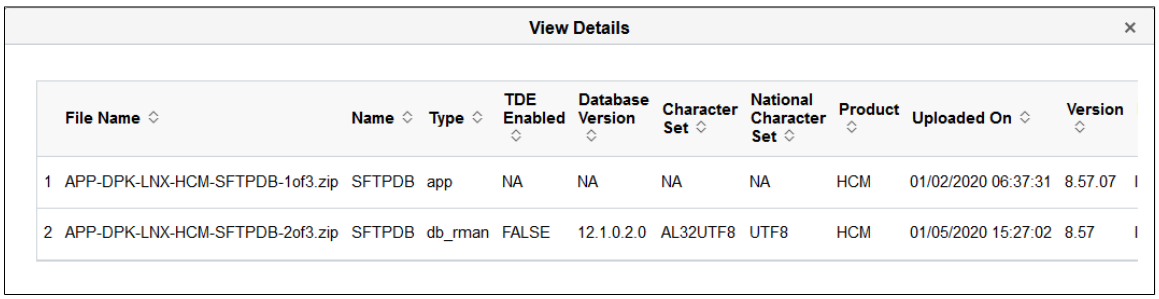

Lift and Shift – Create Environment Wizard

Use the Lift and Shift – Create Environment wizard (ECL_LAS_GENERAL_FL) to perform the Shift operation. The Shift operation facilitates provisioning a new environment using the lifted DPKs.

In Lift and Shift provisioning, you can:

- Select the desired topology based on DB on Oracle Cloud (Compute or DBaaS).
- Modify the sizing and disk space.

Note: The Database type must be set to DEMO on the Lift and Shift template. This field does not appear in the lift and shift provision pages.

Navigation:

Click the Related Action button corresponding to the lifted application. Select Create Environment option. By default, the Lift and Shift - General (New Environment Information) page is displayed.

This example illustrates the fields and controls on the Lift and Shift - New Environment Information page.

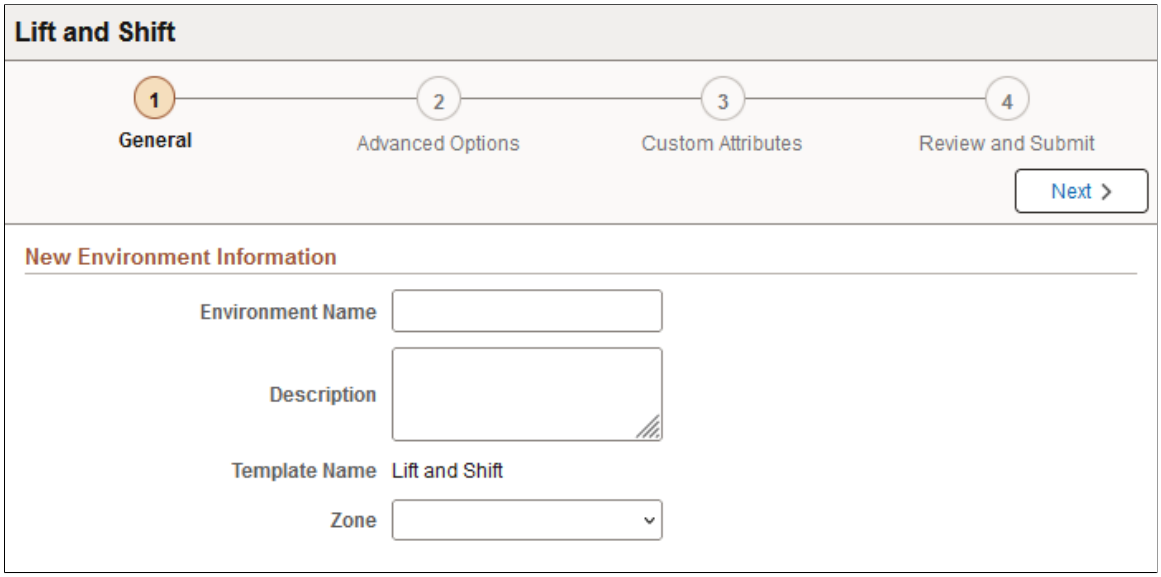

Field or Control	*Description*
Environment Name	Enter the name of the environment which you want to create. **Note:** Length of environment name and identity domain name should not exceed 25 characters.
Description	Enter a description for the environment.
Template Name	Displays the default template to be attached with the environment.

Field or Control	*Description*
Zone	Select the zone on which the environment should be created.

Lift and Shift – Advanced Options Page

Use the Lift and Shift – Advanced Options page (ECL_LAS_ADV_FL) for defining target database details.

Navigation:

Click step 2 or Next at the top of the Lift and Shift guided process.

This example illustrates the fields and controls on the Lift and Shift – Advanced Options page. You can find definitions for the fields and controls later on this page.

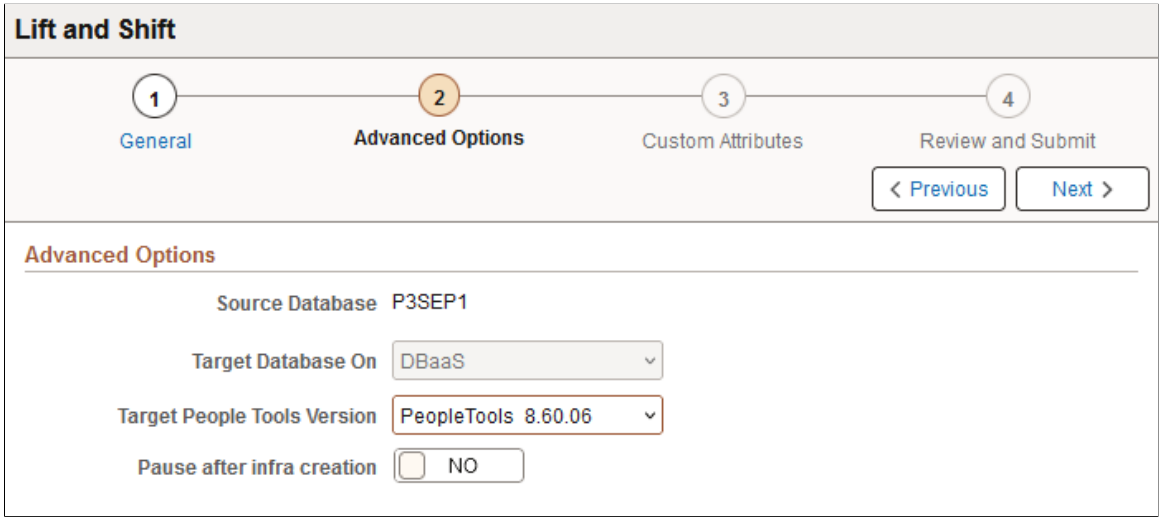

Field or Control	*Description*
Target Database On	The Target Database On is DBaaS for all environments. **Note:** For TDE enabled environments, Target Database On is DBaaS, which cannot be changed.
Target PeopleTools Version	Select the PeopleTools version to be applied on the environment.

Field or Control	*Description*
Pause after infra creation	Select Yes for the environment provisioning to pause after completion of the Infrastructure task. This provides the user the opportunity to do additional setup, actions, or operations on the newly created environment outside of Cloud Manager before proceeding with the PeopleSoft deployment. **Note:** When you are ready to proceed to the PeopleSoft deployment, select Deploy from the related actions on the Environment tile. Select No (default) to continue provisioning the environment when the infrastructure layer is complete.

Lift and Shift – Custom Attributes Page

Use Lift and Shift – Custom Attributes page for defining the custom attributes as per the lifted environment.

Navigation:

Click step 3 or Next at the top of the Lift and Shift guided process.

Enter the custom attributes as per the lifted on-premises environment. It is recommended that the custom attribute values entered on this page match the on-premises configuration.

For details on custom attributes, see Managing Environment Attributes.

This example illustrates the fields and controls on the Lift and Shift – Region and Availability Domains and the Middle Tier Network sections.

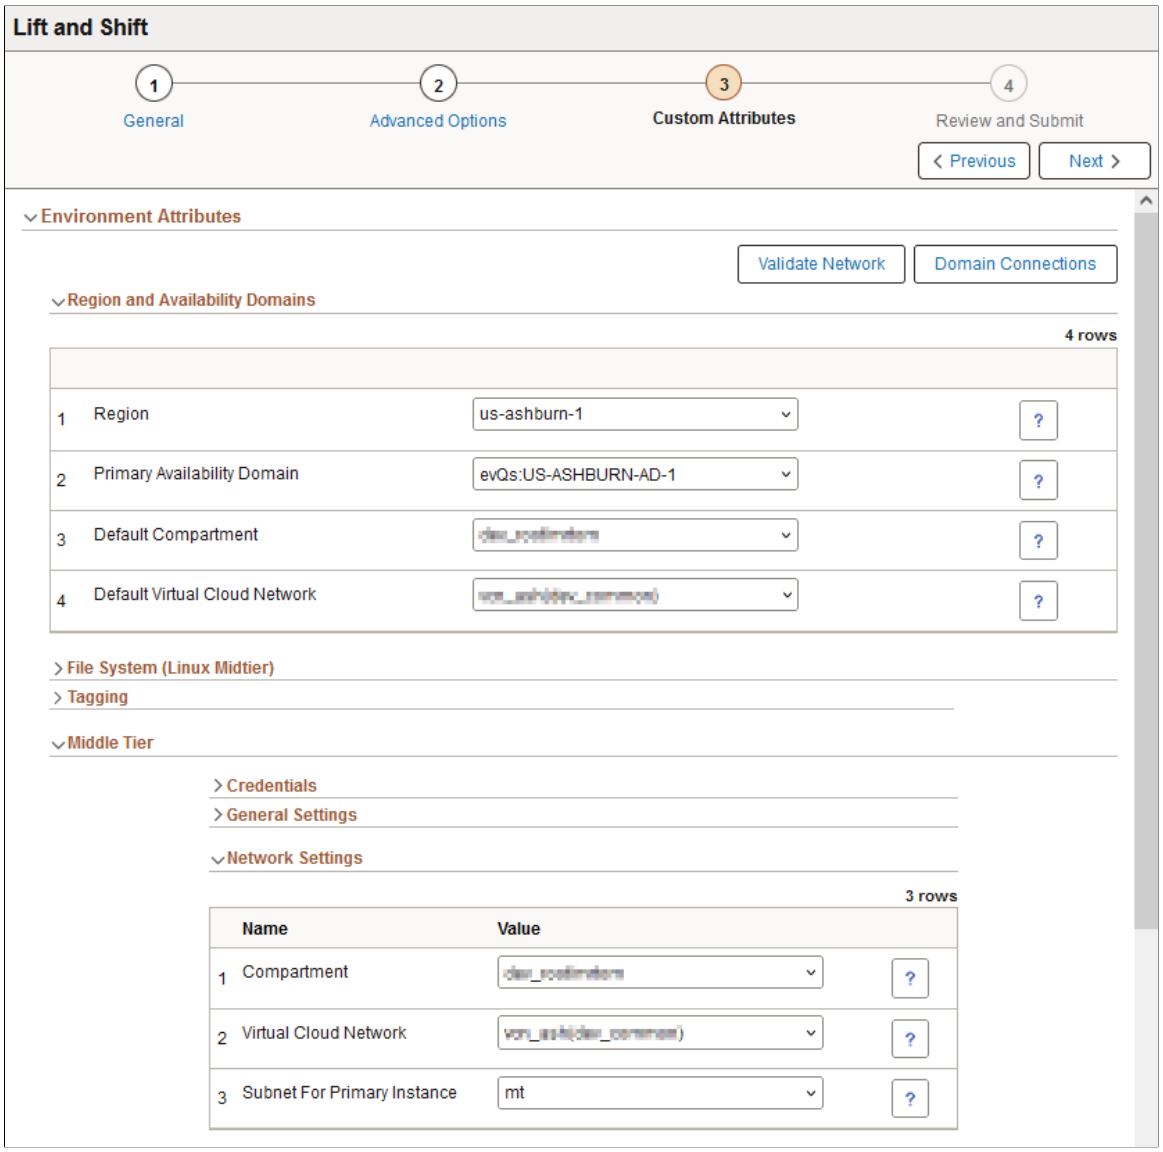

The values in the Network Settings section default to the compartment and VCN in the Region and Availability Domains section; these values can be changed for each tier.

See Configuring Region and Availability Domains, Configuring Network Settings, and Configuring Domain Connections section in Creating An Environment.

File System (Linux Midtier) Section

This example illustrates the fields and controls on the Lift and Shift – File System (Linux Midtier) section.

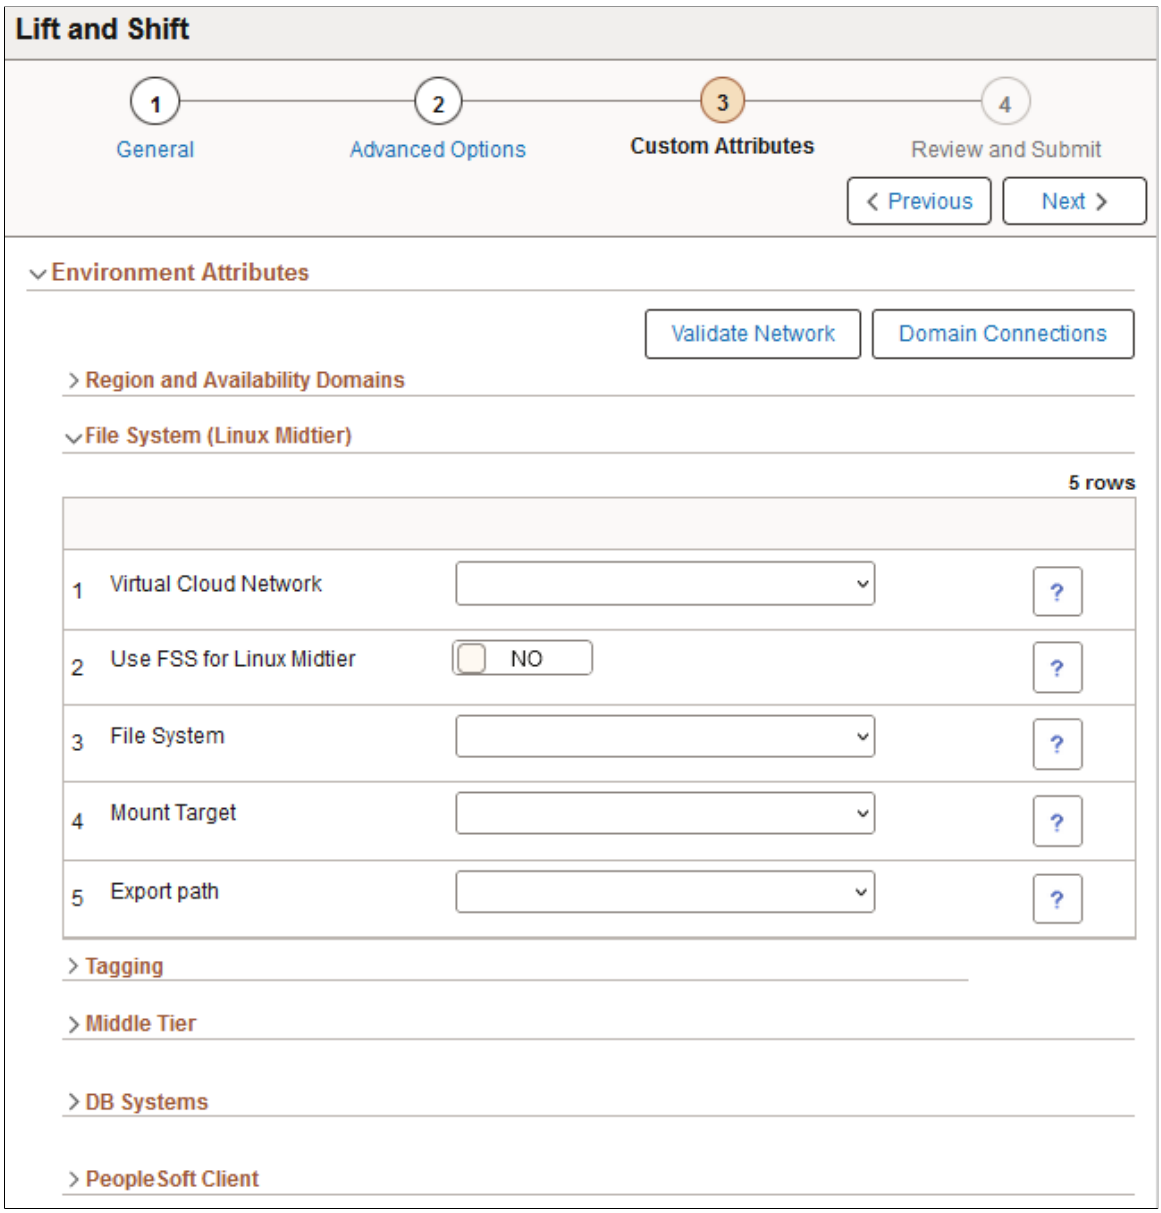

See Using Shared File System for Linux Middle Tier using File Storage Service for a description of the fields and controls.

Middle Tier Section

This example illustrates the fields and controls on the Lift and Shift - Middle Tier section.

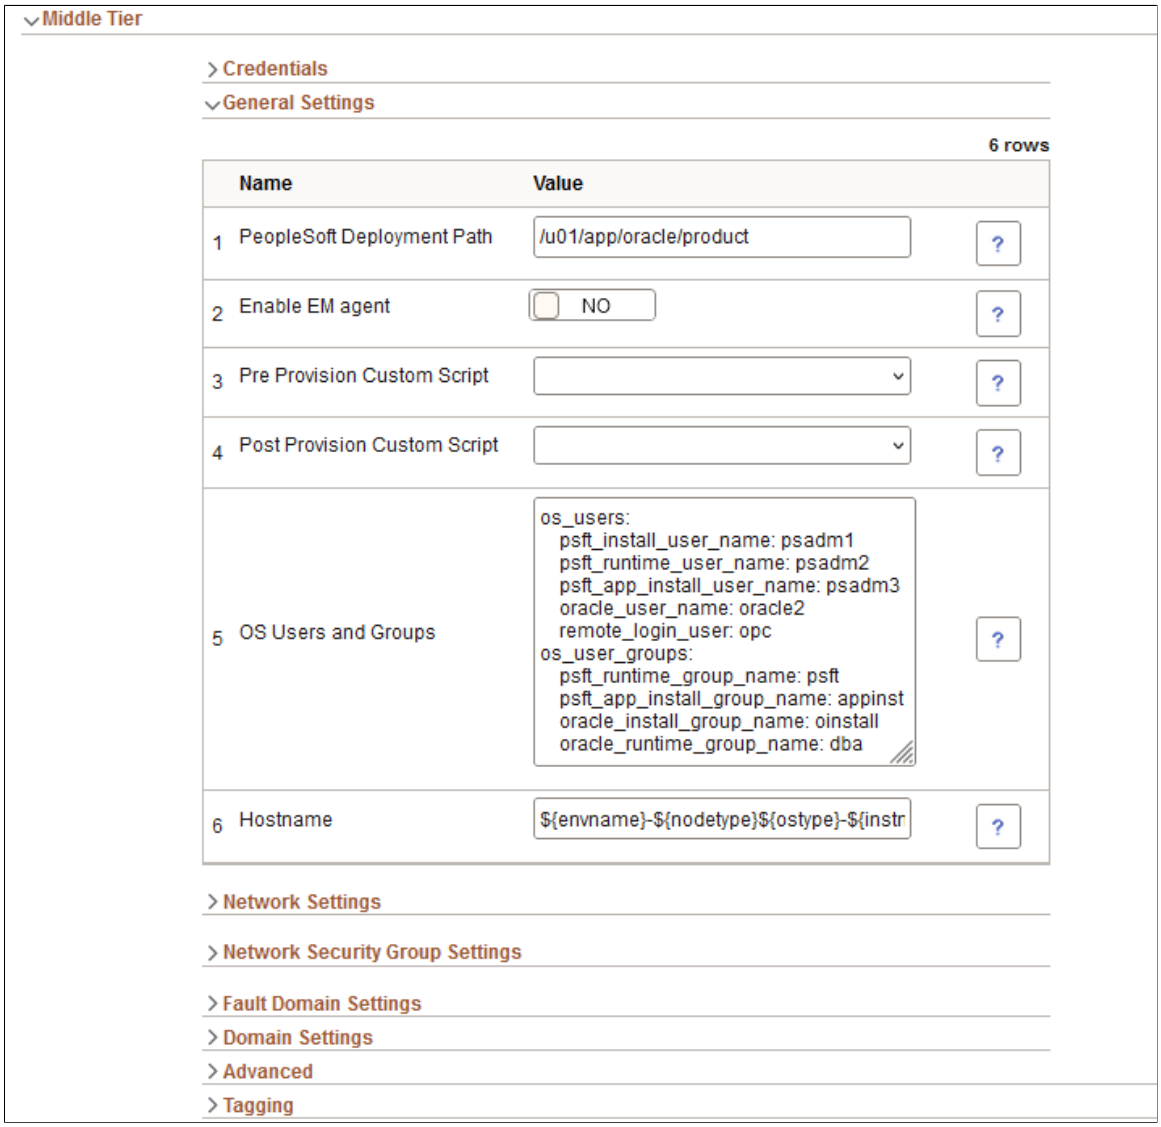

The fields and controls for Middle Tier are described in Configuring Web Server Tier Settings.

This example illustrates the fields and controls on the Lift and Shift – Middle Tier Domain Settings.

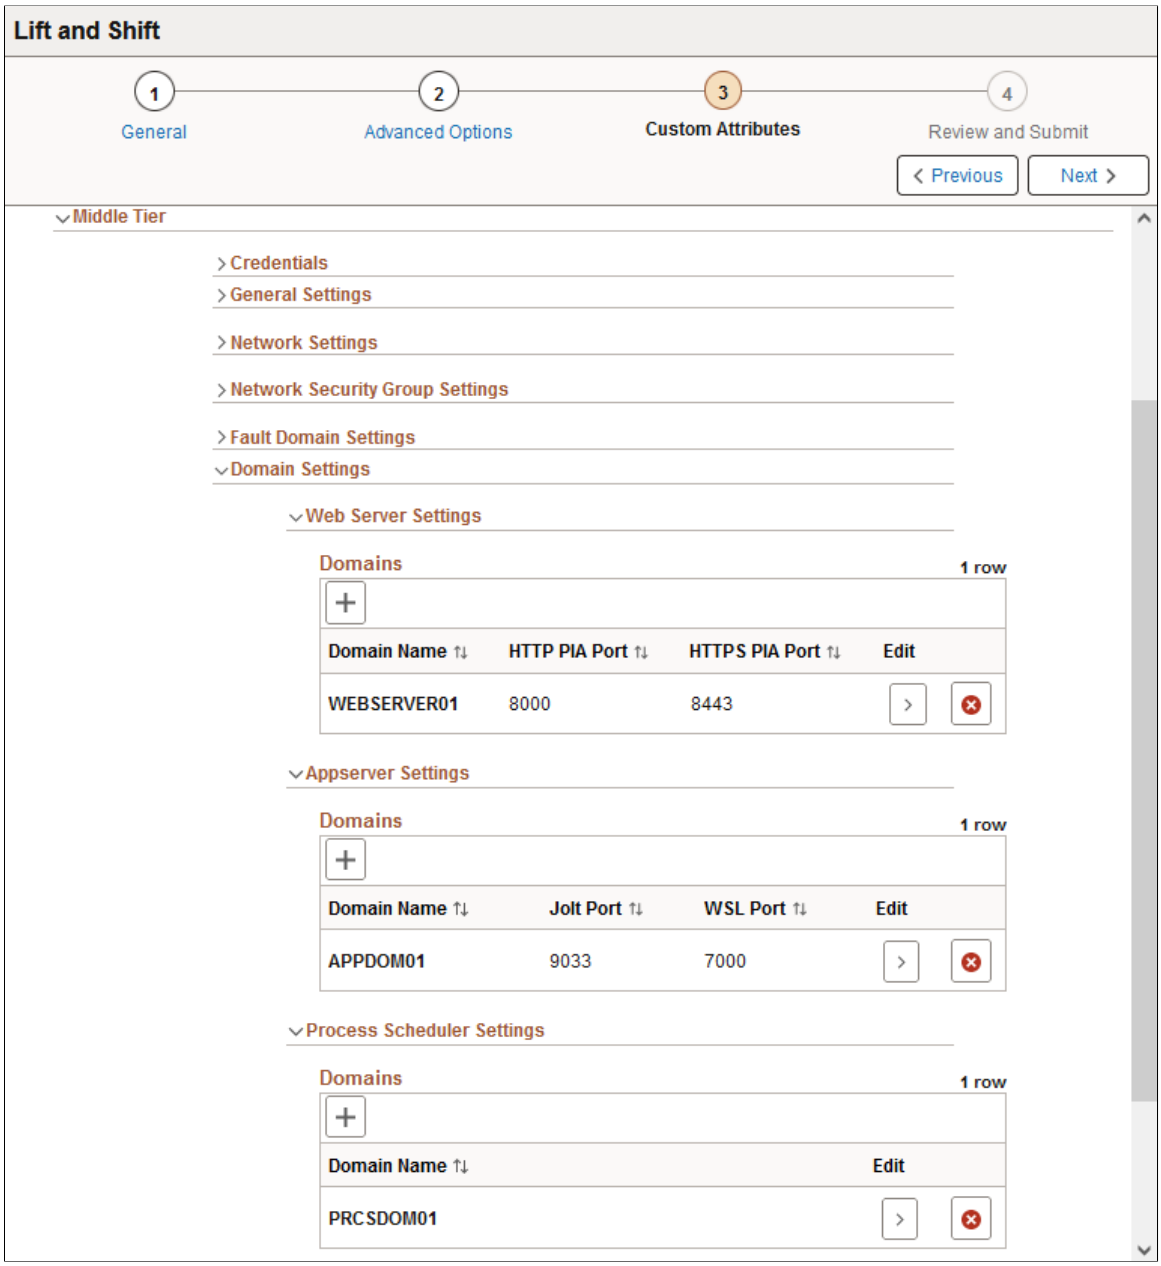

See Domain Settings Section in Configuring Full Tier Template Settings for field definitions and controls.

DB Systems Section

The DB Systems section fields change depending upon the database.

The fields and controls for the DB Systems section are described in Configuring DB Systems Settings.

This example illustrates the fields and controls on the Lift and Shift – DB Systems Credentials section for a TDE-enabled database.

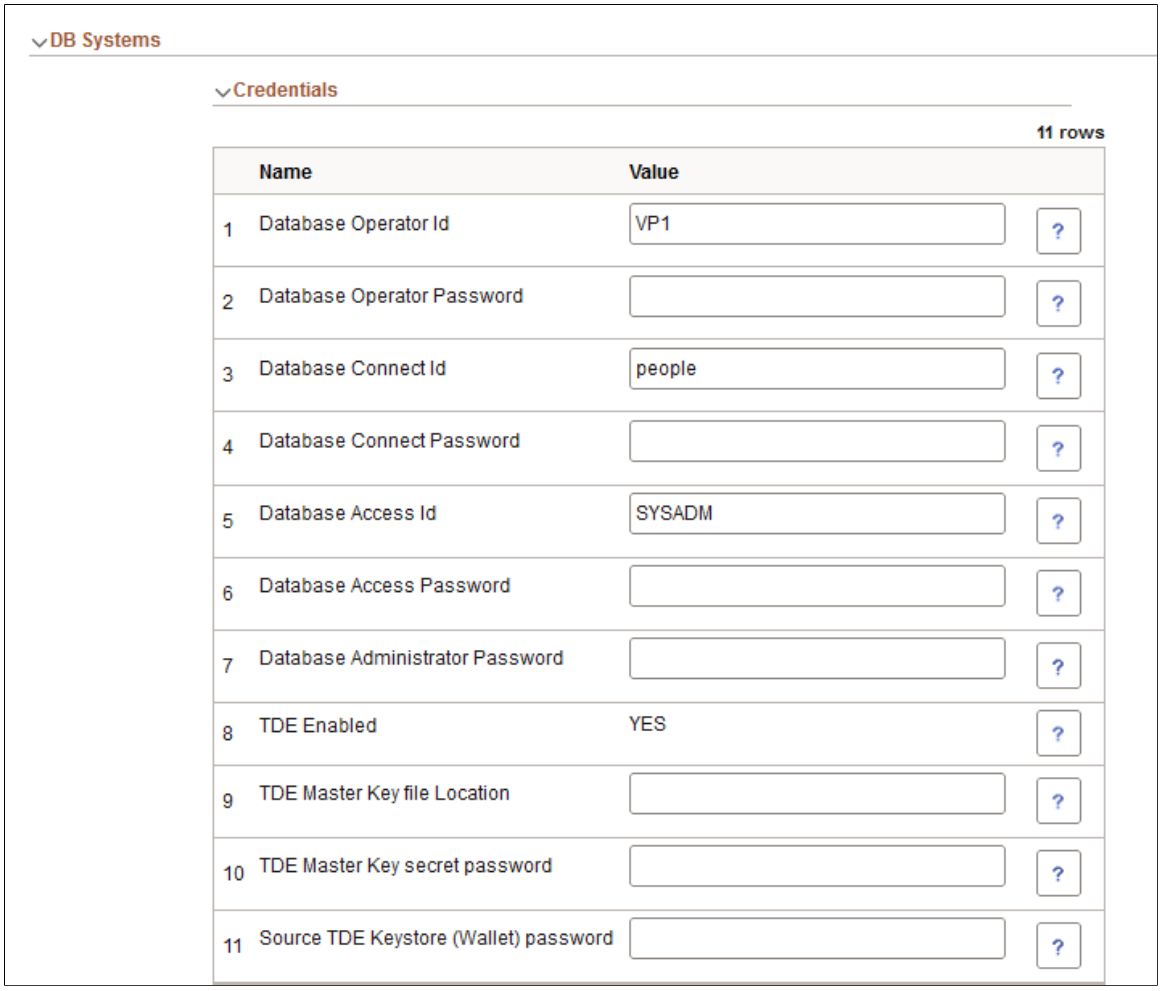

The fields and controls for TDE-enabled databases are described in Shifting TDE Encrypted Database .

This example illustrates the fields and controls on the Lift and Shift – DB Systems Credentials section with RMAN password.

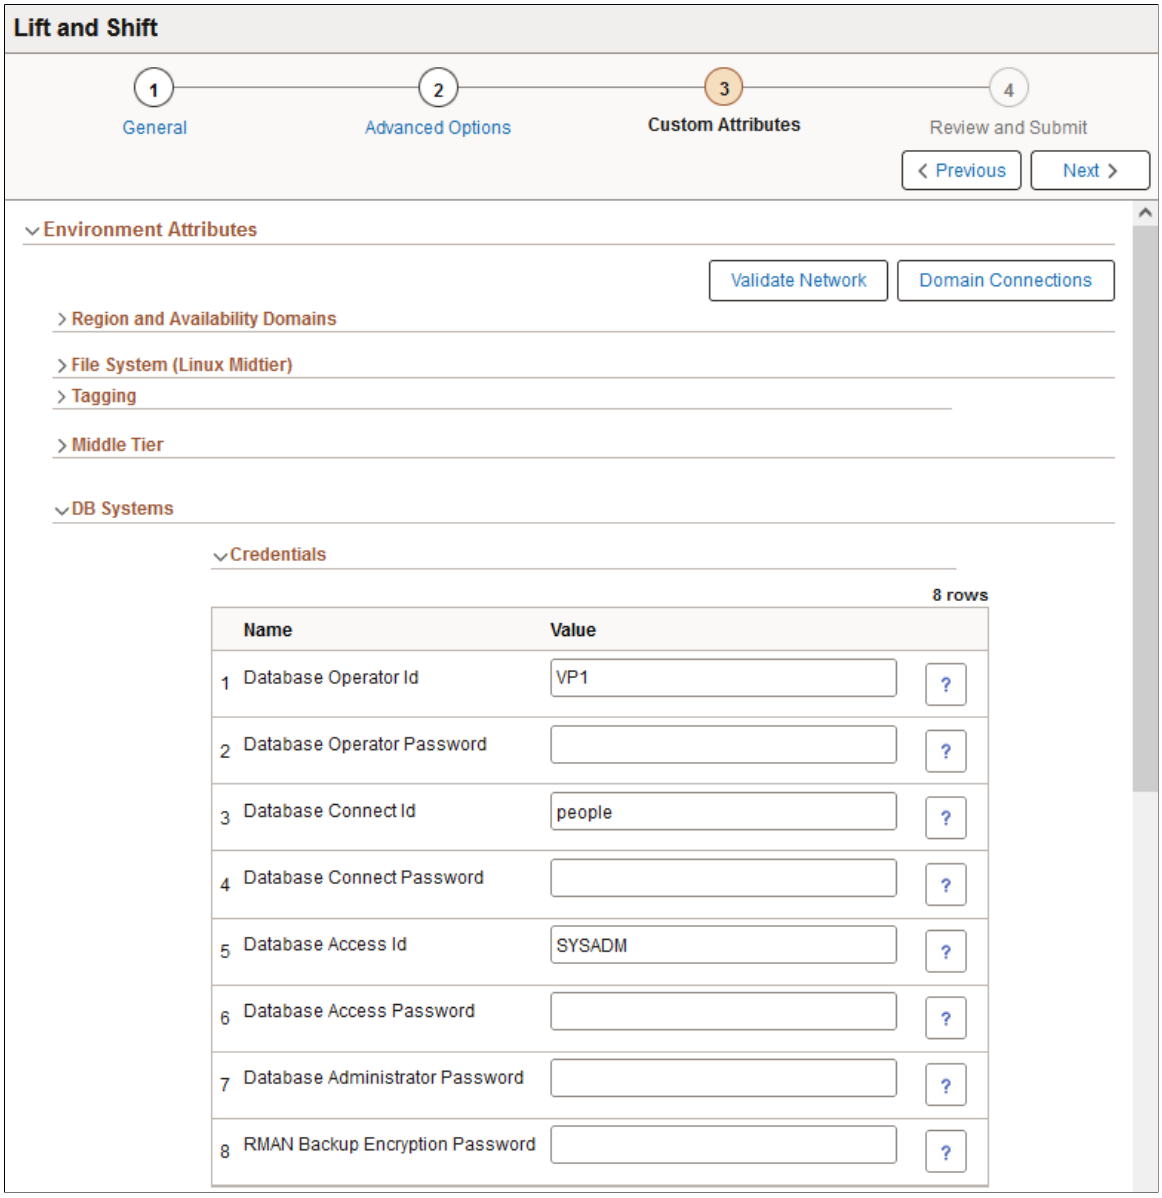

Note: When creating a new shifted environment of a lifted RMAN backup, you must provide a different CDB name than the CDB name of the source database on which the RMAN lift was performed. The shift will fail if the same CDB name is used.
See also Running Lift Using Hot Backup (RMAN).

This example illustrates the fields and controls on the Lift and Shift – DB Systems General Settings section.

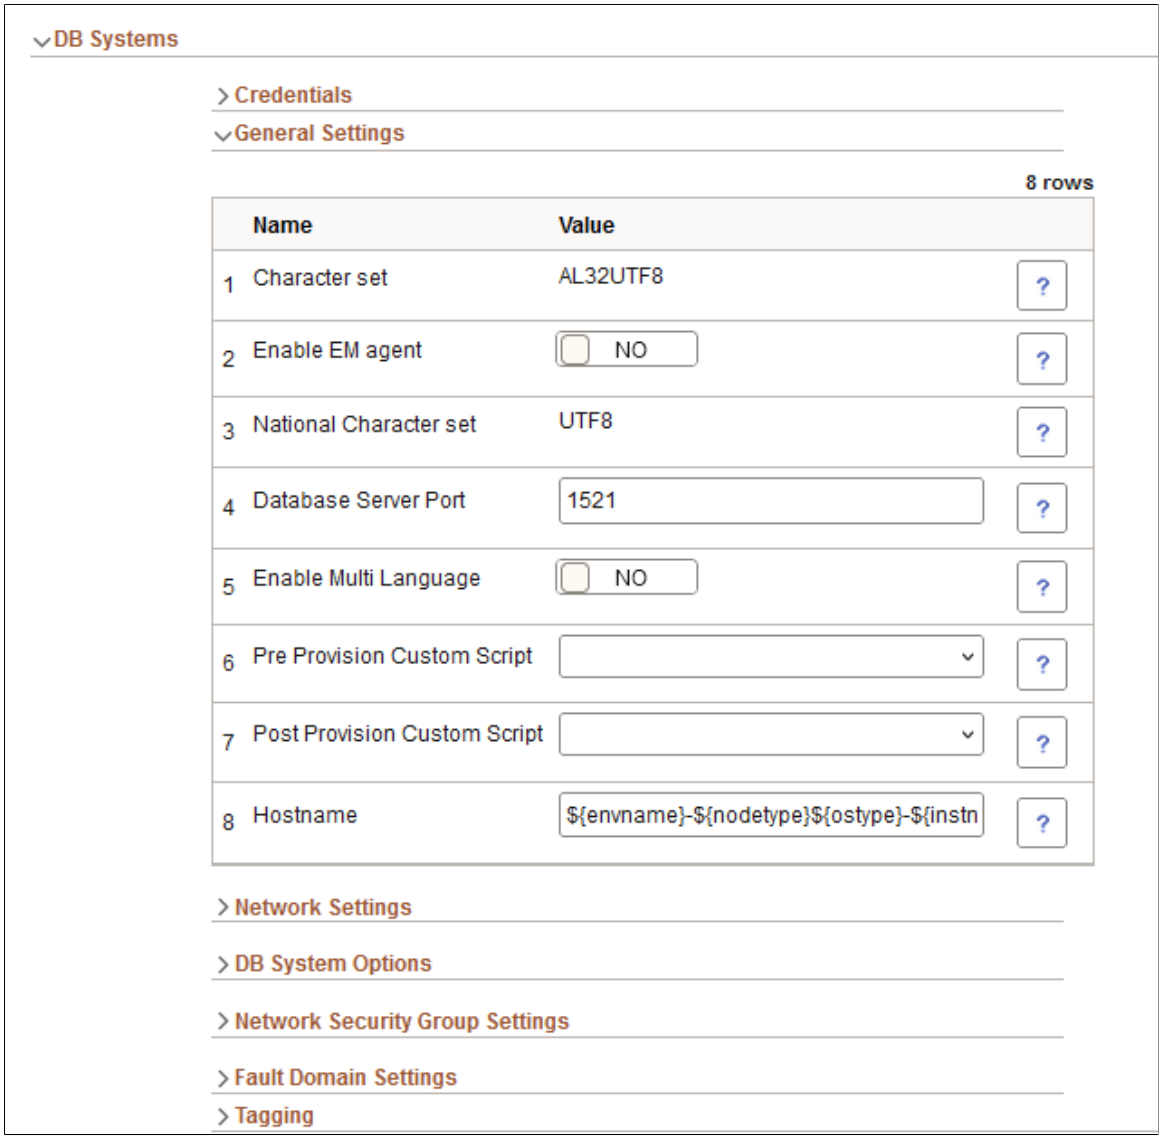

See Configuring DB Systems Settings.

Character Set Attributes

The Character Set and National Character Set attributes are configured with the same values as the on-premises database configuration.

The database character sets to be used for the Shift operation are AL32UTF8 and National Character Set AL16UTF16. Possible values of National Character Set when character set is AL32UTF8 are AL16UTF16 and UTF8. There can be multiple possible values of character set such as UTF8 WE8ISO8859P15. If shifting to DBaaS, you need to modify the character sets based on the database selected.

If you are using Cloud Manager to initiate a DBCS Shift; the “DBaaS Charset” and “DBaaS National Charset” configuration (under the Database Tier section) should match with the “Charset” and “National Charset” of the Database environment where the DB Lift operation is performed.

If there is any mismatch in the Charset data, the DBCS shift will fail.

To find the Charset and National Charset information from the lifted environment, run the following SQL commands on the DB (lifted) environment.

```
select VALUE from nls_database_parameters where parameter='NLS_CHARACTERSET';
select VALUE from nls_database_parameters where parameter='NLS_NCHAR_CHARACTERSET';⇒
```

Output:

```
SQL> SELECT value$ FROM sys.props$ WHERE name = 'NLS_CHARACTERSET' ;
VALUE$
AL32UTF8
SQL> SELECT value$ FROM sys.props$ WHERE name = 'NLS_NCHAR_CHARACTERSET';
VALUE$
UTF8
```

PeopleSoft Client Section

This example illustrates the fields and controls on the Lift and Shift – PeopleSoft Client Credentials section.

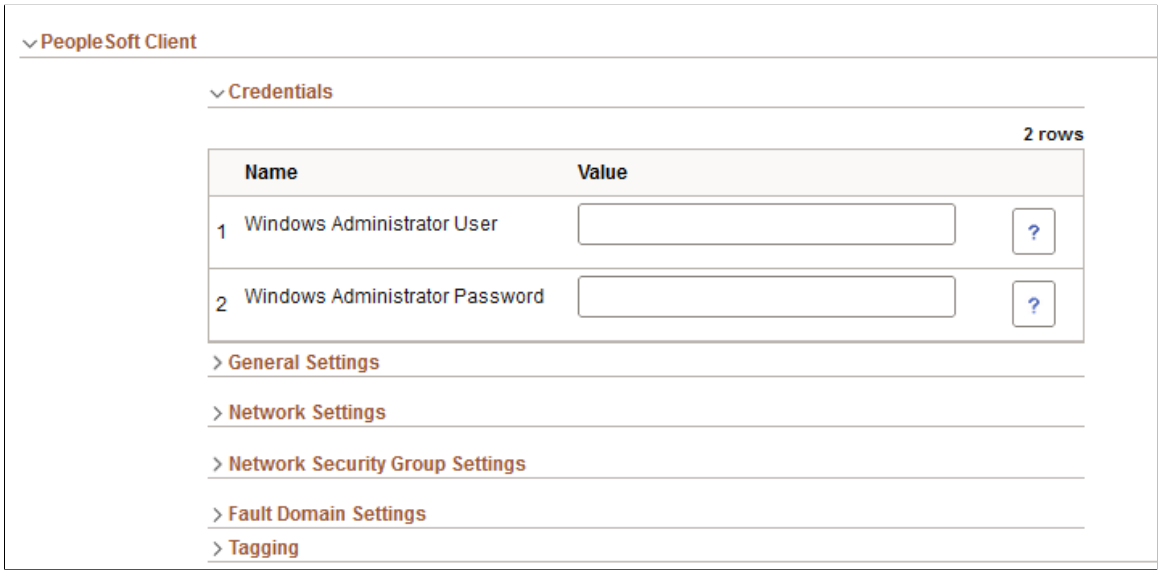

Field or Control	*Description*
Windows Administrator User	The Windows Administrator User is required to access the instance. If you want to specify a user other than the default, opc, you must first create a custom Windows image and add the custom user. The custom user must have administrative privileges. See the tutorial Create a Custom Windows Image for PeopleSoft Cloud Manager in Oracle Cloud Infrastructure (Optional) at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

Field or Control	*Description*
Windows Administrator Password	Enter the password for the Windows administrator.

Lift and Shift – Review and Submit Page

Use the Lift and Shift – Review and Submit page (ECL_LAS_CUSTATR_FL) to review and submit the entered environment details.

Navigation:

Click step 4 or Next at the top of the Lift and Shift guided process.

This example illustrates the fields and controls on the Lift and Shift – Review and Submit page.

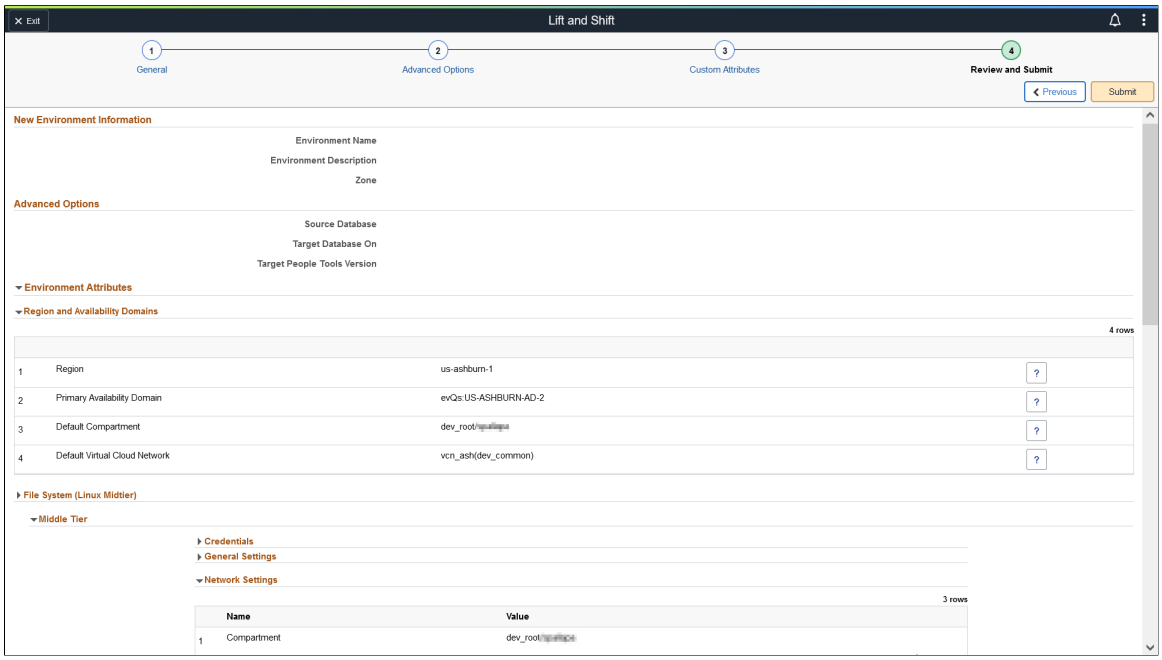

Review the details that were entered for the environment.

Click the Submit button to initiate the creation of a lifted environment in Oracle Cloud based on the details provided.

Once the environment is ready, you will be able to view it under the Environments tile. For details, see Environments Tile

Migrating TDE Enabled Database to Oracle Cloud Using PeopleSoft Cloud Manager

Transparent Data Encryption (TDE) enables customers to encrypt sensitive data, such as Personally Identifiable Information (PII), that are stored in tables and tablespaces.

After the data is encrypted, this data is transparently decrypted for authorized users or applications when they access this data. TDE helps protect data stored on media (also called data at rest) in the event that the storage media or data file is compromised.

This is a graphical representation of the Transparent Data Encryption for Cloud Manager.

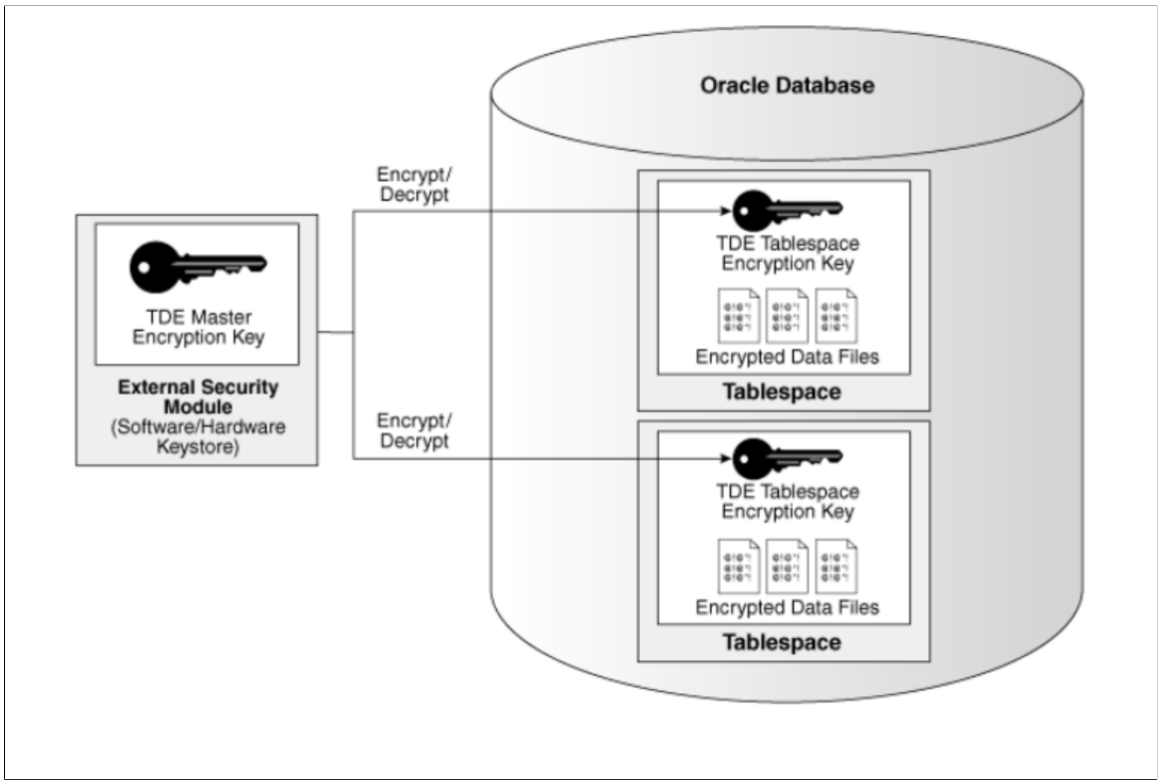

Prerequisites

Below requirements must be satisfied to successfully migrate a TDE enabled database.

- Database being migrated must have TDE enabled and required tablespaces already encrypted.
- Follow the support guidelines for PeopleSoft PeopleTools and PeopleSoft applications on My Oracle Support Certifications and on the PeopleSoft Cloud Manager Home Page, My Oracle Support, Doc ID 2231255.2.
- Database must be an Oracle 19c or later container database.
- Must be a Unicode, non-RAC and non-ASM database.
- Must have a subscription to OCI DBaaS.
- Migration to Compute instance is not supported.
- Remote lift is not supported on TDE.

Lifting TDE Encrypted Database

After the lift process is completed, DPKs are created and the TDE Encryption Keys are exported to a file. This exported file must be securely stored and later provided as input when deploying the lifted DPKs.

1. Download the latest lift utility.
2. Copy and extract the utility on the on-premises environment.
3. Run the lift utility to package database and middle-tier environment into DPKs. The Lift utility when triggered on a TDE Enabled Database prompts for TDE Keystore (Wallet) Password.

   This example illustrates the Lift Utility for TDE enabled database which prompts for the TDE Keystore (Wallet) password.

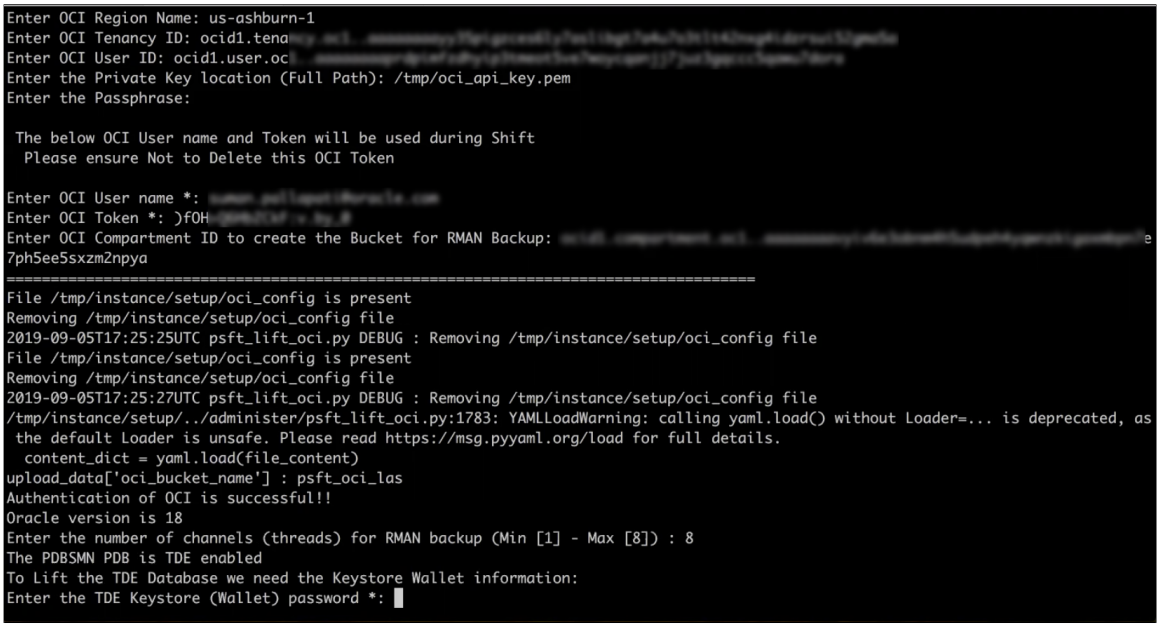

4. Lift utility uploads the DPKs to Oracle Cloud Infrastructure Object Storage.

   See Running Lift Using Hot Backup (RMAN).

5. The TDE encryption wallet directory will be packaged on the on-premises system in a zip file under /<LIFT_UTILITY_PATH>/data/masterkey.zip. The lift log file will have the path to the zip file as shown in the example below. This zip file must be backed up and available when shifting.

   Lift Log File (/<Lift_Utility>/data/psft_lift_session_<PDBNAME>_<SESSIONID>_<PID>.log)

```
2018-01-02T04:36:19EST apputils.py INFO : Master Key zip created successfully
2018-01-02T04:36:19EST apputils.py INFO :
*********************************************************************
2018-01-02T04:36:19EST apputils.py INFO : TDE Key to use during Shift is
packaged within: /mnt/azfs/osl/data/masterkey.zip
2018-01-02T04:36:19EST apputils.py INFO :
*********************************************************************
```

Shifting TDE Encrypted Database

After the lifted DPKs are uploaded to Oracle Cloud Infrastructure Object Storage, navigate to the Lift and Shift page in Cloud Manager and click the button to ‘List Object Store Items’ to refresh the list. Follow below steps to deploy the lifted DPKs.

1. Securely copy the TDE encryption key export file (masterkey.zip, this should be accessible for psadm2 users) to Cloud Manager instance using your favorite SCP tool.

   Note: The length of the path to the zip file must be less than 30 characters.

2. Identify the lifted DPK that must be shifted and initiate shift process by selecting 'Create Environment' in the Actions menu.

3. Provide all the New Environment Information and click Next.

4. In Advanced Options, the **Target Database On** option is set to DBaaS. Compute option is not supported when migrating a TDE encrypted database. Select the PeopleTools patch version and click Next.

   This example illustrates the fields and controls on the Lift and Shift – Advanced Options page.

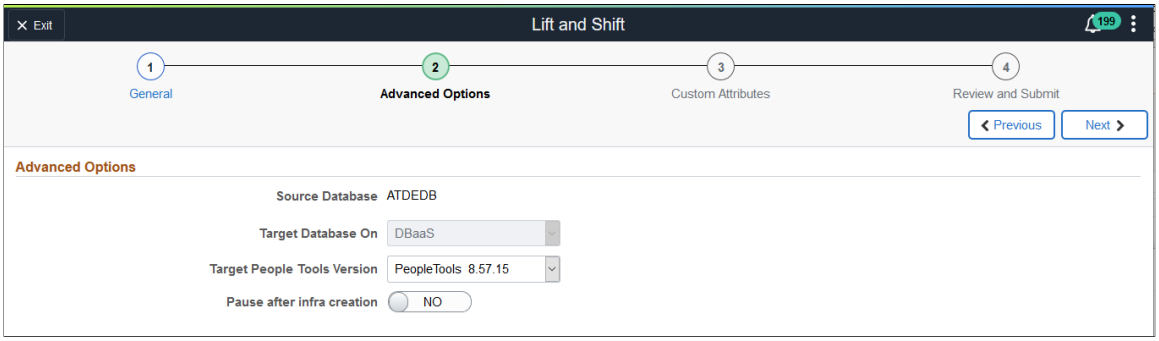

5. In Custom Attributes page, TDE related inputs are listed under DB Systems > Credentials. Provide the path to the masterkey.zip file from step 1 as input to TDE Master Key file location and the secret password. Provide all other required inputs and click Next.

   Note: User is only prompted for TDE Wallet password during Lift, however during Shift the user will be prompted for both TDE Wallet and Master Key secret passwords. Master key secret password is user specific with no restrictions.

   This example illustrates the fields and controls on the TDE Specific Fields in Custom Attributes Page.

6. Finally, review all inputs and submit the request to start provisioning the lifted DPKs.

Shifting to RAC on DBaaS

Cloud Manager only supports shift to RAC on DBaaS.

Before shifting the database to RAC, you must modify the Lift and Shift - DBaaS topology with the required shape, and disk capacity of the database and middle-tier nodes.

Note: The VM shape needs to be supported for RAC (2-node DB system). During Shift, Multi node RAC provisioning needs DHCP Options to be set as "DNS Type: Internet and VCN Resolver" for database subnet.

See Editing an Existing Topology

In order to shift to RAC in OCI, follow the procedure for shifting the migrated environment to the Cloud. See Using the Shift Process to Provision the Migrated Environment on the Oracle Cloud .

- On the Lift and Shift – Advanced Options page, select *DBaaS* for Target Database On field.

  Note: Refer to the Support Matrix for Shift Provisioning on Target Database posted on the PeopleSoft Cloud Manager Home Page.

- On the DB System Options page, the Node Count must be 2.

  The number of nodes in the database system depends on the shape you select. The shape selected in the topology must support 2 nodes.

  This example illustrates the fields and controls on the DB System Options.

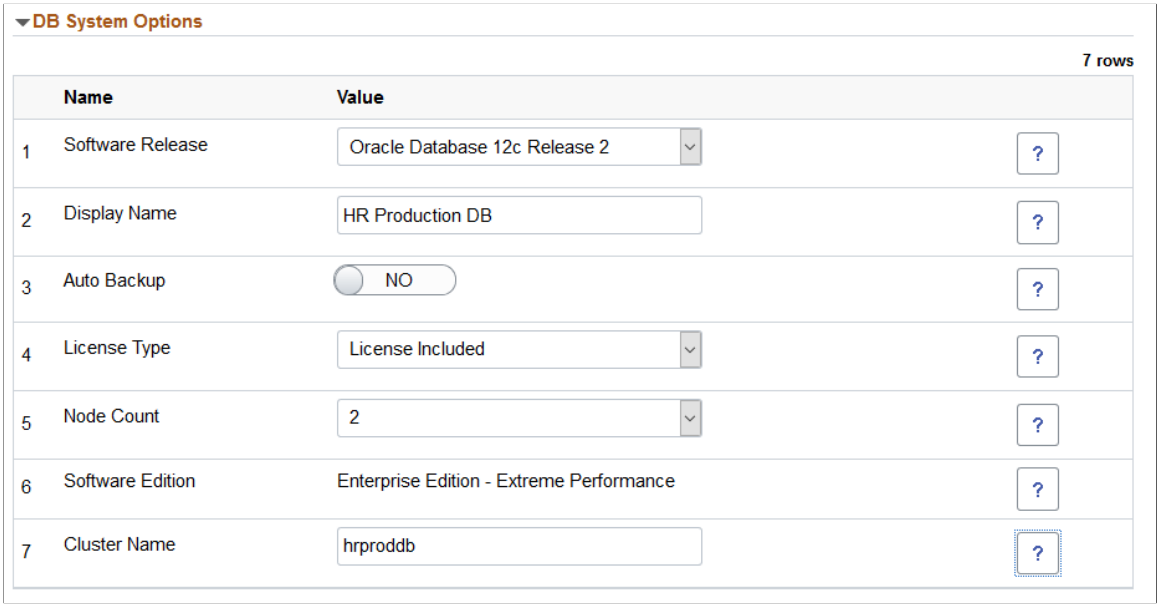

Encrypting Tablespaces Using Transparent Data Encryption

Note: The procedure explained below to encrypt an existing database must be performed on the source environment before performing the lift.

This topic summarizes the procedure to enable Transparent Data Encryption (TDE) Tablespaces Offline Encryption for an Oracle PeopleSoft Applications database. This process is referred to as using the *Fast Offline Conversion* method to convert existing clear data (residing in non TDE encrypted tablespaces) to

TDE encrypted tablespaces. In order to use this feature, the PeopleSoft Applications database requires downtime, as the tablespace(s) to be encrypted need to be temporarily offline. As the encryption is transparent to the application, code does not have to be rewritten, and existing SQL statements work as they are. Transparent also means that any authorized database session can read the encrypted data without any problem: the encryption only applies to data-at-rest, meaning the database data files and any backups of them.

Refer to the information on Transparent Data Encryption in the Oracle Database Advanced Security Guide for your Oracle Database version. See Oracle Database Documentation, https://docs.oracle.com/en/database/oracle/oracle-database/index.html.

Prerequisites

- This procedure can be used with Oracle PeopleSoft Applications Database on supported Oracle Database versions.

  See PeopleSoft Cloud Manager Home Page, My Oracle Support, Doc ID 2231255.2, for support information.

- Understand TDE implications and restrictions and develop a process for maintaining wallets and keys. Refer to the *Oracle Database Advanced Security Administrator's Guide* for further details.

- Ensure the compatible database parameter is set to the appropriate database version, 19c or later.

- Always take a full backup of your database before starting the procedure.

TDE Offline Datafile Encryption Restrictions

The following restrictions apply to implementing Tablespace Encryption using Fast Offline Conversion:

- It can only be performed for application tablespace data files. SYSTEM, SYSAUX, UNDO and TEMP tablespaces cannot be encrypted.

- External Large Objects (BFILEs) cannot be encrypted using TDE tablespace encryption because these files reside outside the database. PeopleSoft applications do not utilize BFILEs.

Procedure to Perform TDE Tablespace Offline

To perform TDE Tablespace Offline Encryption for an Oracle PeopleSoft Applications database, follow the steps below:

1. Shut down application server processes.Shut down all Applications server processes and make sure all jobs are completed cleanly before continuing further. Users should be prevented from using the Applications database until the encryption process is completed.

2. Source your Oracle PeopleSoft Applications Database Oracle Home.

3. Create a wallet by specifying the wallet location in the sqlnet.ora file under the $TNS_ADMIN directory:

   a. Add the following entry to the sqlnet.ora:

   ```
   ENCRYPTION_WALLET_LOCATION = (SOURCE = (METHOD = FILE) (METHOD_DATA = (DI⇒
   ```



```
RECTORY = $ORACLE_HOME/admin/TDE/$ORACLE_SID)))
```

b. Create the corresponding directory manually:

```
$ mkdir -p /$ORACLE_HOME/admin/TDE/$ORACLE_SID
```

c. Check wallet location and status:

```
$ sqlplus / as sysdba;SQL>select * from V$encryption_wallet;
```

4. Create a Keystore in the wallet.

```
SQL>ADMINISTER KEY MANAGEMENT CREATE KEYSTORE '/$ORACLE_HOME/ADMIN/tde/$ORACLE⇒
_SID' IDENTIFIED BY "<Strong password>";
```

5. Open the Keystore create in step 4. As we are in a multitenant environment, we have to specify CONTAINER=ALL in order to set the keystore in all the PDBs:

```
SQL>ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY PASSWORD CONTAIN⇒
ER=ALL;
```

6. Set the master encryption key:

```
SQL>ADMINSTER KEY MANAGEMENT SET KEY IDENTIFIED by "<Strong password>" CONTAIN⇒
ER=ALL;
```

Note: The password must be enclosed in double quotes as shown.

---

7. Bounce the database:

```
SQL> shutdown normal;SQL> exit;
```

8. Start up the database normally, ensuring that the wallet is open:

```
sqlplus "/ as sysdba"SQL>startup;  SQL>ADMINISTER KEY MANAGEMENT SET KEYSTORE ⇒
OPEN IDENTIFIED BY "<Strong password>" CONTAINER=ALL;
```

9. Switch to the PeopleSoft PDB.

```
SQL> ALTER SESSION SET CONTAINER=<PDBNAME>;
```

10. Identify all the temporary and undo tablespaces in the database:

```
SQL>select tablespace_name from dba_tablespaces where contents='TEMPORARY' and⇒
 STATUS='ONLINE';
SQL>select tablespace_name from dba_tablespaces where contents='UNDO' and STAT⇒
US='ONLINE';
```

11. While still in the PDB, generate three scripts, which will be used perform the TDE offline data conversion.

    ALTDATAFILESOFFLINE.SQL

    ALTDATAFILESENCRYPT.SQL

ALTDATAFILESONLINE.SQL

a. Script One takes specific data files offline. Create a script file with the following statements and save file as generatealtdatafilesoffline.sql.

```
sqlplus "/ as sysdba"
SET LINESIZE 256
SET HEADING OFF;
SET TERM OFF;
SET FEED OFF;
SPOOL ALTDATAFILESOFFLINE.SQL
select 'alter database datafile '''||b.file_name|| ''' offline;'
from dba_tablespaces a, DBA_DATA_FILES b
where a.tablespace_name not in ('SYSTEM','SYSAUX','TEMP','PSTEMP','PSGTT0⇒

1')and a.tablespace_name=b.tablespace_name
;
Spool off
Exit
If you call the generation script GENERATEALTDATAFILESOFFLINE.SQL using @⇒

 from SQLPLUS, then you will not have to do any additional editing of the⇒

 generated script.
SQL>alter session set container=<PDBNAME>
System altered.
SQL>@generatealtdatafilesoffline.sql
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2⇒

.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Test⇒

ing options
$
```

b. Script Two offline encrypts data files offline. Create a script file with the following statements and save file as generatealtdatafilesencrypt.sql.

```
sqlplus "/ as sysdba"
SET LINESIZE 256
SET HEADING OFF;
SET TERM OFF;
SET FEED OFF;
SPOOL altdatafilesencrypt.sql
select 'alter database datafile '''||b.file_name|| ''' ENCRYPT;'
from dba_tablespaces a, DBA_DATA_FILES b
where a.tablespace_name not in ('SYSTEM','SYSAUX','TEMP','PSTEMP','PSGTT0⇒

1')and a.tablespace_name=b.tablespace_name
;
Spool off
Exit
If you call the generation script GENERATEALTDATAFILESENCRYPT.SQL using @⇒

 from SQLPLUS, then you will not have to do any additional editing of the⇒

 generated script.
SQL>alter session set container=<PDBNAME>;
System altered.
SQL>@generatealtdatfilesencrypt.sql
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2⇒

.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Test⇒

ing options
$
```


c. Script Three brings data files back online. Create a script file with the following statements and save file as generatealtdatafilesonline.sql.

```
sqlplus "/ as sysdba"
SET LINESIZE 256
SET HEADING OFF;
SET TERM OFF;
SET FEED OFF;
SPOOL altdatafilesoonlineexec.sql
select 'alter database datafile ’’’||b.file_name|| ’’’ online;'
from dba_tablespaces a, DBA_DATA_FILES b
where a.tablespace_name not in ('SYSTEM','SYSAUX','TEMP','PSTEMP','PSGTT0⇒

1')and a.tablespace_name=b.tablespace_name
;
Spool off
Exit

If you call the generation script GENERATEALTDATAFILESONLINE.SQL using @ ⇒

from SQLPLUS, then you will not have to do any additional editing of the ⇒

generated script.

SQL>alter session set container=<PDBNAME>;
System altered.
SQL>@generatealtdatafilesonline.sql
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2⇒

.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Test⇒

ing options
$
```

d. Then get back to root or the CDB level.

```
SQL> ALTER SESSION SET CONTAINER=CDB$ROOT;
```

e. Close the PDB. We want the state to be in 'MOUNT' mode.

```
SQL> ALTER PLUGGABLE DATABASE <PDBNAME> CLOSE IMMEDIATE;
```

f. Switch to the PeopleSoft PDB.

```
SQL> ALTER SESSION SET CONTAINER=<PDBNAME>;
```

12. Bring all the specified tablespaces offline by connecting to SQL*Plus as sysdba, and running the script altdatafilesoffline.sql.

```
$ sqlplus / as sysdbaSQL> @altdatafilesoffline.sql;
```

13. Encrypt your datafiles by running the altdatafilesencrypt.sql offline encryption script from SQL*Plus as sysdba:

```
$ sqlplus / as sysdba
SQL>@altdatafilesencrypt.sql;
```

Note: If you have a large number of data files, you can parallelize their encryption by creating sub-scripts and running the sub-scripts from parallel SQL*Plus sessions.

---

14. Bring all the specified tablespaces online by connecting to SQL*Plus as sysdba, and running the script altdatafilesonline.sql.

```
$ sqlplus / as sysdbaSQL> @altdatafilesonline.sql;
```

Note: Some tablespaces may take time to show as online. These are probably tablespaces that are encrypted.

Check the status of tablespace encryption by connecting to SQL*Plus / as sysdba and running the query shown:

```
$ sqlplus / as sysdba
SQL>select tablespace_name, encrypted from dba_tablespaces;
```

Note: Unless an auto login keystore is created, every time the database is started up, the wallet will need to be opened as in Step 8 above.

To make the wallet auto login, run the following command:

```
$ sqlplus / as sysdba$ administer key management create AUTO_LOGIN keystore fr⇒

om keystore "<Wallet Path>" identified by "<Wallet Password>";
```

Bounce the database.

Chapter 10

Using Zero Downtime Migration to Migrate Environment to Cloud Manager

Understanding Zero Downtime Migration

Zero Downtime Migration (ZDM) gives you a quick and easy way to move on-premises databases and Oracle Cloud Infrastructure Classic instances to Oracle Cloud Infrastructure, Exadata Cloud at Customer, and Exadata Cloud Service without incurring any significant downtime, by leveraging technologies such as Oracle Active Data Guard.

For information on ZDM see Zero Downtime Migration

Using ZDM to migrate your database involves creating a backup of the source database and restoring it to the target database in OCI. The target database must be in DB System.

Once the database is migrated to OCI, you can import the DB System in Cloud Manager and add the middle tier node.

Migrating Environment Using Zero Downtime Migration

To migrate an on-premises environment to Cloud Manager using Zero Downtime Migration (ZDM), perform the following steps:

1. Create a DB System in OCI.
2. Use ZDM to migrate the database to the DB System in OCI.
3. Import the DB System environment in Cloud Manager.
4. Perform an application lift to lift PS_APP_HOME and PS_CUST_HOME.
5. Add middle tier node to the running database environment.

Creating a DB System in OCI

Refer to Creating DB System to create the DB System.

The DB System needs to be accessible to Cloud Manager. Keep the following in mind when creating the DB System:

- Passwords used while creating DB System should be same as source database.
- Container Database (CDB) Name should be same as source database.

- Oracle Database software version should be the same as the source database.
- Cloud Manager must be able to access the DB System.

  It is not necessary for Cloud Manager and the DB System to be in the same VCN, however if they are in different VCNs, then VCN peering has to be done before the database import.

  See VCN Peering.
- Cloud Manager SSH public key must be added to the DB System in order for the import to work.

  Add the SSH key with one of these methods:

  - OCI Console, see Creating DB System.
  - SSH to DB System, see Connecting to a DB System and add the Cloud Manager key.

Using ZDM to Migrate the Database to the DB System in OCI

Prerequisites

- ZDM tool is installed and configured.

  See Zero Downtime Migration 19.2.
- DB System created in OCI.

  See Creating a DB System in OCI
- Source and Target Database environments must be accessible by ZDM tool.

Migrating the Database

Follow the steps in the ZDM documentation for Migrating Your Database with Zero Downtime Migration.

Note: Migration method in the ZDM response file should be "MIGRATION_METHOD=BACKUP_RESTORE_OSS".

Importing the DB System Environment

Import the DB System Environment. See Importing Environment for Add DB System Node.

Perform an Application Lift

Perform an application lift to create the Application DPK, which includes PS_APP_HOME and PS_CUST_HOME.

See Performing Application Lift

Adding Middle Tier Node

To add the middle tier node to the running database environment:

1. Select the Environments tile.
2. Select Manage Node action for your imported DB System environment.
3. Expand the Select Action section and add a middle tier node.

   This example illustrates the fields and controls on the Manage Node page to add a Linux middle tier node.

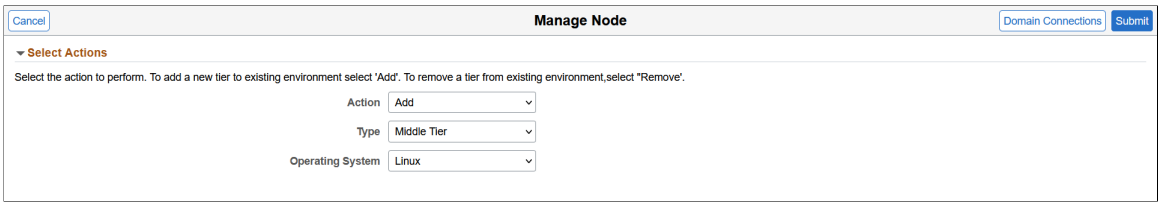

4. Expand the Settings section and select the location for the Application DPK.

   - Object Store

     When you select Object Store, the DPK Object Name drop down list will display all the application DPK files in Object Store.

     This example illustrates the fields and controls on the Manage Node page where the Application DPK location is Object Store.

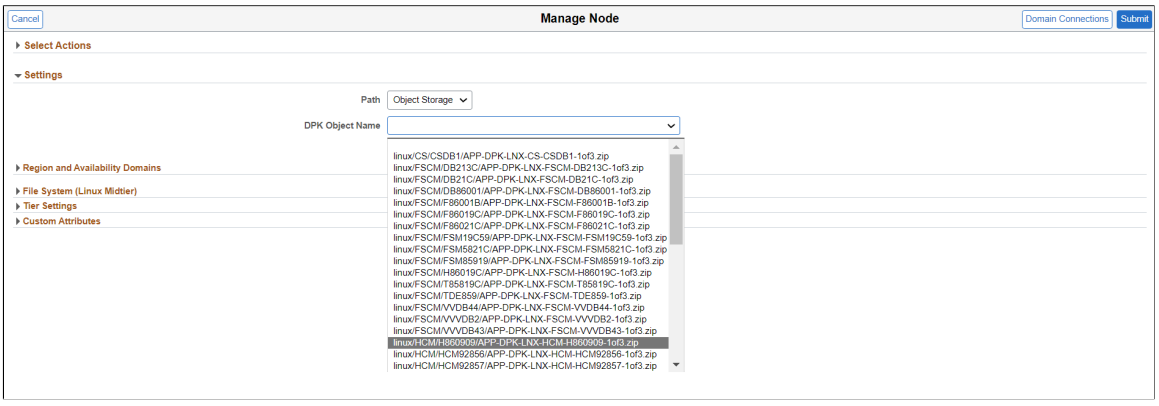

   - File Server

     If you manually copied PS_APP_HOME and PS_CUST_HOME to the file server provide the relative path to the file server.

This example illustrates the fields and controls on the Manage Node page where the Application DPK location is File Server.

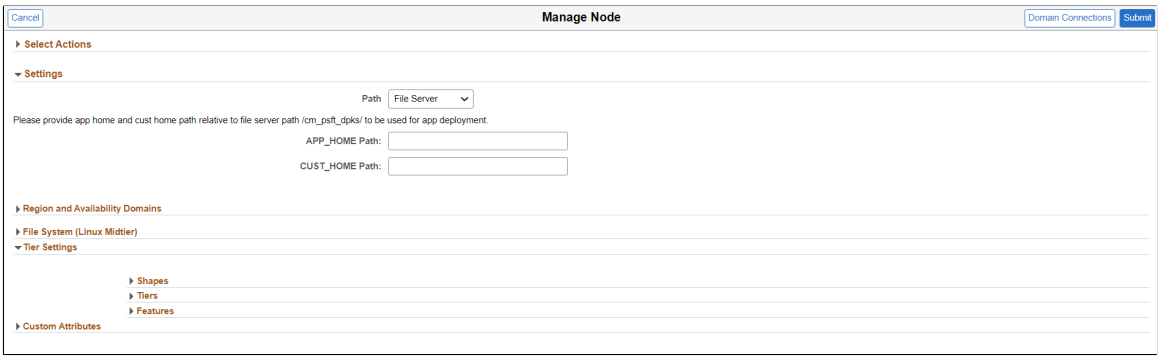

5. Enter the remaining required credentials and submit.

Migrating ADB-Dedicated Environment Using Zero Downtime Migration

Autonomous Database on Dedicated Infrastructure (ADB-D) can be migrated to Cloud Manager.

To migrate a PeopleSoft database environment to Cloud Manager using Zero Downtime Migration (ZDM), perform the following steps:

1. Use ZDM to migrate an on-premises PeopleSoft application database to Autonomous Database – Dedicated (ADB-D) on OCI.

   See Zero Downtime Migration documentation.

2. Import the ADB-D database environment in Cloud Manager.

   See the ADB Instance Type section in Importing Environment.

   Note: Starting with Cloud Manager 13, when the ADB-Dedicated database is imported, the RefreshDBWallet policy is automatically added to the environment.

3. Lift the APP DPK to the object store using the Application Lift (see Performing Application Lift) or copy the APP DPK to a path relative to the FSS of the Cloud Manager Instance.

4. Use Manage Node on the imported environment to add the middle tier node.

   See Managing Nodes.

Note: To support the refresh operation on Cloud Manager, it is advised to reset the passwords to 12 character passwords before migration or after migration to ADB-D.

Migrating ADB-Shared Environment Using Zero Downtime Migration

Autonomous Database on Shared Infrastructure (ADB-S) can be migrated to Cloud Manager using Zero Downtime Migration (ZDM)

Minimum PeopleTools requirements for migrating ADB-S to Cloud Manager are:

- 8.59.01
- 8.58.09
- 8.57.20

Prerequisites for importing ADB-S database include:

1. Backup Bucket must be set before taking backups through OCI (manual backup) or through PeopleSoft Cloud Manager.
2. The following policies are required for ADB-S operations:
    - use vcns for the compartment which the VCN is in
    - use subnets for the compartment which the VCN is in
    - use network-security-groups for the compartment which the network security group is in
    - manage private-ips for the compartment which the VCN is in
    - manage vnics for the compartment which the VCN is in
    - manage vnics for the compartment which the database is provisioned or is to be provisioned in

To migrate a PeopleSoft database environment to Cloud Manager, perform the following steps:

1. Use ZDM to migrate an on-premises PeopleSoft application database to Autonomous Database – Shared (ADB-S) on OCI.

    See PeopleSoft Application with Autonomous Database - Shared, Migration Guide with Oracle ZDM and Zero Downtime Migration documentation.

2. Import the ADB-S database environment in Cloud Manager.

    See the ADB Instance Type section in Importing Environment.

    Note: When the ADB-Shared database is imported, the RefreshDBWallet policy is automatically added to the environment.

3. Lift the APP DPK to the object store using the Application Lift (see Performing Application Lift) or copy the APP DPK to a path relative to the FSS of the Cloud Manager Instance.
4. Use Manage Node on the imported environment to add the middle tier node.

See Managing Nodes

Note: To support the refresh operation on Cloud Manager, it is advised to reset the passwords to 12 character passwords before migration or after migration to ADB-D.

Chapter 11

Enabling Selective Adoption in Cloud Manager

Enabling Selective Adoption in Cloud Manager

Cloud Manager enables customers to take advantage of Selective Adoption by automating creation of PUM source environments and configuration of target databases.

To use Cloud Manager for selective adoption:

- Set up a policy to automatically provision a PUM source environment when a new PUM image is downloaded from My Oracle Support into the Cloud Manager Repository, and keep it current.

  See Adding a Policy to Provision PUM Environments.

- Alternatively, create a PUM source environment as needed.

  See Creating PUM Environments.

- Configure target databases in the PUM source environment.

  See Adding Targets to PUM Sources.

- Access Change Assistant.

  See Accessing Change Assistant in Windows Client

- Follow standard procedure to apply updates to target environments.

For details on the selective adoption process refer to Selective Adoption.

Adding a Policy to Provision PUM Environments

Create an event-driven policy to automatically provision a PUM environment when a new PeopleSoft Update Image (PUM DPKs) is downloaded into the Repository, and migrate PUM metadata to keep the environment current. After you set up the policy, the event is triggered for each new downloaded update image. You can also set up additional optional actions.

To add a policy to provision PUM environments:

1. Set up an Oracle Cloud Infrastructure vault and add secrets for the passwords required for a PUM environment.

   See the information on creating vault resources in the Cloud Manager Installation tutorials at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials

2. Define a password group that includes the passwords required for a PUM environment.

   See Password Groups.

   Here are passwords typically required:

   - Database Administrator Password
   - Database Access Password
   - Database Connect Password
   - Database Operator Password
   - Gateway Administrator Password
   - Weblogic Administrator Password
   - Web Profile Password for user PTWEBSERVER
   - Windows Administrator Password

3. Subscribe to the desired update image in the Repository.

   See Download Subscriptions Page

4. Define a template for a PUM full-tier environment and associate the password group with it.

   Note: A PeopleSoft administrator must enable user access to the newly created PUM source environment template.

   See Creating a Template.

5. Define an event policy with these selections:

Field or Control	***Description***
Policy Object	Repository Artifact
Policy Type	Event
Event Name	PUM DPK Download
Policy Conditions	Select the Product and Release for the PeopleSoft update image to download. For example, HCM and 9.2. **Note:** Only Linux images are supported for PUM environments, so it is not necessary to specify the platform.

Field or Control	*Description*
Policy Action	Provision PUM Migrate PUM Metadata Define and upload target (optional) Delete the old PUM source (optional)

See Configuring Self-Managed Update Images.

6. Use Policy Monitor to review the environment name and process details.

   See Using Policy Monitor.

Creating PUM Environments

To create a new PUM source environment using Cloud Manager:

1. Ensure that the latest required PeopleSoft Update Image is downloaded in the repository.

   See Download Subscriptions Page.

2. Create a new environment template using the latest downloaded PeopleSoft Update Image and PUM full-tier topology.

   Note: A PeopleSoft administrator must enable user access to the newly created PUM source environment template.

   See Creating a Template.

3. Click the **Create Environment** button on the Environments landing page.
4. Enter the required environment attributes inputs.
5. Select the PUM source environment template to deploy.
6. Click **Done**.

   See Creating an Environment

Adding Targets to PUM Sources

After the PUM source environment is deployed and is in a running state, you can add the target database to the PUM source. Use the Manage PUM Connections page (ECL_SA_MANAGEPM_FL) for setting up environments for selective adoption. This page appears only for environments that were deployed using a PeopleSoft Update Image and that have a PeopleSoft Client (Windows Client) as part of the environment.

You can manage target databases for the PUM Source from this page, which will add or remove specified target databases to the PUM source environments. After adding target databases, administrators can use the PIA URL shown on this page to access PUM Dashboard to define change packages. To create and apply change packages, access Change Assistant that is installed on the PeopleSoft Client. You can access the Windows VM with the PeopleSoft Client using remote desktop (RDP).

1. Click the Environments tile available on the Cloud Manager home page.
2. Click the **Related Actions** button corresponding to the PUM source environment.
3. Navigate to the Environment Details page.
4. Select the **Manage PUM Connections** link available on the left panel of the Environment Details page. The Manage PUM Connections page is displayed as shown.

   This example illustrates the fields and controls on the Manage PUM Connections page.

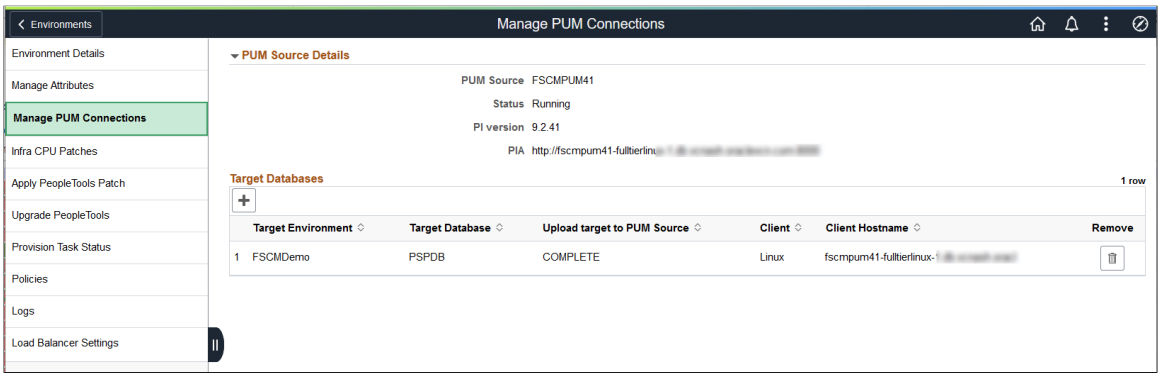

5. Click the **Add target** button to add any environment of the same application type as the PUM source. This displays a modal window for selecting a target database as shown.

   This example illustrates the fields and controls on the Select Target modal window.

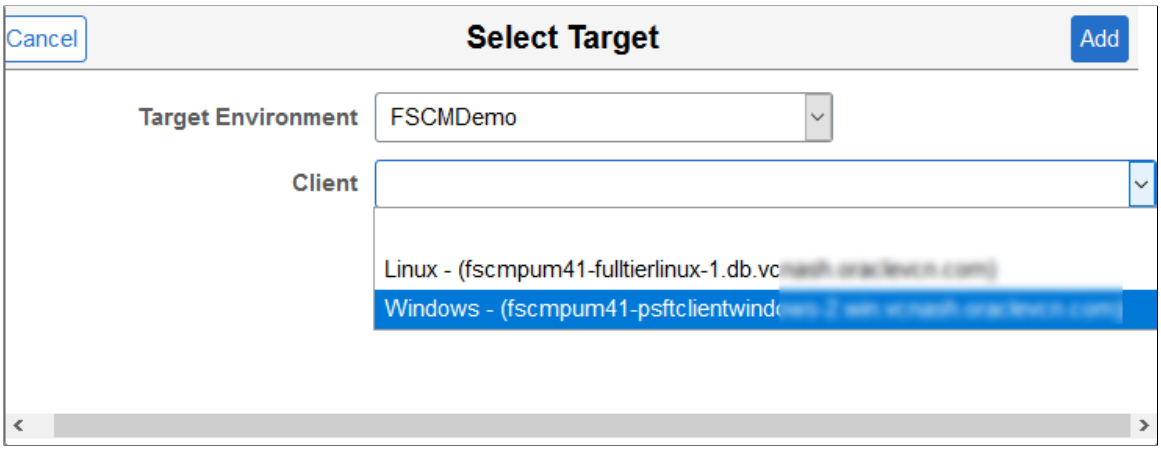

6. Select a target environment.
7. Select the client.
8. Click the **Add** button to add the target database.

Adding the target database takes a few minutes to complete. The target database is configured in Change Assistant and the target database information is uploaded to the PUM source database. The status is displayed as *In Progress* when the job to add the target is running. The status is changed to *Completed* if the target is added successfully, and to *Failed* if the job did not run successfully.

Accessing Change Assistant in Windows Client

Change package can be defined, created, and applied to target environments using the Change Assistant and the PUM source PIA.

To access Change Assistant, perform the following:

1. Determine the IP address or hostname of the PeopleSoft Client that was deployed as part of the PUM source environment from the Environment Details page.

   The IP address and Oracle Cloud name is displayed in PeopleSoft Client section of the Environment Details page.

2. Connect to the Windows Client using remote desktop connection.

3. To apply PRPs to PUM Source environment, you need to copy the downloaded PRPs from the file repository to the Windows Client VM. All downloaded PRPs are accessible to Windows Client VM as a samba share. To access the PRP share on the Windows VM, perform the following:

   - RDP to Windows Client VM.
   - Connect to the samba share using \\<File_Server_IP>\PRP.
   - Copy the required PRPs to D:\psft\pum_download directory on the Windows Client.
   - Use Change Assistant to apply the copied PRPs to the PUM Source environment.

4. Follow the standard selective adoption procedures by:

   - Applying PRPs to the PUM Source environment.
   - Defining the change package by connecting to the PUM source database.
   - Creating and applying the change package.

See the product documentation *PeopleTools: Change Assistant and Update Manager* on Oracle Help Center at https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html.

Chapter 12

Updating Cloud Manager

Updating Cloud Manager Overview

Similar to any PeopleSoft application, Cloud Manager updates are released as PeopleSoft Update Images and PRPs. Cloud Manager updates are available as part of Interaction Hub Update Images and corresponding PRPs. These updates can be applied either using an automated method or manually using selective adoption.

- Automatically Applying Updates using Manage Updates

  See Automatically Applying Updates Using Manage Updates.

- Manually Updating Cloud Manager from N-3 Version

  See Manually Updating Cloud Manager from N-2/N-3 Version to the Latest Version.

Note: Ensure that you back up both the boot volume and block volume of the Cloud Manager instance before applying the updates. See Using Automated Backup and Restore Utility. If the Cloud Manager instance is on a previous version (update image 8 or older), then use the manual method to back up, see Manually Backing Up and Restoring Cloud Manager Using Block Volume Backups for OCI.

Automatically Applying Updates Using Manage Updates

The Automated Cloud Manager Update feature facilitates automatic self-update to the latest Cloud Manager Update Image and automatic PeopleTools update or upgrade when required for the latest Cloud Manager Update Image. This feature is also used to apply PRPs to the Cloud Manager environment.

The update process will:

1. Provision a new PUM Source instance and a Windows Client
2. Apply PRPs (if any) on the PUM Source
3. Upgrade PeopleTools
4. Install and Configure Change Assistant on Windows Client
5. Define a Change Package
6. Create Change Package
7. Apply Change Package
8. Reboot domains (as needed)

Whenever there is a new Interaction Hub (IH) Update Image, or new PRPs are available, Cloud Manager will show a notification on the Cloud Manager Settings tile on the home page, and on the Manage Updates page. You need to click on the Apply button, which will ask for a set of credentials and spin up an IH PUM Source instance. Once the PUM source is up and running, the new updates will be applied to the Cloud Manager instance.

This example illustrates the fields and controls on the Manage Updates page showing the status of update steps. You can find definitions for the fields and controls later on this page.

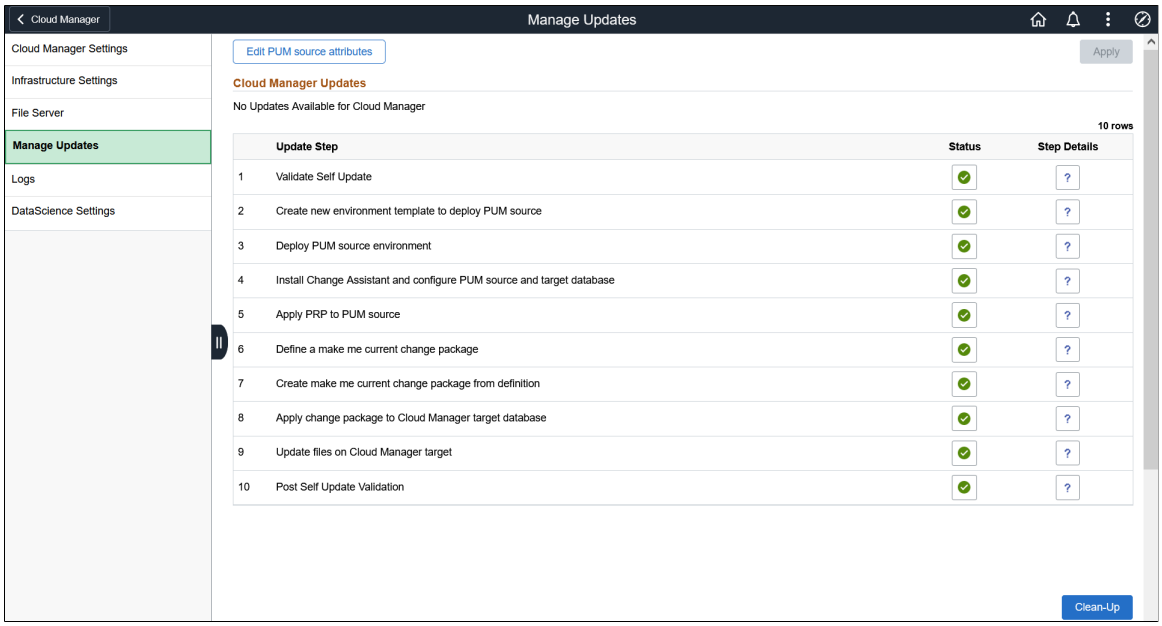

Preparing for Automatic Self-Update

You need to perform the following steps prior to triggering the update:

- Subscribe to the Interaction Hub (IH) download channel. The latest updates for IH must be downloaded before starting the upgrade process.

  See Download Subscriptions Page.

  Important! The latest updates for the IH channel will be downloaded based on the download interval (see Changing Download Interval). To download the updates immediately, unsubscribe and re-subscribe to the IH download channel.

- Subscribe to PeopleTools version of the IH PUM source.

  Note: The Update Image manifest for the IH PUM source image will list the PeopleTools version. See the PeopleSoft Update Manager (PUM) Home Page, My Oracle Support, Doc ID 1641843.2.

- Ensure to take a backup of any customizations done on these permission lists:
  - PACL_001
  - PACL_002

- PACL_003

  These permission lists are applied to the target Cloud Manager as part of the automatic self-update process and the existing instances of these permission lists are overwritten. Back up any customization done on the permission lists before starting the automatic self-update process, and re-apply them after the update is complete.

- Ensure that the PUM Fulltier topology is available and that shapes have been specified for the Full Tier and PeopleSoft Client nodes.
- Ensure a Windows Image is available in your account.
- Configure Windows Image OCID in Cloud Manager Settings page.
- Inform users that Cloud Manager will not be available during the upgrade.
- Ensure to take a backup of Cloud Manager before updating.

Note: If the Cloud Manager update is initiated with jobs currently running, those jobs may fail. The administrator must clean up and resubmit any jobs that failed.

Starting the Automatic Self-Update

To trigger automated Cloud Manager application update, perform the following steps:

1. Log in to Cloud Manager as a user having PACL_CAD user role.
2. Click the Cloud Manager Settings tile, then select Manage Updates.
3. To review the details of the PRPs listed, click the right arrow (>).

   This example illustrates the fields and controls on the Manage Updates page showing that updates are available.

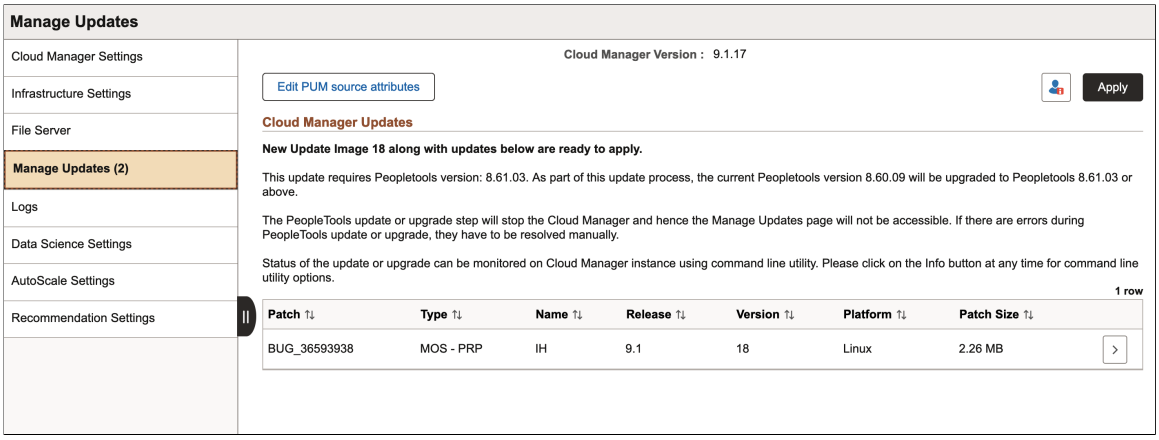

4. Click the Edit PUM source attributes button to input credentials that will be used to deploy a PUM Source environment.

   The fields are defined in Updating Custom Attributes.

Note: The Database Operator Id field value should always be set as VP1.

This example illustrates the fields and controls on the Environment Update Attribute page.

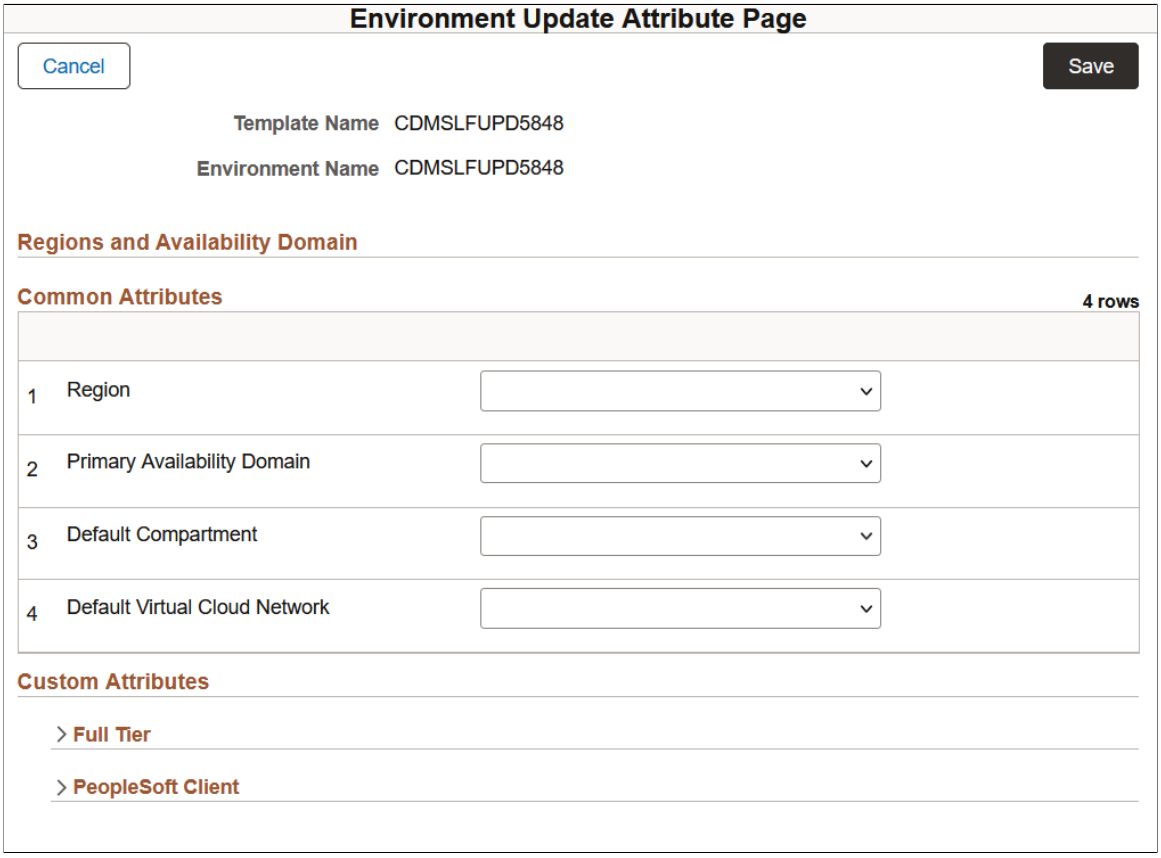

Note: Starting with Cloud Manager Update Image 12, the Virtual Cloud Network for Full Tier and PeopleSoft Client will default to the Default Virtual Cloud Network specified in the Region and Availability Domain section. If VCN peering is set up, the VCN can be changed for the Full Tier or PeopleSoft Client. See the tutorial *Use Custom or Private Network Resources with PeopleSoft Cloud Manager (Optional)*.

5. Click Save.
6. If you are upgrading Cloud Manager to a new update image, along with a PeopleTools upgrade, or if any PUM source attribute is not set, the Environment Updates Attributes page appears again. Select the PeopleTools release for the upgraded Cloud Manager, and specify any missing attributes.
7. Verify the common attributes, and the custom attributes for the Full Tier and PeopleSoft Client.
8. Click Save.
9. Review the results on the Network Validation page.

   Cloud Manager validates port access between subnets that is required during the automatic self-upgrade process. If a warning is shown, click **Cancel** to correct the issue, or click **Ignore and Continue** if you are confident that the network settings are correct.

For information on the ports that Cloud Manager requires, see the tutorial Plan the Virtual Cloud Network for PeopleSoft Cloud Manager (Optional) at https://docs.oracle.com/en/applications/peoplesoft/cloud-manager/index.html#InstallationTutorials.

This example illustrates the Network Validation page.

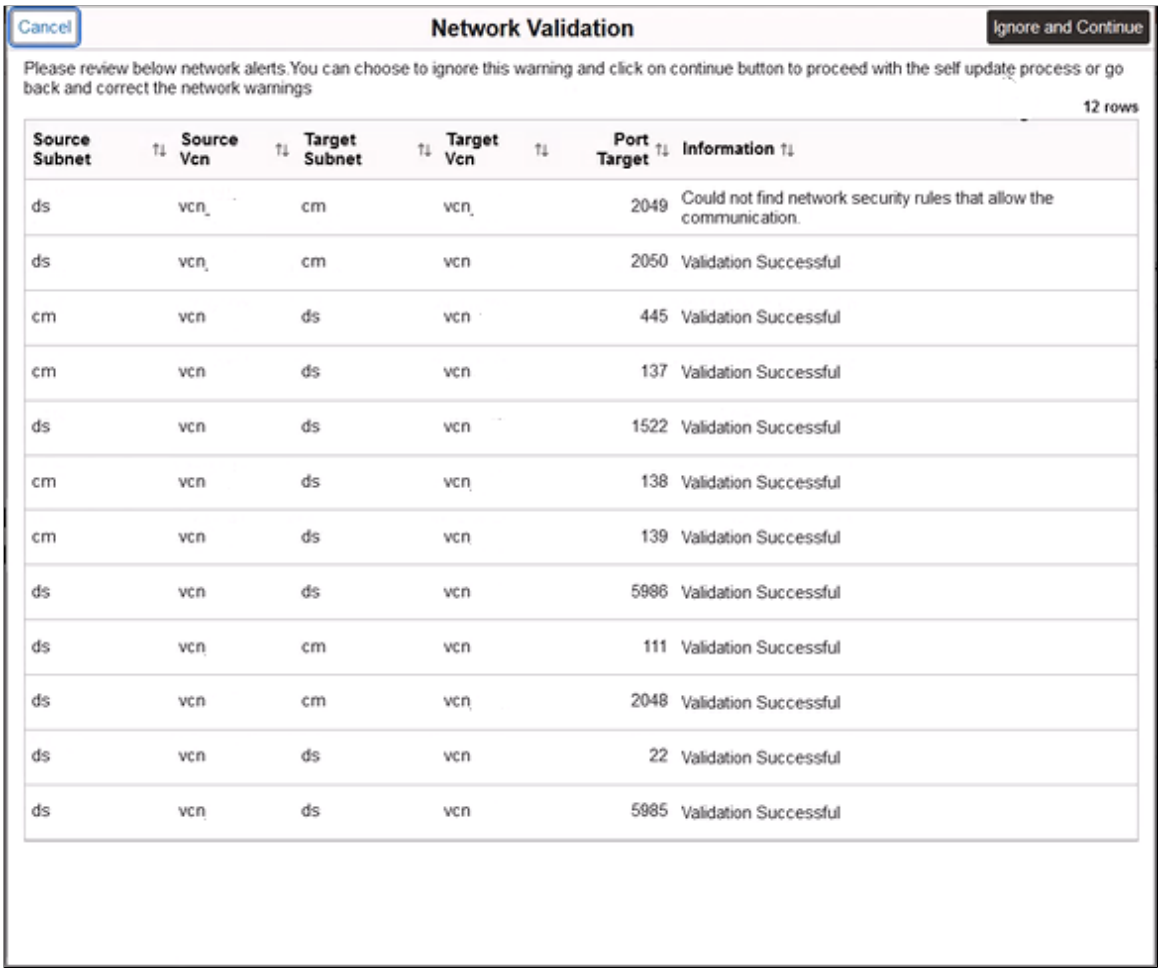

10. Click Apply to initiate Cloud Manager Update. The update steps and status are displayed.

This example illustrates the Manage Updates page showing Cloud Manager Update Steps.

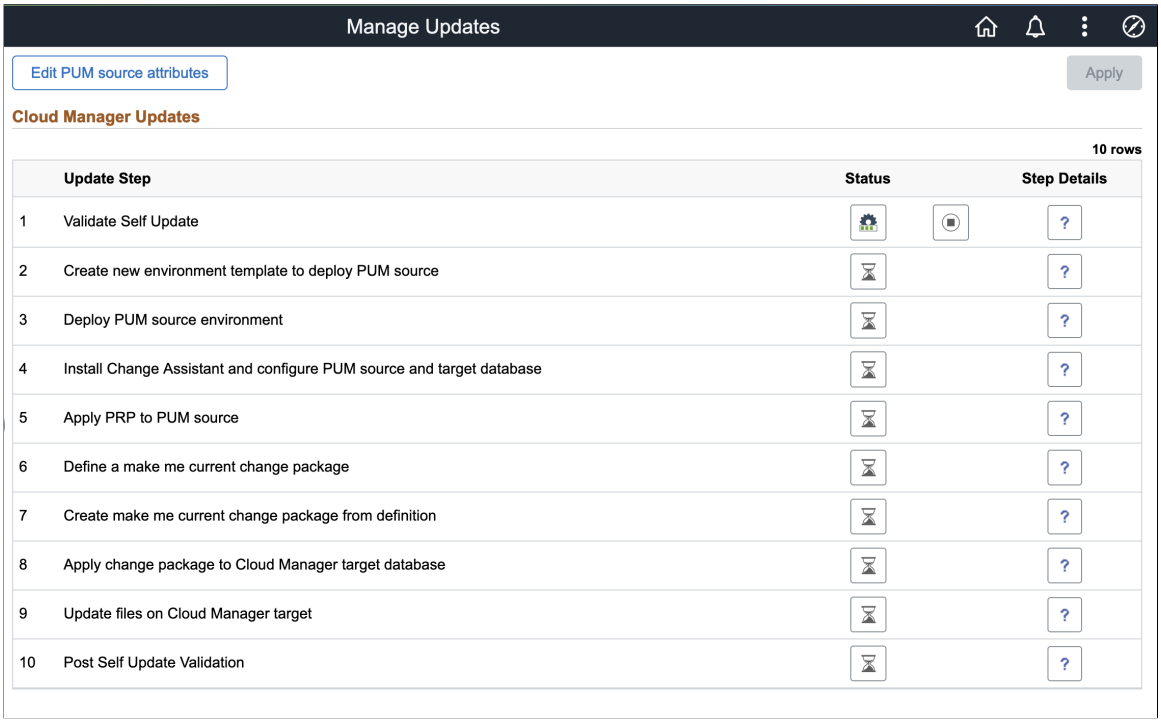

This table lists the update steps:

Step	***Description***
Validate Self Update	This step validates that all of the dependencies for the automation are available and the system is ready to begin the update process.
Create new environment template to deploy PUM source	In this step, a new environment template *CDMSLFUPD<randomly selected number>* is created that will be used to deploy PUM source environment using the latest PeopleSoft Update Image. The template name can be obtained from the Edit Attributes page. If the status is *Success* — A new template was successfully created. If the status is *Failure* — Template creation failed. In this case, the Retry is enabled. You can delete the template if it was created incorrectly and retry the step.

Step	***Description***
Deploy PUM source environment	In this step, a new PUM source environment named *CDMSLFUPD1<randomly selected number>* is created using the template that was created in the previous step. If the status is *Success* — A new PUM source is created and the details are provided in the PUM Source Details section. If the status is *Failure* — Creating a new PUM source environment failed. In this case, the Retry is enabled. Remedial Action — Clean up the failed environment and any instances from both Cloud Manager UI and Oracle Cloud Infrastructure Console that were created and retry the step. The Continue option is disabled until the clean up is complete.
Manual Step	This step is not always present. If present, this step pauses the self update and provides instructions in the task list. After completing the instructions, select to mark the task as Manually Completed.
Pre Tools update or upgrade validation	This step performs validations on the existing PeopleTools.
Install Change Assistant and configure PUM source and target database	This step carries out the processes such as installing Change Assistant on the PeopleSoft Client VM instance, configuring Change Assistant to add source and target database, and uploading target database information to PUM source. If the status is *Success* — Change Assistant is installed and configured with source and target database information. If the status is *Failure* — Failed to install or configure Change Assistant. In this case, the Retry is enabled. Remedial Action — Retry step. Alternatively choose to skip this step after configuring the source and target database manually using Change Assistant and retry.
Apply PRPs on PUM source	In this step, any PRPs that were downloaded are applied and available in Repository on the PUM source. If the status is *Success* — All PRPs were successfully applied on the PUM source environment. If the status is *Failure* — Failed to apply one or more PRPs. In this case, the Retry is enabled. Remedial Action — Retry step. Alternatively choose to skip this step after manually applying all PRPs using Change Assistant. The required PRPs will be available on the PeopleSoft Client VM, if not copy from File Server PRP share.

Step	*Description*
Define make me current change package	In this step, a new change package is defined. If the status is *Success* — Successfully defined a change package which includes all bugs for CM product code. If the status is *Failure* — Failed to define a change package. In this case, the Retry is enabled. Remedial Action — Login to Update Manager PIA of PUM source and delete the change package definition in error and retry step. The name of the definition is in the format CMCHGPKG[n], where n is the sequence number. Alternatively, create the make me current change package definition in Change Assistant. Name the change package CMCHGPKG[n], where n is the sequence number. Mark the failed step as COMPLETED MANUALLY.
Create make me current change package from definition	In this step, a change package using the definition that was created in previous step. If the status is *Success* — Successfully created a change package. If the status is *Failure* — Failed to create a change package. In this case, the Retry is enabled. Remedial Action — Retry step. Alternatively, skip the step after creating the change package manually using Change Assistant with the same name as the definition created in previous step.
Apply change package to Cloud Manager target database	In this step, the change package that was created in the previous step is applied. If the status is *Success* — Successfully applied the change package. If the status is *Failure* — Failed to apply the change package. Remedial Action — Complete the apply step manually using the Change Assistant and continue with next step. ***Warning!*** Reapplying a change package is not recommended as it may apply the fix again.
Update files on Cloud Manager target	In this step, the new and updated files are copied to Cloud Manager target. If the status is *Success* — Successfully copied all file updates. *Failure* — Failed to copy one or more files. In this case, the Retry is enabled. Remedial Action — Retry step.

Step	*Description*
Post Self Update Validation	This step validates that the update is complete and the newly created domains are running.

Updating Custom Attributes

You must update the Full Tier and PeopleSoft Client Settings as part of the Custom Attributes.

This example illustrates the fields and controls in the Full Tier section on Environment Update Attribute page.

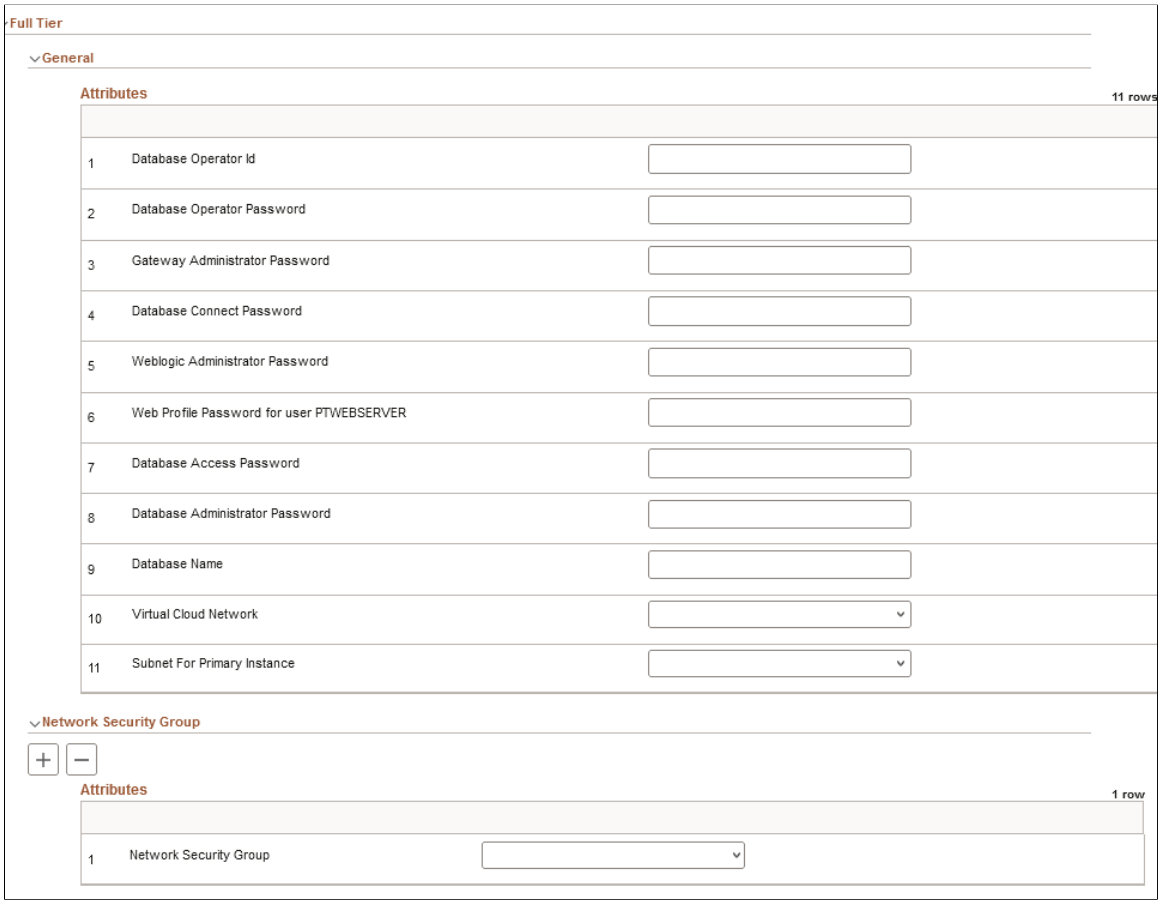

Field or Control	*Description*
Database Operator Id	Enter the database operator ID.
Database Operator Password	Enter the database operator password.
Gateway Administrator Password	Enter the password of gateway administrator.
Database Connect Password	Enter the database connect password.

Field or Control	*Description*
Weblogic Administrator Password	Enter the WebLogic administrator password.
Web Profile Password for user PTWEBSERVER	Enter the web profile password for the user PTWEBSERVER.
Database Access Password	Enter the database access password.
Database Administrator Password	Enter the database administrator password.
Database Name	Enter the database name. The database name must begin with an uppercase letter.
Virtual Cloud Network	Select the Virtual Cloud Network.
Subnet for Primary Instance	Select the subnet within the VCN for primary instance.

This example illustrates the fields and controls in the PeopleSoft Client section on Environment Update Attribute page.

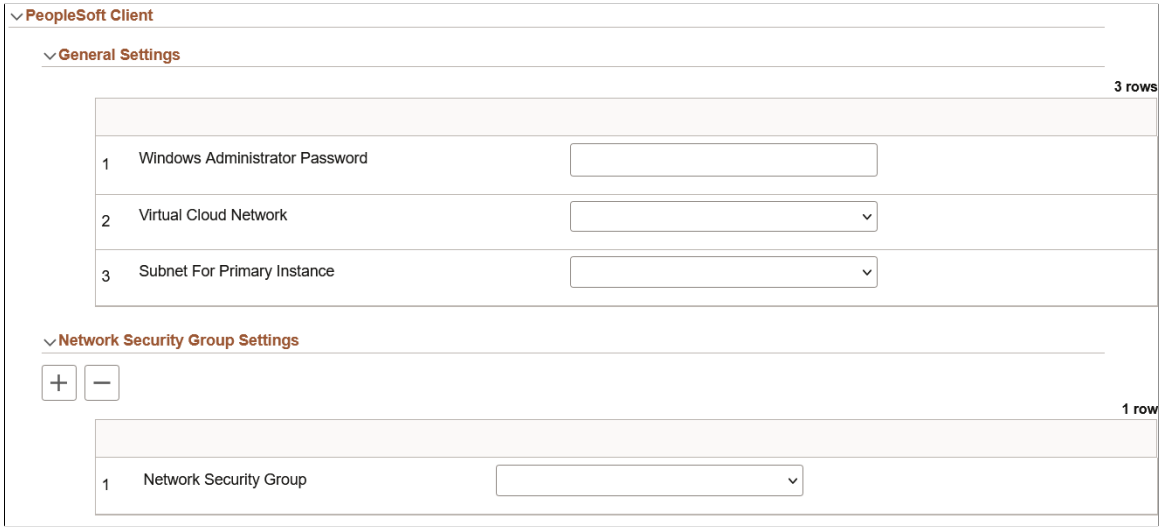

Field or Control	*Description*
Windows Administrator Password	Enter the Windows administrator password.
Virtual Cloud Network	Select the Virtual Cloud Network.
Subnet for Primary Instance	Select the subnet within the VCN for primary instance.

For configuring Network Security Group settings related to Full Tier or PeopleSoft Client, see Configuring Network Security Group Settings.

Running the Post Update Script

When the update is complete, perform the following steps:

Note: Do not run the post update script if you are only applying PRPs to the current Cloud Manager update image. The post update script is only required when updating Cloud Manager to a new Update Image.

1. Run as root the script post_upgrade_script.sh in $PS_APP_HOME/cloud/ .

   ```
   $> ./ post_upgrade_script.sh
   ```

2. Using PSADMIN, restart application server domain, process scheduler domain and web domain on the Cloud Manager instance to ensure the latest updates are running.

3. Navigate to Cloud Manager Settings > Infrastructure Setting and click the Refresh OCI Metadata button. This will update the Fault Domain Settings.

4. The PUM Source environment can then be cleaned up using the **Clean-up** button.

   Important! The post update script must be run prior to selecting Clean-up.

Monitoring Update Steps

The status of the update is displayed on the Manage Updates page.

This example illustrates the fields and controls on the Manage Updates page when the steps are running.

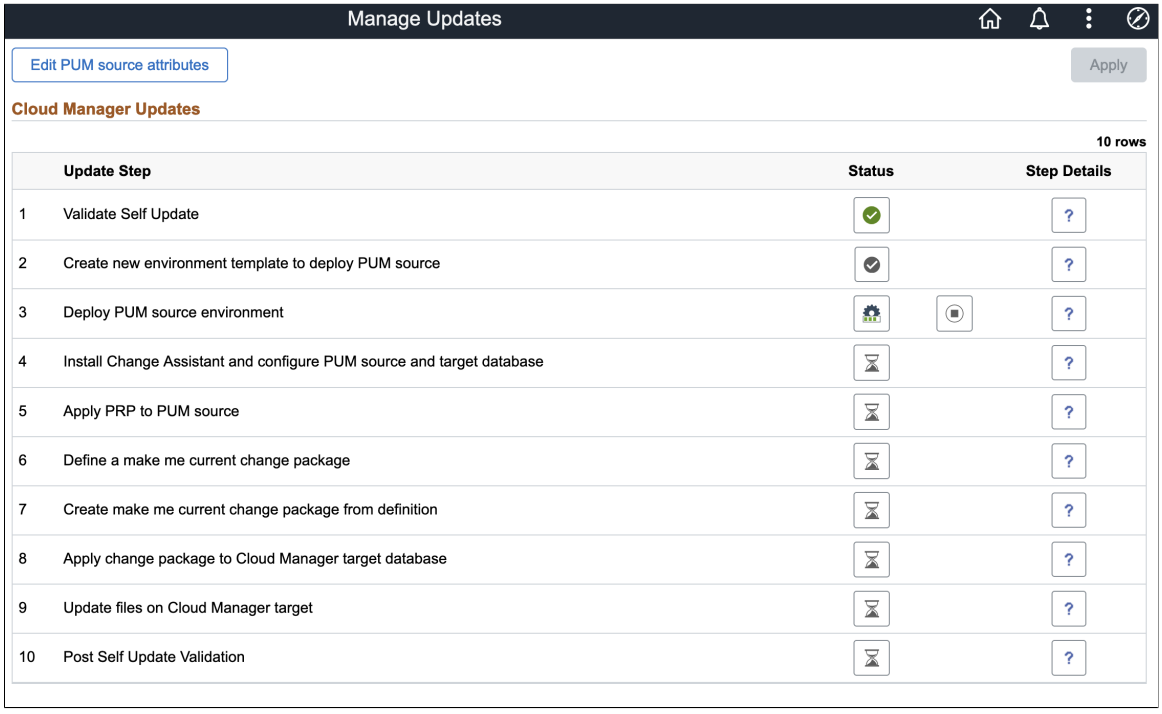

Status Icons

Field or Control	*Description*
	Pending
	Success
	In Progress
	Failed
	Continue
	Abort
	Step details
	Pause

The Abort icon is shown when a step is running. When you click the Abort icon, the modal window is displayed.

This example illustrates the Abort Modal window.

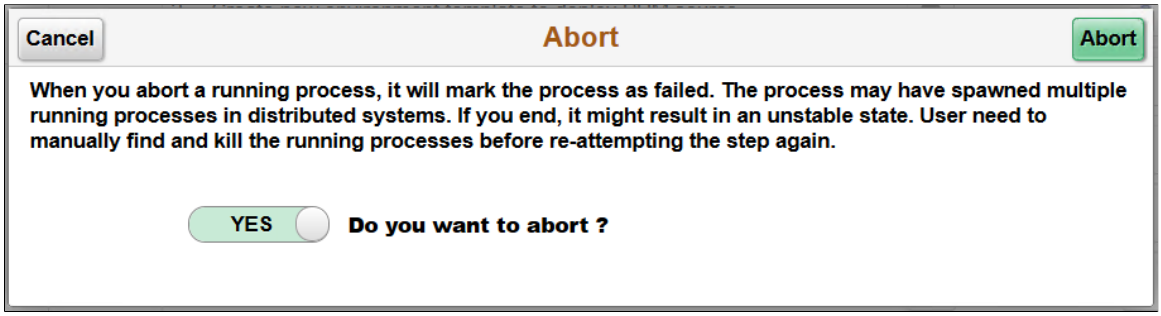

Caution should be used when aborting a step, it is possible that not all processes that were spawned will be aborted. It is recommended to reboot the Cloud Manager instance after aborting a process.

Failed Steps

If a step fails, the process will stop and the Continue icon will appear.

An administrator must resolve the issue and come back to Cloud Manager to continue the update process. For example, if applying change package failed, then the administrator must connect to the Windows Client VM, launch Change Assistant and run the update job to completion. The administrator will then return to the Cloud Manager Update page and continue the automated update process.

All reported errors must be resolved manually by the user. After fixing or manually completing the failed step, click the Continue icon.

This example illustrates a failed step in the Cloud Manager Update process.

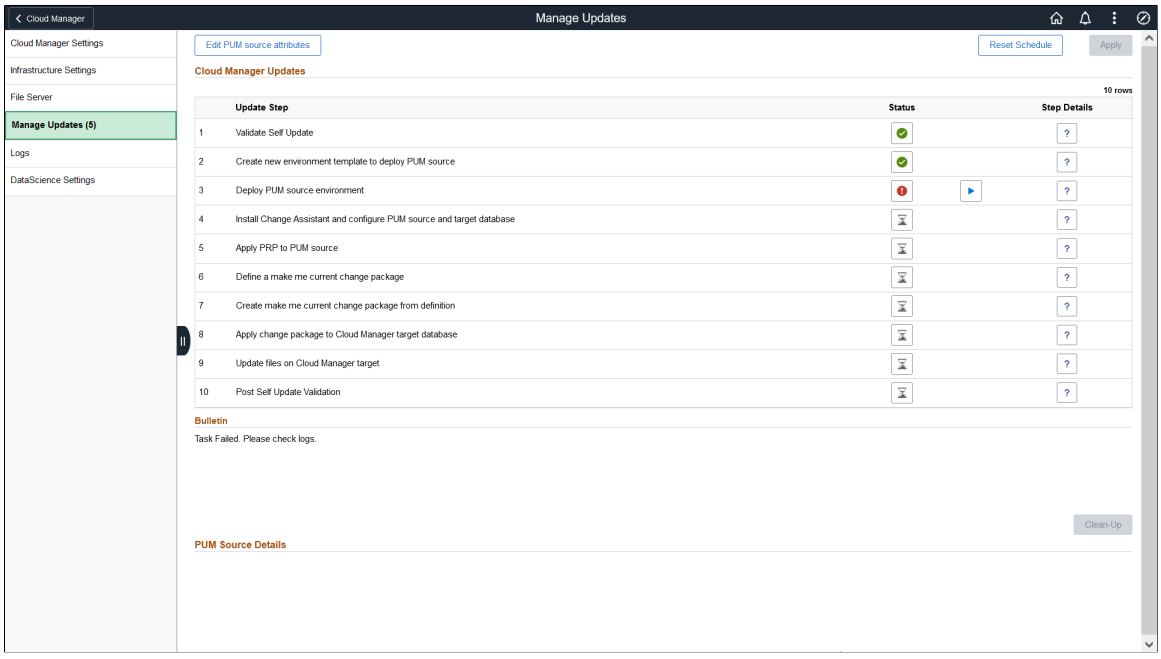

On clicking the Continue button, three options are shown as below:

This example illustrates the options for the failed step. You can find definitions for the fields and controls later on this page.

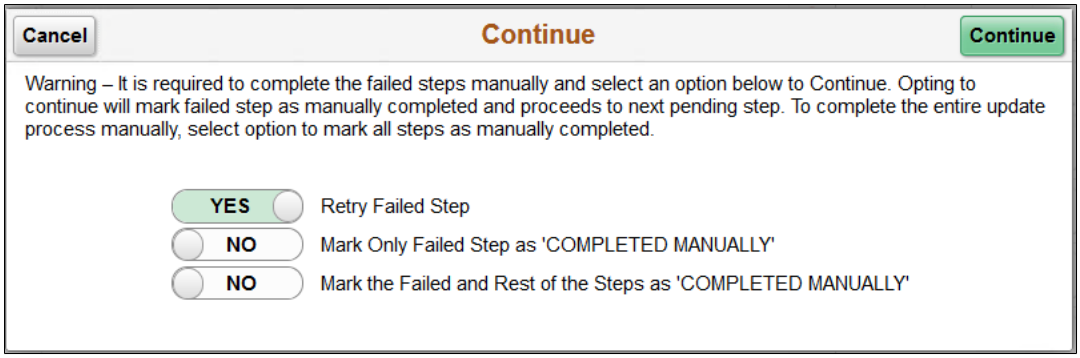

- Retry Failed Step — retry the step again.
- Mark only failed step as 'Completed Manually' — skip the failed step and continue from subsequent step to completion.
- Mark the failed and rest of the pending steps as 'Completed Manually' — skip all steps and set update as complete.

After selecting **Yes** in the Retry Failed Step field, a Continue button is displayed in the top right corner of the Continue modal window.

Monitoring PeopleTools Upgrade

The automated Cloud Manager update process will trigger PeopleTools upgrade when the newly released Cloud Manager (IH) Update Image has a dependency on a particular PeopleTools version. During the update process, when it reaches the step to upgrade PeopleTools, the Cloud Manager PIA user interface will be shut down and the status of PeopleTools upgrade is no longer available on the user interface.

To monitor the PeopleTools upgrade, you need to SSH into the Cloud Manager instance and use the PeopleTools upgrade command line.

The Information icon, at the top right on the Manage Updates page, displays the command line options to use for monitoring the PeopleTools upgrade. This page also appears when you initiate the upgrade.

This example illustrates the fields and controls on the Updating Cloud Manager modal window.

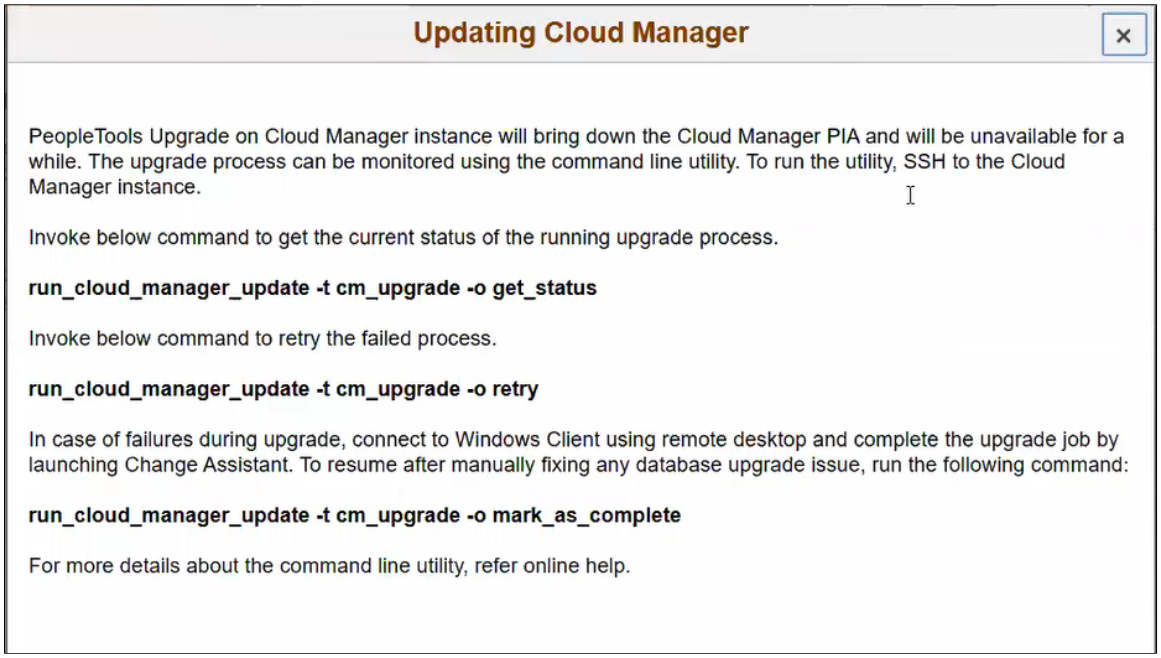

See Command Line Operations for cm_upgrade for additional information.

Manually Updating Cloud Manager from N-2/N-3 Version to the Latest Version

It is highly recommended to update Cloud Manager as soon as a new update image is available. Cloud Manager Update Images up to three versions prior to the current update image can be upgraded to the latest update image.

The self upgrade process can be used to upgrade to the latest version of Cloud Manager (CM_N) from N-1 version. See Automatically Applying Updates Using Manage Updates.

Updating to the latest version from the Cloud Manager version before two or three releases (CM_N-3 or CM_N-2), involves the following two steps.

1. Upgrade PeopleTools version on Cloud Manager N-2/N-3 environment.

   You must upgrade the PeopleTools version first, if there are any major differences in the PeopleTools version of the CM_N-2 or CM_N-3 environment and the CM_N environment. Otherwise, you can skip this step.

   See Upgrading PeopleTools Version on Cloud Manager N-2/N-3 Environment.

2. Perform PUM upgrade on Cloud Manager N-2/N-3 environment.

   For PUM Update process, we are leveraging the existing upgrade process and features in the following:

   a. Cloud Manager (Import, Add Target, Apply PRP)

b. PeopleTools Update Manager and PUM Automated Updates (PAU) features to create Change Package (CP) definition for Cloud Manager (CM) product line

c. Manual Steps/Command-line utility for CP apply and file synchronisation

See Performing PUM Update on Cloud Manager N-2/N-3 Environment.

Upgrading PeopleTools Version on Cloud Manager N-2/N-3 Environment

This example illustrates the PeopleTools Upgrade process.

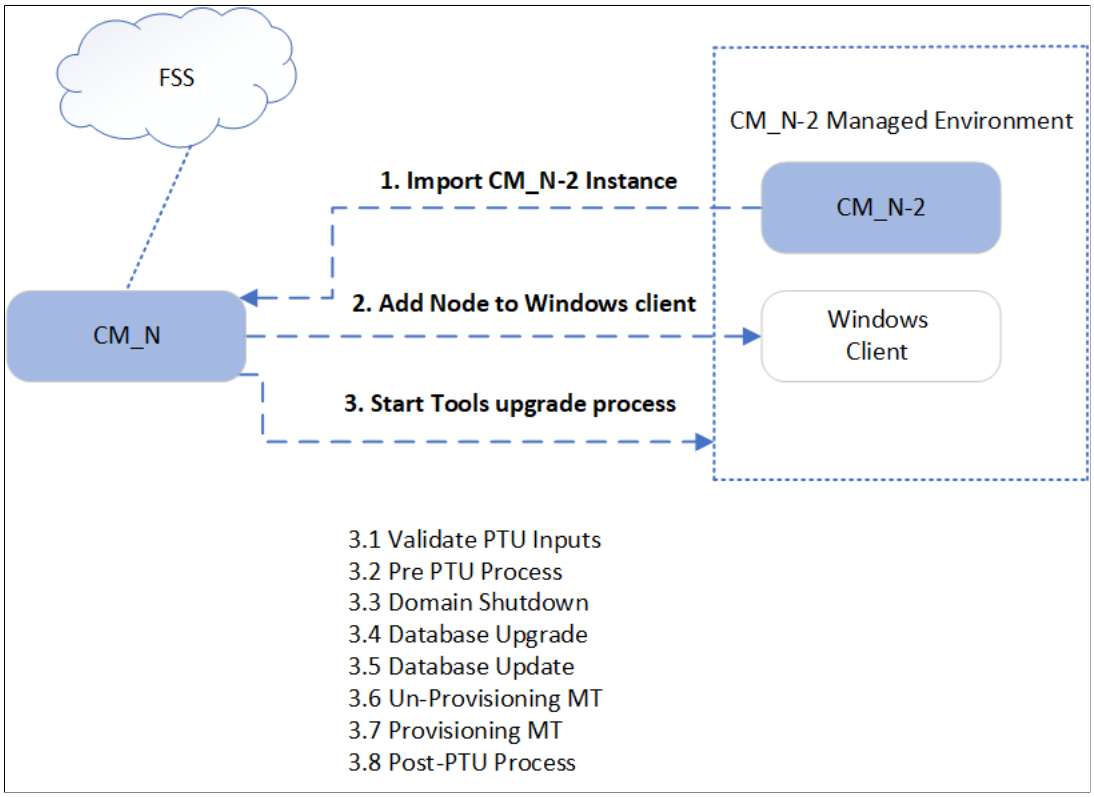

To upgrade the PeopleTools version on Cloud Manager, perform the following steps:

1. Create a new Cloud Manager host instance (CM_HOST) from a Cloud Marketplace update image, named as CM_N, where N is the PI version number. For example, CM_17.

2. Configure File Server and Infrastructure settings on CM_HOST. Ensure that PeopleTools channels, Integration Hub (IH), source PeopleTools version and target PeopleTools version channels are subscribed on CM_HOST.

3. Import the old Cloud Manager, which is the PUM_TARGET, having PI version less than CM_N. For example, CM_15.

4. Add the Windows client node to PUM_TARGET environment so that it becomes FT+Windows client environment. This step is mandatory for performing an upgrade operation in a managed environment.

5. Upgrade PeopleTools in PUM_TARGET. Navigate to PUM_TARGET environment tile context menu > Details > Upgrade PeopleTools, and click Upgrade. This triggers the PUM_TARGET environment

tools upgrade operation. The latest IH PUM DPK tools version is used as *Upgrade to* drop-down value and Report Notification option is disabled.

This example illustrates the Upgrade PeopleTools page.

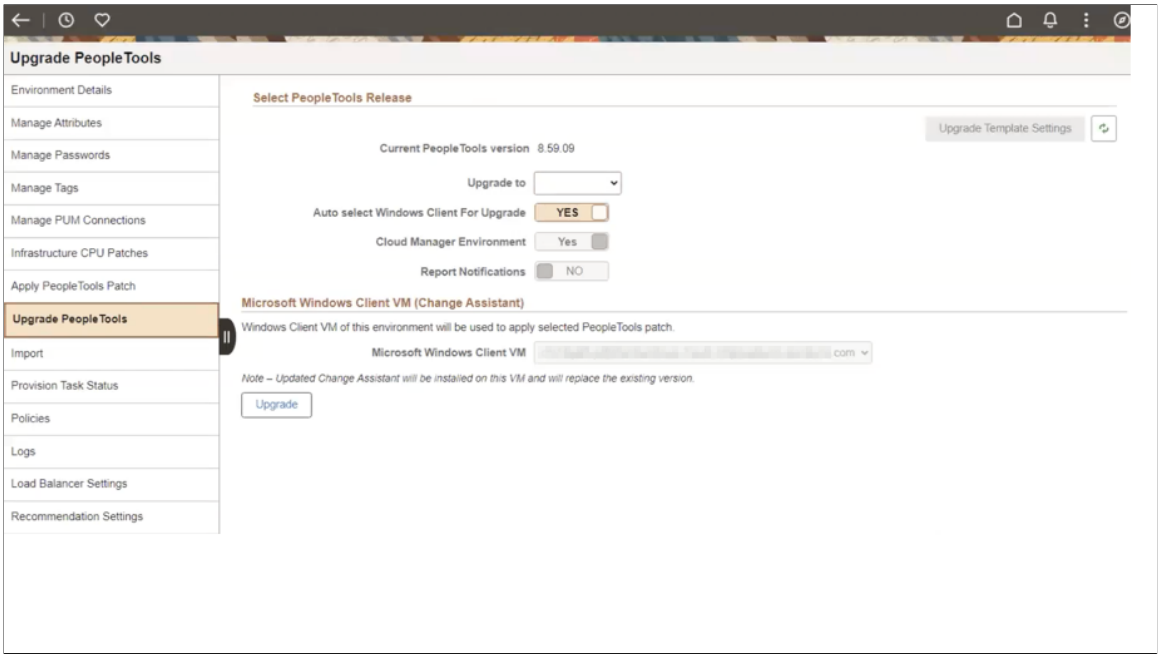

6. Click the Status link in Job grid, which displays the details of the tools upgrade process and current status.

   Three new tasks are added for Cloud Manager PTU process:

   a. Validate PTU Inputs

   b. Pre-PTU Process

   c. Post-PTU Process

This example illustrates the Job Status Information displayed on clicking the Status link.

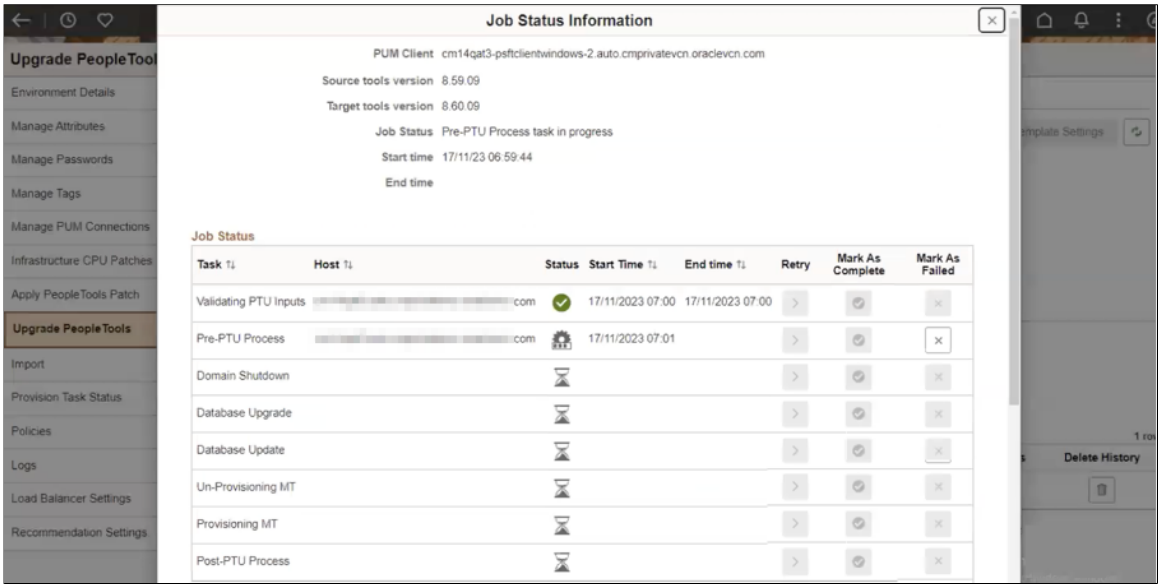

The environment status changes to Applying PeopleTools Upgrade. Once the Upgrade is complete, the status returns to Running state.

This example illustrates the environment status showing that PeopleTools Patch is applied.

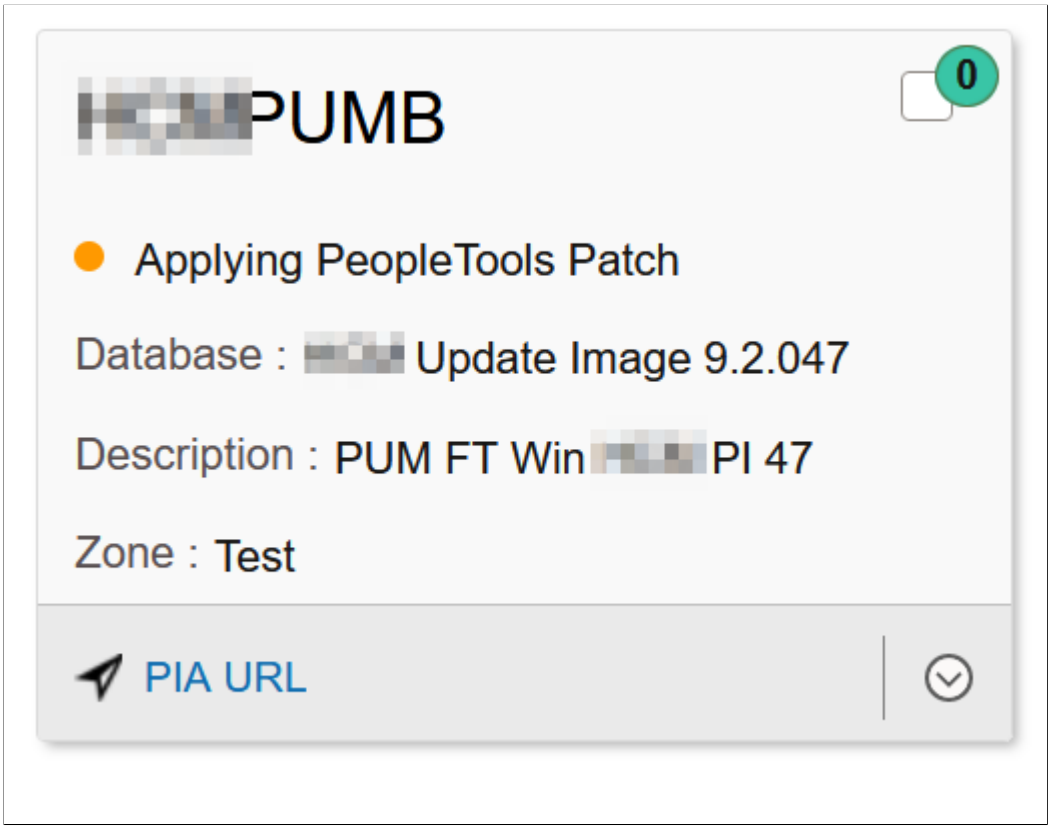

Performing PUM Update on Cloud Manager N-2/N-3 Environment

After you upgrade PeopleTools to the latest version, you must update the PUM version (update image version) of PUM_TARGET.

This example illustrates the PUM Upgrade process.

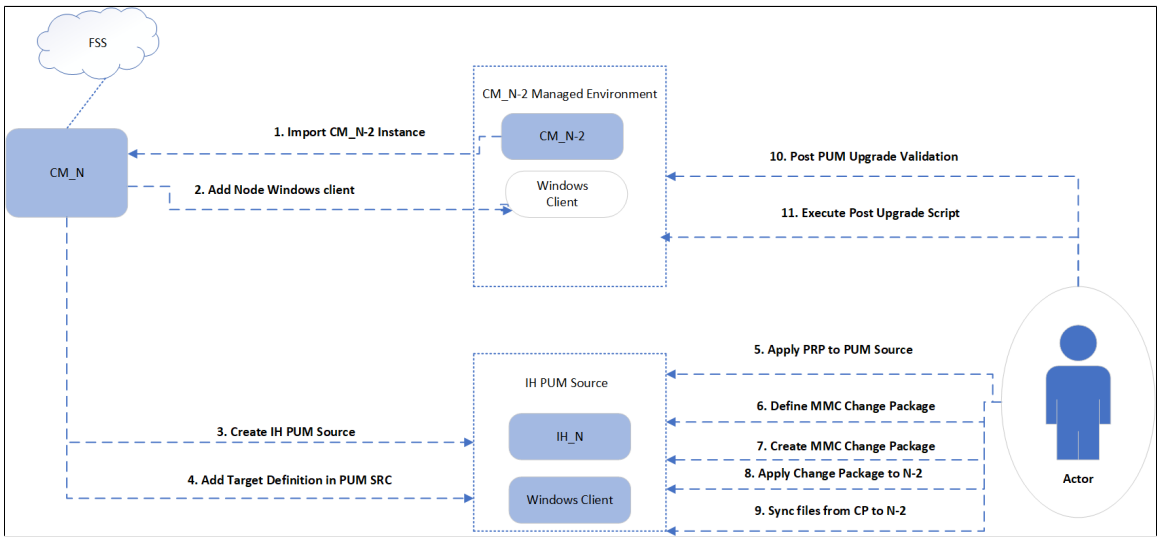

To update Cloud Manager PUM/ Application, perform the following steps:

1. Create a new Cloud Manager host instance (CM_HOST) from Marketplace update image, named as CM_N, where N is the PI version number. For example, CM_17.
2. Configure File Server and Infrastructure settings on Cloud Manager host (CM_HOST). Ensure that PeopleTools channels, Integration Hub (IH), source PeopleTools version and target PeopleTools version channels are subscribed on CM_HOST.
3. Import the old Cloud Manager version, which is the PUM_TARGET, having PI version less than CM_N. For example, CM_15. CM_N-2 or CM_N-3 environment is used as PUM target (PUM_TARGET).
4. Create a new PUM Source (PUM_SOURCE) by provisioning a new IH environment with latest DPK. For example, IH_17. The environment must have a Full Tier instance and a Windows client.
5. Apply PRP to PUM_SOURCE instance using any of the following ways:
    - If the PUM_SOURCE is newly created, all the PRPs will be applied as part of provisioning use case.
    - Configure the PRP-Apply policy in PUM_SOURCE environment template creation. This will make sure all new PRPs will be applied to PUM_SOURCE automatically.
    - Create a new PRP-Apply policy using Policy Editor, which will make sure all the PRPs will be applied to PUM_SOURCE.
    - Manually log in to PUM_SOURCE PIA and use the PUM Automated Updates (PAU) feature for applying the PRPs to PUM_SOURCE.

6. Add PUM_TARGET environment as the Target database by navigating to PUM_SOURCE (IH_17) environment tile > Details > Managed PUM Connections. Only Windows CA support for PUM update.

7. Create a change package definition for CM product line in PUM_SOURCE by making use of Change Assistant Update Manager feature. To create a change package:

   a. Log on to PUM_SOURCE PIA.

   b. Confirm the target database name in Update Manager Tile > Confirm Select Target Database.

   c. Navigate to Define Change Package > Update Change Package > All Other Search Criteria > "Product/Severity/Image" Link > Product Family and select "Cloud Manager". See the topic *Defining an Upgrade Package* in *PeopleTools: Change Assistant and Update Manager*.

   This example illustrates selecting Cloud Manager while defining Change Package.

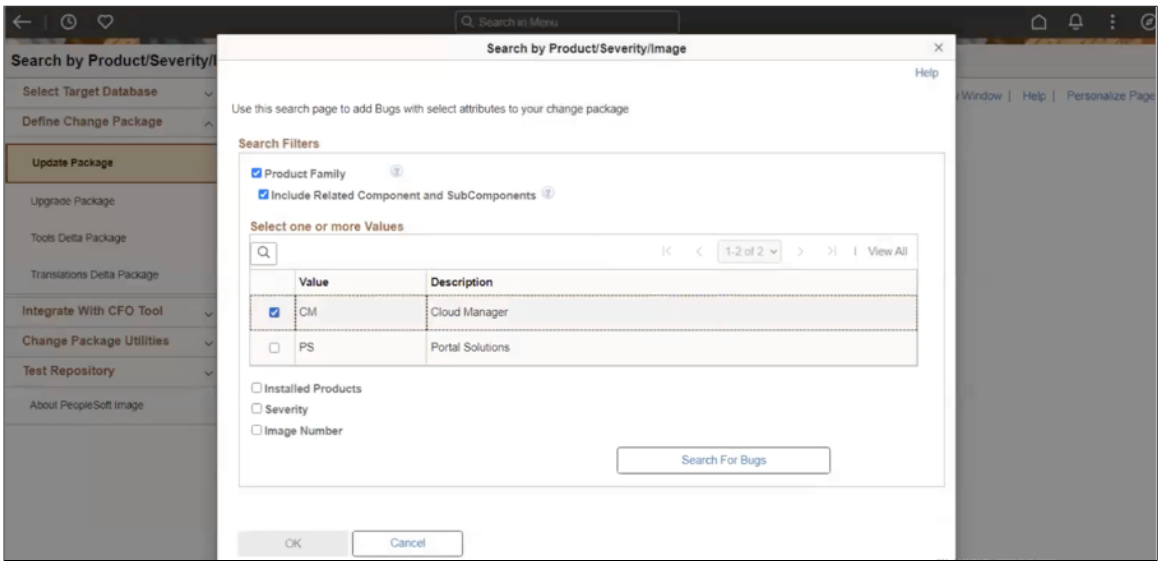

8. Create a change package from PUM source definition by making use of PeopleTools PAU feature. See the product documentation for *PeopleTools: Change Assistant and Update Manager*.

   a. Log on to PUM_SOURCE PIA.

   b. Update Lifecycle roles in Administrator metadata in PUM_TARGET environment by navigating to PeopleTools > Security > User Profile > User Profile (User ID: CLADM) > LifeCycle Tools Roles tab.

   c. Update "Update Settings" information.

   d. Update Scheduler > Define Job > Enable.

      a. Define change package (Select "Search Scope" as "Previously Defined").

      b. Create package.

      c. Apply package.

e. Save and run the job.

This example illustrates the Job Monitor used to view the status of the job.

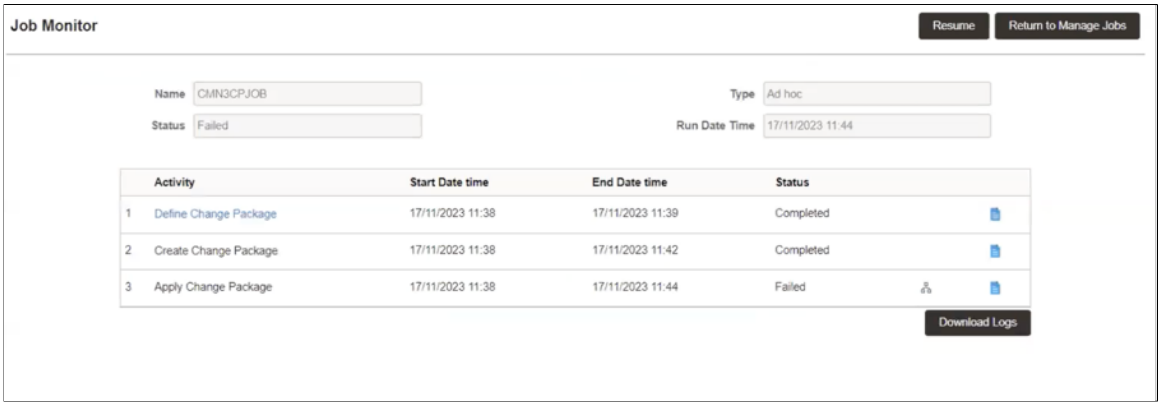

This example illustrates retrieving the change package name by clicking the log icon. Here the name of the change package is CMCP101_20231011.

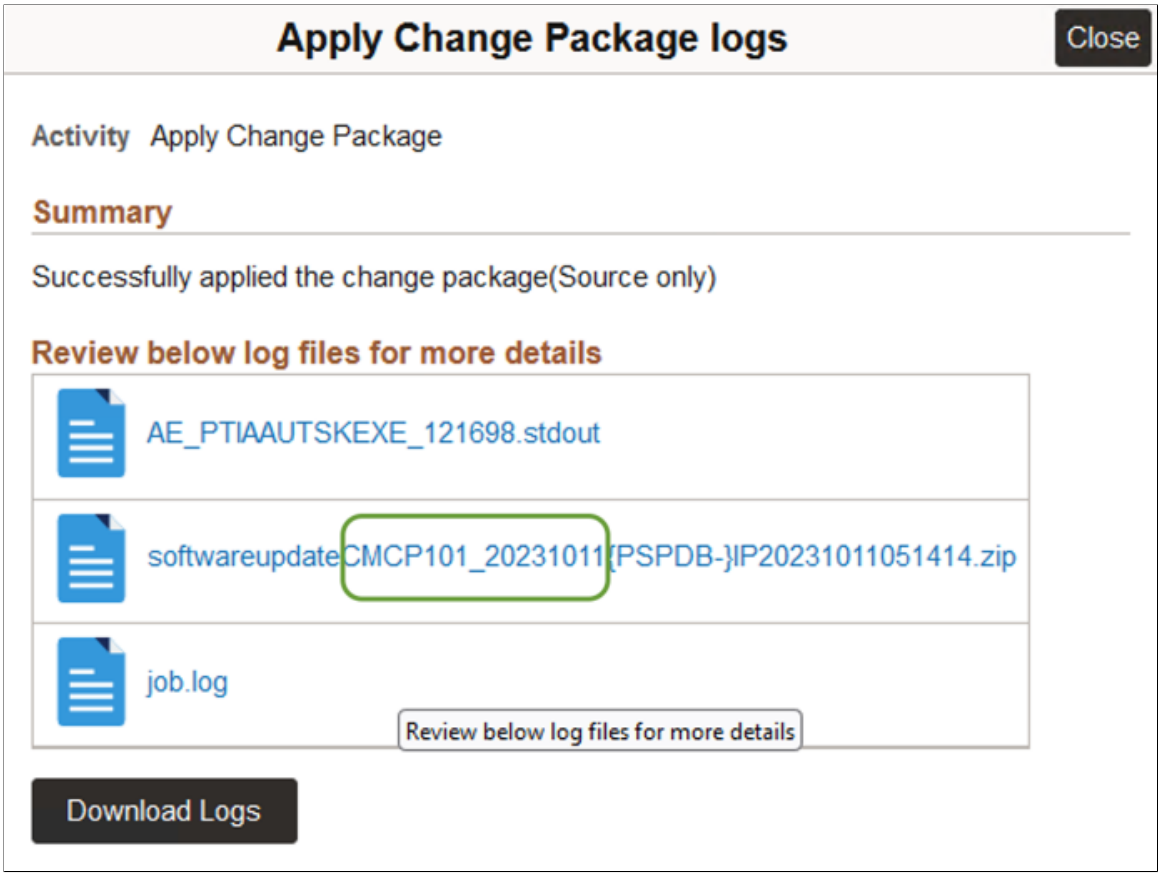

Cross check if the server configuration is applied correctly.

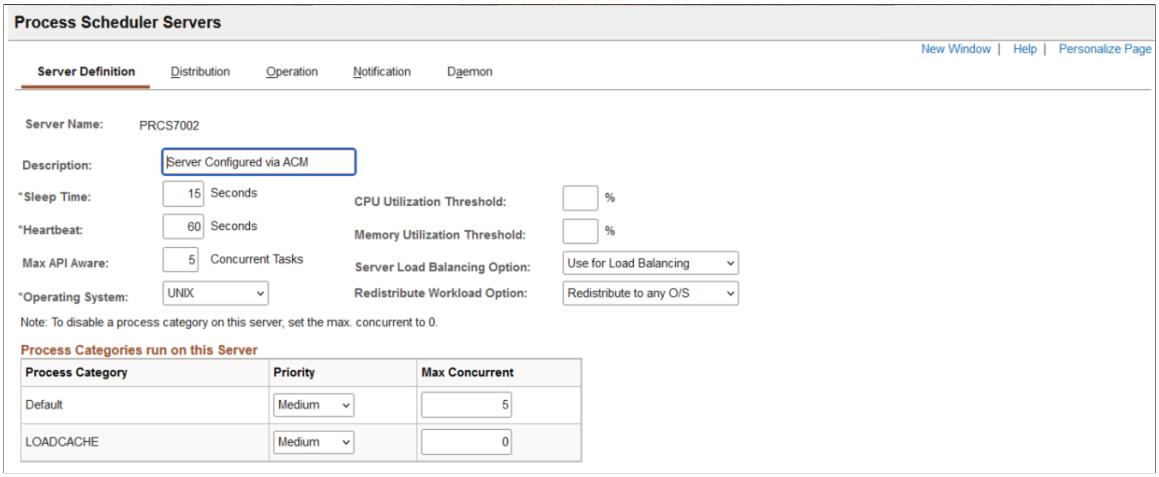

This example illustrates checking the process scheduler server configuration.

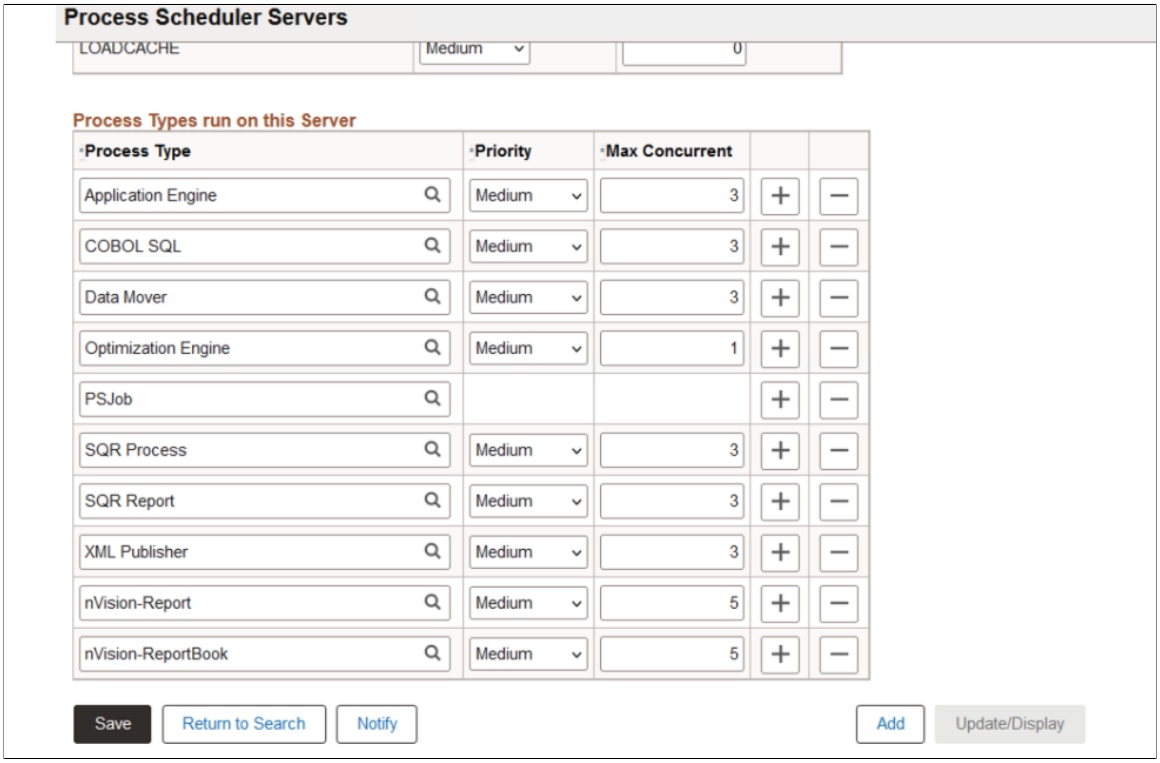

9. Apply change package to Cloud Manager target database and synchronise all change package files in the PUM_TARGET. The task pum_upgrade is used to perform a PUM Upgrade in Cloud Manager instance. The input values are given as response file. The format is:

```
run_cloud_manager_update -t pum_update  -o <operation_name> -r <response_file.⇒
json>
```

See the table provided in step 11 for operation/subtask details.

```
response_file.json:
{
```

```
  "file_server": {
    "mount_path": "<file server export mount path>"
  },
  "pum_source": {
    "windows_client": {
      "private_ip": "<windows_client private ip>",
      "remote_password": "<windows_client password",
      "psft_base": "<windows client psft base directory path>"
    },
    "full_tier": {
      "private_ip": "<full tier private ip>",
      "psft_base": "<full tier psft base directory path>",
      "tools_ver": "<pum source tools version>",
      "pi_number": "<pum source pi version>",
      "env_name": "<pum source environment name>"
    }
  },
  "pum_target": {
    "private_ip": "<full tier private ip>",
    "psft_base": "<full tier psft base directory path>",
    "db_name": "<pum target db name>",
    "tools_ver": "<pum target tools version>",
    "pi_number": "<pum target pi version>",
    "env_name": "<pum target environment name>",
    "access_id": "<access id>",
    "access_password": "<access password>"
  },
  "pum_update_data": {
    "change_package_name": "<change package name in pau"
  }
}
```

Run the following command to get the status:

```
$ run_cloud_manager_update -t PUM_UPDATE -o get_status
```

Run the following command to get more details and log information:

```
$ run_cloud_manager_update -t PUM_UPDATE -o get_status -v
```

The PUM_UPDATE task carries out the following steps:

a. Validate the PUM Update input values in response file.

b. Copy the change package created using PAU process (from PUM_SOURCE FT) to the PUM_SOURCE Windows client computer.

c. Apply the change package DB changes to PUM_TARGET.

d. Synchronise change package files to PUM_TARGET:

   a. Copy the change package files from PUM_SOURCE Windows CA client to PUM_TARGET.

   b. Synchronise CP files to PS_APP_HOME.

   c. Restart PSFT domains.

10. Validate in PUM_TARGET post self update by checking that the UI is working fine in PUM_TARGET.

11. Run the post-upgrade script in Cloud Manager PUM_TARGET. Run the post_upgrade_script.sh as the root user. If all the information is correct on the upgrade configuration details, you can start the post-upgrade process by providing 'y' as input.

Task Name	*Operation Type*	*Sample Command*
PUM Update	execute	run_cloud_manager_update -t pum_update -o execute -r /tmp/pum_update.json
	retry	run_cloud_manager_update -t pum_update -o retry -r /tmp/pum_update.json run_cloud_manager_update -t pum_update -o retry
	mark_as_complete	run_cloud_manager_update -t pum_update -o mark_as_complete -r /tmp/pum_update.json run_cloud_manager_update -t pum_update -o mark_as_complete
	mark_all_step_complete	run_cloud_manager_update -t pum_update -o mark_all_step_complete -r /tmp/pum_update.json run_cloud_manager_update -t pum_update -o mark_all_step_complete
	get_status	run_cloud_manager_update -t pum_update -o get_status run_cloud_manager_update -t pum_update -o get_status -v

Upgrading Cloud Manager PeopleTools Using Command Line

The Cloud Manager command line option can be used for triggering the PeopleTools upgrade and monitoring the upgrade status. This command line utility can be used to upgrade PeopleTools for Cloud Manager instances that are on PeopleTools below 8.58. If self update was used (Cloud Manager Update Image 10 to Update Image 11), the PeopleTools upgrade is included in the process. For Cloud Manager 9 and below, a PeopleTools upgrade is required.

Note: If you are running Cloud Manager Update Image 8 or lower, you must upgrade to Cloud Manager Update Image 11 using selective adoption to obtain the command line utility for automated PeopleTools upgrade.

Note: If you are running Cloud Manager 12 or above, the tools upgrade process is automatically triggered as part of "Automatically Applying Updates Using Manage Updates". Use the command line for checking the status and retrying the failed step, see Monitoring PeopleTools Upgrade.

To use the command line utility to upgrade PeopleTools:

1. Subscribe to the PeopleTools download channel. Ensure that the download is complete.

   For example: Interaction Hub PI 9.1.11 has PeopleTools version 8.58.03, therefore you must subscribe to PeopleTools 8.58 Linux download channel with a minimum patch version of 8.58.03. To determine the PeopleTools version on the IH PUM Source, SSH into the IH PUM Source, switch user to psadm2 and run **psadmin -v** command. See Accessing Provisioned Environments

2. Log into the Cloud Manager instance using SSH.

3. Create a response file.

   See Creating Response File

4. Back up Cloud Manager instance. See Using Automated Backup and Restore Utility.

   Important! The backup will be available in case of any failures in the PeopleTools upgrade process, allowing you to restore to this point.

5. Run the command line to carry out the upgrade.

   ```
   run_cloud_manager_update -t <task_name>  -o <operation_name > -r <response_fil⇒
   e.json>
   ```

   Example: run_cloud_manager_update -t cm_upgrade -o execute -r /tmp/cm_update_rsponse_file.json

   See Command Line Operations for cm_upgrade

Understanding the Command Line

The command line has the following hierarchy:

```
Command > Task > Operation > Subtask >Activity
```

- Command

  The command is **run_cloud_manager_update**.

- Task

  For Cloud Manager Update Image 9, the only supported task is **cm_upgrade** [-t option].

- Operation

  Multiple operations [-o option] are available for a task. For cm_upgrade operations see Command Line Operations for cm_upgrade

- Subtasks

Each operation contains one or more subtasks. See Subtasks for PeopleTools Upgrade [PTU] in Cloud Manager Instance.

- Activity

  Each subtask may have multiple activities.

Command Line Operations for cm_upgrade

The task *cm_upgrade* is used to perform a PeopleTools Upgrade [PTU] in the Cloud Manager instance. The input values are given as a response file. The format is:

```
run_cloud_manager_update -t cm_upgrade  -o <operation_name> -r <response_file.json>⇒
```

This table lists the operation/subtask details.

Operation Type [-o]	***Description/Details***
execute	This operation is used for running the current task. If the response file is present it will read the response file and that will be used for creating the input values for the task implementation.
get_status	This operation is used to display the current task status in the console. The status will show the implementation status at the activity level. -v [verbose] option can be used for the detailed status. The activity status will be updated, only after the task implementation completes. The statuses supported by activities are: • PENDING/UNKNOWN • SUCCESS • FAILURE

Operation Type [-o]	*Description/Details*
retry	This operation is used for rerunning failed activities. This can be run with or without a response file. The response file should be used when the task implementation failed due to input value error and the response file has been updated. Retry will: • Check if a response file is provided. If provided the response file is used for creating input data. If the response file is not provided, the utility will get the input data from the previous or failed run. • Automatically determine the current activity name in failure state. • Run the rollback step for the current failed activity and skip all other successful activities and subtasks. • Run the activity from the failed step forward.. **Note:** In the status summary RETRY is displayed for the activity that was in a failed state. In verbose summary, the activity status will be the current failed status. After the activities finish running, the activity status and task status will be updated properly.
mark_as_complete	This operation is used to mark an activity as manually fixed. The user manually fixes the failed activity and then runs this operation. This operation will then skip the current failed activity and mark it as MANUAL_SUCCESS, then resume to next subtask. **Note:** The status summary will show the MANUAL_SUCCESS status for the particular sub task. In verbose status the activity status is not changed. This is because the task or activity is not rerun, therefore the verbose status display the old status.
mark_all_steps_complete	This operation is used to skip all the subtask processing in case of failure. In this case the user can manually fix all pending and failed subtasks. **Note:** The status summary will show the MANUAL_SUCCESS status for the particular subtask. In verbose status the activity status will be the current status [in FAILED / PENDING/ SUCCESS state]. Since the subtask or activity is not rerun, the verbose status displays the old status.

Subtasks for PeopleTools Upgrade [PTU] in Cloud Manager Instance

The **cm_upgrade** task is used to perform a PeopleTools Upgrade [PTU] in Cloud Manager instance. The input values should given as response file. This task will upgrade PeopleTools on the Cloud Manager instance with no manual stop.

The following sub tasks are performed.

1. Validate the Cloud Manager PTU response file input values.

   The validation activity (PsftCMUpdateValidationActivity) will validate:

   • WinRm connectivity to the Windows client instance

   • Windows client user id and password

- File server validation
- Current PeopleTools version
- PeopleTools DPK for the new PeopleTools version
- Database information in TNS entry
- PIA port - http and https
- psftserver values and jolt port in wls config
- WLS port and Jolt port in Application domain
- Operator id/password
- WLS admin is/password
- Web profile user id/password
- DB admin password
- Operator id and password
- Connect id and password
- Access id and password
- All user input mandatory values

2. Take a backup of current CM_PS_HOME and PS_CFG_HOME.

   The activity to prepare for the update (PsftPrepareCMUpdateActivity) will:

   - Take a backup of PS_HOME and PS_CFG_HOME
   - Take a backup of PTU specific files.
   - Copy the PTC and ODC DPK from CM DPKs to the file server.

3. Stop the Cloud Manager psft domains.

   The activity to shut down PSFT domains (PsftCMDomainRestartActivity) will run,

4. Cloud Manager PTU upgrade process.

   This subtask contains multiple activities:

   - Copy the Cloud folder from Cloud Manager file server to Windows client.
   - Install the PeopleTools client for current CM PeopleTools version.
   - Install the PeopleTools client for new CM PeopleTools version.
   - Configuration Change Assistant for CM PeopleTools Upgrade process.

- Trigger the Change Assistant PTU command line process.

5. Cloud Manager PTP update process.

   This subtask contains multiple activities:

   - Copy the Cloud folder from Cloud Manager file server to Windows client.
   - Install the PeopleTools client for new CM PeopleTools version.
   - Configuration Change Assistant for CM PeopleTools Update process.
   - Trigger the Change Assistant PTP command line process.

6. Uninstall the Cloud Manager middle tier.

   The activity to uninstall the middle tier will:

   - Uninstall middle tier from the Cloud Manager instance.
   - Call the puppet clean up command using the root user.

7. Re-provision the Cloud Manager middle tier.

   This subtask contains multiple activities:

   - Provision middle tier:
     - Installation of new PeopleTools middle tier in the Cloud Manager middle tier.
     - Restore old configuration files in Cloud Manager.
     - Restore log4j, open ssl, puppet config and processing scripts.
     - Will do Cloud Manager specific configuration in Cloud Manager puppet yaml file and will trigger the Cloud Manager specific puppet profiles.
     - Run puppet apply using root user for recreating the middle tier.
   - Place holder for doing the post mid-tier creation steps, if any.

8. Cloud Manager PTU post update settings.

   The activity for PTU post update settings will:

   - Post the upgrade task.
   - Run the application engine program for notifying the UI about the processing status.

9. Restarting the Cloud Manager domains.

The activity will restart all Cloud Manager psft domains.

Creating Response File

Create a json file that includes all the mandatory values. The other values will be discovered in job processing. If the mandatory values are not included in the response file, running the job results in an error.

This is an example of the response file containing the mandatory values:

```
{
  "pum_source": {
    "windows_client": {
      "private_ip": "<windows_client private IP>",
      "remote_password": "<Windows Client Password>"
    }
  },
  "pum_target": {
    "psft": {
      "access_pwd": "<Access Password>",
      "opr_pwd": "<Operator Password>",
      "admin_pwd": "<DB Admin Password>",
      "connect_pwd": "<Connect Password>",
      "gw_admin_user_pwd": "<Gateway User Password>",
      "webprofile_user_pwd": "<Web Profile User Password>",
      "weblogic_admin_pwd": "<Webserver Admin Password>"
    }
  }
}
```

The response file can also contain additional values, however the mandatory values must be included.

This is an example of a response file with additional values:

```
{
  "file_server": {
    "hostname": "hostname.example.com"
  },
  "pum_source": {
    "windows_client": {
```



```
      "host_name": "win_client_hostname",
      "private_ip": "<Windows Client Private IP>",
      "remote_password": "<Windows Client Password>"
    }
  },
  "pum_target": {
    "psft": {
      "opr_id": "CLADM",
      "connect_id": "people",
      "access_id": "SYSADM",
      "weblogic_admin_user": "system",
      "gw_admin_user": "administrator",
      "db_name": "CMPSDB",
      "db_service_name": "CMPSDB",
      "db_host": "hostname.example.com",
      "psft_server": "hostname.example.com",
      "wsl_port": "7000",
      "jolt_port": "9033",
      "pia_https_port": "8443",
      "pia_http_port": "8000",
      "db_port": "1522",
      "tools_version": "8.56.12",
      "new_tools_version": "8.57.03",
      "pi_number": "8",
      "access_pwd": "<Access Password>",
      "opr_pwd": "<Operator Password>",
      "admin_pwd": "<DB Admin Password>",
      "connect_pwd": "<Connect Password>",
      "gw_admin_user_pwd": "<Gateway User Password>",
      "webprofile_user_pwd": "<Web Profile User Password>",
      "weblogic_admin_pwd": "<Webserver Admin Password>"
    }
  }
}
```


Getting Status of PeopleTools Upgrade Job

Use the following command to get the current status of the PeopleTools Upgrade:

```
run_cloud_manager_update -t CM_UPGRADE -o get_status
```

When the PeopleTools Upgrade begins, the status will show in progress and you can see what step is running.

This example illustrates Cloud Manager PeopleTools Upgrade Status.

```
-----------------------------------------
  CLOUD MANAGER PEOPLETOOLS UPGRADE STATUS
-----------------------------------------
  ORCHESTRATION ID:              CM_ORCH_PTU_20190501_154931  [Time:  2019-05-01
 15:50:26:158869]
  JOB SCHEDULE STATUS:           SUCCESS
  JOB STATUS:                    IN_PROGRESS

  1.  Validating Input Data:    IN_PROGRESS
  2.  Cloud Manager Pre-Update Settings:        PENDING
  3.  Stopping the Cloud Manager PSFT Domains:  PENDING
  4.  Cloud Manager PeopleTools Upgrade [PTU]:  PENDING
  5.  Cloud Manager PeopleTools Update [PTP]:   PENDING
  6.  Un-Provisioning the Middle Tier of Cloud Manager:         PENDING
  7.  Re-Provisioning the Middle Tier of Cloud Manager:         PENDING
  8.  Cloud Manager Post Update Settings:       PENDING
```

If a step fails, the status will show failure. You will need to correct the failure and either retry or mark the step as complete in order to continue.

This example illustrates Cloud Manager PeopleTools Upgrade Status where a step has failed.

```
-----------------------------------------
  CLOUD MANAGER PEOPLETOOLS UPGRADE STATUS
-----------------------------------------
  ORCHESTRATION ID:              CM_ORCH_PTU_20190429_082823  [Time:  2019-04-29
 08:29:17:737108]
  JOB SCHEDULE STATUS:           SUCCESS
  JOB STATUS:                    FAILURE

  1.  Validating Input Data:    FAILURE
  2.  Cloud Manager Pre-Update Settings:        PENDING
  3.  Stopping the Cloud Manager PSFT Domains:  PENDING
  4.  Cloud Manager PeopleTools Upgrade [PTU]:  PENDING
  5.  Cloud Manager PeopleTools Update [PTP]:   PENDING
  6.  Un-Provisioning the Middle Tier of Cloud Manager:         PENDING
  7.  Re-Provisioning the Middle Tier of Cloud Manager:         PENDING
  8.  Cloud Manager Post Update Settings:       PENDING
```

This example illustrates Cloud Manager PeopleTools Upgrade Status where all steps are successful.

```
-----------------------------------------
  CLOUD MANAGER PEOPLETOOLS UPGRADE STATUS
-----------------------------------------
  ORCHESTRATION ID:              CM_ORCH_PTU_20190501_160902  [Time:  2019-05-01
 16:09:57:122185]
  JOB SCHEDULE STATUS:           SUCCESS
  JOB STATUS:                    SUCCESS

  1.  Validating Input Data:    SUCCESS
  2.  Cloud Manager Pre-Update Settings:        SUCCESS
  3.  Stopping the Cloud Manager PSFT Domains:  SUCCESS
  4.  Cloud Manager PeopleTools Upgrade [PTU]:  SUCCESS
  5.  Cloud Manager PeopleTools Update [PTP]:   SUCCESS
  6.  Un-Provisioning the Middle Tier of Cloud Manager:         SUCCESS
  7.  Re-Provisioning the Middle Tier of Cloud Manager:         SUCCESS
  8.  Cloud Manager Post Update Settings:       SUCCESS
```

Verbose

To display the verbose status, add the -v verbose option:

```
run_cloud_manager_update -t CM_UPGRADE -o get_status -v
```

This is an example of the verbose Cloud Manager PeopleTools Upgrade Status:

```
-----------------------------------------
  CLOUD MANAGER PEOPLETOOLS UPGRADE STATUS
-----------------------------------------
  ORCHESTRATION ID:              CM_ORCH_PTU_20190501_160902  [Time:  2019-05-01 16⇒

:09:57:122185]
  JOB SCHEDULE STATUS:           SUCCESS
  JOB STATUS:                    SUCCESS
  JOB DETAILS:                   Successfully completed
  JOB LOG:                       /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MAN⇒

AGER_INSTANCE/CM_ORCH_PTU_20190501_160902/cm_job_psft_cm_ptu_job_0_20190501_160903_⇒

out.log
  JOB SCHEDULER LOG:             /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MAN⇒

AGER_INSTANCE/CM_ORCH_PTU_20190501_160902/out.log
  JOB STATUS CHECK LOG:          /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MAN⇒

AGER_INSTANCE/CM_DATA/out_20190501.log

  1.  Validating Input Data:    SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_PTU_VALIDATE_0_20190501_160906/out.log
      1.1  CM Update Validation Process:        SUCCESS
           Details: Successfully completed

  2.  Cloud Manager Pre-Update Settings:        SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_PRE_UPDATE_0_20190501_160912/out.log
      2.1  Cloud Manager Prepare Update Process:        SUCCESS
           Details: Successfully completed
```

```
  3.  Stopping the Cloud Manager PSFT Domains:  SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_DOMAIN_STOP_0_20190501_160918/out.log
      3.1  CM Domain Restart/Stop Process:      SUCCESS
           Details: Successfully completed

  4.  Cloud Manager PeopleTools Upgrade [PTU]:  SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_PTU_UPGRADE_0_20190501_160924/out.log
      4.1  Performing PsftActivityWindowsCloudFolderSync:       SUCCESS
           Details: Successfully completed

      4.2  Performing Cloud Manager Windows Client Install Activity [Source PeopleT⇒

ools]:       SUCCESS
           Details: Successfully completed

      4.3  Performing Windows Client Install Activity:  SUCCESS
           Details: Successfully completed

      4.4  Adding Environment Info to CA:       SUCCESS
           Details: Successfully completed

      4.5  CA PTU Apply Process:        SUCCESS
           Details: Successfully completed

  5.  Cloud Manager PeopleTools Update [PTP]:   SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_PTP_UPDATE_0_20190501_160930/out.log
      5.1  Performing PsftActivityWindowsCloudFolderSync:       SUCCESS
           Details: Successfully completed

      5.2  Performing Windows Client Install Activity:  SUCCESS
           Details: Successfully completed

      5.3  Adding Environment Info to CA:       SUCCESS
           Details: Successfully completed

      5.4  CA PTP Apply Process:        SUCCESS
           Details: Successfully completed

  6.  Un-Provisioning the Middle Tier of Cloud Manager:         SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_UNPROV_MT_0_20190501_160936/out.log
      6.1  Performing Unprovision MT:   SUCCESS
           Details: Successfully completed

  7.  Re-Provisioning the Middle Tier of Cloud Manager:         SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒

GER_INSTANCE/CM_REPROV_MT_0_20190501_160942/out.log
      7.1  Performing Reprovision MT:   SUCCESS
           Details: Successfully completed

      7.2  Cloud Manager Post MT Reprovision Process:   SUCCESS
           Details: Successfully completed

  8.  Cloud Manager Post Update Settings:       SUCCESS
      TASK DETAILS:              Successfully completed
      TASK LOG:                 /home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANA⇒
```



```
GER_INSTANCE/CM_POST_UPDATE_0_20190501_160948/out.log
      8.1  Cloud Manager Post Update Process:   SUCCESS
            Details: Successfully completed
```

Troubleshooting PeopleTools Upgrade Failures

Several logs are available to assist the users in troubleshooting errors that may occur during the upgrade process.

Log	***Location/Example log file name***
JOB LOG	/home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANAGER_INSTANCE/CM_ORCH_PTU_20190308_054432/cm_job_psft_cm_ptu_job_0_20190308_054432_out.log
JOB SCHEDULER LOG	/home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANAGER_INSTANCE/CM_ORCH_PTU_20190308_054432/out.log
JOB STATUS CHECK LOG	/home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANAGER_INSTANCE/CM_ORCH_PTU_20190308_054607/out.log
TASK LOG	/home/psadm2/psft/data/cloud/cmlogs/envs/CLOUD_MANAGER_INSTANCE/CM_POST_UPDATE_0_20190308_054441/out.log

Note: The log file names are only examples; the file name in a customer environment will vary.

Chapter 13

Cloud Manager Logs

Understanding PeopleSoft Cloud Manager Logs

Logs contain useful information for analyzing any environment related issues or failures that may occur in the system.

The flow diagram below illustrates an overview of logging process in Cloud Manager.

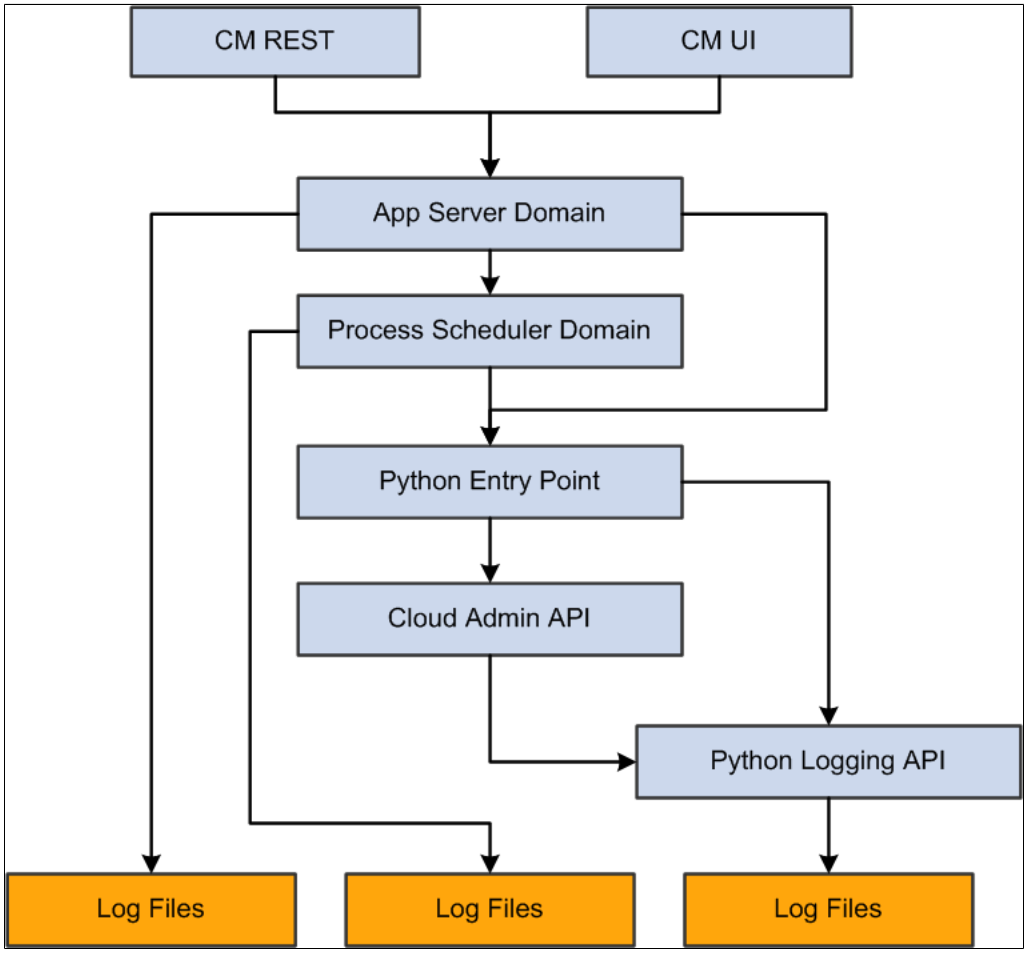

Cloud Manager logs include:

- Python logs
- Environment Action logs
- Download Manager logs
- Patching logs

- App Server Domain logs
- Process Scheduler Domain logs
- Puppet logs in Provisioned VMs
- Terraform logs

Describing Cloud Manager Logs

Python Logs

- All Python environment action logs will be under the following folder: <CM Python Log Root>/envs/
- All logs related to a particular environment <env name> will be under: <CM Python Log Root>/envs/<env name>/. The path of <CM Python Log Root> is /home/psadm2/psft/data/cloud/cmlogs.
- All logs related to the action <Type> on the environment denoted by <env name> will be under: <CM Python Log Root>/envs/<env name>/<Type>_TimeStamp
- The actions can be:
  - CREATE
  - DEPLOY (Only for OCI)
  - REMOVE
  - ACTIONS (Start, Stop, and so on)
  - ADD_TARGET
  - UPGRADE
  - BACKUP
  - RESTORE
  - CLONE
  - REFRESH

Download Manager Logs

Log files generated by the download manager are available in the following folder: <CM Python Log Root>/dm/

Note: A contextual logs UI that can be accessed from the environment details page is available in Cloud Manager for administrator and end users while debugging issues in their environments.

Since the number of folders and files under cmlogs will grow over time, an archiving process for older files is included in Cloud Manager.

Note: When the user deletes an environment, the log files are automatically moved to an archive directory, for example: CM Python Log Root>/envs/archive_dir

Machine Learning Logs

- Log files generated for data upload are available in the following folder: <CM Python Log Root>/cmlogs/mltraining/OCI_DATA_UPLOAD_<TimeStamp>/
- Log files generated for Model training are available in the following folder: <CM Python Log Root>/cmlogs/mltraining/OCI_DS_MODEL_TRAINING_<TimeStamp>/
- Log files generated for Model Prediction run are available in the following folder: <CM Python Log Root>/cmlogs/mlprediction/
- To view Model Training OCI Logs in OCI console:
  - Navigate to Data Science > Project > Jobs.
  - Select Current Job Run.
  - Select logs in Job run details.

Application Server Domain Logs

Application Server Domain logs are written in the default application server domain logs directory. $PS_CFG_HOME/appserv/APPDOM/LOGS

PeopleSoft Cloud Manager Log Levels

The log levels that can be configured by the customer are:

- Critical
- Error
- Warning
- Debug

Note: Logging formats and levels are controlled using Python Logging configuration. The default log level is Debug.

Attribute name	***Format***	***Description***
asctime	%(asctime)s	Human-readable time when the LogRecord was created. By default this is of the form ‘2003-07-08 16:49:45,896’ (the numbers after the comma are millisecond portion of the time).
created	%(created)f	Time when the LogRecord was created (as returned by time.time()).

Attribute name	*Format*	*Description*
filename	%(filename)s	Filename portion of pathname.
funcName	%(funcName)s	Name of function containing the logging call.
levelname	%(levelname)s	Text logging level for the message ('DEBUG', 'INFO', 'WARNING', 'ERROR', 'CRITICAL').
levelno	%(levelno)s	Numeric logging level for the message (DEBUG, INFO, WARNING, ERROR, CRITICAL).
lineno	%(lineno)d	Source line number where the logging call was issued (if available).
message	%(message)s	The logged message, computed as msg % args. This is set when Formatter. format() is invoked.
module	%(module)s	Module (name portion of filename).
msecs	%(msecs)d	Millisecond portion of the time when the LogRecord was created.
name	%(name)s	Name of the logger used to log the call.
pathname	%(pathname)s	Full pathname of the source file where the logging call was issued (if available).
process	%(process)d	Process ID (if available).
processName	%(processName)s	Process name (if available).
relativeCreated	%(relativeCreated)d	Time in milliseconds when the LogRecord was created, relative to the time the logging module was loaded.
thread	%(thread)d	Thread ID (if available).
threadName	%(threadName)s	Thread name (if available).

Configurable Log Root: /home/psadm2/psft/data/cloud/cmlogs will be the Cloud Manager Python Log Root.

Changing Log Levels

The customer can edit psc_constants.py and pca_int.conf files to set the log level.

The default logging level is *Debug*. To customize it to another level, modify the following entry in the file:

```
<PS_APP_HOME>\cloud\psc_cloud\psc_utils\psc_constants.py logging_level = logging.DE⇒
BUG
```

Note: You do not need to restart the domains after the changing the log levels.

Important! In OCI, for Python logging configuration, two locations have to be configured.
cloud/pca_int.conf - This controls the log level in Download Manager and Terraform handler.
cloud/psc_cloud/psc_utils/psc_constants.py - This controls the log level in PSFT deployment code.

Terraform Logs for OCI

When Cloud Manager is used for provisioning environments, the provisioning of infrastructure is the first task that is run. The Terraform log files generated during the provisioning task can be found under the logs directory for the environment: /home/psadm2/psft/data/cloud/cmlogs/envs/<Environment Name>/CREATE_<Time Stamp>/

Term	***Definition***
Log File Type	Description
tf.out	This is the Output Log, which contains the Terraform's stdout stream output.
tf.err	This is the Error Log, which contains the Terraform's stderr stream output.
out.log	This is the Driver Output generated by the Cloud Manager module that invokes Terraform.
console.log	This contains the uncaught exceptions.

Terraform Input and Output Files

The Terraform input/output files used by Cloud Manager for provisioning an environment can be found under: /home/psadm2/psft/data/cloud/ocihome/envs/<Environment Name>/

Term	***Definition***
Log File Type	Description
terraform.tf.json	The .json file contains the specification of the VMs, storage volumes, database systems etc.

Term	*Definition*
variables.tf	This file contains the tenancy OCID, user OCID, API key paths, finger print etc.
tf.result.json	This file contains a summary of the resources that were successfully created by Terraform.

Chapter 14

Backing Up and Restoring Cloud Manager

Understanding Cloud Manager Backup and Restore

There are two methods for backing up and restoring Cloud Manager.

1. Automated Backup and Restore Utility.

   See Using Automated Backup and Restore Utility.

2. Back up and restore using Block Volume Backups for OCI.

   See Manually Backing Up and Restoring Cloud Manager Using Block Volume Backups for OCI

Important! If your Cloud Manager instance was created using a custom fqdn, the PRESERVE_HOSTINFO field must be set to 0 before doing CM backup. If this value is not set as 0, there is a chance that the Cloud Manager restore will fail as OCI will not be able to configure the new instance with new IP details in /etc/hosts file. See tutorial *Install the PeopleSoft Cloud Manager Stack with Resource Manager*.

Warning! Backing up or restoring Cloud Manager will shut down all services. Please ensure that no provisioning or lifecycle jobs are running. Any running jobs will be abruptly ended and may result in an unstable or unusable state.

Important! Ensure the process scheduler does not have any jobs in its queue. If there are pending jobs in the queue, those may get scheduled to run whenever a backup is restored.

Using Automated Backup and Restore Utility

The automated backup and restore utility provides the ability to create a backup of the Cloud Manager instance. This backup/restore utility provides the following options:

- Backup and restore Cloud Manager Instance block volume.
- Backup Cloud Manager Instance boot volume.
- Delete backups.
- List Cloud Manager backups in OCI.
- Create an OCI config file.

To run the automated backup and restore utility:

1. Log into Cloud Manager instance using SSH as user opc.

2. Change the directory to */home/opc/bootstrap/cm_backup_and_restore*.
3. Copy 2 files (cm_backup_restore.py and cm_backup_restore.sh) from */opt/oracle/psft/pt/ps_app_home/cloud/psc_cloud/psc_utils/* to */home/opc/bootstrap/cm_backup_and_restore*.
4. Run the shell script **cm_backup_restore.sh**.

Understanding the Backup and Restore Shell Script

For a summary of the usage and optional arguments for shell script **cm_backup_restore.sh**, use the -h option.

This is an example of the help.

This example illustrates running cm_backup_restore.sh -h.

```
[opc@smncm12r1 cm_backup_and_restore]$ sh cm_backup_restore.sh -h
usage: cm_backup_restore.py [-h] -o option [-n backupname] [-t {BLOCK,BOOT}]
                            [-c DB_ACCESS_ID DB_ACCESS_PWD PDB_NAME]

Utility to Backup/Restore storage volume of Cloud Manager Instance

optional arguments:
  -h, --help            show this help message and exit
  -o option             Provide an option to Backup/Restore/Delete/Display

                        backup       (for taking backup of the Cloud Manager Instance)
                        restore      (for restoring block volume of Cloud Manager Instance)
                        delete       (for deleting backup from OCI)
                        list         (to display list of Cloud Manager Instance backups in OCI)
                        createconfig (to create oci config file)

  -n backupname         Custom Backup Name [Optional]
                        Default format:'<Vol Type>_VOLBKP_<CM Instance Name><timestamp(ddmmyyyy)>_<timestamp(HHMMSS)>'
                        (this field is mandatory for Restore and Delete operations)

  -t {BLOCK,BOOT}       Backup Volume Type [Optional]
                        Default : BLOCK
                        (this field is required for Backup operation)

  -c DB_ACCESS_ID DB_ACCESS_PWD PDB_NAME
                        Provide Database Access Id, Password and PDB Name
                        Example <DB_ACCESS_ID> <DB_ACCESS_PWD> <PDB_NAME> (separated by space)
```

Creating Config File

You need to create a config file if:

- This is the first time you are running the automated backup and restore utility.
- You want to run the utility with different OCI user credentials.

To create the Config file:

1. Run the utility with the createconfig option and provide the Cloud Manager database Access ID and password.

   ```
   sh cm_backup_restore.sh -o createconfig -c <DB_ACCESS_ID> <DB_ACCESS_PWD> <PDB⇒
   _NAME>
   ```

2. This option will read OCI User ID, Tenancy ID, Fingerprint and Private Key File Path from database and display on screen.
3. Check the displayed values.
   - If the values are correct, press *y* to confirm and the configuration file will be created.

- If you want to change the values, press *n* and provide Tenancy ID, User ID, Fingerprint, Private Key File Path, Region details to create config file.

This example illustrates accepting default values.

```
[opc@smncm12r1 cm_backup_and_restore]$ sh cm_backup_restore.sh -o createconfig -c EMDBO Water123 CMPSDB
Please check whether below OCI details are correct

Tenancy OCID : 'ocid1.tenancy.oc1..aaaaaaaayy35pigzces6ly7aslibgt7a4u7o3tlt42nxg4idzrsui52gma5a'
User OCID: 'ocid1.user.oc1..aaaaaaaaprdpimfzdhyip3tmeot5ve7woycqanjj7juz3gqccc5qawu7dora'
Fingerprint of Public Key : '49:4f:45:64:2a:1c:d1:11:2f:61:6a:09:95:97:b4:74'
Private Key Path : '/home/psadm2/psft/data/cloud/ocihome/keys/oci_api_key.pem'
Region : 'us-ashburn-1'

 Enter y if above values are correct : y
2021-02-10 13:35:06 -     INFO -  406 : Created config file 'config.json'
[opc@smncm12r1 cm_backup_and_restore]$
```

This example illustrates modifying default values.

```
[opc@smncm12r1 cm_backup_and_restore]$ sh cm_backup_restore.sh -o createconfig -c EMDBO Water123 CMPSDB
Please check whether below OCI details are correct

Tenancy OCID : 'ocid1.tenancy.oc1..aaaaaaaayy35pigzces6ly7aslibgt7a4u7o3tlt42nxg4idzrsui52gma5a'
User OCID: 'ocid1.user.oc1..aaaaaaaaprdpimfzdhyip3tmeot5ve7woycqanjj7juz3gqccc5qawu7dora'
Fingerprint of Public Key : '49:4f:45:64:2a:1c:d1:11:2f:61:6a:09:95:97:b4:74'
Private Key Path : '/home/psadm2/psft/data/cloud/ocihome/keys/oci_api_key.pem'
Region : 'us-ashburn-1'

 Enter y if above values are correct : n
Enter Tenancy OCID : ocid1.tenancy.oc1..aaaaaaaayy35pigzces6ly7aslibgt7a4u7o3tlt42nxg4idzrsui52gma5a
Enter User OCID : ocid1.user.oc1..aaaaaaaaprdpimfzdhyip3tmeot5ve7woycqanjj7juz3gqccc5qawu7dora
Enter Fingerprint of Public Key : 49:4f:45:64:2a:1c:d1:11:2f:61:6a:09:95:97:b4:74
Enter region : us-ashburn-1
Enter Private Key File Path : /home/psadm2/psft/data/cloud/ocihome/keys/oci_api_key.pem
2021-02-10 13:37:00 -     INFO -  406 : Created config file 'config.json'
[opc@smncm12r1 cm_backup_and_restore]$
```

Note: If config.json is already present, this utility will replace the existing one.

Creating Backups

When creating a backup of Cloud Manager, backup both block (data) volume and the boot volume. This will ensure creating a pair of time consistent backups that can be restored together when restoring the entire Cloud Manager instance in case of boot volume issues or failures.

Backups can be created using the default name which is *<Vol Type>_VOLBKP_<CM Instance Name>_<timestamp(ddmmyyyy)>_<timestamp(HHMMSS)>* or you can provide a custom name for the backup. Use the optional argument *-n <Custom Backup Name>* to create a backup with a custom name.

The default backup type is BLOCK.

Cloud Manager data is saved in two locations: Oracle Database files on data volume and configuration files on boot volume. Both must be backed-up to restore to a consistent point. The backup utility creates a full backup of both data volume as well as files on boot volume automatically.

Creating the Block Volume Backup

To create the block volume backup:

1. Run the utility with the backup option.

   ```
   sh cm_backup_restore.sh -o backup
   ```


Note: Default value for volume type (-t) is BLOCK.

2. Optional arguments include name (-n) and volume type (-t).

   ```
   sh cm_backup_restore.sh -o backup -n SMNCM_CST_BKP_BLK
   ```

   In the above example the backup name is SMNCM_CST_BKP_BLK.

3. User will be prompted that the backup task will stop Database, Application Server, Process Scheduler and PIA services. Press *y* to continue.

4. User will be prompted to enter a passphrase for the private API signing key.

   Enter the passphrase if it exists or press enter.

5. The backup automation process will:

   - Shut down the database and domains.
   - Backup files on boot volume and save them on data volume as /opt/oracle/psft/dpks/cm_boot_vol_files.tar.gz. Files backed up include:
     - All files under PS_CFG_HOME
     - /home/psadm1/.bashrc
     - /home/psadm2/.bashrc
     - /home/psadm3/.bashrc
     - /home/oracle2/.bashrc
     - /home/esadm1/.bashrc
     - /etc/profile
     - /etc/bashrc
     - /usr/lib/systemd/system/psft-db-<DB_NAME>.service
     - /usr/lib/systemd/system/psft-appserver-APPDOM.service
     - /usr/lib/systemd/system/psft-prcs-PRCSDOM.service
     - /usr/lib/systemd/system/psft-pia-peoplesoft.service
   - Create a backup of data volume.
   - Restart the database and domains.

Example:

```
[opc@smncm cm_backup_and_restore]$ sh cm_backup_restore.sh -o backup
Backup task will stop Database, Application Server, Process Scheduler and PIA servi⇒

ces
Do you want to continue (y/n) ? y
```

```
Please Enter Private Key Pass Phrase :
2019-04-19 08:36:22 -     INFO -  408 : Compressing BOOT volume directories/files '⇒

/home/psadm1/.bashrc,/home/psadm2/.bashrc,/home/psadm3/.bashrc,/home/oracle2/.bashr⇒

c,/home/esadm1/.bashrc,/etc/profile,/etc/bashrc,/etc/init.d/psft-db,/etc/init.d/psf⇒

t-appserver,/etc/init.d/psft-prcs,/etc/init.d/psft-pia' and moving to '/opt/oracle/⇒

psft/dpks/cm_boot_vol_files.tar.gz'
2019-04-19 08:36:22 -     INFO -  418 : $PS_CFG_HOME directory '/home/psadm2/psft/p⇒

t/8.57'
2019-04-19 08:36:22 -     INFO -  427 : Compressed BOOT volume directories/files '/⇒

home/psadm1/.bashrc,/home/psadm2/.bashrc,/home/psadm3/.bashrc,/home/oracle2/.bashrc⇒

,/home/esadm1/.bashrc,/etc/profile,/etc/bashrc,/etc/init.d/psft-db,/etc/init.d/psft⇒

-appserver,/etc/init.d/psft-prcs,/etc/init.d/psft-pia,/home/psadm2/psft/pt/8.57' an⇒

d moving to '/opt/oracle/psft/dpks/cm_boot_vol_files.tar.gz'
2019-04-19 08:37:00 -     INFO -  598 : Cloud Manager Instance Backup Volume Id :oc⇒

id1.volume.oc1.iad.xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
2019-04-19 08:37:00 -     INFO -  448 : Changing DB service state to 'Stop'
2019-04-19 08:37:27 -     INFO -  481 : DB service is in 'Stop' state
2019-04-19 08:37:27 -     INFO -  448 : Changing APPSERVER service state to 'Stop'
2019-04-19 08:37:43 -     INFO -  481 : APPSERVER service is in 'Stop' state
2019-04-19 08:37:43 -     INFO -  448 : Changing PRCS service state to 'Stop'
2019-04-19 08:42:35 -     INFO -  481 : PRCS service is in 'Stop' state
2019-04-19 08:42:35 -     INFO -  448 : Changing PIA service state to 'Stop'
2019-04-19 08:42:54 -     INFO -  481 : PIA service is in 'Stop' state
2019-04-19 08:42:54 -     INFO -  875 : Created input.json for BLOCK Volume Backup ⇒

: BLOCK_VOLBKP_SMNCM_20190419_084254
2019-04-19 08:42:54 -     INFO -  876 : Creating BLOCK Volume Backup : BLOCK_VOLBKP⇒

_SMNCM_20190419_084254
2019-04-19 08:44:15 -     INFO -  897 : Created BLOCK volume backup 'BLOCK_VOLBKP_S⇒

MNCM_20190419_084254'
2019-04-19 08:44:15 -     INFO -  448 : Changing DB service state to 'Start'
2019-04-19 08:44:24 -     INFO -  481 : DB service is in 'Start' state
2019-04-19 08:44:24 -     INFO -  448 : Changing APPSERVER service state to 'Start'
2019-04-19 08:44:49 -     INFO -  481 : APPSERVER service is in 'Start' state
2019-04-19 08:44:49 -     INFO -  448 : Changing PRCS service state to 'Start'
2019-04-19 08:46:45 -     INFO -  481 : PRCS service is in 'Start' state
2019-04-19 08:46:45 -     INFO -  448 : Changing PIA service state to 'Start'
2019-04-19 08:46:46 -     INFO -  481 : PIA service is in 'Start' state
```

Creating Boot Volume Backup

Boot volume backups must be restored manually using OCI console. This backup can be used if the CM instance becomes unusable and inaccessible.

Note: When restoring a boot volume backup ensure to create the Cloud Manager instance with the same IP addresses as the original instance. The original instance must be terminated before creating a new instance. If unable to reuse the same IP address, then all references to old IP addresses in Cloud Manager application and domains must be manually updated to reflect the new IP address of the instance.

Warning! If the backup is not restored correctly, the Cloud Manager instance may not come up with the proper network configuration which could result in losing the ability to manage already provisioned environments.

To create a boot volume backup:

1. Run the utility with the backup option.

   ```
   sh cm_backup_restore.sh -o backup -t BOOT
   ```

2. Optionally you can provide a custom backup name using the -n argument..

   ```
   sh cm_backup_restore.sh -o backup -n CUSTOM_NAME -t BOOT
   ```

   In the above example the backup name is CUSTOM_NAME.

3. User will be prompted that the backup task will stop Database, Application Server, Process Scheduler and PIA services. Press *y* to continue.

4. User will be prompted to enter a passphrase for the private API signing key.

   Enter the passphrase if it exists or press enter.

5. The backup process will:

   - Shut down the database and domains.
   - Create the backup.
   - Restart the database and domains.

Example:

```
[opc@smncm cm_backup_and_restore]$ sh cm_backup_restore.sh -o backup -t BOOT
Backup task will stop Database, Application Server, Process Scheduler and PIA servi⇒

ces
Do you want to continue (y/n) ? y
Please Enter Private Key Pass Phrase :
2019-04-19 09:58:54 -     INFO -  598 : Cloud Manager Instance Backup Volume Id :oc⇒

id1.bootvolume.oc1.iad.abuwcljrhcuapoaveqp2yf2oof2pc6lsurxx4xu2fmjarjjrz4qyfkif3wpq
2019-04-19 09:58:54 -     INFO -  448 : Changing DB service state to 'Stop'
2019-04-19 09:59:14 -     INFO -  481 : DB service is in 'Stop' state
2019-04-19 09:59:14 -     INFO -  448 : Changing APPSERVER service state to 'Stop'
2019-04-19 09:59:30 -     INFO -  481 : APPSERVER service is in 'Stop' state
2019-04-19 09:59:30 -     INFO -  448 : Changing PRCS service state to 'Stop'
2019-04-19 10:04:33 -     INFO -  481 : PRCS service is in 'Stop' state
2019-04-19 10:04:33 -     INFO -  448 : Changing PIA service state to 'Stop'
2019-04-19 10:07:35 -     INFO -  481 : PIA service is in 'Stop' state
2019-04-19 10:07:35 -     INFO -  875 : Created input.json for BOOT Volume Backup :⇒

 BOOT_VOLBKP_SMNCM_20190419_100735
2019-04-19 10:07:35 -     INFO -  876 : Creating BOOT Volume Backup : BOOT_VOLBKP_S⇒

MNCM_20190419_100735
2019-04-19 10:09:17 -     INFO -  897 : Created BOOT volume backup 'BOOT_VOLBKP_SMN⇒

CM_20190419_100735'
2019-04-19 10:09:17 -     INFO -  448 : Changing DB service state to 'Start'
2019-04-19 10:09:26 -     INFO -  481 : DB service is in 'Start' state
2019-04-19 10:09:26 -     INFO -  448 : Changing APPSERVER service state to 'Start'
2019-04-19 10:09:51 -     INFO -  481 : APPSERVER service is in 'Start' state
2019-04-19 10:09:51 -     INFO -  448 : Changing PRCS service state to 'Start'
2019-04-19 10:11:47 -     INFO -  481 : PRCS service is in 'Start' state
2019-04-19 10:11:47 -     INFO -  448 : Changing PIA service state to 'Start'
2019-04-19 10:11:47 -     INFO -  481 : PIA service is in 'Start' state
```

Listing Existing Backups

To list existing block and boot volume backups:

1. Run the utility with the list option:

   ```
   sh cm_backup_restore.sh -o list
   ```

2. User will be prompted to enter a passphrase for the private API signing key.

   Enter the passphrase if it exists or press enter.

3. The list of backups is displayed.

Example:

```
[opc@smncm cm_backup_and_restore]$ sh cm_backup_restore.sh -o list
Please Enter Private Key Pass Phrase :
  BLOCK volume backup list :
   * BLOCK_VOLBKP_SMNCM_20190410_110419
  BOOT volume backup list :
   No BOOT volume backups found
```

Restoring from a Backup

Both data volume and a set of files on boot volume must be replaced from a backup to restore Cloud Manager instance to a certain backup point.

First step is to restore the data volume. To restore a block (data) volume backup the user must provide a backup name. User can list the existing backups and then select the backup to restore.

To restore a block (data) volume backup:

1. While you are still on the latest version, before restoring the data volume, generate the list of IP addresses of all managed instances provisioned by Cloud Manager. This is required later in the restore process to synchronize the '/cm_psft_dpks/cloud' directory to all managed nodes in case the files are not up-to-date. This step is required whenever restoring from a latest Cloud Manager version to an older version. For example, restoring back from Cloud Manager Image 10 to Cloud Manager Image 9.

   - SSH to Cloud Manager
   - Switch user to psadm2

     ```
     $ sudo su - psadm2
     ```

   - Change directory to PS_APP_HOME/cloud

     ```
     $ cd /opt/oracle/psft/pt/ps_app_home/cloud
     ```

   - Generate the list of IP addresses into file /home/psadm2/managedenvironments.txt

     ```
     $ get_managed_envs.sh $PS_CFG_HOME
     ```

   - Take backup of /home/psadm2/managedenvironments.txt

     ```
     $ cp /home/psadm2/managedenvironments.txt  <backup_path>
     ```

2. Run the utility with the restore option.

   ```
   sh cm_backup_restore.sh -o restore -n <backup_name>
   ```

3. User will be warned with a message that backup automation will stop Database, Application Server, Process Scheduler and PIA services. Press *y* to continue.

4. User will be prompted to enter a passphrase for the private API signing key.

   Enter the passphrase if it exists or press enter.

5. The restore process will:

   - Shut down the database and domains.
   - Restore from the identified volume backup .
   - Restart the database and domains.

   Example:

```
[opc@smncm cm_backup_and_restore]$ sh cm_backup_restore.sh -o restore -n BLOCK⇒

_VOLBKP_SMNCM_20190419_084254
Restore task will stop Database, Application Server, Process Scheduler and PIA⇒

 services
Do you want to continue (y/n) ? y
Please Enter Private Key Pass Phrase :
2019-04-19 09:14:20 -     INFO -  658 : Getting BLOCK volume backup 'BLOCK_VOL⇒

BKP_SMNCM_20190419_084254' details from OCI
2019-04-19 09:14:20 -     INFO -  666 : BLOCK volume backup 'BLOCK_VOLBKP_SMNC⇒

M_20190419_084254' Id : 'ocid1.volumebackup.oc1.iad.abuwcljryiuqzxut5ldxtb2hma⇒

xgj5w65ykxukza6ztmz7odk323yiip3r7a'
2019-04-19 09:14:20 -     INFO -  824 : Creating storage volume 'StorageVol_sm⇒

ncm_19APRIL2019_0914' from Backup 'ocid1.volumebackup.oc1.iad.abuwcljryiuqzxut⇒

5ldxtb2hmaxgj5w65ykxukza6ztmz7odk323yiip3r7a'
2019-04-19 09:15:01 -     INFO -  853 : Created storage volume Name:'StorageVo⇒

l_smncm_19APRIL2019_0914',
Id:'ocid1.volume.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllreleoh⇒

sqe2vcq'
2019-04-19 09:15:01 -     INFO -  598 : Cloud Manager Instance Backup Volume I⇒

d :ocid1.volume.oc1.iad.abuwcljr7k4bq4eqw3nfoyl7vxdsykzyrrthzgrco7jnwiyg5f6zga⇒

ghlmva
2019-04-19 09:15:01 -     INFO -  448 : Changing DB service state to 'Stop'
2019-04-19 09:15:22 -     INFO -  481 : DB service is in 'Stop' state
2019-04-19 09:15:22 -     INFO -  448 : Changing APPSERVER service state to 'S⇒

top'
2019-04-19 09:15:34 -     INFO -  481 : APPSERVER service is in 'Stop' state
2019-04-19 09:15:34 -     INFO -  448 : Changing PRCS service state to 'Stop'
2019-04-19 09:20:41 -     INFO -  481 : PRCS service is in 'Stop' state
2019-04-19 09:20:41 -     INFO -  448 : Changing PIA service state to 'Stop'
2019-04-19 09:20:57 -     INFO -  481 : PIA service is in 'Stop' state
2019-04-19 09:20:57 -     INFO -  497 : Kill some of the running 'psadm' proce⇒
```

```
sses
2019-04-19 09:21:02 -     INFO -  530 : Unmounted block volume '/opt/oracle/ps⇒

ft'
2019-04-19 09:21:02 -     INFO -  688 : Running ISCSI command to logout of sto⇒

rage volume iqn.2015-12.com.oracleiaas:22afb62d-a893-4ab4-a8fe-6dd14053c61b
2019-04-19 09:21:02 -     INFO -  698 : Logged out of storage volume iqn.2015-⇒

12.com.oracleiaas:22afb62d-a893-4ab4-a8fe-6dd14053c61b
2019-04-19 09:21:02 -     INFO -  700 : Running ISCSI command to delete storag⇒

e device iqn.2015-12.com.oracleiaas:22afb62d-a893-4ab4-a8fe-6dd14053c61b
2019-04-19 09:21:02 -     INFO -  709 : Deleted storage device iqn.2015-12.com⇒

.oracleiaas:22afb62d-a893-4ab4-a8fe-6dd14053c61b
2019-04-19 09:21:02 -     INFO -  793 : Initiated REST Call to Detach storage ⇒

volume ocid1.volumeattachment.oc1.iad.abuwcljrfnxyjammza3scounohq7qzawelzslmqi⇒

accvef7cm6w7m2wxjxuq
2019-04-19 09:21:23 -     INFO -  802 : Detached storage volume ocid1.volumeat⇒

tachment.oc1.iad.abuwcljrfnxyjammza3scounohq7qzawelzslmqiaccvef7cm6w7m2wxjxuq
2019-04-19 09:21:23 -     INFO -  915 : Attaching storage volume 'ocid1.volume⇒

.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllreleohsqe2vcq'
2019-04-19 09:21:23 -     INFO -  918 : Creating input.json for attaching bloc⇒

k volume 'ocid1.volume.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxll⇒

releohsqe2vcq'
2019-04-19 09:21:23 -     INFO -  920 : Created input.json for attaching block⇒

 volume 'ocid1.volume.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllr⇒

eleohsqe2vcq'
2019-04-19 09:21:23 -     INFO -  928 : Initiated REST Call to attach storage ⇒

volume ocid1.volume.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllrel⇒

eohsqe2vcq
2019-04-19 09:22:04 -     INFO -  935 : Attached storage volume 'ocid1.volume.⇒

oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllreleohsqe2vcq'
2019-04-19 09:22:04 -     INFO -  598 : Cloud Manager Instance Backup Volume I⇒

d :ocid1.volume.oc1.iad.abuwcljrf2hcy3ff5mxhbumnldvftljsp5yz66n73wzsxllreleohs⇒

qe2vcq
2019-04-19 09:22:04 -     INFO -  718 : Running ISCSI Attach Commands
2019-04-19 09:22:04 -     INFO -  733 : Running ISCSI command to add new node ⇒

iqn.2015-12.com.oracleiaas:f6f08414-d96f-4fd2-a852-69c0208ca8b2
2019-04-19 09:22:04 -     INFO -  739 : Added new node iqn.2015-12.com.oraclei⇒

aas:f6f08414-d96f-4fd2-a852-69c0208ca8b2
2019-04-19 09:22:04 -     INFO -  741 : Running ISCSI command to start node on⇒

 boot iqn.2015-12.com.oracleiaas:f6f08414-d96f-4fd2-a852-69c0208ca8b2
2019-04-19 09:22:04 -     INFO -  747 : Updated node iqn.2015-12.com.oracleiaa⇒

s:f6f08414-d96f-4fd2-a852-69c0208ca8b2 settings to start on boot
2019-04-19 09:22:04 -     INFO -  749 : Running ISCSI command to login to node⇒

 iqn.2015-12.com.oracleiaas:f6f08414-d96f-4fd2-a852-69c0208ca8b2
2019-04-19 09:22:04 -     INFO -  755 : Logged in to node iqn.2015-12.com.orac⇒
```



```
leiaas:f6f08414-d96f-4fd2-a852-69c0208ca8b2
2019-04-19 09:22:04 -     INFO -  543 : Mounted block volume '/opt/oracle/psft⇒

'
2019-04-19 09:22:13 -     INFO -  448 : Changing DB service state to 'Start'
2019-04-19 09:22:25 -     INFO -  481 : DB service is in 'Start' state
2019-04-19 09:22:25 -     INFO -  448 : Changing APPSERVER service state to 'S⇒

tart'
2019-04-19 09:22:56 -     INFO -  481 : APPSERVER service is in 'Start' state
2019-04-19 09:22:56 -     INFO -  448 : Changing PRCS service state to 'Start'
2019-04-19 09:24:52 -     INFO -  481 : PRCS service is in 'Start' state
2019-04-19 09:24:52 -     INFO -  448 : Changing PIA service state to 'Start'
2019-04-19 09:24:52 -     INFO -  481 : PIA service is in 'Start' state
2019-04-19 09:24:52 -     INFO -  947 : Getting storage volume name for Id : '⇒

ocid1.volume.oc1.iad.abuwcljr7k4bq4eqw3nfoyl7vxdsykzyrrthzgrco7jnwiyg5f6zgaghl⇒

mva'
2019-04-19 09:24:52 -     INFO -  957 : Storage volume name for Id 'ocid1.volu⇒

me.oc1.iad.abuwcljr7k4bq4eqw3nfoyl7vxdsykzyrrthzgrco7jnwiyg5f6zgaghlmva' is 'S⇒

torageVol_smncm_18APR2019_1044'
2019-04-19 09:24:52 -     INFO - 1061 : Please remove detached volume 'Storage⇒

Vol_smncm_18APR2019_1044' from OCI manually
```

6. Manually restore boot volume files from the backup /opt/oracle/psft/dpks/cm_boot_vol_files.tar.gz

Next restore files on boot volume. To restore, run the following set of commands and set appropriate ownership on restored files.

1. SSH to Cloud Manager instance
2. Switch user to root.

```
$ sudo bash
```

3. Uncompress and extract boot volume files backup that were saved as /opt/oracle/psft/dpks/cm_boot_vol_files.tar.gz to /tmp/CMbkup

```
$ mkdir /tmp/CMbkup
$ cd /tmp/CMbkup
$ tar -xvf /opt/oracle/psft/dpks/cm_boot_vol_files.tar.gz
```

4. Restore PS_CFG_HOME files. The PS_CFG_HOME path varies if the PeopleTools version before backup was different. Ensure to use the right path.

```
$ cp -r  /tmp/CMbkup/home/psadm2/psft/pt/* /home/psadm2/psft/pt/
$ chown -R psadm2:oinstall /home/psadm2/psft/pt/
$ chmod 755 -R /home/psadm2/psft/pt/
```

5. Copy profile files for users psadm1, psadm2, psadm3, oracle2 and esadm1.

```
$ cp /tmp/CMbkup/etc/profile /etc/profile
$ cp /tmp/CMbkup/etc/bashrc /etc/bashrc
$ chown root:root /etc/profile /etc/bashrc
$ chmod 644 /etc/profile /etc/bashrc

$ cp /tmp/CMbkup/home/psadm1/.bashrc /home/psadm1/.bashrc
$ chown psadm1:oinstall /home/psadm1/.bashrc
$ chmod 644 /home/psadm1/.bashrc
```

```
$ cp /tmp/CMbkup/home/psadm2/.bashrc /home/psadm2/.bashrc
$ chown psadm2:oinstall /home/psadm2/.bashrc
$ chmod 644 /home/psadm2/.bashrc

$ cp /tmp/CMbkup/home/psadm3/.bashrc /home/psadm3/.bashrc
$ chown psadm3:appinst /home/psadm3/.bashrc
$ chmod 644 /home/psadm3/.bashrc

$ cp /tmp/CMbkup/home/oracle2/.bashrc /home/oracle2/.bashrc
$ chown oracle2:oinstall /home/oracle2/.bashrc
$ chmod 644 /home/oracle2/.bashrc

$ cp /tmp/CMbkup/home/esadm1/.bashrc /home/esadm1/.bashrc
$ chown esadm1:oinstall /home/esadm1/.bashrc
$ chmod 644 /home/esadm1/.bashrc
```

6. Copy init scripts.

```
$ cp /tmp/CMbkup/usr/lib/systemd/system/psft-db-<DB_NAME>.service /usr/lib/sys⇒

temd/system/
$ cp /tmp/CMbkup/usr/lib/systemd/system/psft-appserver-APPDOM.service /usr/lib⇒

/systemd/system/
$ cp /tmp/CMbkup/usr/lib/systemd/system/psft-prcs-PRCSDOM.service /usr/lib/sys⇒

temd/system/
$ cp /tmp/CMbkup/usr/lib/systemd/system/psft-pia-peoplesoft.service /usr/lib/s⇒

ystemd/system/
$ chown root:root /usr/lib/systemd/system/psft-db-<DB_NAME>.service /usr/lib/s⇒

ystemd/system/psft-appserver-APPDOM.service /usr/lib/systemd/system/psft-prcs-⇒

PRCSDOM.service /usr/lib/systemd/system/psft-pia-peoplesoft.service
$ chmod 755 /usr/lib/systemd/system/psft-db-<DB_NAME>.service /usr/lib/systemd⇒

/system/psft-appserver-APPDOM.service /usr/lib/systemd/system/psft-prcs-PRCSDO⇒

M.service /usr/lib/systemd/system/psft-pia-peoplesoft.service
```

7. Sync files to FS cloud directory.

```
$ mv /cm_psft_dpks/cloud /cm_psft_dpks/cloud_upgbkup
$ cp -r /opt/oracle/psft/pt/ps_app_home/cloud /cm_psft_dpks
$ chown -R root:root /cm_psft_dpks/cloud
$ chmod 755 -R /cm_psft_dpks/cloud
```

8. Verify the permission and ownership of files using below command.

```
$ ls -l /home/psadm1/.bashrc /home/psadm2/.bashrc /home/psadm3/.bashrc /home/o⇒

racle2/.bashrc /home/esadm1/.bashrc /etc/profile /usr/lib/systemd/system/psft-⇒

db-<DB_NAME>.service /usr/lib/systemd/system/psft-appserver-APPDOM.service /us⇒

r/lib/systemd/system/psft-prcs-PRCSDOM.service /usr/lib/systemd/system/psft-pi⇒

a-peoplesoft.service
-rw-r--r--. 1 root    root     3182 May  2 05:43 /etc/bashrc
-rwxr-xr-x. 1 root    root     1908 Apr 29 15:30 /usr/lib/systemd/system/psft-⇒

appserver-APPDOM.service
-rwxr-xr-x. 1 root    root     6891 Apr 29 15:24 /usr/lib/systemd/system/psft-⇒

db-<DB_NAME>.service
-rwxr-xr-x. 1 root    root     1773 Apr 29 15:36 /usr/lib/systemd/system/psft-⇒
```



```
pia-peoplesoft.service
-rwxr-xr-x. 1 root    root     1900 Apr 29 15:33 /usr/lib/systemd/system/psft-⇒

prcs-PRCSDOM.service
-rw-r--r--. 1 root    root     2354 May  2 05:43 /etc/profile
-rw-r--r--. 1 esadm1  oinstall  974 Apr 29 15:12 /home/esadm1/.bashrc
-rw-r--r--. 1 oracle2 oinstall  370 Apr 29 15:12 /home/oracle2/.bashrc
-rw-r--r--. 1 psadm1  oinstall  878 Apr 29 15:12 /home/psadm1/.bashrc
-rw-r--r--. 1 psadm2  oinstall 1097 Apr 29 15:12 /home/psadm2/.bashrc
-rw-r--r--. 1 psadm3  appinst   929 Apr 29 15:12 /home/psadm3/.bashrc
```

9. Start Cloud Manager using below commands or use the psadmin utility.

```
$ sudo systemctl start psft-db-<DBName>.service
$ sudo systemctl start psft-appserver-APPDOM.service
$ sudo systemctl start psft-prcs-PRCSDOM.service
$ sudo systemctl start psft-pia-peoplesoft.service
```

10. Restore the /cm_psft_dpks/cloud folder to all managed instances. This step is required if restoring from a latest Cloud Manager version to an older version. For example, restoring back to Cloud Manager Image 8 from Cloud Manager Image 9.

    Restoring to Linux instances

    a. SSH into Cloud Manager.

    b. Switch user to psadm2.

    ```
    $ sudo su - psadm2
    ```

    c. Securely copy the cloud directory from Cloud Manager to a managed node.

    ```
    $ scp -i /home/psadm2/psft/data/cloud/ocihome/keys/cm_adm_pvt_key -r /cm_⇒

    psft_dpks/cloud/  opc@< Instance1 IPADDRESS>:/home/opc/cloud
    ```

    Where <Instance1 IPADDRESS> is the first item in each row having unix as second field in backed up file /home/psadm2/managedenvironments.txt.

    d. Repeat the above copy step for all IP addresses tagged as unix in the file managedenvironments.txt.

    Restoring to Windows instances

    a. RDP into any Windows instance in the same VCN as the Cloud Manager.

    b. From the above Windows machine, RDP into each Windows instance listed in the backed up file managedenvironments.txt. The password is also captured in the same file.

    c. Access fileserver machine by opening the share \\<file_server_IP>\u01\app\oracle\product.

    d. Copy cloud folder from fileserver into D:\cloud.

Important! Delete the managedenvironments.txt file after completing the cloud folder restores on all nodes.

Deleting Backup

To delete a block volume backup, the user must provide the backup name. User can list the existing backups then select the backup name to delete.

To delete a block (data) volume backup:

1. Run the utility with the restore option.

   ```
   sh cm_backup_restore.sh -o delete -n <backup_name>
   ```

2. User will be prompted to enter a passphrase for the private API signing key.

   Enter the passphrase if it exists or press enter.

3. The backup volume will be deleted.

Example:

```
[opc@smncm cm_backup_and_restore]$ sh cm_backup_restore.sh -o delete -n BLOCK_VOLBK⇒

P_SMNCM_20190410_110419
Please Enter Private Key Pass Phrase :
2019-04-19 10:16:02 -     INFO -  658 : Getting BLOCK volume backup 'BLOCK_VOLBKP_S⇒

MNCM_20190410_110419' details from OCI
2019-04-19 10:16:02 -     INFO -  666 : BLOCK volume backup 'BLOCK_VOLBKP_SMNCM_201⇒

90410_110419' Id : 'ocid1.volumebackup.oc1.iad.abuwcljre6zjpqrmk6ceqim3vu5jnmwt5zw4⇒

zb5vostgq55umw2bij5tt7cq'
2019-04-19 10:16:02 -     INFO - 1146 : Deleting volume backup : BLOCK_VOLBKP_SMNCM⇒

_20190410_110419
2019-04-19 10:16:02 -     INFO - 1157 : Deleted BLOCK volume backup 'BLOCK_VOLBKP_S⇒

MNCM_20190410_110419'
```

Manually Backing Up and Restoring Cloud Manager Using Block Volume Backups for OCI

Using block volume backup feature in OCI, the Cloud Manager data can be backed up and restored on demand.

Backing Up Cloud Manager

To back up the Cloud Manager instance for OCI using block volumes, perform the following:

1. To create a consistent backup, shut down the database, application server, web server (PIA), and Process Scheduler domains.

   Note: Ensure that there are no provisioning or lifecycle jobs running. If there are any such jobs, they will be abruptly ended and may result in environments in an unstable or unusable state.

2. Access the Cloud Manager instance with SSH and run the following commands or use the psadmin utility.

```
$ sudo systemctl stop psft-pia-peoplesoft.service
$ sudo systemctl stop psft-prcs-PRCSDOM.service
$ sudo systemctl stop psft-appserver-APPDOM.service
$ sudo systemctl stop psft-db-<DBName>.service
```

3. Back up the set of files on boot volume that are listed below on to local file system on CM instance or any remote instance.

Note: Use the environment variable PS_CFG_HOME to determine the exact path. Make a note of this in case the path gets modified during PeopleTools upgrade.

- All files under PS_APP_HOME/cloud (/opt/oracle/psft/pt/ps_app_home/cloud)
- All files under PS_CFG_HOME
- /home/psadm1/.bashrc
- /home/psadm2/.bashrc
- /home/psadm3/.bashrc
- /home/oracle2/.bashrc
- /home/esadm1/.bashrc
- /etc/profile
- /etc/bashrc
- /usr/lib/systemd/system/psft-db-<DB_NAME>.service
- /usr/lib/systemd/system/psft-appserver-APPDOM.service
- /usr/lib/systemd/system/psft-prcs-PRCSDOM.service
- /usr/lib/systemd/system/psft-pia-peoplesoft.service

4. On the OCI console, navigate to Compute | Instances | Cloud Manager instance.
5. Navigate to Cloud Manager Instance Details page.
6. Scroll down to the Attached Block Volumes section. Click on the attached volume name which will have a name in the format StorageVol_<CMinstance>_<timestamp>. This volume is available as disk /dev/sdb in Cloud Manager instance. It is mounted on /u01/app/oracle/product, where Cloud Manager application is installed.
7. This will bring up the volume details. On this page, click on 'Create Backup'.
8. Optionally, create a backup of the boot volume in similar way.
9. Provide a name for the backup and click 'Create Backup'.
10. After few minutes a backup is created.

11. Start the database, pia, app and prcs domains. Use below commands or psadmin utility.

```
$ sudo systemctl start psft-db-<DBName>.service
$ sudo systemctl start psft-appserver-APPDOM.service
$ sudo systemctl start psft-prcs-PRCSDOM.service
$ sudo systemctl start psft-pia-peoplesoft.service
```

Restoring Cloud Manager

To restore a backup using block volumes, perform the following:

1. If restoring to an older version of Cloud Manager from a newer version, then generate the list of IP address of all managed instances that were provisioned by Cloud Manager. Follow step 1 in Restoring from a Backup.
2. On the OCI console, navigate to Storage | Backups.
3. Select the backup to restore and click 'Create Block Volume' using menu on the right.
4. Enter a name for the block volume and choose the Availability Domain in which the volume will be created. Ensure to choose the same Availability Domain where Cloud Manager instance is deployed.
5. A new volume is created in few seconds.
6. Access the Cloud Manager instance with SSH and shut down database, pia, app and prcs domains using commands below or psadmin utility.

```
$ sudo systemctl stop psft-pia-peoplesoft.service
$ sudo systemctl stop psft-prcs-PRCSDOM.service
$ sudo systemctl stop psft-appserver-APPDOM.service
$ sudo systemctl stop psft-db-<DBName>.service
```

7. Clean up any running processes that might be using the data volume that needs to be restored.

```
$ ps -ef | grep psadm
psadm2 2969 1 0 Feb01 ? 00:00:19 rmiregistry 10100
psadm2 3495 1 0 Feb01 ? 00:00:20 rmiregistry 10200
$ sudo kill 2969 3495
```

8. Unmount /dev/sdb which is mounted on /opt/oracle/psft.

```
$ sudo umount /opt/oracle/psft
```

9. Navigate to OCI | Compute | Instances | Cloud Manager instance. Scroll down to the Attached Block Volumes. Select the volume to be restored and click Detach.
10. On the Detach Block Volume page, copy all DETACH COMMANDS.
11. Run the detach commands on the Cloud Manager instance.
12. Click 'Continue Detachment' (from step 9) and confirm detachment.
13. Verify in OCI UI for the instance that the volume is now removed.
14. Now restore the volume backup. Click Attach Block Volume. Select ISCSI attachment type. Select the block volume compartment where the backup volume was restored and select the restored volume. Select read-write access mode.
15. Click Attach to attach the restored volume to Cloud Manager instance.

16. After the status shows Attached. Retrieve the iSCSI commands that must be run on the instance to attach the volume in the OS. Click the Actions icon (Actions icon) next to the volume, and then click iSCSI Commands and Information. Copy all ATTACH COMMANDS.

17. Access the Cloud Manager instance with SSH and run the copied attach commands.

18. Verify the disk is attached using "sudo fdisk –l" command. There should now be an entry for /dev/sdb.

```
Disk /dev/sdb: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk identifier: 0x00000000
```

19. Restore below set of files that were backed up from boot volume. Ensure to restore the PS_CFG_HOME files to the right path in case the backup contains files from an older PeopleTools release.

    - Restore /cm_psft_dpks/cloud/ from backup of PS_APP_HOME/cloud
    - All files under PS_CFG_HOME
    - /home/psadm1/.bashrc
    - /home/psadm2/.bashrc
    - /home/psadm3/.bashrc
    - /home/oracle2/.bashrc
    - /home/esadm1/.bashrc
    - /etc/profile
    - /etc/bashrc
    - /usr/lib/systemd/system/psft-db-<DB_NAME>.service
    - /usr/lib/systemd/system/psft-appserver-APPDOM.service
    - /usr/lib/systemd/system/psft-prcs-PRCSDOM.service
    - /usr/lib/systemd/system/psft-pia-peoplesoft.service

20. Run 'mount -a' command on the CM instance and reboot the instance. Check status of Cloud Manager domains using following commands.

```
$ sudo systemctl status psft-db-<DBName>.service
PeopleSoft Container Database CDBHCM Status is Up
PeopleSoft Pluggable Database PSPDB Status is Open
PeopleSoft Database Listener is Up
$ sudo systemctl status psft-prcs-PRCSDOM.service
PeopleSoft Process Scheduler Domain PRCSDOM is Up
$ sudo systemctl status psft-appserver-APPDOM.service
PeopleSoft Application Server Domain APPDOM is Up
$ sudo systemctl status psft-pia-peoplesoft.service
PeopleSoft PIA Domain peoplesoft is Up
```

If database and domains do not come up automatically then start them using the following commands. Reboot only if necessary.

```
$ sudo systemctl start psft-db-<DBName>.service
$ sudo systemctl start psft-appserver-APPDOM.service
$ sudo systemctl start psft-prcs-PRCSDOM.service
$ sudo systemctl start psft-pia-peoplesoft.service
```

If the database or domains don't start successfully, then the restored backup may have issues, In such scenario, there are two options at this point:

a. Restore the original volume. Follow steps 5 to 18 described under 'How to restore a backup' section.

b. Troubleshoot the reason for failures and bring up the database or domains manually.

21. SSH into Cloud Manager and remove the directory /home/psadm2/psft/data/cloud/dm/cache/.

22. Before accessing the restored Cloud Manager PIA URL, you need to clear the application domain cache. To clear cache:

    a. SSH into Cloud Manager instance.

    b. Switch user to psadm2.

    ```
    sudo su - psadm2
    ```

    c. Start psadmin.

    d. Select 1) Application Server.

    e. Select 1) Administer a domain.

    f. Select 1) APPDOM.

    g. Select 8) Purge Cache.

23. If restoring to an older version from a newer version of Cloud Manager, copy the restored /cm_psft_dpks/cloud to all the managed instances. Follow step 10 in Restoring from a Backup.

Note: If you want specific downloads to begin, unsubscribe and then subscribe to the required download channels. Otherwise, the downloads will begin at the next scheduled time.
