

**Oracle Utilities Cloud Services**  
Administration Guide  
For 21A Releases  
**F39181-01**

April 2021

Copyright © 2017, 2021 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

---

---

# Contents

## Chapter 1

Introduction .....	1-1
--------------------	-----

## Part One

### Identity and Access Management

## Chapter 2

Overview .....	2-1
About This Book .....	2-2
Identity Cloud Service Tenancy.....	2-3

## Chapter 3

Quick Start Guide .....	3-1
Activate Security Administrator Account .....	3-2
Evaluate Federated Single Sign-On Requirements.....	3-2
Adjust the Default Oracle Identity Cloud Service Settings.....	3-2
Prepare User Community.....	3-2
Setup Process Summary.....	3-3

## Chapter 4

Security Administrator Account.....	4-1
Setting Up the Security Administrator Account .....	4-2
Navigating to the Identity Cloud Service Admin Console.....	4-3
Accessing via Cloud Account Portal .....	4-3
Accessing Identity Cloud Services Admin Console Directly.....	4-3
Verifying Security Administrator Identity Cloud Service Access.....	4-4
Verifying Subscription Contents .....	4-5
Exploring the Applications .....	4-5
Verifying Access to Object Storage .....	4-6
Verifying Security Administrator Access to Service.....	4-6

## Chapter 5

User Management Procedures .....	5-1
User Onboarding - My Services Portal .....	5-2
Setting Up a New User.....	5-2
Setting Up a New Security Administrator.....	5-3
Updating or Removing a User .....	5-4
Defining User Group Membership.....	5-4
Managing Groups.....	5-4
Advanced User and Access Management - Identity Cloud Service Admin Console.....	5-5
Managing Users .....	5-5
Managing Groups.....	5-6
Managing Applications .....	5-6
Bulk Upload and Download.....	5-7
Updating Settings .....	5-8

Updating Security Privileges .....	5-8
Sign-On Policies for Online Access .....	5-8
Available Reports .....	5-11

## Chapter 6

<b>User Provisioning for Oracle Utilities Cloud Services .....</b>	<b>6-1</b>
Overview .....	6-2
Pre-Defined Application Roles.....	6-3
Configuring Just in Time Provisioning .....	6-3
Setting Up Groups for Provisioning - Identity Cloud Service .....	6-4
Configuring User Provisioning Rules - OUAF .....	6-4
Creating and Provisioning Users.....	6-6
Setting Up an OUAF Security and Access Administrator.....	6-6
Setting Up an Online Application User.....	6-6
Setting Up an Integration User for REST/SOAP Web Services .....	6-6
Setting Up a User with Access to BI Publisher and Analytics Visualization .....	6-7
Setting Up a User with Access to Utilities Testing Accelerator.....	6-7
Setting Up a User Authorized to Execute Ad-hoc SQL Queries .....	6-7
Cloud Service Implementation User.....	6-7

## Chapter 7

<b>User Provisioning for Oracle Utilities Analytic Insights.....</b>	<b>7-1</b>
Overview .....	7-2
Pre-Defined Application Roles.....	7-2
Setting Up Application Users .....	7-2

## Chapter 8

<b>User Management for Oracle Utilities Analytic Insights.....</b>	<b>8-1</b>
Locating Provisioned Users in Oracle Utility Analytics Insights .....	8-2
Assigning Groups and Roles in Oracle Utility Analytic Insights .....	8-2
Assigning and Removing User Group Permissions .....	8-3
Assigning User Group Permissions .....	8-3
Removing User Group Permissions.....	8-3
Assigning User Role Permissions .....	8-3
Removing User Role Permissions .....	8-4

## Chapter 9

<b>Using Federated Single Sign-On .....</b>	<b>9-1</b>
Overview .....	9-2
Setup External Identity Provider.....	9-3
Service Access for Federated Users.....	9-4
Just In Time Provisioning for Federated Users .....	9-4

## Part Two

### Object Storage Setup

## Chapter 10

<b>Introduction .....</b>	<b>10-1</b>
---------------------------	-------------

## Chapter 11

<b>Object Storage Management .....</b>	<b>11-1</b>
Object Storage Structure .....	11-2
Compartments .....	11-2
Object Storage Buckets .....	11-2
Security and Access Management .....	11-3
Accessing the Cloud Infrastructure Console .....	11-3
First Time Login.....	11-4
Managing Users .....	11-4
Managing Groups.....	11-5
Managing Policies.....	11-6

Tenant Information.....	11-7
Regions.....	11-7
API Access.....	11-8
<b>Chapter 12</b>	
<b>Connecting to Oracle Cloud Object Storage .....</b>	<b>12-1</b>
Object Storage Connection Configuration .....	12-2
API Key Management.....	12-3
Registering the API Key.....	12-3
Referencing Files on Object Storage .....	12-4
Using the Bucket Name Prefix.....	12-4
<b>Chapter 13</b>	
<b>Recommended Object Storage Structure for a New Implementation.....</b>	<b>13-1</b>
Security Considerations .....	13-1
Compartments .....	13-1
Users.....	13-2
User Groups.....	13-2
Policies .....	13-2
Recommended Setup for a Single Cloud Service .....	13-3
Oracle Cloud Infrastructure - IAM and Object Storage .....	13-3
Example: Oracle Utilities Customer Cloud Service.....	13-4
Recommended Setup for Multiple Cloud Services.....	13-5
<b>Chapter 14</b>	
<b>Initial Testing of Object Storage Connectivity.....</b>	<b>14-1</b>
<b>Chapter 15</b>	
<b>Cross-Region Disaster Recovery Considerations.....</b>	<b>15-1</b>
Home and Disaster Recovery (DR) regions.....	15-2
Preparing your Disaster Recovery Region.....	15-3
Copying Your Object Storage Bucket Structure .....	15-3
Copying Your Object Storage Data .....	15-3
Recovering from a Disaster.....	15-5
Switching To Your Disaster Recovery Region.....	15-5
Switching Back To Your Home Region .....	15-5
Copying Back Your Object Storage Data.....	15-5
<b>Part Three</b>	
<b>Cloud Operations</b>	
<b>Chapter 16</b>	
<b>Cloud Operations Service Requests .....</b>	<b>16-1</b>
Backup and Restore Policy Overview .....	16-2
Requesting a Full Ad-Hoc Backup.....	16-3
Service Level Objective .....	16-3
Restoring from a Full Ad-Hoc Backup .....	16-3
Service Level Objective .....	16-4
Requesting a Database Restore Point.....	16-4
Service Level Objective .....	16-5
Rolling Back Environment to Restore Point .....	16-5
Service Level Objective .....	16-5
Request for DNS Address List.....	16-6
Service Level Objective .....	16-6
Request to Execute Data Correction Scripts.....	16-6
Service Level Objective .....	16-7
Request for Temporary Additional Threads .....	16-7
Service Level Objective .....	16-8
Request to Export Cloud Environment Schema for On-Premises Use .....	16-8

---

Service Level Objective .....	16-9
Refreshing Data (Cloning).....	16-9
Cloning Process Overview.....	16-9
Submitting a Request for Cloning .....	16-11
Service Level Objective .....	16-12

# Chapter 1

---

## Introduction

This document provides information related to ongoing administration and maintenance of the following Oracle Utilities cloud services:

- Oracle Utilities Billing Cloud Service
- Oracle Utilities Customer Cloud Service
- Oracle Utilities Meter Solution Cloud Service
- Oracle Utilities Operational Device Cloud Service
- Oracle Utilities Rate Cloud Service
- Oracle Utilities Work and Asset Cloud Service

The topics described in this document include:

- [Part One: Identity and Access Management](#)
- [Part Two: Object Storage Setup](#)
- [Part Three: Cloud Operations](#)

# Part One

---

## Identity and Access Management

This section provides instructions for Security Administrators to set up user accounts for Oracle Utilities cloud services, managing end-to-end lifecycle of user identity, and governing user authentication in multiple business applications. User management tasks include creation of the user record, a setup of the application roles, establishing user's membership in the role and, granting user an access to the target business applications, which includes:

- [Chapter 2: Overview](#)
- [Chapter 3: Quick Start Guide](#)
- [Chapter 4: Security Administrator Account](#)
- [Chapter 5: User Management Procedures](#)
- [Chapter 6: User Provisioning for Oracle Utilities Cloud Services](#)
- [Chapter 7: User Provisioning for Oracle Utilities Analytic Insights](#)
- [Chapter 8: User Management for Oracle Utilities Analytic Insights](#)
- [Chapter 9: Using Federated Single Sign-On](#)



# Chapter 2

---

## Overview

End user provisioning involves creating user records and granting appropriate access for users of Oracle Utilities cloud services.

This guide provides instructions for Security Administrators to set up user accounts for Oracle Utilities cloud services, managing the end-to-end lifecycle for user identity, and governing user authentication in multiple business applications. Identity management tasks include creation of the user and user group records, granting users and groups an access to the target business applications, and managing various security settings.

This guide also provides an introduction to working with Oracle Identity Cloud Service (). Identity Cloud Service is provisioned to customers with subscriptions to Oracle Utilities cloud services. Customers receive an instance of Identity Cloud Service (also referred to as Identity Cloud Service tenancy). The tenancy is managed exclusively by the customer (see **Identity Cloud Service Tenancy** on page 2-3 for more information).

This chapter includes the following:

- [About This Book](#)
- [Identity Cloud Service Tenancy](#)

# About This Book

This guide contains the following:


- [Chapter 2: Overview](#) (this chapter)
- [Chapter 3: Quick Start Guide](#) provides an overview of the end user provisioning process, with references to additional information in the following chapters.
- [Chapter 4: Security Administrator Account](#) describes how to set up a security administrator account for user provisioning.
- [Chapter 5: User Management Procedures](#) describes general procedures related to managing users and groups.
- [Chapter 6: User Provisioning for Oracle Utilities Cloud Services](#) describes specific tasks related to user provisioning for Oracle Utilities cloud services.
- [Chapter 7: User Provisioning for Oracle Utilities Analytic Insights](#) describes specific tasks related to user provisioning for Oracle Utilities Analytic Insights.
- [Chapter 8: User Management for Oracle Utilities Analytic Insights](#) describes specific tasks related to user management for Oracle Utilities Analytic Insights.
- [Chapter 9: Using Federated Single Sign-On](#) describes tasks required when using an external identity management system to provide authentication for the application instances within your cloud subscription.

# Identity Cloud Service Tenancy

Identity Cloud Service tenancy is provided to the customer as part of the service subscriptions.

The following configurations are defined in Identity Cloud Service:

- **Application:** In Oracle Utilities cloud services the application represents a single environment, Production or non-Production. Applications are created by the subscription provisioning process.
- **Application Role:** In Oracle Utilities cloud services the Application Role represents an entitlement to access a component within the environment. Assigning user to an Application Role provides this user with access to this component. Application Roles are created by the subscription provisioning process.
- **User:** Users represent a human or non-human entity that is accessing the environment. User records are created and managed by the Security Administrator.
- **Group:** Groups comprise of one or more users. Groups are created and managed by the Security Administrator.



# Chapter 3

---

## Quick Start Guide

This chapter provides an overview of the initial set up of your cloud server user community including:

- [Activate Security Administrator Account](#)
- [Evaluate Federated Single Sign-On Requirements](#)
- [Adjust the Default Oracle Identity Cloud Service Settings](#)
- [Prepare User Community](#)
- [Setup Process Summary](#)

---

# Activate Security Administrator Account

Access the Oracle Identity Cloud Service (IDCS) Admin console and perform the verification of the provisioned environments. Follow the steps described in [Chapter 4: Security Administrator Account](#).

## Evaluate Federated Single Sign-On Requirements

If you are using IDCS as your only identity management system, proceed with adjusting the IDCS cloud settings and perform the user community setup. Otherwise if the user identities are managed by an existing enterprise identify management system then evaluate any Federated Single Sign On (SSO) requirements. If the federation is required for all user accounts, including the implementation team, immediately proceed with the federated SSO setup as described in [Chapter 9: Using Federated Single Sign-On](#). Otherwise, if the federation is required for the actual production users only, it may be configured in the later stages.

## Adjust the Default Oracle Identity Cloud Service Settings

Modify Oracle Identity Cloud Service (IDCS) settings as follows:

- Define your user naming conventions: decide whether the primary email address will be also used as a user name (login).
- You may want to include user name in the communication emails. Update notification(s) accordingly (see **Notification Update Example: Welcome Email** on page 5-8 for an example of updating notifications).
- Update the notifications further to include additional details, for example the contact information of the technical support team.
- Evaluate the default Password Policy and amend according to your organization's requirements.
- Customize the look of the IDCS login page with your company's branding elements (optional).

See **Updating Settings** on page 5-8 for more information about updating settings.

## Prepare User Community

Explore the Users list. Beside the Security Administrator account you may find a Process Automation user. This is created as part of the service provisioning and usually linked to the Security Administrator's email address. Process Automation is an internal user for inter-domain communications.

Take advantage of the IDCS's user import feature to quickly establish user access to the provisioned environments, using the following steps:

- Determine the list of users who'll be accessing the provisioned environment(s):
  - Provide access to the non-production environments for key members of the implementation team

- Provide access to the production environment users
- Define IDCS Group(s) for Just-In-Time Provisioning (if required). See **Setting Up Groups for Provisioning - Identity Cloud Service** on page 6-4 for more information).
- Browse the IDCS Applications and determine Application Roles that users will be assigned to.
- Download the bulk upload template files from IDCS and create import files for:
  - Users
  - Groups
  - Application Roles

See **Bulk Upload and Download** on page 5-7 for more detailed information about uploading and downloading template files.

## Setup Process Summary

Note: The following assumes the Security Administrator account has been activated.

- If you wish to delegate the just-in-time provisioning and access/authorization setup, assign the IDCS administrator role to at least one user per environment (see **Updating Security Privileges** on page 5-8).
- Access the environment and configure Just-In-Time provisioning according to the product's specifications (see **Configuring User Provisioning Rules - OUAF** on page 6-4).
  - For example setup the IDCS Integration Master Configuration for CCS. Make sure the IDCS Groups are the same Groups that were used for the User/Group import files.
- Perform import of Users, Groups and Application Roles using the import files prepared above (see **Bulk Upload and Download** on page 5-7).
- Setup at least one integration (non-human) user per environment of CCS and communicate the credentials to the implementation team (see **Setting Up an Integration User for REST/SOAP Web Services** on page 6-6)
- Setup access to production environment for those users who are responsible for legacy data migration

# Chapter 4

---

## Security Administrator Account

This chapter how to set up a security administrator account for user provisioning, including:

- [Setting Up the Security Administrator Account](#)
- [Navigating to the Identity Cloud Service Admin Console](#)
- [Verifying Security Administrator Identity Cloud Service Access](#)
- [Verifying Subscription Contents](#)
- [Exploring the Applications](#)
- [Verifying Access to Object Storage](#)
- [Verifying Security Administrator Access to Service](#)

---

# Setting Up the Security Administrator Account

The account for the Security Administrator is created during the tenancy provisioning. The customer provides the name and the email address of the intended security administrator as part of the service order.

Once the order is completed the Security Administrator receives a cloud account activation email.

The activation email contains:

- Activation URL
- The user name and the temporary one-time password

Security administrators should use the following procedure the first time logging into the **Oracle Cloud Account Portal**:

1. Press the activation link or copy the link into the internet browser's address.

You will be redirected to the login page.

2. Enter the user name and the temporary password.
3. Follow the prompts to create a new permanent password.

Finally, you will be redirected to the **Oracle Cloud Account Portal** dashboard.



---

# Navigating to the Identity Cloud Service Admin Console

The Identity Cloud Service Admin Console can be accessed either directly or via Cloud Account Portal.

## Accessing via Cloud Account Portal

On the **Oracle Cloud Account Portal** dashboard click **Users** in the top right corner of the screen.

On the **Users** tab click **Identity Console**. You'll be redirected to the **Identity Cloud Services** console.

Click the menu icon at the left top corner to expand the left-side navigation pane

## Accessing Identity Cloud Services Admin Console Directly

After navigating to the Admin console for the first time you can copy the URL from the internet browser address bar and bookmark it for the further use.

The URL is structured as follows:

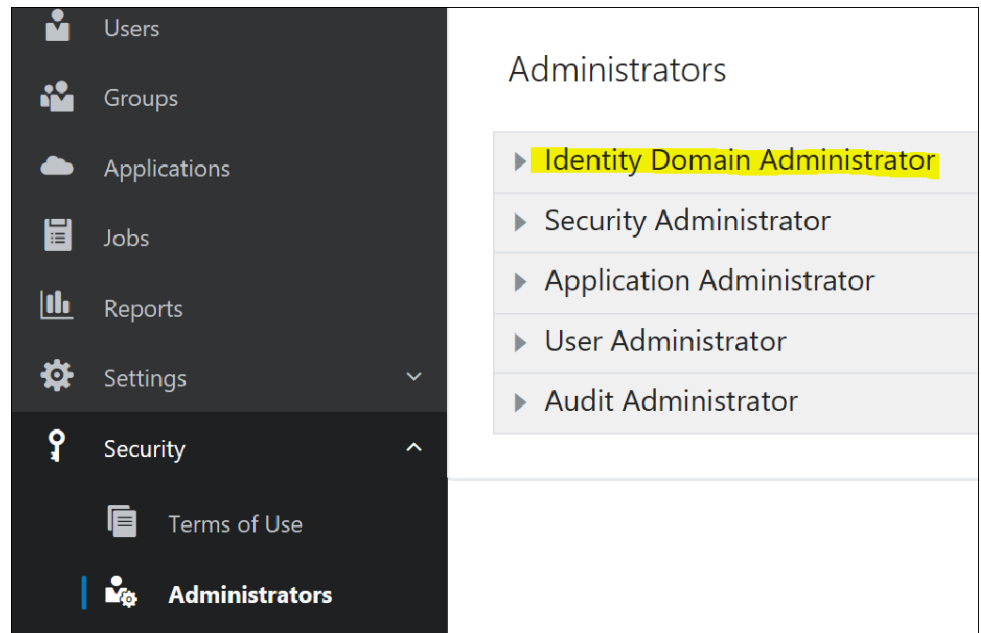
```
https://< tenancy>/ui/v1/adminconsole
```

where < tenancy> represents the instance of the IDCS that belongs to the customer's subscription.

In this scenario the user is re-directed to the **Identity Cloud Services** admin console dashboard. Use the menu icon on the left top corner to expand the navigation pane. You can also browse various help topics listed in the upper section of the page.

# Verifying Security Administrator Identity Cloud Service Access

Expand the **Security** topic on the navigation pane and click **Administrators**.



On the page, click **Identity Domain Administrator** and verify that your name is on the list of Identity Domain Administrators.

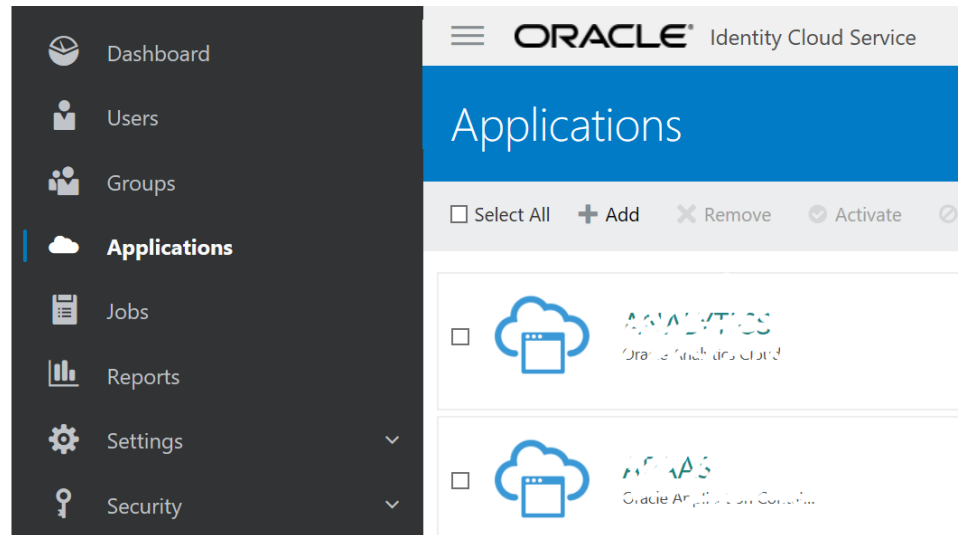
## Verifying Subscription Contents

Click **Oracle Cloud Services** (**Applications** in earlier versions) on the navigation pane. The main panel displays a list of available applications.

Each Application in represents an environment, for example Production or Test.

**Note:** A typical subscription includes one Production environment, and at least one Development and one Test environment. The number of environments depends on specific customer requirements and may include multiple Development and/or Test instances

The list of applications may also include an instance of Oracle Cloud Object Storage.



## Exploring the Applications

Click on one of the applications on the list and display the single application. Most of the information is system-generated and read-only. Users and Groups should be assigned to Application Roles within the application in order to gain access to the environment.

Click the **Application Roles** tab and review available Application Roles.

While the application represents a single environment, the different Application Roles represent different components within the environment. In order to authorize user's access to a certain component the user has to be assigned to a corresponding Application Role. Application Roles include:

- Online Application Access
- Web services REST/SOAP API
- Access to supporting Applications such as BI Publisher and SQL Developer Web

Application Roles also used to support coarse-grained authorization in the target component, for example the BI Content Author versus an ordinary BI Consumer.

## Verifying Access to Object Storage

Refer to the *Oracle Utilities Cloud Services Object Storage Setup Guide* in your service's documentation library for more information about object storage.

## Verifying Security Administrator Access to Service

As part of the service activation notifications, the security administrator is provided with URL-s for all components within Production and Non-Production environments.

Perform the following steps to verify the access:

- Assign the security administrator user to online-related Application Roles in all environments (Application Role description indicates whether the access is given for online or for the REST/SOAP API)
- Try to access the URL-s for the online applications

See Oracle Utilities Integration documentation for more details on how to verify API access.

# Chapter 5

---

## User Management Procedures

This chapter general procedures related to managing users and groups, including:

- [User Onboarding - My Services Portal](#)
- [Advanced User and Access Management - Identity Cloud Service Admin Console](#)
- [User Onboarding - My Services Portal](#)

# User Onboarding - My Services Portal

The basic user access management operations can be performed directly on the **My Services** portal on the **Oracle Cloud Account Portal**.

The link to the **User Management** portal is located on the upper navigation bar.

## Setting Up a New User

Click **Add** on the **Users** tab of the **User Management** portal to set up a new user.

### Add User Details

Enter the minimum required information:

- Last Name
- First Name
- Email address

**Note:** By default the email address is used as the user name. Uncheck **Use Email as User Name** to enter the User Name manually

- User Name

Click **Next** to set up the user's Service Access

### Define Access to Service

The **Service Access** page displays a list of environments and services.

Locate the environment in the list or use Search to filter out a specific environment.

To add one or more roles for an environment, click on the field beneath the environment's name.

To add all available roles at once click **Add User Roles**.

Click **Finish** button to complete the setup.

The new user appears on the **User Management** portal.

Note: Additional product-specific setup may be required in order to provide user authorization and Just In Time provisioning. See [Chapter 6: User Provisioning for Oracle Utilities Cloud Services](#) and [Chapter 7: User Provisioning for Oracle Utilities Analytic Insights](#) for more information.

## Setting Up a New Security Administrator

The new security administrator is configured as follows:

- Add new user record as shown above
- Grant administrative role(s) to the new user.


### Cloud Account Administrator

The Cloud Account administrator's is able to manage every aspect of the subscription including but not limited by Identity Cloud Service administration.

The Security Administrator is assigned this role.

To grant the same privileges to the new user:

- Filter the services list Cloud Account service on the list
- Select the Cloud Account Administrator role.



### Identity Administrator

Identity administration roles authorize users to manage configurations and administer Identity Cloud Service. There are various level of access:

- User Administrators are allowed to create and manage users and groups.
- The Application Administrator role is limited to the Application configuration and lifecycle.
- Audit and Security Administrator roles provide access to basic security settings and Identity-related reports.
- The Identity Domain Administrator role includes all of the above.

In order to grant user the administrative role in Identity Cloud Service:

- Filter the services list and locate an Identity Cloud service.
- Select one or more roles from the list or click **Add Admin Roles** to add all available roles at once.

## Updating or Removing a User

User records are displayed on the **User Management** portal.

### Update User Details

To update details for a user, double-click the user or click on the action menu icon to open the user record and update the user information as appropriate.

**Note:** First and Last names are editable. The email address is editable only if not used as user name (login)

### Update Access to Service

To update the access a user has to services, modify the existing user's access to services by adding or removing roles.

### Remove User

To remove a user, click **Remove** from the menu.

**Note:** Removing a user is irreversible.

## Defining User Group Membership

Select one or more user records and the multi-record actions became available:

- **Add to Group:** Adds selected users to an existing group
- **Create New Group:** Creates a new Group and adds the selected users to it
- **Clear Selections:** Deselects selected users



## Managing Groups

Click the **Groups** tab on the **User Management** portal. The portal displays a list of all available group.



## Add New Group

To add a new group click **Add**.

Enter the **Group Name** and **Description** and save the new group.

## Add Users

To add users to a group, click on the group name on the list or use the Edit menu action. The portal displays the selected group record.

Click the **Users** tab, then click **Add Users** to add one or multiple users to the group.

## Group Access to Service

Click the **Roles** tab.

The access setup steps are similar to setting up an individual user's access.

The portal displays a list of available environments and services. Filter the list and assign group to one or more roles.

In order to setup a group with administrator privileges, locate Identity Cloud on the list and add one or more administrative roles to the group.

# Advanced User and Access Management - Identity Cloud Service Admin Console

Use the **Identity Cloud Service** admin console to manage applications, perform user management and administer general and security settings also view basic reports.

## Managing Users

Users can be added and maintained via **Identity Cloud Service** admin console. Access the **Users** portal from the **Identity Cloud Service** admin console dashboard or from the navigation bar.

To add a new User, click the **+Add** button and populate the required user details. On the next screen you can also immediately add user to one or more Groups

In addition to add and remove, the following multi-record actions are available on the **User** page:

- Resend Invitation
- Reset Password
- Activate/Deactivate User
- Update User information and preferences (on individual User record)
- Unlock User (on individual User record)

## Resend Invitation to Service

The initial email invitation to access the service is sent to the user immediately upon user record creation. This invitation is expired after certain period of time.

## Reset Password

Resets a single, multiple, or all passwords. Users will receive a password reset email notification immediately

## Activate/Deactivate User

User can be temporarily activated or deactivated. The email notification is sent to the user immediately.

If the deactivation lasts longer than the password rotation period the activation will cause password reset.

## Update User Information and Preferences

Updates details for individual users. In addition to the minimum required information provided during user creation the following details can be updated:

- Title
- Time Zone and Address including Country
- Preferred language
- Alternative email and contact information

## Unlock User

Unlocks a locked user account. The user's account may be locked for various reasons for example after too many unsuccessful login attempts.

Select **Unlock User** from the **More** menu to unlock the locked account.

## Managing Groups

Users and groups can be added and maintained using the Identity Cloud Service admin console.

Access the **Groups** portal from the Identity Cloud Service admin console dashboard or from the navigation bar.

Select one or more entries from the list.

In addition to add and remove, the following actions are available:

- Import Groups
- Export Groups

## Managing Applications

The applications that represent the provisioned services are pre-created during the service order processing. The Application Roles are also pre-configured.

The administrator is authorized to activate or deactivate certain applications, assign users to Application Roles and also perform import and export of application role's members.

## Bulk Upload and Download

IDCS supports import and export of users, groups and application roles membership. The bulk identity data operations may be required for the fast user onboarding or as part of the federated single sign on setup.

The **Import** and **Export** actions are available on multiple Admin Console pages:

- **Users** page:
  - Import all or a selected set of users
  - Export information for one or more users
- **Groups** page:
  - Import all or a selected set of groups and their member users
  - Export one or more groups and their member users
- **Application > Application Roles** page:
  - Import all or a selected set of application role's membership (users and groups)
  - Export one or more application role's membership (users and groups)

### Importing

1. Navigate to the **Users, Groups, or Applications (Application Roles tab)** page as appropriate.
2. Click **Import** on the top actions bar.
3. Download the sample file.
4. Review the sample file. Note that you can provide different type of information:
  - Users
  - Groups
  - Application Roles Membership
5. Populate the file with user's data and save.
6. Import the file into **Identity Cloud Service**.

### Exporting

1. Navigate to the **Users, Groups, or Applications (Application Roles tab)** page as appropriate.
2. Select entries for the export.
3. Click **Export** on the top actions bar

A notification email is sent as soon as the export job is completed and the file is available for the download.

## Updating Settings

Use the navigation bar to expand the **Settings** topic. The following settings can be modified:

- **Default Settings:** Used to manage default time zone, language and audit setup
- **Session Settings:** Used to manage session expiration
- **Password Policy:** Used to amend the default password policy according to your requirements
- **Notifications:** Used to modify the default email notification templates provided with Identity Cloud Service

### Notification Update Example: Welcome Email

The email notification templates are provided for multiple identity management-related events. The default content of these notifications can be amended to reflect customer's business requirements.

For example, there are two approaches to user account creation: using email address as a user name as opposed to using a manually defined user name. The former means the user knows what to specify on the login screen (email address). The later means the user name that is created manually by the security administrator has to be communicated to the user. In order to communicate the **user name** in the **Welcome** email perform the following steps:

- Select **Notification** on the left-side navigation bar
- Click on the **Email Templates** tab
- Expand the **Welcome** template:  
In the email body the greeting line reads: Hello \${user.displayName}
- Modify the greeting to include the user name (login) as follows:  
Hello \${user.displayName} (\${user.userName})

Note that other substitution variables are also available for use in the notifications. To explore the variables available to a specific template click the **Email Variables** link above the email body editor.

## Updating Security Privileges

Use side navigation panel to expand the **Security** topic. Use **Administrators** link to add or remove administrative privileges from the users.

## Sign-On Policies for Online Access

IDCS supports the ability to restrict web-browser-based access to the applications based on set of conditions including the user's client IP addresses. Both IP "blocklist" and "allowlist" approaches are supported.

- A blocklist defines a set of IP addresses that are blocked from the access. This approach should be used when the "bad" IP-s are well-known and permanent and the list is not expected to change very often.

- An allowlist defines the set of IP addresses that are permitted to access the application while everybody else is denied access.

In addition to IP addresses the following can be allowed or blocked:

- Specific users
- Groups
- User's administrative role in IDCS
- User being authenticated by a specific external identity provider(s)

**Note:** Sign-On Policies are applied ONLY when user attempts to authenticate to IDCS using a web browser. They are not applicable for requests submitted via REST/SOAP API.

## Setup a Network Perimeter

A Network Perimeter represents a set of IP addresses, and can be defined as:

- A list of one or more IP addresses
- A range of IP addresses
- One or more IP addresses in IPv4 CIDR notation, which encompass all IP addresses belonging to a subnet. You can also use the IPv4 CIDR notation to refer to the entire internet: 0.0.0.0/0.

Create Network Perimeters:

- Use side navigation panel to expand the Security Topic
- Locate Network Perimeters
- Add one or more Network Perimeters that define "blocklist" and/or "allowlist" IP addresses

## Setup Sign-On Policies

Sign-on policies define the set of rules used for granting the access to the applications. The out-of-box default policy contains a single default rule that grants the access to every authenticated user. You can either modify the default policy or create a new one(s).

Sign-on policy rule definition includes multiple optional conditions to filter the users and an action to **allow** or **deny** the access:

- By authenticating the Identity Provider: Denying/allowing access for users authenticated by specific external IP in case of a federated SSO
- By group membership: Denying/allowing access for specific set of groups
- By being or not being an IDCS administrator
- By being one of the explicit list of users
- By the user client's IP address being in one or more of the Network Perimeters

The rules on the policy are evaluated top-to-bottom. The first result halts the evaluation. Meaning if the user satisfies the rule's condition, the rule's action (**allow** or **deny** access) is applied and evaluation ends.

**Note:** the default rule on the default policy cannot be deleted, therefore it has to be modified first.

**Example:**

Let's assume that the requirement is to:

- Allow access from IP addresses on the company's intranet
- In addition, allow certain administrators to connect from their personal home computers
- Block anyone else

To configure this example:

- Create two new Network Perimeters:
  - **NP1-Company** to represent the intranet: specify the an entire subnet using CIDR notation, like, for example, 10.10.0.1/24, which means all addresses in 10.10.0 subnet
  - **NP2-Admins:** specify one or more IP addresses, comma-separated
- Configure Default Sign-On Policy:
  - Modify Default Rule:
    - Set the rule's "*and the user's client IP address is*" condition to "*in one or more of these network perimeters*" and specify **NP1-Company**
    - Set the rule's action to "**Allowed**"
  - Add new Rule:
    - Set the rule's "*And is an administrator*" condition to "true"
    - Set the rule's "*and the user's client IP address is*" condition to "*in one or more of these network perimeters*" and specify **NP2-Admins**
    - Set the rule's action to "**Allowed**"
  - Add new Rule
    - Set the rule's "*and the user's client IP address is*" condition to "*Anywhere*"
    - Set the rule's action to "**Denied**"

**Sample Sign-in Scenarios:**

**Scenario 1:** An employee is trying to login from the office computer that is connected to the intranet.

- The first rule (the default rule) is evaluated first. The user's IP satisfies the condition by being on the NP1-Company perimeter. The rule's action ("*Allowed*") is applied and the user is allowed to sign in.

**Scenario 2:** The administrator is trying to login with admin's user name from a personal computer whose IP is listed in NP2-Admins perimeter.

- The first rule (the default rule) is evaluated first. The user's IP does not satisfy the condition by being on the NP1-Company perimeter
- The second rule is evaluated. The user's IP does satisfies both conditions: being and administrator and being on the Np2-Admins perimeter.

- The rule's action ("*Allowed*") is applied and the user is allowed to sign in

**Scenario 3:** The employee is trying to connect from the home computer.

- The first rule (the default rule) is evaluated first. The user's IP does not satisfy the condition by being on the NP1-Company perimeter.
- The second rule is evaluated. The user's IP does not satisfy any of the conditions: being neither an administrator nor being on the Np2-Admins perimeter.
- The third rule is evaluated. The user's IP satisfies the "Anywhere" condition.
- The rule's action ("*Denied*") is applied and the sign in is blocked. The IDCS login error message: "Sign-on policy denies access." is displayed.

Refer to IDCS documentation for the detailed instructions regarding Sign-On Policy and Network Perimeter setup.

## Available Reports

The following **Identity Cloud Service** reports are available for review and download:

- Successful Login Attempts
- Unsuccessful Login Attempts
- Application Access
- Granted and Revoked Application Roles

# Chapter 6

---

## User Provisioning for Oracle Utilities Cloud Services

This chapter describes user provisioning for Oracle Utilities cloud services, including:

- [Overview](#)
- [Pre-Defined Application Roles](#)
- [Configuring Just in Time Provisioning](#)
- [Creating and Provisioning Users](#)
- [Cloud Service Implementation User](#)

The information in this chapter applies to the following cloud services:

- Oracle Utilities Customer Cloud Service
- Oracle Utilities Meter Solution Cloud Service
- Oracle Utilities Work and Asset Cloud Service
- Oracle Utilities Operational Device Cloud Service
- Oracle Utilities Rate Cloud Service
- Oracle Utilities Billing Cloud Service



---

# Overview

Each Oracle Utilities cloud service environment included in the subscription contains multiple components:

- Business Applications that run on the Oracle Utilities Application Framework, (OUAF)

Oracle Utilities Application Framework supports fine-grained authorization to access various features within the Business Application. It stores users and user groups.

For each user authorized to access Oracle Utilities Application Framework the corresponding application user is created in the Oracle Utilities Application Framework.

For the online application access, Oracle Utilities Application Framework users are created through Just in Time Provisioning flow.

**Important Note:** Identity Cloud Service user names (login id) are **case-sensitive**. This information should be communicated to users in order to avoid authorization and authentication issues.

- Supplemental components such as BI Publisher that don't maintain their own user records and support role-based authentication and authorization.

## Pre-Defined Application Roles

The following roles are pre-defined in the Applications that represent Oracle Utilities service environments. Each role represents an entitlement within the environment and grants user an access to a certain component:

Application Role	Authorized Access
Online Application User	The user assigned to this role may access online application.
Web Services Access	The user assigned to this role is authorized to access the REST/SOAP API-s.
BI Consumer	The user assigned to this role may access the BI Publisher within the environment and view and execute pre-defined reports.
BI Content Author	The user assigned to this role may access the BI Publisher within the environment and author new and view/execute existing reports.
Test Administrator	The user assigned to this role may access the Utilities Testing Accelerator within the environment and perform administrative tasks as well as develop and approve other's work.
Test Approver	The user assigned to this role may access the Utilities Testing Accelerator within the environment and develop and approve Test Components and Test Flows.
Test Developer	The user assigned to this role may access the Utilities Testing Accelerator within the environment and develop Test Components and Test Flows.
SQL Developer Web Online User	The user assigned to this role may access the SQL Developer Web online and query the database to retrieve the information from both production and conversion schema
REST Enabled SQL	The user assigned to this role may use cURL utility to invoke REST services and query the database to retrieve the information from both production and conversion schemas.

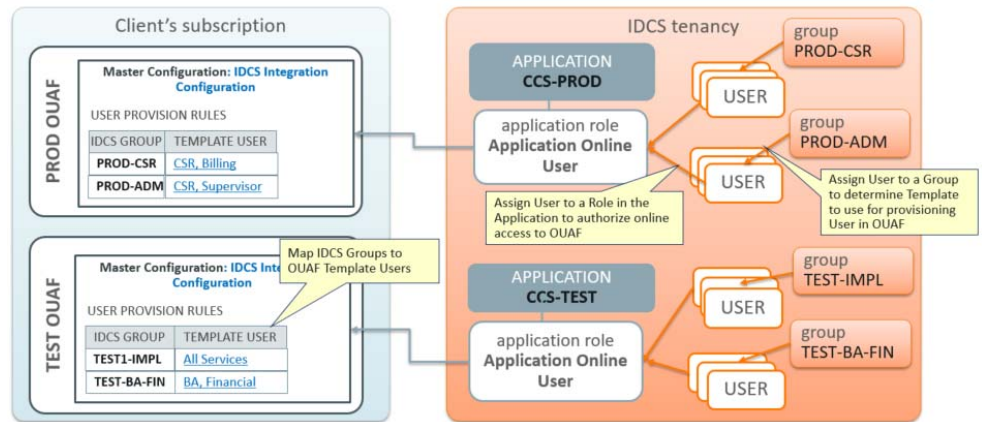
**Note:** Additional pre-defined roles for Oracle Utilities Analytics Visualization may be provided with specific cloud services. Examples include *CustomerContentCreator*, *CustomerContentConsumer* among others.

## Configuring Just in Time Provisioning

Just In Time provisioning is a process that creates application user record in the OUAF-based business applications upon first successful login.

The new user is created in the business application based on a pre-defined OUAF Template User.

The Template User is determined from the mapping between Groups and OUAF Template Users defined in **Integration Configuration**.



Steps to configure Just In Time provisioning:

- Assign Security Administrator user to a Online Application Role in the environment (this is required to access the OUAF with access administrator privileges)
- [Setting Up Groups for Provisioning - Identity Cloud Service](#)
- [Configuring User Provisioning Rules - OUAF](#)

## Setting Up Groups for Provisioning - Identity Cloud Service

Create Groups in Identity Cloud Service that represent broad functional areas and/or authorization level in the service. For example:

- For Non-Production (Development and Testing) environments:
  - Implementers
  - Business Analysts
  - QA Team
  - Security Testing
  - Functional Testing
- For Production environments:
  - Call Center
  - Call Center Supervisor
  - Business Administrator
  - Accounting

## Configuring User Provisioning Rules - OUAF

To configure Identity Cloud Service Integration in OUAF:

- Create Template Users that represent various level of access authorization
- Review existing Template Users.

If your intention is to use a Template User to provision integration (non-human) users you might have to assign Default Access Group to the Template User.

- Map the Groups created above to the Template Users in OUAF in the IDCS Integration Master Configuration.

If the IDCS Integration Master Configuration is not configured at the time the user record is created, the user will be provisioned with K1MINACS (default minimal access)

---

# Creating and Provisioning Users

This section describes steps involved in creating users and providing access to the cloud service's various components.

## Setting Up an OUAF Security and Access Administrator

Perform the following steps:

1. Create a new user or search for and select an existing user.
2. Assign this user to the User Administrator role. See **Setting Up a New Security Administrator** on page 5-3 for more details.
3. After first login to OUAF this user will be provisioned with Template User K1SCRADM (security administrator).

## Setting Up an Online Application User

Perform the following steps:

1. Create a new user or search for and select an existing user.
2. Assign the user to the group that represents the appropriate level of authorization for the environment.
3. Locate the application that is corresponding to the environment. Assign the user to the Online Application User role in the environment.

## Setting Up an Integration User for REST/SOAP Web Services

REST/SOAP API doesn't perform Just-In-Time provisioning. Users for web services must be created manually in both and OUAF applications.

An email address must be provided as part of user creation:

- It is recommended that this email address is used for non-human user setup only
- All email notifications concerning user account are sent to this email address
- Security administrator must have an access to this email account

Perform the following steps:

1. Create a new user or search for and select existing user
  - Specify the email address allocated for the integration/non-human users.
  - When the activation email is received, reset the user's password and communicate the email address and password to the integration team.
2. Assign the User to the REST/SOAP Web Services role in the Application that represents the environment.
3. Login to OUAF and create a new User with Login ID = User Name. Assign the user to user groups that provide access to all or selected application services, according to the business requirements.

## Setting Up a User with Access to BI Publisher and Analytics Visualization

Perform the following steps:

1. Create a new user or search for and select an existing user.
2. Locate the application that is corresponding to the environment. Assign the user to one of the Application Roles available in the environment:
  - **BI Publisher:** Choose one (or both) of the following application roles:
    - BI Consumer
    - BI Content Author
  - Choose one or more product-specific application roles related to Oracle Utilities Analytics Visualization features, such as CustomerContentConsumer. See [Administer the Service](#) in the *Oracle Utilities Analytics Visualization User Guide* for more information about application roles used with Oracle Utilities Analytics Visualization.

## Setting Up a User with Access to Utilities Testing Accelerator

Perform the following steps:

1. Create a new user or search for and select an existing user.
2. Locate the application that is corresponding to the environment. Assign the user to one of the Utilities Testing Accelerator roles available in the environment:
  - Test Administrator
  - Test Approver
  - Test Developer

## Setting Up a User Authorized to Execute Ad-hoc SQL Queries

Perform the following steps:

1. Create a new user or search for and select an existing user.
2. Locate the application that is corresponding to the environment. Assign the user to one of the following roles:
  - **SQL Developer Web Online User:** Provides access to the online web-based interface that enables user to execute queries
  - **Rest Enabled SQL:** Provides the ability to execute REST calls using cURL command

## Cloud Service Implementation User

The environment provisioning process creates an internal (non-human) user account named "K1IPROCESS" that is used by cloud service implementation tools and processes, including configuration migration between environments.

# Chapter 7

---

## User Provisioning for Oracle Utilities Analytic Insights

This chapter describes user provisioning for Oracle Utilities Analytic Insights, and includes the following:

- [Overview](#)
- [Pre-Defined Application Roles](#)
- [Setting Up Application Users](#)

## Overview

Oracle Utilities Analytic Insights features internal user access management that grants users the access to various features within the application. It stores and maintains users.

For each user that authorized to access Oracle Utilities Analytic Insights the corresponding application user is created in OUAL.

The Oracle Utilities Analytic Insights user is created thru Just in Time Provisioning flow.

## Pre-Defined Application Roles

The following roles are pre-defined in the Applications that represent Oracle Utilities Analytic Insights environments:

Application Role	Authorized Access
Online Application User	The user assigned to this role may access online application
Online Application Admin	The user assigned to this role may access the online application and also perform the administrative tasks in Oracle Utility Analytic Insights.
BI User	The user assigned to this role may access the Oracle Business Intelligence Enterprise Edition (OBIEE) component within the environment

In addition to the above listed roles, BIConsumer, BIContentCreator, and Analytics Visualization roles may also be provisioned.

## Setting Up Application Users

Perform the following steps to provision new user:

1. Create a new user or search for and select an existing User.
2. Locate the application that is corresponding to the Oracle Utilities Analytic Insights environment.
  1. Assign the user to one of the Online Application roles.
  2. Assign the user to the BI User role



# Chapter 8

---

## User Management for Oracle Utilities Analytic Insights

Once an Identity Cloud Service user has been created and assigned to the application role(s), the user may access Oracle Utilities Analytic Insights online.

Upon first successful login, the user record is created in the Oracle Utilities Analytic Insights user list, and is ready to be configured and assigned to Oracle Utilities Analytic Insights groups and roles. Groups and roles determine the user features and functionality available to each user.

This chapter outlines specific steps that need to be performed to configured users in Oracle Utility Analytic Insights including:

- [Locating Provisioned Users in Oracle Utility Analytics Insights](#)
- [Assigning Groups and Roles in Oracle Utility Analytic Insights](#)
- [Assigning and Removing User Group Permissions](#)

# Locating Provisioned Users in Oracle Utility Analytics Insights

Once a user has been created in Identity Cloud Service and logged into the application for the first time, it will appear in the Oracle Utility Analytic Insights user list. See [Chapter 7: User Provisioning for Oracle Utilities Analytic Insights](#).

**Note:** You need to have customer administration rights to complete this task. Verify that your user has been assigned to the Application Admin role in the application that is corresponding to the environment in Identity Cloud Service.

Use the following procedure to locate a user.

1. Log in to Oracle Utility Analytic Insights.
2. Select **Administer**, then **Security**, then **Users** to navigate to the **Administer Users** page.
3. Search for the user you created in Identity Cloud Service by completing one of the user information fields and then clicking **Get Users**.

You may search for a user by any data entered when creating the user.

The data table will return with the user information and links that allow you to assign their user environment. See the *Oracle Utilities Analytic Insights User Guide* ([https://docs.oracle.com/cd/E72219\\_01/documentation.html](https://docs.oracle.com/cd/E72219_01/documentation.html)) for additional information about the **Administer User** page.

4. Click **Edit** in the user row to open the **Manage Users** dialog box. The **Manage Users** dialog box allows you to modify group and role permissions.

## Assigning Groups and Roles in Oracle Utility Analytic Insights

User environment access is managed through the Add Group and Add Role functions located in the **Manage Users** pane.

- **Add Group:** Determines general user interface characteristics (for example, the menus that are displayed) and, consequently, which pages are accessible to the user and sets of users.
- **Add Role:** Assigns user roles and determines the features that are available on the pages made available by the user's group privileges.

Roles are associated with modules. Assigning a role automatically associates the user to a module. The following table provides an example of possible user role to module associations. See the *Oracle Utilities Analytic Insights User Guide* ([https://docs.oracle.com/cd/E72219\\_01/documentation.html](https://docs.oracle.com/cd/E72219_01/documentation.html)) for more information.

Module	Role
Meter to Bill	AMI Deployment Billing Meter Operations Safety

Revenue Protection	Revenue Protection
Distribution Planning and Operations	Distribution Planning
Demand Response and Energy Efficiency	Demand Response and Energy Efficiency

Most end users have access to environments with Explore and Export functionality based on their group assignment. The features available for the user on the Explore and Export pages are determined by their role. For example, a user with a Billing role in the Meter to Bill module has different algorithms and panels on the Explore page than a user assigned to the Distribution Planning role in the Distribution Planning and Operations module.

**Note:** The user interface features defined for groups and roles are determined by licensing and implementation. They are not configurable by the customer.

## Assigning and Removing User Group Permissions

The **Group** options in this section are examples only. Your environment may have different group types or group names.

### Assigning User Group Permissions

Use the following procedure to assign a user to a group:

1. Locate the user in Oracle Utility Analytic Insight sand open the dialog box for the user. See **Locating Provisioned Users in Oracle Utility Analytics Insights** on page 8-2.
2. Click **Assign Additional Group**.
3. Select the appropriate group from the **Add Group** drop-down list, and click **Save**. The **Manage Users** dialog box will update the **Group** field with the assigned group.
4. If a user needs permissions for multiple groups, repeat the previous steps for each additional group.
5. Click **Cancel** or any area outside of the dialog box to close the dialog box.

### Removing User Group Permissions

Use the following procedure to remove group assignments:

1. Open the **Manage User** dialog box for the user. See **Locating Provisioned Users in Oracle Utility Analytics Insights** on page 8-2.
2. Click **Remove** next to the group name you want to remove.

### Assigning User Role Permissions

Users must be assigned roles in order to access the environment. Once assigned roles, a user will be able to choose from the modules that correspond to their assigned roles.

Use the following procedure to assign user role permissions:

1. Open the **Manage User** dialog box for the user. See **Locating Provisioned Users in Oracle Utility Analytics Insights** on page 8-2.
2. Click **Assign Additional Role**. The dialog will update with a drop-down list of the available roles based on the modules licensed to the customer.
3. Select the appropriate role from the list and click **Save**. The **Manage Users** dialog box will update the **Role** field with the newly assigned role.
4. If the user needs permissions for multiple roles, repeat the steps for each additional role.
5. Click **Cancel** or any area outside of the dialog box to close the dialog box.

## Removing User Role Permissions

Use the following procedure to remove user role permissions:

1. Open the **Manage User** dialog box for the user. See **Locating Provisioned Users in Oracle Utility Analytics Insights** on page 8-2.
2. Click **Remove** next to the role.

# Chapter 9

---

## Using Federated Single Sign-On

This chapter describes tasks required when using an external identity management system to provide authentication for the application instances within your cloud subscription, including:

- [Overview](#)
- [Setup External Identity Provider](#)
- [Service Access for Federated Users](#)
- [Just In Time Provisioning for Federated Users](#)

---

# Overview

Federated Single Sign-On (SSO) allows your organization to use an external identity management system to provide online authentication for the application instances within your cloud subscription.

- The configuration and verification of the Federated Single Sign On should be available after the subscription is live.
- The Federated Single Sign-On only concerns online access; it is not applicable for the integration and other non-human accounts.
- The option to configure federation with existing Identity and Access Management is included with Identity Cloud Service subscriptions as part of Oracle Utilities cloud services.

# Setup External Identity Provider

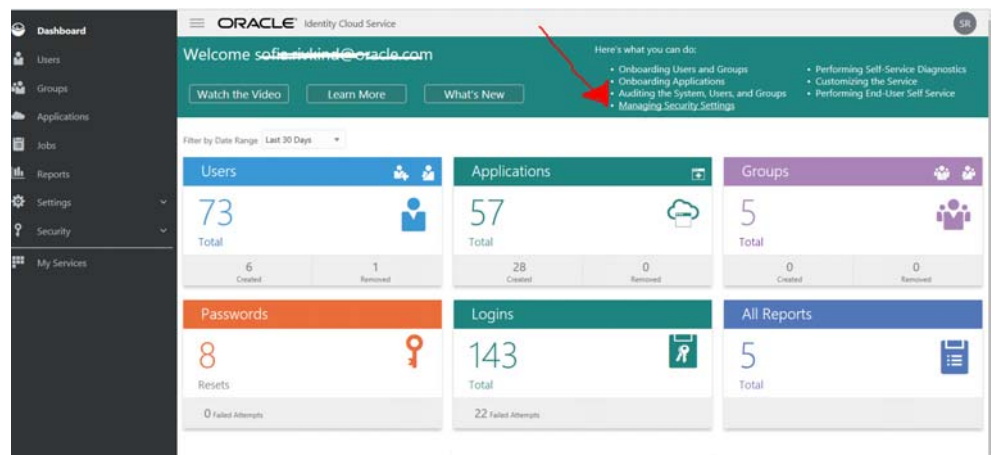
Configure a Security Assertion Markup Language (SAML) 2.0 external identity provider such as Active Directory Federation Services (AD FS) for federated SSO to Oracle Identity Cloud Service.

Configuration steps include:

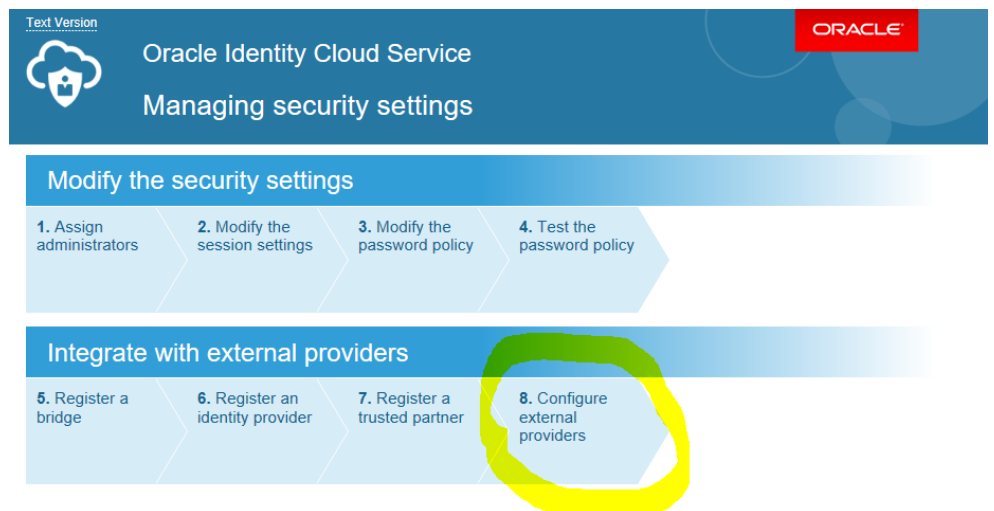
- Configure Microsoft Active Directory Bridge or implement user data synchronization via REST SCIM API or flat file import.
- Setup the Security Assertion Markup Language 2.0 Identity Provider.
- Verify Federated Single Sign-On.

To access detailed configuration instructions provided by Identity Cloud Service:

- Navigate to the Identity Cloud Service console dashboard and select **Managing Security Settings** to access online Identity Cloud Service tutorials.



- Follow the instructions under **Configure External Provider**.



Note: Federated authentication is enabled by default. This configuration means the user credentials will be validated against a configured Identity Provider. When configuring Identity Bridge define the federated authentication as follows:

- To continue validate credentials and maintain passwords and password rules in the external identity management system leave the **Federated Authentication** checkbox checked
- To validate credentials and manage passwords in Identity Cloud Service uncheck the **Federated Authentication** checkbox. Identity Cloud Service will generate the password for the users and send the notification by email (the email attribute must be filled in Microsoft Active Directory and mapped to the Oracle Identity Cloud Service).

## Service Access for Federated Users

Users created in Identity Cloud Service via federation should be granted access to the environments within the subscription the same way as the users created directly in Identity Cloud Service.

See **Update Access to Service** on page 5-4, **Setting Up an Online Application User** on page 6-6, and **Setting Up Application Users** on page 7-2 for the instructions on how to assign user to the online access application roles.

Possible approaches:

- Process users one by one: locate user in Identity Cloud Service and assign to the application roles
- Process multiple users:
  - Export users from directly or from the group (see **Exporting** on page 5-7 for more details).
  - Copy the information into Application Role import file and import users and/or groups to the Application Role (see **Importing** on page 5-7 for more details).

## Just In Time Provisioning for Federated Users

In the federated SSO scenario the Identity Cloud Service users and groups are imported from the external identity provider's data repository.

- Evaluate the groups created in Identity Cloud Service as a result of sync with external Identity Provider and determine whether to use them for Just In Time provisioning purpose.
- Login to the OUA-based application and set up Template Users that represent authorization levels corresponding to the Identity Cloud Service groups synchronized from the external provider.
- Configure the Identity Cloud Service Group - Template User mapping in the Master Configuration.

See **Configuring Just in Time Provisioning** on page 6-3 for more detailed configuration instructions.



# Part Two

---

## Object Storage Setup

This section describes the tasks that are required for connecting the system to Object Storage and the basic administration that is needed for implementation stages and beyond that:

- [Chapter 10: Introduction](#)
- [Chapter 11: Object Storage Management](#)
- [Chapter 12: Connecting to Oracle Cloud Object Storage](#)
- [Chapter 13: Recommended Object Storage Structure for a New Implementation](#)
- [Chapter 14: Initial Testing of Object Storage Connectivity](#)
- [Chapter 15: Cross-Region Disaster Recovery Considerations](#)

# Chapter 10

---

## Introduction

Oracle Cloud Object Storage is a part of Oracle Cloud Infrastructure Storage Services and it is a required service for Oracle Utilities Cloud Services, including Oracle Utilities Customer Cloud Service (CCS)

These cloud services uses Oracle Cloud Object Storage as the vehicle to exchange data files with customers during an implementation and in production.

Oracle Infrastructure Services get provisioned separately from Oracle Utilities Cloud Services but are grouped together under the same customer Cloud Account.

Access and administration of Oracle Cloud Infrastructure Services is done via the Oracle Cloud Infrastructure Console that can be accessed from the Oracle Cloud Account.

This document describes the tasks that are required for connecting the system to Object Storage and the basic administration that is needed for implementation stages and beyond that.

For more information on Oracle Cloud Object Storage (including concepts, security best practices, and more), please refer to Oracle documentation about Oracle Cloud Infrastructure Services at: <https://cloud.oracle.com/iaas>.

This section provides information about setup and configuration of object storage for use with Oracle Utilities Cloud services, including:

- [Object Storage Management](#)
- [Connecting to Oracle Cloud Object Storage](#)
- [Recommended Object Storage Structure for a New Implementation](#)
- [Initial Testing of Object Storage Connectivity](#)
- [Cross-Region Disaster Recovery Considerations](#)

# Chapter 11

---

## Object Storage Management

This chapter outlines the basic administration tasks of Oracle Cloud Infrastructure related to Object Storage, including:

- [Object Storage Structure](#)
- [Security and Access Management](#)
- [Tenant Information](#)
- [API Access](#)

---

# Object Storage Structure

This section provides an overview how object storage is structured, including:

- [Compartments](#)
- [Object Storage Buckets](#)

## Compartments

All cloud infrastructure resources are organized in Compartments.

A tenancy can include several compartments. A compartment is a logical grouping of resource types. For object storage, compartments help manage the structure of objects that are stored in the cloud.

Compartments can have child-compartments which support multi-level hierarchy of resource grouping.

Each compartment is identified by a unique Oracle Cloud ID (OCID).

When connecting the system to object storage, the compartment identification is part of the required connection configuration information.

There are no hard requirements as to the structure or number of compartment that should be created. A recommended setup is described later in this document and has reference to compartments as well.

### Root Compartment

The Root Compartment is created for each account and is the top level of the compartments hierarchy. The name of that compartment includes the string "(root)" in it.

## Object Storage Buckets

Oracle Cloud Object Storage is organized in buckets. A bucket is like a folder or a directory that stores one or more objects. Objects can be any file and can include documents, images, etc.

Each compartment can have one or more buckets. Buckets cannot include other buckets.

An example of Object Storage structure can be:

- Root Compartment
  - Compartment A
    - Child Compartment A1
      - Bucket A1-1
      - Bucket A1-2
      - Bucket A1-3
    - Bucket A1
  - Compartment B
    - Bucket B1
    - Bucket B2

Bucket names are unique within a tenancy which means that the same bucket name cannot be used in different compartments. Compartments have a unique identifier (OCID) so they are in fact unique within the tenancy.

The system can be configured to connect to any compartment and bucket that you define. This configuration is described in the next chapter.

## Security and Access Management

Oracle Utilities Cloud Services security is managed by an Oracle Identity Cloud Service (IDCS) instance that gets created when that services are provisioned. Oracle Cloud Infrastructure security is managed by Oracle Identification and Access Management (IAM).

These two identity management system are linked together and synchronized to allow easy access and security administration tasks.

This document includes only the information needed for the security administration of Oracle Cloud Infrastructure services. For information about security management of Oracle Utilities Cloud Services (that is done using IDCS), refer to the *User Provisioning Guide* document that is included with the service.

## Accessing the Cloud Infrastructure Console

Access to the console can be done by selecting **Open Service Console** from the small action menu on the lower right side of the **Compute** tile on Oracle Cloud Account. In addition, the URL for the console can be found on the **My Admin Accounts** tab when selecting the **Account Management** box in the **Oracle Cloud Account** page. The URL for the console will appear next to the **Compute (OCI) Users** account type.

**Note:** if you don't see a tile called **Compute**, click the **Customize Dashboard** tile on the dashboard and select to show the **Compute** service from the list under the **Infrastructure** category. If you cannot see that service or it is not available yet, please contact your Oracle support representative.

## Authentication and Access Management: Federated and Non-Federated Users

When accessing Oracle Cloud Infrastructure, authentication can be Federated or Non-Federated:

- **Federated** users are defined in Oracle Identity Cloud Service (IDCS), they are synchronized with IAM and are authenticated by IDCS when logging into Oracle Cloud Infrastructure.
- **Non-Federated** users are defined only in IAM and are authenticated by IAM only.

The initial security administration user is created as BOTH a Federated and Non-Federated user. That means that this administration user can login into Oracle Cloud Infrastructure from the Cloud Account Portal without the need to provide their credentials again.

## First Time Login

Since the security administrator has users definitions that are both Federated and Non-Federated, they can login into Oracle Cloud Infrastructure for the first time in several ways:

- Login from their Oracle Cloud Account (using the Open Service Console option on the Compute tile): this automatically logs the user into Oracle Cloud Infrastructure without the need to provide any credentials.
- Login directly to Oracle Cloud Infrastructure (using the direct URL): when using this option the user is presented with two authentication options:
  - Login using Single Sign On (SSO): this requires Federated user credentials. If the user is already logged into their Cloud Account, they will not need to provide their credentials.
  - Login directly into Oracle Cloud Infrastructure: this requires Non-Federated user credentials. In the case of a first login, the temporary password that was assigned to the federated user will be the same for the non-federated user.

## Managing Users

There are two types of users that should have access to infrastructure services (Object Storage being one of these): UI Access users and API Access users.

UI Access users should typically include administrator level personnel that use the Infrastructure Console to manage security and the various infrastructure services (such as Object Storage). These users are typically Federated (although they can also be Non-Federated) and therefore should be created in IDCS (refer to the *Oracle Utilities Cloud Services End User Provisioning Guide* for more information).

**Note:** UI Access users that should not have administrator access to Object Storage but are only involved in business operations (for example: uploading files to an Object Storage Bucket) should have Non-Federated users with non-administration security access setup.

API Access users are applications that use the API to access the various services but do not have access to the console user interface. These users can be Federated or Non-Federated. However, the instructions below refer to Non-Federated users only!

The recommended setup outlined later in the document includes details about both types of users.

### Adding a New User:

1. In order to add a new user, use the upper left menu in the infrastructure console, select **Identity**, then **Users**. Click **Create User** to create a new user.
2. After saving the new user information (name and description are sufficient in this case) you should be able to see the new name in the list of users.

API Access users do not need a password since they are identified via API keys. API Key management is described later in the document.

**Note:** When looking at the users defined for Oracle Cloud Infrastructure you will be able to see Federated and Non-Federated

---

users. Federated users will typically have a name in a format similar to "oracleidentitycloudservice/username...".

## Creating or Resetting User Password

Note: Initial password setup is required for Non-Federated UI Access users.

1. From the **User** list in the console, select the user name to go to the user details page.
2. Click **Create/Reset Password** to create an initial password for the user. The new temporary password can be emailed to the customer for them to login. They will be required to change the password on their first login.

## User Identification

A User is identified by an OCID key that is displayed underneath the user name. That key is used to identify users when connecting to Object Storage via API calls.

## User API Keys

API Access users that use API calls to connect to object storage should generate an encryption key pair (private/public) in PEM format and register the public key for the appropriate user (that is used in the API call).

### To register a public key for a User:

1. From the **User** list in the console, select the **User** name to go to the **User** details page.
2. Select the **API Keys** option from the **Resource List** on the left for that User.
3. Click **App Public Key**.
4. Copy and paste the public key content into the page and click **Add**.

## Managing Groups

Security management is done in Oracle Cloud Infrastructure by User Groups. Oracle Cloud Infrastructure includes an Administrator User Group that is predefined and contains the initial administrator user.

### Adding a New User Group:

1. In order to add a new user group, use the upper left menu in the **Infrastructure Console**, select **Identity**, then **Groups**. Click **Create Group** to create a new group.
2. Provide a **Name** and a **Description** for the group. Tags are optional and are not covered in this document.

### Adding Users to a User Group:

Users can be added to user groups in two ways:

1. When editing a user group record, you can add a user from the **Group Members** section by clicking **Add User to Group**.

2. When editing a user record, select the **Groups** option from the **Resource** list on the left for that user and click **Add User to Group** on the **Groups** section that is shown for that user.

## Managing Policies

Policies can be used to enforce access rights for Users that are a part of a User Group. Policies are defined in IAM using the **Identity > Policies** menu.

Using policy definitions, you can define the access rights to your infrastructure services, for example, Object Storage. You can define what compartment or bucket user groups have access to, and the type of access (read, write, etc).

Policies can apply to specific compartments or the root compartment, in which case it will apply to all of the compartments. A policy is a collection of statements with specific syntax that describe access rights to resources. For example, in a policy, you can define that a certain user group has access to create and delete buckets and objects in a certain compartment.

Refer to *Oracle Cloud Infrastructure* documentation for Identify and Access Management to find out more about policies.



---

## Tenant Information

Information about the tenancy is displayed when selecting **Administration**, then **Tenancy Details** from the upper left menu in the **Infrastructure Console**.

The information displayed is important for connecting the system to that Object Storage instance, and includes:

- The OCID key of the tenancy: This is the tenancy identification.
- Home Region: This is the main data region selected for this tenancy. Additional data regions added to this tenancy can be defined.
- Object Storage Namespace: This identification is pre-generated and is needed for the connection of the system to Object Storage.

## Regions

When a cloud account is created, a Home Region is assigned to it. This is the main data region that is linked to that account. Additional data regions can be subscribed to for the tenancy if access to regions outside the home regions are required.

The list of all available regions is displayed under the **Regions** section of the **Tenancy Details** page. Clicking **Subscription** for a region will add that to the list of available regions for this tenancy. All administration tasks will be conducted at the home region but will be synced to the other regions automatically. Please note that when connecting the system to object storage the region has to be identified as well.

---

# API Access

Oracle Cloud Object Storage can be accessed via the **Infrastructure Console** or via three types of APIs:

- Command Line Interface (CLI)
- REST calls
- Java SDK

The system connects to Object Storage using REST calls to the Object Storage endpoints that are documented for each of the data regions to which your cloud service has access.

For more information about Object Storage APIs, please refer to *Oracle Cloud Infrastructure Object Storage* documentation (go to: <https://cloud.oracle.com/storage> and select the **Documentation** tab).

# Chapter 12

---

## Connecting to Oracle Cloud Object Storage

The system supports and manages connections to Object Storage via metadata configuration. The system can connect to any number of Object Storage locations and Tenancies.

REST API calls issued by the system, to interact with the Cloud Object Storage, require API key signature. The system is designed to have a unique private/public key pair for each environment that connects to Object Storage. This means that each system environment should have a unique user defined in IAM with a registered unique API Key.

Currently the system supports accessing files on Object Storage via batch processing. Referencing a file location as Object Storage is done using a special notation.

This chapter includes the following:

- [Object Storage Connection Configuration](#)
- [API Key Management](#)
- [Referencing Files on Object Storage](#)

For additional information refer to **External File Storage** help topic in the cloud service online help.

# Object Storage Connection Configuration

Each connection configuration is represented in the system via the File Storage Configuration extendable lookup (F1-FileStorage). Each value for that extendable lookup should contain the information described below.

In order to configure a new connection, go to the Extendable Lookup portal by selecting **Admin**, then **General**, then **Extendable Lookup**, then **Search**, and search for "File Storage Configuration". After selecting it, click **Add** to add a new value.

When adding a new value, select the Oracle Cloud Object Storage file adapter and provide the following information:

- **User:** the User Identification (OCID Key) that is used for that connection.

A unique user ID should be defined for each system environment (e.g. Dev, Test, Prod) that is connecting to that object storage tenancy. It is strongly recommended that this user ID is not used for other purposes.

If one system environment is required to connect to multiple object storage tenancies, there should be a different user ID for each of these tenancies.

- **Tenancy:** the tenancy ID (OCID Key) of the object storage tenancy.
- **Compartment:** the compartment ID (OCID Key) of the compartment for that connection.

Each compartment needs a separate connection configuration.

- **Namespace:** the Namespace of the object storage tenancy.
- **Key Ring:** the Key Ring name that was created in the system. See [API Key Management](#) for more information.
- **Region:** the region of the object storage tenancy for that connection. Reminder: object storage tenancies can have multiple regions if additional subscription was done.
- **Bucket Name Prefix:** a name prefix that will be added to the bucket name of file paths referencing object storage (see **Referencing Files on Object Storage** on page 12-4 for more information).

---

# API Key Management

Secured access to Object Storage is accomplished by using API Signature Key. Each configured connection to Object Storage includes a Key Ring.

A key ring is an object that hold a set of private/public encryption key pairs. Object Storage connections can share the same key ring and even the same key in the key ring for the same system environment.

For example, key ring A can be defined and used in all the system environments: Dev, Test and Prod. However, the key pairs inside the ring have to be different in each of the environments. The connections defined for Object Storage can all use the same key ring A in all the environments since the actual key pair that is used in each environment, is different.

To create a new key ring, select **Admin**, then **Security**, then **Add Key Ring**. Make sure to generate a key pair in that ring after creating it.

## Registering the API Key

Once a key ring has been created with an active key pair, click **View** for the Public Key of that key pair to copy the public key content. That content should be pasted into the User API Key in IAM (see the **User API Keys** section in the **Security and Access Management** section of the **Managing Object Storage** chapter).

# Referencing Files on Object Storage

Reference to Object Storage can be used anywhere that a file location reference is allowed in the system.

The format is: `file-storage://<File Location>/<Bucket>/<Filename.ext>`

where:

- **<File-Location>**: The File Storage Configuration extendable lookup value defined for that file. This will include the compartment identification.
- **<Bucket>**: The object storage bucket in the compartment that is defined as part of the File Storage Configuration extendable lookup value.
- **<Filename.ext>**: The name of the file.

For example, the "payment\_info.dat" file in the "Payment-Upload" bucket in a compartment that is referenced in the "AB-Payments" File Storage Configuration extendable lookup value can be referenced as:

`"file-storage://AB-Payments/Payment-Upload/payment_info.dat"`.

## Using the Bucket Name Prefix

If you set the Bucket Name Prefix in the File Storage configuration, any file path referencing this configuration will be automatically revised at runtime, adding the name prefix to the bucket name.

This allows you to define different name prefix for buckets for each environment (or for production vs non-production environments) and keep your file paths for your batch jobs the same in each environment.

For example:

- You can create all your non-production buckets with a "NP-" name prefix, and all your production buckets without a name prefix.
- You can then define a File Storage configuration named "OS-APP" in each of your environments and set the **Bucket Name Prefix** to:
  - "NP-" in all of the non-production environments
  - Blank in the production environment
- When you will use a file path reference on your batch jobs, for example "file-storage://OS-APP/AB-Payments" then:
  - When the job related to that file runs in a non-production environment it will reference the payment files in the "NP-AB-Payments" bucket
  - When the job runs in the production environment it will reference the "AB-Payments" bucket

# Chapter 13

---

## Recommended Object Storage Structure for a New Implementation

This chapter describes a recommended configuration and structure for your Object Storage tenancy for your service implementation. Using the recommended setup can simplify the initial implementation and testing activities of your new service but they are not mandatory. Furthermore, you can start with the recommended setup and adjust it per your implementation needs.

Refer to the following topics in the Cloud Service Foundation online help:

- Object Storage
- Process Automation Tool
- Data Conversion.

### Security Considerations

The system connection to Oracle Cloud Object Storage is governed by a combination of User, User Group (optional) and Access Policies that are defined in IAM (see the Managing Object Storage chapter for more information). As a reminder, the User ID details are provided as part of the File Storage Extendable Lookup value in the system.

### Compartments

It is recommended to divide your resources amongst several compartments:

- **Production Compartment:** This compartment includes all the production resources (such as object storage buckets and objects that store production data).
- **Non-Production Compartment:** This compartment includes all the non-production resources used during the implementation and testing phases.
- **Shared Compartment:** This compartment is used to hold resources that are used by special activities or processes and can be accessed by production and non-production users. A good example of that can be configuration data (that can be exported from a testing environment and moved to the production environment when ready, using the Configuration Migration Assistant) or

---

conversion data that can be used in both production and non-production environments (during the implementation phases).

## Users

It is recommended that each system environment uses a unique user ID in IAM so that access rights to production vs non-production files or objects can be enforced for that tenancy. Each user will have its own API Key registered and should be a part of a user group, which will simplify the security access definitions.

## User Groups

It is recommended to assign the users to several groups, for example:

- **Application Access User Group for Production:** This group includes the user assigned to the production system environment and other users that will need access to object storage production information via API calls.
- **User Access User Group for Production:** This group includes all the users that will need access to object storage production information via the **Infrastructure Console**.
- **Application Access User Group for Non-Production:** This group includes the users assigned to the non-production system environments and other users that will need access to object storage non-production information via API calls.
- **User Access User Group for Non-Production:** This group includes all the users that will need access to object storage non-production information via the **Infrastructure Console**.

These groups can be referenced when defining the security policies for production and non-production access.

## Policies

It is recommended to create Policies to control access to resources based on:

- **Production vs Non-Production:** For example, it is recommended to restrict access to production resources only to production users.
- **System vs Human Users:** For example, it is recommended to restrict certain operations from system users (such as ability to delete objects or buckets).



# Recommended Setup for a Single Cloud Service

If you are using a single Oracle Utilities cloud service (such as Customer Cloud Service) consider the following recommended setup:

## Oracle Cloud Infrastructure - IAM and Object Storage

### Compartments and Buckets

- Root Compartment
  - CCS-Prod (Compartment)
  - CCS-Non-Prod (Compartment)
  - CCS-Shared (Compartment)
    - CMA-Files (Bucket)  
[for the system Configuration Migration Assistant]
    - CONV-Upload (Bucket)  
[for Data Conversion]
    - CONV-Output (Bucket)  
[for Data Conversion]

### Application Users and User Groups for Object Storage Access

- CCSDEV (for the Development environment)  
[part of User Group CCSObjectStorageAppNonProdAccess]
- CCSTEST (for the Testing environment)  
[part of User Group CCSObjectStorageAppNonProdAccess]
- CCSPROD (for the production environment)  
[part of User Group CCSObjectStorageAppProdAccess]

Additional environments will each have their own unique User with the "CCS" prefix and will be a part of the CCSObjectStorageAppNonProdAccess User Group.

### Policies for Object Storage

- Policy for application access to object storage in the Production Compartment:
  - Defined under the root compartment.
  - Open only to production user groups.
  - Allows read, create and modify access to buckets and objects in the Production Compartment and the Shared Compartment.
- Policy for application access to object storage in the Non-Production Compartment
  - Defined under the root compartment.

- Open only to non-production user groups.
- Allows read, create and modify access to buckets and objects in the Non-Production Compartment and the Shared Compartment.

## Example: Oracle Utilities Customer Cloud Service

The following example references the setup in the Customer Cloud Service (CCS) application outlined above.

### File Storage Configuration

The following File Storage Configuration extendable lookup values should be defined to correspond to the cloud infrastructure setup above:

- OS-SHARED: This value will point to the Shared Compartment:
  - The user ID will be different in each environment (CCSDEV, CCSTEST, CCSPROD)
  - The key ring can be the same in all environment but each environment key ring will have different key pairs (generated separately in each environment).
- Additional values can be defined based on the file location your specific processes will need to access, for example:
  - OS-Payment: for Payment upload interface
  - OS-MR-Up: for Meter Reads upload interface
  - OS-MR-Dl: for Meter Reads download interface
  - The Extendable Lookup values (the name) will be the same in each environment but some of the information that is defined for them will be different in each environment:
    - User ID, compartment (Prod vs Non-Prod) and keys.

---

## Recommended Setup for Multiple Cloud Services

If you are using multiple Oracle Utilities Cloud Services (for example Customer Cloud Service and Work and Asset Cloud Service) and you are still using a single Oracle Cloud Infrastructure tenancy (and therefor single Object Storage tenancy), then:

- Duplicate the Cloud Infrastructure setup (compartments, buckets, users, groups, policies, etc), one set with the CCS name prefixed and one set with the WACS name prefix.
- The setup in the Utilities Cloud Service (CCS or WACS) would be identical for both. The differences will be in the references to the various Cloud Infrastructure resources prefixed with CCS or WACS, for example:
  - OS-SHARED in CCS will point to CCS-Shared Compartment with User CCSDEV/TEST/PROD.
  - OS-SHARED in WACS will point to WACS-Shared Compartment with User WACSDEV/TEST/PROD.

# Chapter 14

---

## Initial Testing of Object Storage Connectivity

This chapter contains step by step instructions for initial testing of your connection between your cloud service and your object storage. The instructions represent a simple setup for testing the connection to object storage. These instructions do not represent the complete recommended setup that was described in previous chapter.

1. Log into Oracle Cloud Infrastructure Console using credentials provided to you by your security administrator:
  - a. In the **Identity** menu, select **Users**:
    - i. Create a new user named "INIT-TEST" (Take note of the user OCID). (This will be a Non-Federated user.)
    - ii. Add that user to the Administrator user group.
  - b. In the **Identity** menu, select **Compartments**:
    - i. Create a new compartment named "INIT-TEST" (take note of the compartment OCID).
  - c. In the **Object Storage** menu, select **Object Storage**:
    - i. Select the INIT-TEST compartment in the **Compartment** field under the **List Scope** section.
    - ii. Create the following buckets under the INIT-TEST compartment:
      1. CMA-Files
  - d. In the **Administration** menu, select **Tenancy Details**:
    - i. Take note of the tenancy OCID (under **Tenancy Information**)
    - ii. Take note of the namespace (**Name** field under **Tenancy Information**)
    - iii. Take note of the home region
2. Log into the Utility Cloud Service development environment (DEV), using credentials provided to you by your security administrator:
  - a. Go to the **Key Ring** portal (use the Menu Search option):
    - i. Add a new Key Ring named "INIT-TEST"
    - ii. After creating the new Key Ring, click **Generate Key**.
    - iii. In the **Key Pair** section, choose the **Activate** action for the new generated Key Pair.

- iv. Click **View** to get the public key portion of the key pair.
        - v. Copy the full content of the public key displayed in a popup window, save it in a text document. You will use this later.
      - b. Go to the File Storage Configuration extendable lookup and search for a value of OS-SHARED.
      - c. Edit that value and enter the following information:
        - i. **User**: the user OCID of INIT-TEST User from step #1.
        - ii. **Tenancy**: the tenancy OCID from step #1.
        - iii. **Compartment**: the compartment OCID of INIT-TEST Compartment from step #1.
        - iv. **Namespace**: the namespace noted in step #1.
        - v. **Key Ring**: search for the INIT-TEST key ring created above and select it.
        - vi. **Region**: the home region noted in step #1.
        - vii. Click **Save**.
      - d. Go to the **Master Configuration** portal (use the Menu Search option):
        - i. Look for the **Migration Assistant Configuration** master configuration.
        - ii. Make sure that the **Import** and **Export** directories have the following value:

```
"file-storage://OS-SHARED/CMA-Files"
```
3. Log back into Oracle Cloud Infrastructure Console using credentials provided to you by your security administrator:
  - a. From the **Identity** menu, select Users:
    - i. Select the INIT-TEST user created earlier.
    - ii. In the **API Keys** section, click **Add Public Key**.
    - iii. In the popup window paste the public key value saved in previous step (the public key portion of the key pair generated in the Utilities Cloud Service application), and click **Add**.
4. You are ready to test the object storage connectivity. Log back into the Utility Cloud Service development environment (DEV), using credentials provided to you by your security administrator:
  - a. Go to the **Migration Request** portal (use the Menu Search option).
  - b. Search for a Migration Request named Users (F1-Users).
  - c. Click **Export** for that request (Users).
    - i. In the popup window enter the file name "init\_test" (for example)
    - ii. Click **Save**. You will be directed to the **Migration Data Set Export** page.
  - d. Go to the **Batch Job Submission** portal and submit a job with the F1-MGDPR batch code. When the job ends, go back to the **Migration Data Set Export** portal and check the status:

- i. If the status changed to Exported, log into the Oracle Cloud Infrastructure Console, navigate to the CMA-Files object storage Bucket under the INIT-TEST Compartment and check that there is a file called `init_test.cma` there.
  - ii. If the file exist, the test is successful!
5. If the connectivity test was successful, proceed with the overall setup of the Object Storage and your Cloud Service application per the recommended setup above.

# Chapter 15

---

## Cross-Region Disaster Recovery Considerations

This chapter outlines the considerations for Object Storage connection and configuration in case the cross-regional disaster recovery option has been enabled for your system.

This chapter includes the following:

- [Home and Disaster Recovery \(DR\) regions](#)
- [Preparing your Disaster Recovery Region](#)
  - [Copying Your Object Storage Bucket Structure](#)
  - [Copying Your Object Storage Data](#)
- [Recovering from a Disaster](#)
  - [Switching To Your Disaster Recovery Region](#)
  - [Switching Back To Your Home Region](#)
  - [Copying Back Your Object Storage Data](#)

---

## Home and Disaster Recovery (DR) regions

Your system has a Home Region, which is the data region that it was initially provisioned at. This will be referred to as the System Home Region. When cross regional disaster recovery is enabled for your system it will have a designated disaster recovery (DR) region. The disaster recovery region is the data region that your system will be switched to in case your home region is no longer available. This will be referred to as the System Disaster Recovery Region.

Your Oracle Cloud Infrastructure (where your Object Storage resides) has also a home region, that will be referred to as the Object Storage Home Region. If your system has a designated disaster recovery region, it will make sense for your object storage to have a designated disaster recovery region as well, which will be referred to as the Object Storage Disaster Recovery Region.

In most cases the System Home Region will be the same as the Object Storage Home Region but it could be different if it was chosen to be different. The same is true for the System Disaster Recovery Region and the Object Storage Disaster Recovery Region. Selecting an Object Storage Disaster Recovery Region will be covered in the next section.

**Note:** If the Object Storage Home Region is different than the System Home Region, you can skip this chapter since the cross region disaster recovery procedures will not affect your object storage and will not affect your system connection to object storage.



# Preparing your Disaster Recovery Region

If cross-region disaster recovery was enabled for your system, it will be automatically set up to be ready for a disaster event in terms of availability of resources on your System Disaster Recovery Region, according to your service level agreements.

It is your responsibility to make sure that your object storage is ready as well.

Since Object Storage is a regional service, there is no automatic disaster recovery for that. Assuming your Object Storage Home Region is identical to your System Home Region, you need to plan for the eventuality that this region might become unavailable and so you will need to have your object storage available on another region.

The first thing you will need to do is to subscribe to an additional data region to be your Object Storage Disaster Recovery Region.

In order to subscribe to an additional region you should do the following:

1. In the Oracle Cloud Infrastructure Console, select **Administration**, then **Manage Regions** and look at the list of additional available data regions. Select the data region to designate as the Object Storage Disaster Recovery Region (is it recommended to have it identical to your System Disaster Recovery Region, if possible).
2. Your request for subscription to a new data region will be processed and when it is completed, you will see your new region in the list of available regions.
3. You will also be able to switch to this data region in your Oracle Cloud Infrastructure Console via the **Region** drop-down list.

## Copying Your Object Storage Bucket Structure

Your Oracle Identification and Access Management (IAM) definitions (i.e. users, groups, policies and compartments) are all maintained in your Object Storage Home Region and these definitions are replicated automatically to all the other regions to which you are subscribed.

Object Storage Buckets are region dependent which means that each data region can have its own set of buckets.

In order for your system to continue to work properly once it is switched to your System Disaster Recovery Region (for functions that require access to object storage), your object storage bucket structure should exist in your Object Storage Disaster Recovery Region.

Therefore we recommend that you synchronize your bucket structure periodically between your Object Storage Home Region and Object Storage Disaster Recovery Region. This means, at a minimum, that buckets created in your Object Storage Home Region should be also added to your Object Storage Disaster Recovery Region.

## Copying Your Object Storage Data

You may also choose to periodically copy the objects inside your buckets from your Object Storage Home Region to your Object Storage Disaster Recovery region.

Please note that copying data from one region to another will result in the use of additional object storage space, which in turn can lead to additional cost per billing period.

Refer to [Using Replication](#) in the **Object Storage** section of the Oracle Cloud Infrastructure Documentation for more information about configuring data replication policies to copy data between buckets in different regions.

If you have the ability to re-create lost data when a disaster occurs, then you might not need to copy your data across regions in advance, for example:

- Most files generated by your system via batch jobs can be regenerated if necessary
- 3rd party applications that load files into object storage may also be able to reproduce these files upon request

---

# Recovering from a Disaster

A disaster is defined as an event that will cause your System Home Region to become unavailable.

When a disaster occurs, your system will automatically be switched to your System Disaster Recovery Region, based on your service level agreements. When that happens you are responsible to tell the system what object storage region to connect to instead of the current one that is was linked to when the disaster happened (if that region has also become unavailable).

This section covers what you should do during a disaster and after it is resolved.

## Switching To Your Disaster Recovery Region

Once your system has been switched to its System Disaster Recovery region, you will need to point it to a different data region for object storage access:

1. Log into each of the system environments.
2. In each environment look at all of your current File Storage Configurations.
3. Edit each File Storage Configuration and change the region field to your Object Storage Disaster Recovery Region.
4. Save your changes.

## Switching Back To Your Home Region

When your home region has been recovered and data was restored, the system will be switched back to your System Home Region. At this point you will need to point it back to your Object Storage Home Region for object storage access:

1. Log into each of the system environments.
2. In each environment look at all of your current File Storage Configurations.
3. Edit each File Storage Configuration and change the region field to your Object Storage Home Region.
4. Save your changes.

## Copying Back Your Object Storage Data

When you are switched back to your Object Storage Home Region, you may need to copy back some of the data that was created in your Object Storage Disaster Recovery Region. This may also include changes in bucket structure that you may have done while working in your disaster recovery regions.

- Changes in object storage bucket structure can be repeated in your home region manually after that region has been recovered.
- If you need to copy data back to your home region, refer to [Using Replication](#) in the **Object Storage** section of the Oracle Cloud Infrastructure Documentation for guidance.

# Part Three

---

## Cloud Operations

This section provides details regarding service requests (SR) provided by the Oracle Utilities Cloud Operations team for customers during implementation and operation of Oracle Utilities Cloud Services:

- [Chapter 16: Cloud Operations Service Requests](#)

# Chapter 16

---

## Cloud Operations Service Requests

This document provides details regarding service requests (SR) provided by the Oracle Utilities Cloud Operations team for customers during implementation and operation of the following Oracle Utilities Cloud Services:

- Oracle Utilities Billing Cloud Service
- Oracle Utilities Customer Cloud Service
- Oracle Utilities Meter Solution Cloud Service
- Oracle Utilities Rate Cloud Service
- Oracle Utilities Work and Asset Cloud Service
- Oracle Utilities Operational Device Cloud Service

The service requests described in this chapter include:

- [Backup and Restore Policy Overview](#)
- [Requesting a Full Ad-Hoc Backup](#)
- [Restoring from a Full Ad-Hoc Backup](#)
- [Requesting a Database Restore Point](#)
- [Rolling Back Environment to Restore Point](#)
- [Request for DNS Address List](#)
- [Request to Execute Data Correction Scripts](#)
- [Request for Temporary Additional Threads](#)
- [Request to Export Cloud Environment Schema for On-Premises Use](#)
- [Refreshing Data \(Cloning\)](#)

Service request descriptions in this document include the following information:

- A brief description of the service to be performed
- **Service Level Objective (SLO):**
  - **Advanced Notice:** The amount of advanced notice (in days) the Cloud Operations team needs to perform the service request, starting from the first full business day after the service request was created.

A future dated request for performing any of the above mentioned activities on end of business day, Friday will be executed in the next week on the first

working day. For example, if the service request states to take a back up on January 7th 2021 after 6PM PST, it will be executed in the week of January 11th 2021 on the first working day, during agreed maintenance hours.

Service requests will not be executed on weekend days since weekends are reserved for scheduled planned maintenance activities like upgrades.

- **Acknowledge/Schedule:** The amount of time the Cloud Operations team would take to acknowledge and schedule the request.
- **Execution Time:** The amount of time (in days) it will take to perform the service request

**Note:** Times provided in this document are based on the data center where the cloud service is deployed. For example, a reference to "8:00 AM" means 8:00 AM in the time zone of the data center, not necessarily the time zone of the customer.

## Backup and Restore Policy Overview

### Introduction

Oracle Cloud Infrastructure (OCI) team is responsible for configuring and managing the backup and restore for all Oracle Databases. The OCI team has adopted by default the regime known as the Oracle Suggested Backup Strategy.

### Scheduling

The backups for the Oracle-Recommended Backup Strategy are scheduled to run by default as follows:

- Weekly **full database** backup on Sunday at 2AM
- Daily **incremental backups** on every other day at 2AM.
- Daily **archive log backups** are taken three times per day at 6AM, 2PM, and 10PM

**Note:** These time are based on the Local Data Center Time.

### Encryption

All backups are encrypted at two levels:

- OCI Storage
- RMAN (Oracle Recovery Manager)

### Monitoring and Alerting

Monitoring of the backup activity is done through Oracle Enterprise Manager. Alerts are sent to the Oracle Cloud Infrastructure database administration team if backup jobs fail so that remedial action can be taken.

### Backup Retention

The retention period for backups defaults to **60 days**.

# Requesting a Full Ad-Hoc Backup

As specified in [Backup and Restore Policy Overview](#) section, weekly full backups are taken on every Sunday at 2 AM data center time. The customer or system integrator can request an ad-hoc full backup if they want to keep a copy of the database at a particular state so they can later restore the database back to that copy.

## Request Specifications

- The request backs up the database, not object storage files or application code.

## Customer Obligations

- The customer needs to utilize the backup within 60 days from the date the backup has been taken or the backup will be lost.
- When creating the service request, the customer needs to specify the source environment, the target environment, and the date and time (specifying the time zone) when back up is needed.
- The customer should note the version of code in the backup, as it will only be able to be restored to a target environment on the same code version.

## Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 3 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1 business day

# Restoring from a Full Ad-Hoc Backup

The customer or system integrator will request the rollback to a previously requested full backup by opening a Service Request with Oracle Support.

## Request Specifications

- The backup being requested to be restored, must be of same environment.
- The backup must be from the same version of code as the target environment to be used for a restore.
- The restore request can also be from a scheduled weekly full or daily incremental backups.
- The restore impacts the database, not object storage files or application code.
- The production environment may not be restored to a backup unless the environment is not yet used for production purposes.

### Customer Obligations

- The customer needs to utilize the backup within 60 days or else the backup will be lost.
- When creating the service request, the customer needs to specify details related to the source environment, the source environment backup date, the target environment, and the date and time (including time zone) the restoration has to be performed.
- The customer should note the version of code in the backup, as it will only be able to be restored to a target environment on the same code line.

### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 3 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1 business day

## Requesting a Database Restore Point

A Restore Point is a snapshot of data that is saved at a specific point in time. The customer or system integrator will request a Restore Point by opening a Service Request with Oracle Support.

### Request Specifications

- The restore point impacts the database, not object storage files or application code.
- The restore point may not be restored to a backup unless the environment is not yet used for production purposes.

### Customer Obligations

- The customer needs to utilize the restore within **3 calendar days** or else the restore will be lost.
- When creating the service request, the customer needs to specify details related to the source environment, the source environment backup date, the target environment, and the date and time (including time zone) the restoration has to be performed.
- The customer should note the version of code in the backup, as it will only be able to be restored to a target environment on the same code line.

### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request



- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 1 business day
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1 business day

## Rolling Back Environment to Restore Point

The customer or system integrator will request the rollback to a previously saved Restore Point by opening a Service Request with Oracle Support.

### Request Specifications

- The restore point must be of the same environment to be used for a restore.
- The restore point impacts the database, not object storage files or application code.
- The restore point must be the same version of code as the target environment.
- The restore point may not be restored to a backup unless environment is not yet used for production purposes.

### Customer Obligations

- The customer needs to utilize the restore within 3 days since when the restore point was taken or else the restore will be lost.
- When creating the service request, the customer needs to specify details related to the source environment, the source environment backup date, the target environment, and the date and time (including time zone) the restoration has to be performed.
- The customer should note the version of code in the backup, as it will only be able to be restored to a target environment on the same code line.

### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** Within 2 calendar days from the time the restore point was created
- **Acknowledge/Schedule:** 1 business day

- **Execution Time:** 1 business day

## Request for DNS Address List

The customer or system integrator can request a DNS (Domain name service) name to be added in the allow list for outbound interface communication. An allow list provides access to specified DNS addresses that the Oracle network would otherwise prevent access to. For Oracle Utilities cloud services, a customer or system integrator must request a DNS to be added to allow list for outbound communication to all external systems. This section describes the process to for customers or system integrators to request a DNS allow list.

### Request Specification:

- A single request to allow a named DNS or a list of named DNS will ensure allowing the same DNS in all customer environments. Customer need not to request per environment (Dev, Test, Prod).

### Customer Obligations

- The customer must provide the named DNS OR URL along with the justification for its allowance.
- The SSL port 443 must be used.
- The hosts must have a valid TLS certificate signed by a trusted public certificate authority (Digicert, Symantec/VeriSign, Comodo, etc.)

### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 7 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 2 business days

## Request to Execute Data Correction Scripts

The customer or system integrator can request execution of data correction scripts.

Note: In most cases, the customer or system integrator should be able to do this by creating a Plug-In Driven batch control. If a situation should occur where this cannot be done please follow this procedure.

### Request Specification:

- The customer must ensure there are no commits in the scripts and provide the count of the records that are going to be updated by the given script. The Cloud Operations team will commit upon successful execution of the script.
- The customer must specify the schema name where the script to be executed. ex: CISADM, STGADM and SCHADM
- The execution of data correction scripts may or may not require a maintenance window depending on the nature of the scripts being executed

#### Customer Obligations

- The customer must provide a justification for executing data correction scripts.
- The customer must provide the environment name where the script has to be executed.
- The customer must provide the exact SQL they request to be run.
- The customer must first request SQL be run in a non-production environment before executing in production.

#### Oracle Utilities Product Development Team Obligations

- Oracle Utilities Product Development team must review and approve the script, before the Cloud Operations team can execute it against the environment
- Oracle Utilities Product Development team will also suggest if the script would require outage on the environment

#### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Coordinate with Infrastructure teams (if needed) for completion of the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 3 business days from the time the script is approved by Oracle Utilities Product Development
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1 business day

## Request for Temporary Additional Threads

The customer or system integrator can request additional threads to be added to an environment temporarily.

#### Request Specification:

- The customer must specify a percentage increase/ or absolute number of threads in Batch thread pool workers to be increased with justification.
- The customer must specify duration for which additional threads are required (Default duration is 1 Month maximum).

- Cloud Operations team will assess capacity in order to accept request
- The customer must specify the environment where threads to be added.

**Customer Obligations**

- Refer to the Request Specification above.

**Oracle Cloud Operations Team Obligations**

- Acknowledge and review the additional threads requested.
- Assess the capacity and seek management approval
- Schedule the execution of service request and communicate the schedule to customer
- Communicate the status upon completion of the request

**Service Level Objective**

- **Advanced Notice:** 7 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1 business day

## Request to Export Cloud Environment Schema for On-Premises Use

The customer or system integrator can request export of a cloud environment schema (CISADM) for on-premises use. This schema export is limited to be used for reporting purposes and is intended for initial extract only. The subsequent change is expected to be loaded via Generalized Data Extracts batch jobs

**Request Specification:**

- The customer must request the export of the cloud environment application schema providing details of the source environment.
- The customer should also provide details of the Object Storage Location and access details (including credentials to access object storage bucket) where the export files are to be copied

**Customer Obligations**

- The customer or system integrator must specify the environment from which the application schema export is to be taken, including Date and Time of the export.
- This request requires 7 business days lead time from the time of creation of service request

**Oracle Cloud Operations Team Obligations**

- Acknowledge and validate access to Object Storage
- Schedule the execution of the service request
- Execute the activities to complete the service request
- Communicate the status upon completion of the service request

## Service Level Objective

- **Advanced Notice:** 7 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** 1-5 business days, depending on the size of the database

Database Size	Duration of Export
Up to 0.5 TB	1-2 days
Between 0.5 TB and 1 TB	2-3 days
Between 1 TB and 5 TB	3-4 days
Larger than 5 TB	5+ days

NOTE: The customer will receive daily status updates on the service request for the duration of the export activity.

## Refreshing Data (Cloning)

### Cloning Process Overview

#### What is environment cloning?

Environment cloning is a process of copying data from source to target environment. The most frequent use cases include:

- During implementation, to clone a 'gold' configuration to another environment.
- During implementation, to clone an environment with converted data to another environment for System Integration testing or User Acceptance Testing.
- Once live, to clone Production data to a Test environment.

#### Pre-requisites for requesting environment clones

- In order to clone the environment, the source and target environments must be on the same version and patch set level.
- Source environments must maintain data integrity, and also should not have unusable indexes and if found, the Cloud Operations team would notify the customer to fix it and reschedule the cloning activity.
- During the implementation phase, a customer can request:
  - A Dev-sized environment to be cloned to another Dev, Test or Prod environment.
  - A Test-sized environment can be cloned to another Test or Prod environment.
  - A Prod environment can only be cloned to a Test environment.

Once the customer is live, a non-prod environment will not be cloned to Prod environment.

#### Frequency of Clone requests

There is a limit to number of data refresh (cloning) requests that a customer can make during the implementation phase and after go-live. Refer to the **Usage Limits** section in the [Utilities Global Business Unit Service Descriptions and Metrics](#) document.

### How to submit request for Cloning

Please refer to the [Submitting a Request for Cloning](#) section below for instructions on how to submit an environment clone request.

### Process of Cloning

Cloning in Oracle Utilities cloud environments is done using Oracle Data Pump. Oracle Data Pump technology enables very high-speed movement of data and metadata from one database to another.

The following data would be truncated and replaced with data from source environment

- Metadata
- Admin Data
- Master Data
- Transactional Data
- Conversion Staging and Scheduler data is optional and is copied upon request

In addition, the table structures/partitions would match the source environment.

The following data would be preserved in the target environment

- Users and User Groups  
Users and User Groups from the source will be appended to the target application tables. This means that Users who could previously access the target environment can still access it after the clone. However, note that Identity Cloud Service (IDCS) Users are not migrated from source to target, so any users that exist as IDCS Users with rights to access the source environment will not automatically have IDCS rights to access the target environment. Customer may need to adjust this manually as needed (please refer to the *Oracle Utilities Cloud Services End User Provisioning Guide* for details).
- External system configuration
- Process automation configuration
- Configuration related to how the target environment is identified by name is preserved during the clone. For example, if the target environment is called "DEV01"
- Configuration related to Oracle Object Storage that are environment specific are preserved during the clone and not overwritten by source data.
- Outbound integration references specified on external system application tables are preserved during the clone and not overwritten by the source data. (eg: message senders, external system references).

### Business Impact

There is no impact on the source environment during the data pump export process and no outage is caused. For all tables in the export to be consistent to the same point in time, export utility parameter FLASHBACK\_TIME is used.

The target environment would require outage for the duration of data pump import process depending on the size of the database being cloned.

Database Size	Duration of Import
Up to 0.5 TB	1 - 2 Days
Between 0.5 TB and 1 TB	2 - 3 Days
Between 1 TB and 5 TB	4 - 6 Days
Larger than 5 TB	6+ Days

Note: The customer will receive daily status update on the service request for the duration of the cloning activity.

### Post Cloning Steps

There are some post-cloning steps that the customer needs to perform from the application user interface on the target environment before it can be released to end-users. Oracle Support would communicate these steps via the service request upon completion of request.

## Submitting a Request for Cloning

The customer or system integrator can request to refresh data from one environment to another environment(s) belonging to the customer.

### Request Specifications

- The customer can request a data refresh from their Production environment to Non-Production environments, or from one non-production environments to other non-production environments.
- The source environment can't be larger than target environment.
- The customer needs to specify the date/ time of the cut of the source environment.

### Customer Obligations

- The customer needs to understand and follow post cloning procedure to complete the clone.
- The customer needs to request if cloning needs to be taken as soon as possible or at a specific date and time (specifying time zone).

### Oracle Cloud Operations Team Obligations

- Acknowledge and schedule the execution of the service request
- Execute the activities to complete the service request
- Provide daily status update if the outage during clone is more than a business day.
- Verify if the table row counts matches between source and target.
- Communicate the status upon completion and provide post cloning steps

## Service Level Objective

- **Advanced Notice:** 7 business days
- **Acknowledge/Schedule:** 2 business days
- **Execution Time:** Based on the size of the source environment. Please see the table in [Business Impact](#) section above.