Oracle Utilities Cloud Services

Object Storage Setup Guide

For 20C Releases

F35475-01

December 2020

ORACLE

Oracle Utilities Customer Cloud Services 20C Object Storage Setup Guide

Contents

Chapter 1

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 guide 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

Chapter 2

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 includes 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 Fon-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 3

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:

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 3-4 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 4

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 simply 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 accesses 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 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 5

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. 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 6

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

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 designed 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 that 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 became 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.