

Oracle® Cloud

Configuring an Oracle-Managed Disaster Recovery Solution for Oracle Integration 3



G10051-10
March 2025



Oracle Cloud Configuring an Oracle-Managed Disaster Recovery Solution for Oracle Integration 3,
G10051-10

Copyright © 2024, 2025, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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®, Java, MySQL, and NetSuite 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

Preface

Audience	iv
Documentation Accessibility	iv
Diversity and Inclusion	iv
Related Resources	iv
Conventions	v

1 Introduction to an Oracle-Managed Disaster Recovery Solution

Introduction to Oracle-Managed Disaster Recovery	1-1
What's Supported?	1-4
What's Not Supported?	1-5

2 Set Up and Perform Disaster Recovery

Perform Preinstallation Tasks	2-1
Install and Configure Oracle Integration for Disaster Recovery	2-2
Perform Failover Prerequisite Tasks	2-6
Understand Your Connectivity Agent Responsibilities During a Failover	2-6
Configure Email Delivery If Using Your Own Email Tenancy	2-6
Configure File Server for Disaster Recovery	2-7
Fail Over to the Other Instance	2-7
Configure Email Notification Settings After Failover	2-10

Preface

This document describes how to configure a disaster recovery solution for Oracle Integration.

Topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Diversity and Inclusion](#)
- [Related Resources](#)
- [Conventions](#)

Audience

This document is intended for personnel who are responsible for configuring a disaster recovery solution for Oracle Integration.

Documentation Accessibility

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

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://support.oracle.com/portal/> or visit [Oracle Accessibility Learning and Support](#) if you are hearing impaired.

Diversity and Inclusion

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

Related Resources

For more information, see the Oracle Integration documentation in the [Oracle Cloud Library on the Oracle Help Center](#).

Conventions

The following text conventions are used in this document.

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Introduction to an Oracle-Managed Disaster Recovery Solution

Oracle Integration is available in an Oracle Cloud Infrastructure region governed by service-level agreements (SLAs). This guide details the procedure to install and use a cross-region, Oracle-managed disaster recovery solution for Oracle Integration, specifically for the Integrations and File Server features in Oracle Integration.

Topics:

- [Introduction to Oracle-Managed Disaster Recovery](#)
- [What's Supported?](#)
- [What's Not Supported?](#)

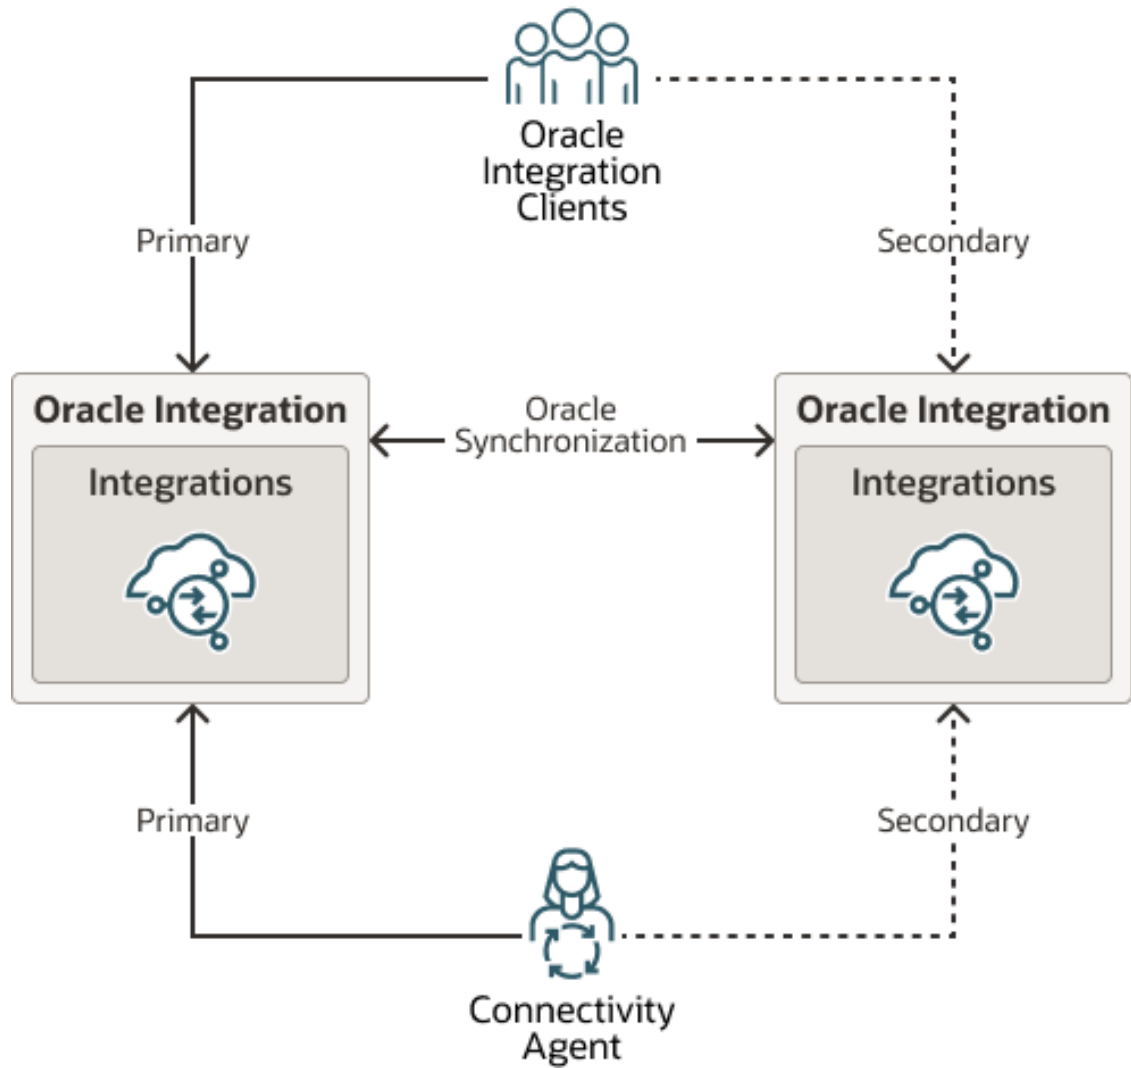
Introduction to Oracle-Managed Disaster Recovery

Failover is the process in which a secondary (standby) instance takes over when the primary working instance fails. Oracle provides a disaster recovery solution that allows you to fail over quickly from natural or human disasters and provide business continuity in your secondary instance. You can also use this solution for planned migrations and switch between instances periodically. Oracle manages nearly all disaster recovery responsibilities automatically for you. Your administrative responsibilities are minimal.

Note:

Oracle-managed disaster recovery is a paid feature. Consult with your sales representative for details.

You don't need to worry about managing DNS changes, load balancing, design-time data synchronization between instances, object storage buckets, and other responsibilities. All message traffic is automatically forwarded to the correct instance. All messaging is bidirectional, meaning you can fail over from one instance to another, and back. Data synchronization between the two instances occurs automatically in near real time to minimize data loss.



At a high level, the Oracle-managed disaster recovery solution works as follows:

1. You work in your primary instance (for example, in the Ashburn instance), which then fails and becomes unreachable.
2. Your administrator logs in to the Oracle Cloud Infrastructure Console for your secondary instance (for example, in the Phoenix instance) and selects to fail over from the Ashburn instance to the Phoenix instance. No other administrator-initiated tasks are required for failover to complete.
3. Once failover completes, you are prompted to log in to the new primary instance in a different region with the global (regionless) URL and resume work. Data synchronization has occurred in near real time between the two instances since you installed your disaster recovery solution. For this reason, any data loss is minimized.
4. You work in the instance in Phoenix, which has become the primary instance, until the original primary instance in Ashburn is restored.
5. Your administrator logs in to the Oracle Cloud Infrastructure Console in either the Ashburn instance or the Phoenix instance and selects to fail over from the Phoenix instance back to the Ashburn instance. No other administrator-initiated tasks are required for failover to complete.

- Once failover completes, you are prompted to log in to the instance in the Ashburn instance (which once again becomes the primary instance) and resume work. Because of data synchronization in near real time between the two instances, the data changes you made in the Phoenix instance are visible in the restored Ashburn primary instance.

User Responsibilities

Because Oracle handles nearly all disaster recovery management tasks, your responsibilities as an administrator are kept to a minimum. You only have several major responsibilities in an Oracle-managed disaster recovery environment:

Task	Description	See..
Subscribe to the secondary region	Subscribe to the secondary region to ensure that secondary instance creation is successful and configure the necessary policies for any default or defined tags you are using.	Perform Preinstallation Tasks
Install primary and secondary instances	Select the Enable disaster recovery toggle when installing an instance in the Oracle Cloud Infrastructure Console. This action creates primary and secondary instances in separate, predetermined regions. Data synchronization between the two instances is automatically configured and occurs in near real time.	Install and Configure Oracle Integration for Disaster Recovery
Perform prerequisites prior to failover	Review the following prerequisites to determine if they apply to your setup: <ul style="list-style-type: none"> If your connectivity agent is installed in an Oracle Cloud Infrastructure Compute instance that fails, you must have a plan in place that allows for a quick recovery of the connectivity agent. If you are using your own email tenancy, you must manually maintain your email notification details in your primary and secondary instances. If you are using File Server, connections to it must use the port and hostname, rather than the port and IP address. 	Perform Failover Prerequisite Tasks
Fail over and fail back between instances in different regions	Two types of failover are supported: <ul style="list-style-type: none"> If your primary instance is unreachable, you click Start Failover in the Oracle Cloud Infrastructure Console of the secondary instance to fail over to that instance. Once the original primary instance is restored, you can then fail back to that instance. If you want to perform a planned migration between instances periodically, you click Start Failover in the Oracle Cloud Infrastructure Console of either the primary or secondary instance to fail over to the secondary instance. 	Fail Over to the Other Instance
Configure email notification settings after failover	After failover occurs, you must configure email notification settings on the Notifications page of the new primary instance.	Configure Email Notification Settings After Failover

Understand Failover Behavior

- Bidirectional data synchronization (replication) is regularly performed in near real time between the two instances to reduce the chance of data loss after failover.

- Failover is a one instance-to-one instance replication, meaning you can only fail over to a second instance. You cannot fail over to multiple instances.
- When a failover is performed, the secondary instance takes over the responsibility of providing all features of the primary instance.
- The primary instance goes into standby mode and becomes a passive listener when the secondary instance becomes active.
- All traffic that was originally sent to the initial primary instance is forwarded to the new primary instance.
- The life cycle operations in the standby instance are disabled in the Oracle Cloud Infrastructure Console with the exception of performing a failover.
- There are no changes in OAuth credentials after failover.
- Only design-time metadata is synchronized. Runtime tracking data such as that shown in the activity stream, Instance page, and other observability pages is not synchronized with the secondary instance.
- You log in to primary and secondary instances with a global URL that does not include a region name.
- If you delete the primary instance, the secondary instance is also deleted.
- If you start and stop the primary instance, this has no impact on the secondary instance, which simply remains a passive listener.

What's Supported?

The Oracle-managed disaster recovery solution provides support for the following:

- The Integrations and File Server components.
- The Enterprise edition and Healthcare edition. See [Install and Configure Oracle Integration for Disaster Recovery](#).
- Active-passive topologies. Only one instance of the disaster recovery configuration can be active and processing transactions at a time. Active-passive topologies mean that only one instance processes all the load even though both instances are expected to be up and running. An active instance is determined by which instance the traffic is directed to and is enabled to execute all functions, and not by its start or stop status. It is recommended that you not turn off your disaster recovery instance.
- Oracle-managed synchronization (replication) of design-time metadata only.
- Support for disaster recovery instances in the following regions. Each region is paired for failover with another region.

This Region...	Is Paired with this Region...
Ashburn (IAD)	Phoenix (PHX)
Germany Central (FRA)	Netherlands Northwest (AMS)
London UK (LHR)	Newport UK (CWL)
Singapore (SIN)	Japan East (Tokyo) (NRT)
Toronto (YYZ)	Montreal (YUL)

- Support for *new* instances only. You cannot configure a disaster recovery solution for an existing instance.

- A recovery point objective (RPO) period of one hour. The RPO is the period after a disaster occurs during which the service can tolerate lost data before the disaster begins to affect the business.
- A recovery time objective (RTO) period of one hour. The RTO is the target time within which your service must be restored to the secondary system after a failover request.
- Support for user-initiated failover (not automatic). When your current primary instance becomes unreachable, you manually select the **Start failover** option in the Oracle Cloud Infrastructure Console of the secondary instance to fail over to that instance. See [Fail Over to the Other Instance](#).
- Extended data retention.
If you have data retention set to more than 30 days on the primary instance and then fail over to the secondary instance, those settings are retained on the secondary instance.

**Note:**

Integrations with polling triggers require special handling. See [Deactivate and Reactivate Integrations with Polling Triggers](#).

What's Not Supported?

The Oracle-managed disaster recovery solution does not support the following:

- The following components:
 - Oracle Cloud Infrastructure Process Automation
 - B2B for Oracle Integration
 - Visual Builder
 - Robotic process automation (RPA)
- Changing the compartment of either the primary instance or the secondary instance.
- Standard edition and development shape when performing an installation. See [Install and Configure Oracle Integration for Disaster Recovery](#).
- Event-based integrations.
After a failover, an event-based integration does not work until the event definition is manually recreated and the integration is reactivated.
- Custom endpoints and private endpoints.
- Network access (ACLs).
- Synchronization (replication) of runtime data between instances (for example, audits and tracking data).
- In-flight transactions and asynchronous integrations.
These types are aborted upon the failover request. They must be retried after failover to the secondary instance is successful. Scheduled integrations continue to work on the secondary instance.

2

Set Up and Perform Disaster Recovery

This section describes how to set up and perform disaster recovery.

Topics:

- [Perform Preinstallation Tasks](#)
- [Install and Configure Oracle Integration for Disaster Recovery](#)
- [Perform Failover Prerequisite Tasks](#)
- [Fail Over to the Other Instance](#)
- [Configure Email Notification Settings After Failover](#)

Perform Preinstallation Tasks

Subscribe to the secondary region to ensure that secondary instance creation is successful. If you are using a non-default identity domain, you must configure data replication for the secondary region. If you are using default or defined tags, ensure that you configure the necessary policies.

- [Subscribe to the Secondary Region](#)
- [Perform Data Replication Across Non-Default Domains](#)
- [Configure Policies for Default or Defined Tags](#)

Subscribe to the Secondary Region

If you have not subscribed to the region in which to install your secondary instance, you must do so to ensure that instance creation is successful.

See [Managing Regions](#).

Perform Data Replication Across Non-Default Domains

Replication is always enabled for the Default identity domain. The Default identity domain always replicates to all regions to which the tenant is subscribed. However, if you are using a *non-default* identity domain in which to install disaster recovery, data is *not* automatically replicated to the secondary region. You must explicitly enable replication for disaster recovery installation to succeed. See [Replicating an Identity Domain to Multiple Regions](#).

Configure Policies for Default or Defined Tags

If you are using default or defined tags, ensure that you configure the necessary policies. Add a policy to use tags for tenancies requiring default defined tags set up by the tenancy owner.

1. From the Home page, select **Identity & Security**, then **Policies**.
2. Click **Create Policy** and add a policy as follows:
If you are using default defined tags or defined tags:

- a. If there is only a default domain, add a tag-based policy in the default domain:

```
allow group group_name use tag-namespaces in tenancy
```

- b. If there is a secondary domain, add a tag policy in the secondary region domain to allow the use of defined tags:

```
allow group domain-secondary-region/group_name use tag-namespaces in tenancy
```

See [Required Permissions for Working with Defined Tags](#).


Install and Configure Oracle Integration for Disaster Recovery

You can select to enable disaster recovery when provisioning an Oracle Integration instance. This action installs a primary instance in one region and a secondary instance in another region.

1. Sign in to the Oracle Cloud Infrastructure Console and note your selected region. See [Sign In to the Oracle Cloud Infrastructure Console in Provisioning and Administering Oracle Integration 3](#).



2. Open the navigation menu and click **Developer Services**. Under **Application Integration**, click **Integration**.
3. From the **Compartment** list, click through the hierarchy of compartments and select the one in which to create the instance. You may need to expand the + icon to find the compartment to use. Compartments can contain other compartments.



The page is refreshed to show any existing instances in that compartment.

4. Click **Create instance**.
5. Enter the following details.

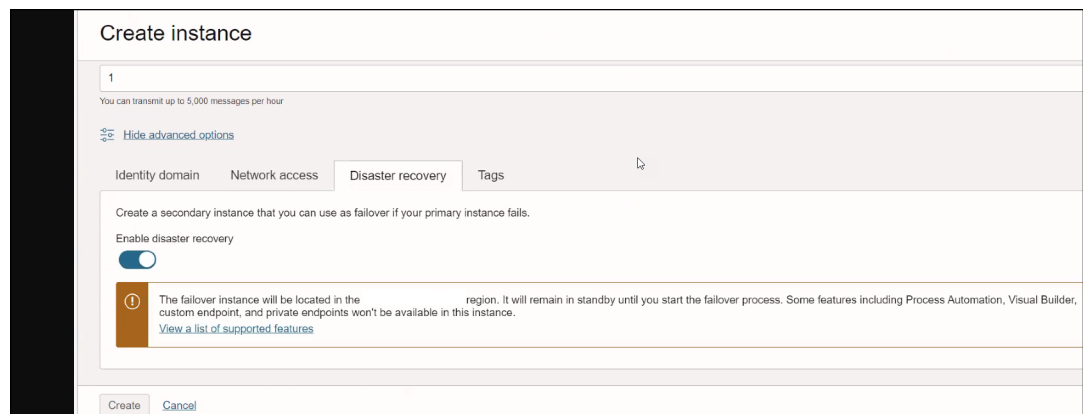
Field	Description
Display Name	Enter the display name for the instance.
Edition	Select from the following supported editions for disaster recovery. <ul style="list-style-type: none"> • Enterprise • Healthcare <p>Note: If you select the Standard edition, an error is displayed when you click Create.</p> <p>See Oracle Integration Editions in <i>Provisioning and Administering Oracle Integration 3</i> to see what's licensed in each supported edition.</p>
Shape	Only the Production shape is available for disaster recovery instances. <p>Note: The Development shape is not supported.</p>
License Type	Select the license type. See Choose a License Type in <i>Provisioning and Administering Oracle Integration 3</i> .
Message Packs	Enter the message pack number. See Choose a Message Pack Number in <i>Provisioning and Administering Oracle Integration 3</i> .

The **Use advanced options** link appears at the bottom of the Create instance page.

6. Click **Use advanced options** and select the **Disaster recovery** tab.
7. Select the **Enable disaster recovery** toggle.

A message is displayed.

The failover instance will be located in the `region_name` region. It will remain in standby until you start the failover process. . Some features including Process Automation, Visual Builder, custom endpoints, and private endpoints won't be available in this instance.



 **Note:**

Do *not* select to configure the options under the **Network access** tab. This feature is not supported in disaster recovery-enabled environments.

8. Click **Create**.

The **Work requests** section shows installation progress. As installation progresses, the **% Complete** value changes.

Work requests				
Operation	Status	% Complete	Accepted	Started
Create integration instance	In progress	35	Fri, Jul 19, 2024, 17:14:59 UTC	Fri, Jul 19, 2024, 17:15:13 UTC

 **Note:**

Creation of the primary and secondary instances can take time to complete due to DNS configuration.

9. Click **Create integration instance** to view details about installation progress. When creation of the primary and secondary instances in different regions completes successfully, the **Finished** field is updated. You do not receive a popup message when installation completes.

WR

Create integration instance

Work request information

Percent complete: 100

OCID: ...y7jp3e5ida [Show](#) [Copy](#)

Accepted: Wed, Sep 25, 2024, 16:54:04 UTC

Started: Wed, Sep 25, 2024, 16:54:21 UTC

Finished: Wed, Sep 25, 2024, 16:58:31 UTC

Log messages

Message	Accepted
Integration instance create operation is in progress.	Wed, Sep 25, 2024, 16:54:21 UTC
Successfully created integration instance.	Wed, Sep 25, 2024, 16:58:31 UTC

Showing

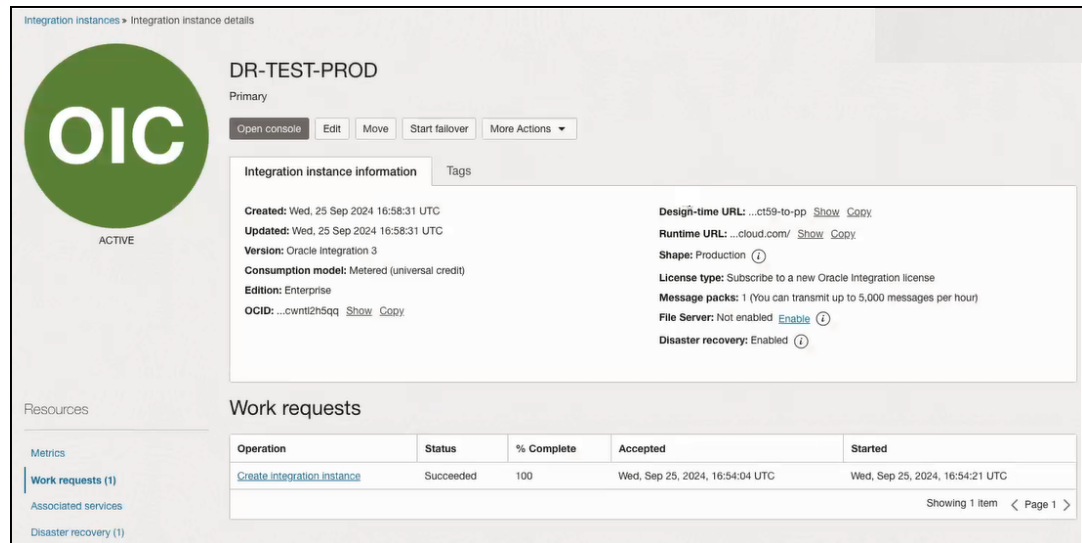
Resources

[Log messages \(0\)](#)

[Error messages \(0\)](#)

SUCCEEDED

The installed primary instance includes a green circle with the label **ACTIVE**. The word **Primary** appears below the instance name to indicate that this is the primary instance.



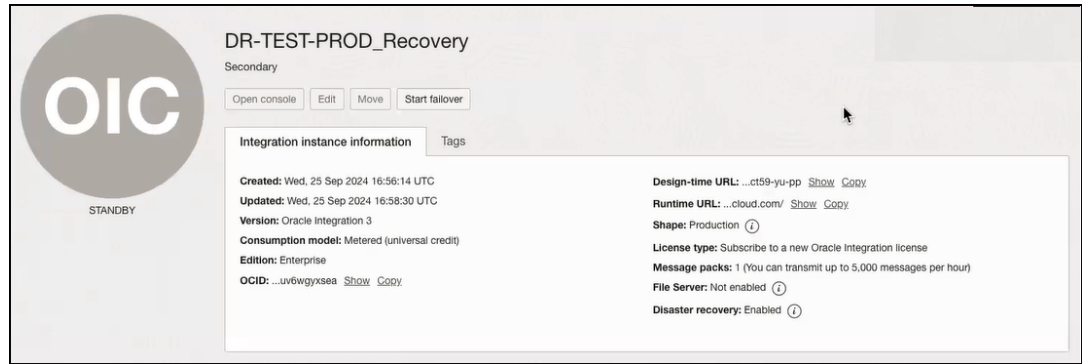
10. Under **Resources** in the left navigation pane, click **Disaster recovery**.



11. Click the secondary instance.

The details page for the secondary instance in the other region is displayed.

- The word **_Recovery** is appended to the end of the secondary instance name. The **_Recovery** word always appears as part of this instance name, even when the secondary instance becomes the primary instance after a failover.
- The word **Secondary** appears below the instance name. When you fail over from the initial instance to the secondary instance, **Secondary** is replaced with **Primary**.
- A status of **STANDBY** appears below the circle labeled **OIC**.
- The **Start failover** button is enabled for use. This is the only life cycle action that can be performed from the secondary instance.



12. Start designing integrations in your primary instance. As soon as you start, the synchronization of design-time metadata begins between the primary and secondary instances in near real time. This reduces the chance for data loss when a failover is required.

Perform Failover Prerequisite Tasks

You must perform several additional tasks while configuring Oracle Integration for disaster recovery. Complete these tasks prior to initiating a failover.

Topics:

- [Understand Your Connectivity Agent Responsibilities During a Failover](#)
- [Configure Email Delivery If Using Your Own Email Tenancy](#)
- [Configure File Server for Disaster Recovery](#)

Understand Your Connectivity Agent Responsibilities During a Failover

The location of your connectivity agent installation during a disaster determines if manual intervention is required. In most cases, the connectivity agent fails over and automatically redirects traffic to the secondary instance with no need for manual intervention. Review the following table to determine your responsibilities.

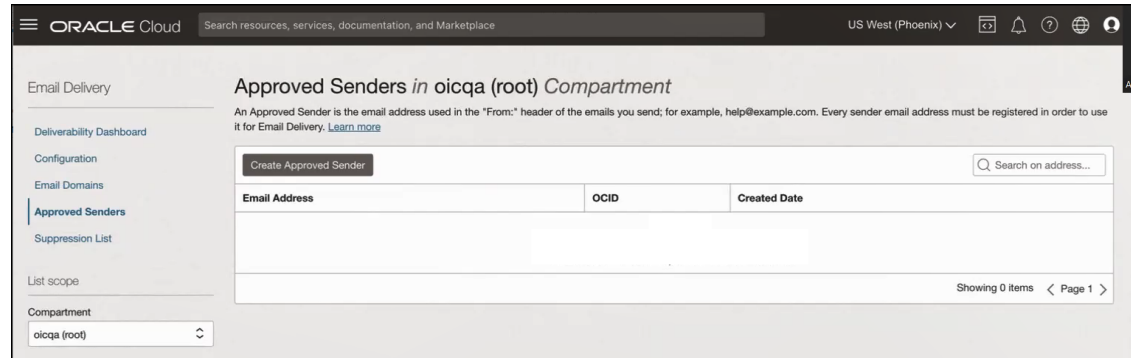
Scenario	Your Responsibilities
If the connectivity agent is installed in Oracle Cloud Infrastructure Compute, and the compute is impacted because of a disaster in that region or for another reason.	You must have a plan in place that allows for a quick recovery of the connectivity agent. See <i>Recover the Connectivity Agent in Case of Host Failure</i> in <i>Using Integrations in Oracle Integration 3</i> .
For all other scenarios not impacted by Oracle Cloud Infrastructure Compute failure (for example, the connectivity agent is installed in your on-premises environment).	A second connectivity agent installation is <i>not</i> required. After failover, the existing connectivity agent automatically redirects traffic to the secondary site and continues to work without any manual intervention required.

Configure Email Delivery If Using Your Own Email Tenancy

If you are using your own email tenancy, you must manually maintain your email notification details (for example, approved sender email addresses, email domains, and suppression list) in *both* the primary and secondary instances in the Oracle Cloud Infrastructure Console. You

must perform these tasks *prior* to failover. These tasks are not required if you are using the default Oracle Integration email tenancy.

The configuration pages for email domains, approved sender email addresses, and suppression list are visible in the Email Delivery section of the Oracle Cloud Infrastructure Console. For configuration details, see [Email Delivery](#).



The approved (default) sender and other email addresses must match with those specified on the Notifications page of the Oracle Integration instance. See *Configure Notification Emails in Using Integrations in Oracle Integration 3*.

Configure File Server for Disaster Recovery

All connections to File Server must use its port and hostname, rather than its port and IP address. Otherwise, the failover for File Server doesn't occur.

1. Update FTP Adapter connections to File Server.
 - a. In all integrations, identify all connections to File Server that are based on the FTP Adapter.
 - b. Review the connections to determine whether they connect to File Server using its port and hostname.
 - c. For any connections that use the port and IP address, update them so that they use the port and hostname.
2. Update all connections from SFTP clients.
 - a. Identify all the SFTP clients that connect to File Server.
 - b. Review the connections to determine whether they connect to File Server using its port and hostname.
 - c. For any connections that use the port and IP address, update them so that they use the port and hostname.

Fail Over to the Other Instance

You manually initiate a failover from the primary instance or from the secondary instance (if the primary instance is unreachable). You can perform failover during an actual disaster recovery occurrence or to test this functionality per your business or legal requirements.

- [Perform a Failover](#)

- [Deactivate and Reactivate Integrations with Polling Triggers](#)
- [Change the Built-in API Calls to Reflect the New Hostname and Integration Instance Name](#)

Perform a Failover



Note:

No notification is sent by Oracle Integration when an instance fails. However, Oracle Cloud Infrastructure provides functionality for setting alarms when issues occur. See [Managing Alarms](#).

1. Go to the primary instance in the Oracle Cloud Infrastructure Console. For this example, failover is performed from the primary instance in one region to the secondary instance in a different region. The steps are the same when performing a failover from the secondary instance.
2. Click **Start failover**.

Integration instances > Integration instance details

DR-TEST-PROD
Primary

Open console Edit Move Start failover More Actions

Integration instance information Tags

Created: Wed, 25 Sep 2024 16:58:31 UTC
Updated: Wed, 25 Sep 2024 16:58:31 UTC
Version: Oracle Integration 3
Consumption model: Metered (universal credit)
Edition: Enterprise
OCID: ...cwnt12h5sq Show Copy

Design-time URL: ...ct59-to-pp Show Copy
Runtime URL: ...cloud.com/ Show Copy
Shape: Production ⓘ
License type: Subscribe to a new Oracle Integration license
Message packs: 1 (You can transmit up to 5,000 messages per hour)
File Server: Not enabled Enable ⓘ
Disaster recovery: Enabled ⓘ

Resources

Metrics

Work requests (1)

Associated services

Disaster recovery (1)

Operation	Status	% Complete	Accepted	Started
Create integration instance	Succeeded	100	Wed, Sep 25, 2024, 16:54:04 UTC	Wed, Sep 25, 2024, 16:54:21 UTC

Showing 1 item < Page 1 >

3. Click **Failover** when prompted to fail over to the secondary instance in the other region. The process to fail over to the instance in the other region begins. A second tab opens in your browser for failover status about the secondary instance.

Because data synchronization between the two instances has occurred in near real time since the completion of disaster recovery installation, the failover process takes approximately the same amount of time regardless of the amount of design-time metadata data in your instance.
4. Follow failover progress in the **Work Requests** section.
5. When failover completes, the previous primary instance goes into standby mode. The word **Secondary** now appears below the instance name. The status changes to **STANDBY** below the circle labeled **OIC**.
6. Click the tab in your browser to access the new primary instance. The status is shown as **ACTIVE** below the circle labeled **OIC** and the word **Primary** now appears below the instance name. The word **_Recovery** remains appended to the end of the new primary name.

If you were logged in to the primary instance of Oracle Integration during failover, you receive the following message indicating that the primary instance is now in a standby state.



7. Click the global (regionless) URL specified in the message to log in to the new primary instance. The global URL does not mention any region name.
8. Continue working as you were prior to receiving this message. Activated integrations in the original primary instance are displayed as activated in the new primary instance. As you work in the new primary instance, your changes to data are synchronized with the original primary instance automatically.
9. When the original primary instance is restored, click **Failover** in the Oracle Cloud Infrastructure Console of either instance *if* you want to fail back to the original primary instance.

Deactivate and Reactivate Integrations with Polling Triggers

When you initiate a failover for integration instances running with polling endpoints (for example, integrations involving Oracle Cloud Infrastructure Streaming, databases, JMS, and others), the integration transfers to the failover instance but is not automatically activated. After failover, you must manually identify and reactivate these integrations. Once reactivated, message flow resumes as expected.

Change the Built-in API Calls to Reflect the New Hostname and Integration Instance Name

If you use the Oracle Integration built-in APIs, the hostname and integration instance name change in the API call after failover completes.

For example:

- Pre-failover API call:

```
https://mydesign-pp-integration-prod-gen3.integration.us-  
ashburn-1.ocp.oraclecloud.com/ic/api/integration/v1/integrations/  
?integrationInstance=sysqa-drtest-cgs-inst2-bxffubagcv29-to-pp
```

- Post-failover API call:

```
https://mydesign-pp-integration-prod-gen3.integration.us-  
phoenix-1.ocp.oraclecloud.com/ic/api/integration/v1/integrations/  
?integrationInstance=sysqa-drtest-cgs-inst2-remote-bxffubagcv29-yu-pp
```

This is expected behavior because the global (regionless) URL is used only for runtime. You must change the design-time hostname and integration instance name in the API call to reflect the post-failover values.

Configure Email Notification Settings After Failover

If using your own email tenancy, the SMTP configuration section is empty in the new primary instance. After every failover occurs, you must go to the Notifications page in the new primary instance and manually re-enter the SMTP configuration settings.



Note:

If you are using the default Oracle Integration email tenancy, no manual configuration is necessary. Configuration is automatic after failover.

1. Go to the new primary Oracle Integration instance.
2. In the navigation pane, click **Settings**, then **Notifications**. The Notifications page is displayed.
3. In the **SMTP configuration** section, click to enable the customer tenancy mode.
4. Specify the SMTP user name and password and the default sender email address (that is, the `from` address). These three fields are required. You can also optionally specify a default address for system notifications. The settings you enter must match those settings entered prior to failover in [Configure Email Delivery If Using Your Own Email Tenancy](#).

SMTP configuration

If you use your own SMTP server, you must manage approved senders and the suppression list in Oracle Cloud Infrastructure. Before using your own configuration, you must generate the SMTP credentials and save them. [Open Oracle Cloud Infrastructure](#)

Use my SMTP configuration Test

Configure SMTP details

Prior to saving, test your configuration which validates it by sending an email to the specified recipient(s).

SMTP username Required

SMTP password Required

Default sender email address Required

Sender email address for system notification

5. Click **Test** to validate the SMTP credentials.
6. Click **Save**. Configuration is complete.