Oracle® Cloud

Using Oracle Cloud Infrastructure Full Stack Disaster Recovery

F32206-10
May 2024

Oracle Cloud Using Oracle Cloud Infrastructure Full Stack Disaster Recovery,

F32206-10

Primary Author: Roopam Jain

Contributors: Shekhar Borde, Praveen Sampath, Gregory King, Jean-Francois Verrier, Padmaja Potineni, Rama Vijjapurapu, Glen Hawkins, Suraj Ramesh, Mahesh Desai, Santhosh Shankaramanchi, Aravind Kadiyala, Chandra Sekhar Atla, Saurabh Sachdev, Harsh Patel

Contributing Authors: Prakash Jashnani, Subhash Chandra, Ramya P

Contents

Preface

1 Overview of Full Stack Disaster Recovery

2 Get Started with Full Stack Disaster Recovery

3 Manage Disaster Recovery Protection Groups

A Reference

Preface

Topics

Audience

This document is intended for Disaster Recovery (DR) administrators responsible for defining, creating, and implementing DR strategies for enterprise applications, databases, and infrastructure deployed in Oracle Cloud.

Documentation Accessibility

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

Access to Oracle Support

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

Related Resources

See these Oracle resources:

- Get Started with Oracle Cloud
- Oracle Public Cloud

  https://cloud.oracle.com

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.

Convention	Meaning
italic	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 Overview of Full Stack Disaster Recovery

Learn about terminology, concepts, and benefits of Full Stack Disaster Recovery.

- About Full Stack Disaster Recovery
  Full Stack DR is an Oracle Cloud Infrastructure (OCI) disaster recovery orchestration and management service that provides comprehensive disaster recovery capabilities for all layers of an application stack, including infrastructure, middleware, database, and application.
- Benefits of Full Stack Disaster Recovery
  Full Stack DR provides multiple benefits in the area of business continuity.
- Full Stack Disaster Recovery Terminology and Concepts
  Before using Full Stack DR, familiarize yourself with the following key terms and concepts.

About Full Stack Disaster Recovery

Full Stack DR is an Oracle Cloud Infrastructure (OCI) disaster recovery orchestration and management service that provides comprehensive disaster recovery capabilities for all layers of an application stack, including infrastructure, middleware, database, and application.

> **Note:**
>
> The terms *Oracle Cloud Infrastructure Full Stack Disaster Recovery*, *OCI Full Stack Disaster Recovery*, *Oracle Cloud Infrastructure Full Stack Disaster Recovery Service*, *Full Stack DR*, and *Disaster Recovery* are used interchangeably throughout this documentation. All the terms refer to the same service.

Figure 1-1 Overview of region map

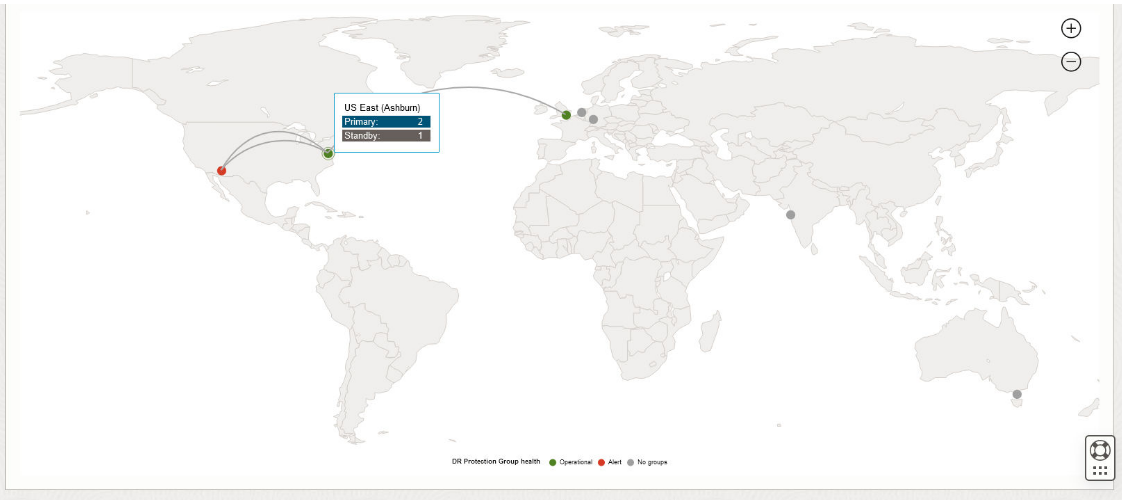

Full Stack Disaster Recovery assures comprehensive business continuity from a variety of data center outages to ensure that organizations have a minimal impact from region-wide outages or Availability Domain (AD) outages.

Full Stack DR is flexible enough to easily integrate with various Oracle platforms, non-Oracle applications, and infrastructure. Full Stack DR generates, runs, and monitors disaster recovery plans for services and applications deployed in your tenancy. Full Stack DR operates at the service level, so there is no impact on other services running in your tenancy. Based on your specific needs, you can customize the disaster recovery plans generated by Full Stack DR.

You can actively monitor the progress of Full Stack DR operations and take corrective actions if there are errors during an operation.

You can also validate and monitor business continuity readiness and compliance by periodically running Full Stack DR Prechecks.

Benefits of Full Stack Disaster Recovery

Full Stack DR provides multiple benefits in the area of business continuity.

- **Provides Full Application Recovery**: Full Stack DR provides recovery for the entire application stack (business system) instead of recovery of individual components such as databases or compute instances. Business continuity depends on recovering the entire application stack instead of a few select components.
- **Intelligent Plan Generation**: Full Stack DR evaluates the topology and deployment characteristics of a variety of application stacks and automatically creates a dedicated Disaster Recovery (DR) workflow for recovering all the components of the application stack in another region.
- **Minimizes Disaster Recovery Time**: Full Stack DR eliminates the need to perform manual disaster recovery for individual resources in an application stack, such as application components, middleware, databases, and infrastructure components. You can quickly bring up the entire application stack in a different region using fully automated DR workflows.
- **Validates Disaster Recovery Workflows and Configurations**: A common problem with disaster recovery workflows is that topology changes, configuration changes, or environment drift can make disaster recovery workflows unusable because they do not match the environment they are protecting. Full Stack DR provides a comprehensive series of prechecks that you can use to continuously validate the conformity of DR workflows with the application environment you configure them to protect.
- **Reduces Human Errors**: Disaster Recovery workflows are error-prone as they require a complicated orchestration of many precise, interconnected steps to be performed by humans. Full Stack DR automates and sequences these steps to provide fast and seamless disaster-recovery operations across regions without human intervention, thus reducing Disaster Recovery (DR) workflow errors.
- **Flexible and Customizable DR Workflows**: Disaster Recovery workflows are highly flexible and customizable. You can adapt them to perform DR for any application stack or a set of IT assets. You can add your own callouts and custom logic to any workflow and precisely configure the action of these customizations.
- **Eliminates the Need for Special Skills**: The operators and administrators do not require any special skills or domain expertise in areas such as applications, and storage replication. You can create and test DR workflows ahead of time. This process eliminates the need for experts when performing DR operations.

- **Single Pane of Glass Monitoring and Management**: Full Stack DR provides a single pane of glass monitoring and management capability for all disaster recovery needs. You can create DR protection configurations, monitor DR readiness, execute DR workflows, and manage errors using the OCI console.

Full Stack Disaster Recovery Terminology and Concepts

Before using Full Stack DR, familiarize yourself with the following key terms and concepts.

- **Disaster Recovery (DR)** – The process of restoring some or all parts of a business system (a service) after an outage. The recovery of this business system can occur in the same geographical region or in another geographical region.
- **Full Stack** – A term used to collectively refer to all the functional layers of a business system or application or software service. An application can be comprised of different functional layers or tiers such as: application layer, middleware layer, database layer, and infrastructure layer.
- **Recovery Point Objective (RPO)** – The RPO defines the maximum amount of data loss that can be tolerated as part of the DR restoration. RPO is typically expressed in units of time.
- **Recovery Time Objective (RTO)** – The RTO defines the maximum amount of time that the application or service under DR protection can be unavailable until service is restored. RTO is typically expressed in units of time.
- **Primary** – The production version of an application or service that is currently in use. Full Stack DR refers to the Primary version of an application as having a **Primary** role.
- **Standby** – The reserved version of an application or service. **Standby** is also used to refer to the alternate region in which the application or service will be restored. Full Stack DR refers to the Standby version of an application as having a **Standby** role.
- **Warm Standby** - A DR model in which some or all of the components of an application or service are pre-deployed in the standby region to prepare for a future DR transition. This model involves higher operating costs but a lower RTO.
- **Cold Standby** - A DR model in which very few or none of the components of an application or service need to be pre-deployed in the standby region in preparation for a future DR transition. The application components are deployed as part of the DR transition. This model involves lower operating costs but a higher RTO.
- **Role** – Specifies whether an application and its region is currently the Primary (production) version or the Standby (reserved) version. An application's and its region's role changes as a result of a DR transition.
- **Resource** – A resource is a component in OCI that can be used and managed independently. Examples of OCI resources are: compute instances, block volumes, databases, load balancers, etc. Examples of resources provided by Full Stack DR are: DR Protection Groups, DR Plans, and DR Plan Executions.
- **DR Protection Group** – A resource type used by Full Stack DR. A DR Protection Group represents a consistency grouping defined for the purposes of disaster recovery. It is a collection of different OCI resources that comprise an application and must be treated as a combined group when performing disaster recovery operations. For example, a DR Protection Group may consist of application servers (compute instances), associated block storage (grouped as volume groups), and databases.
- **Association** – A pair relationship defined between two DR Protection Groups. DR Protection Groups in Full Stack DR must be associated (paired) in a Primary and Standby relationship before they can be used to implement DR services. An association between

two DR Protection Groups is exclusive, that is, a DR Protection Group can only be associated with one other DR Protection Group.

- **DR Plan** – A resource type used by Full Stack DR. A DR Plan represents a DR workflow associated with a pair of DR Protection Groups. A DR Plan is represented as a sequence of Plan Groups. These Plan Groups in turn consist of Plan Steps. A DR Plan can only be created at the Standby DR Protection Group.
- **DR Plan Execution** – A resource type used by Full Stack DR. A DR Plan Execution represents an execution (a running instance) of a DR Plan. A DR Plan Execution can only be created (launched) at a Standby DR Protection Group.
- **Plan Group** – A group of steps in a DR Plan. A DR Plan consists of one or more Plan Groups that execute sequentially. All steps in a Plan Group execute in parallel.
- **Plan Step** – A single indivisible unit of execution in a DR Plan. A Plan Step must belong to a Plan Group.
- **Built-In Groups or Steps** – A type of Plan Group or Step that is generated automatically by Full Stack DR when a DR Plan is created. Examples of Built-in Plan Steps are: Launch Compute Instance, Switchover Database, etc.
- **User-Defined Groups or Steps**– A type of Plan Group or Step that is added by the user to a DR Plan after the DR plan is created by Full Stack DR.
- **Precheck** – A predetermined set of checks associated with a DR Plan. A Precheck for a DR Plan performs a set of checks to validate that a DR Plan is compliant with the members and configuration of the DR Protection Groups with which the DR Plan is associated. Prechecks are used to perform ongoing DR Plan validation (DR readiness checks) to ensure that the DR Plan (DR workflow) stays aligned with the topology it protects.
- **Switchover** – A type of DR Plan that performs a planned transition of services from the Primary DR Protection Group to the Standby DR Protection Group. Switchover plans perform an orderly transition by shutting down the application stack in the primary region and then bringing it up in the standby region. Therefore, a switchover plan requires that application stack components and other required OCI services be available in both regions.
- **Failover** – A type of DR Plan that performs an unplanned transition of services to the Standby DR Protection Group. Failover plans usually perform an immediate transition by bringing up the application stack in the standby region, without attempting to shutdown service in the primary region. Hence, a failover plan only requires that OCI services be available in the standby region. Failover plans are typically used to perform DR transitions when an outage or disaster affects the primary region.
- **DR Drill -** Performing a DR Drill for a pair of associated DR Protection Groups brings up a replica of the application stack in the standby DR Protection Group. This replica stack can be used to test and validate the effectiveness of the DR processes. A **Start DR Drill** plan execution creates the application stack replica in the standby, and a **Stop DR Drill** plan execution terminates this application stack replica.

2
Get Started with Full Stack Disaster Recovery

Learn about accessing Full Stack DR, prerequisites, and understand the workflow for using the service.

- How Full Stack Disaster Recovery Works?
  You can use the Full Stack DR service to create dedicated Disaster Recovery (DR) configurations for each of your application stacks that requires DR protection.
- How to Access Full Stack Disaster Recovery?
  You can access Oracle Cloud Infrastructure Full Stack DR using the Oracle Cloud Infrastructure Console (a browser based interface), REST APIs, Oracle Cloud Infrastructure SDKs, command-line interface, and DevOps tools.
- Prerequisites for Full Stack Disaster Recovery
  Before you start using Full Stack DR, learn about the prerequisites like preparing compute instances, block storage, applications, and more.

How Full Stack Disaster Recovery Works?

You can use the Full Stack DR service to create dedicated Disaster Recovery (DR) configurations for each of your application stacks that requires DR protection.

Create a dedicated Full Stack DR configuration for an application stack using the following steps:

1. Identify all the components and dependencies of that application stack.
2. Create a consistency grouping of these application components and dependencies.
3. Create automated DR plans (workflows) to execute planned and unplanned DR migrations.
4. Customize the DR plans to incorporate additional DR requirements.
5. Test and validate the DR plans to ensure that they are error-free.
6. Execute the pretested DR workflows to perform planned or unplanned DR operations.

Full Stack Disaster Recovery Service currently supports disaster recovery for the following OCI resource types:

- Compute Instances
- Boot and Block Volumes (Volume Groups)
- Oracle Exadata Database Service
- Oracle Base Database Service
- Oracle Autonomous Database Serverless
- File Systems
- Load Balancers
- Network Load Balancers

> **Note:**
>
> Full Stack DR can support cross-region as well as intra-region DR configurations. However, Oracle recommends using cross-region DR configurations to protect against region-wide outages.

How to Access Full Stack Disaster Recovery?

You can access Oracle Cloud Infrastructure Full Stack DR using the Oracle Cloud Infrastructure Console (a browser based interface), REST APIs, Oracle Cloud Infrastructure SDKs, command-line interface, and DevOps tools.

- **Using the Oracle Cloud Infrastructure Console (UI)**

  To access Full Stack Disaster Recovery using the Console:

  1. Sign in to your Oracle Cloud account at cloud.oracle.com. See Signing in to Oracle Cloud in the *Oracle Cloud Infrastructure* documentation.
  2. Select the region.
     If your tenancy is subscribed to multiple regions, select your region from the **Region** menu at the top of the page.
  3. Open the navigation menu, click **Migration & Disaster Recovery**, and select **Disaster Recovery** or **DR Protection Groups**.
     The Disaster Recovery (DR) Protection Groups page is displayed. All existing DR protection groups in your compartment are displayed. Use the **Compartment** list in the panel on the left to change the compartment.

- **Using the Command-Line Interface (CLI)**

  If you are using a command-line interface (CLI), then you can access the Full Stack DR service by installing the Oracle Cloud CLI and listing all the available commands by running the command `oci disaster-recovery --help` at the shell prompt.

  For details on how to install and use the Oracle Cloud CLI, see: Command Line Interface (CLI).

- **Using the REST API**

  REST API users can access the Full Stack DR service by connecting to one of the regional REST API endpoints listed in Full Stack Disaster Recovery API and then invoking an HTTP method such as GET, PUT, POST, or DELETE.

  For more information, see: REST APIs.

- **Using a Software Development Kit (SDK)**

  You can access Full Stack DR using numerous language-specific Software Development Kits (SDKs). Full Stack DR supports SDKs for the following languages:

  - Java
  - Python
  - Typescript and Javascript
  - .NET
  - Go
  - Ruby

For more information, see Software Development Kits and Command Line Interface.

- **2.2.5 Using other DevOps Tools**

  You can use the following DevOps tools to integrate with the Full Stack DR service:

  - Terraform Provider
  - Windows Powershell

  For more information, see: DevOps Tools and Plug-ins.

Prerequisites for Full Stack Disaster Recovery

Before you start using Full Stack DR, learn about the prerequisites like preparing compute instances, block storage, applications, and more.

- Who can use Full Stack Disaster Recovery?
  Full Stack DR is available to all Oracle Cloud customers using **Universal Credits** and **Pay As You Go**.
- Preparing Object Storage Buckets for Operation Logs
  Full Stack DR configurations use Object Storage to store Disaster Recovery (DR) operation logs.
- Preparing Compute Instances for Full Stack Disaster Recovery
  If the disaster recovery topology contains any compute instances, use the following steps to prepare each compute instance for Full Stack DR:
- Preparing Block Storage for Full Stack Disaster Recovery
  If the disaster recovery topology contains any block storage volumes attached to compute instances, or any stand-alone block storage volumes, use the steps in this topic to prepare the block storage volumes for Full Stack DR.
- Preparing Oracle Databases for Full Stack Disaster Recovery
  If the disaster recovery topology contains any databases related to Oracle Base Database Service, Oracle Exadata Database Service on Dedicated Infrastructure, use the steps in this topic to prepare the databases for Disaster Recovery (DR).
- Preparing Oracle Autonomous Database Serverless for Full Stack Disaster Recovery
  If the disaster recovery topology contains any Oracle Autonomous Database Serverless, use the steps in this topic to prepare the databases for Full Stack DR.
- Preparing Applications for Full Stack Disaster Recovery
  If the disaster recovery topology contains any Oracle applications or custom applications, use the steps in this topic to prepare the applications for Full Stack DR.
- Preparing Other Full Stack Disaster Recovery Workflow Customizations
  If your Disaster Recovery (DR) topology requires additional user-defined custom steps to be run as part of the DR workflow, then use the steps described in this topic to prepare such customizations.
- Creating Default NFS Client Export Option for File System
  You can add the IP addresses from the standby region and those are copied over during the switchover of the file system to the other region.
- Setting Up Policies for Full Stack Disaster Recovery and Other Services
  In addition to preparing various components for disaster recovery, you must set up IAM policies that allow Full Stack DR to manage these components during the disaster recovery process.

Related Topics

- Signing In to the Console
- Setting Up Your Tenancy
- VCNs and subnets
- How Policies Work

Who can use Full Stack Disaster Recovery?

Full Stack DR is available to all Oracle Cloud customers using **Universal Credits** and **Pay As You Go**.

Full Stack DR is not available for users of the Oracle Cloud Free Tier. The service is currently available in the following regions:

Table 2-1 Full Stack Disaster Recovery Regions

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
Australia East (Sydney)	ap-sydney-1	Sydney, Australia	SYD	OC1	1
Australia Southeast (Melbourne)	ap-melbourne-1	Melbourne, Australia	MEL	OC1	1
Brazil East (Sao Paulo)	sa-saopaulo-1	Sao Paulo, Brazil	GRU	OC1	1
Brazil Southeast (Vinhedo)	sa-vinhedo-1	Vinhedo, Brazil	VCP	OC1	1
Canada Southeast (Montreal)	ca-montreal-1	Montreal, Canada	YUL	OC1	1
Canada Southeast (Toronto)	ca-toronto-1	Toronto, Canada	YYZ	OC1	1

Table 2-1 (Cont.) Full Stack Disaster Recovery Regions

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
Chile (Santiago)	sa-santiago-1	Santiago, Chile	SCL	OC1	1
Chile West (Valparaiso)	sa-valparaiso-1	Valparaiso, Chile	VAP	OC1	1
Colombia Central (Bogota)	sa-bogota-1	Bogota, Colombia	BOG	OC1	1
France Central (Paris)	eu-paris-1	Paris, France	CDG	OC1	1
France South (Marseille)	eu-marseille-1	Marseille, France	MRS	OC1	1
Germany Central (Frankfurt)	eu-frankfurt-1	Frankfurt, Germany	FRA	OC1	3
India South (Hyderabad)	ap-hyderabad-1	Hyderabad, India	HYD	OC1	1
India West (Mumbai)	ap-mumbai-1	Mumbai, India	BOM	OC1	1
Israel Central (Jerusalem)	il-jerusalem-1	Jerusalem, Israel	MTZ	OC1	1
Italy Northwest (Milan)	eu-milan-1	Milan, Italy	LIN	OC1	1

Table 2-1 (Cont.) Full Stack Disaster Recovery Regions

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
Japan Central (Osaka)	ap-osaka-1	Osaka, Japan	KIX	OC1	1
Japan East (Tokyo)	ap-tokyo-1	Tokyo, Japan	NRT	OC1	1
Mexico Central (Queretaro)	mx-queretaro-1	Queretaro, Mexico	QRO	OC1	1
Mexico Northeast (Monterrey)	mx-monterrey-1	Monterrey, Mexico	MTY	OC1	1
Netherlands Northwest (Amsterdam)	eu-amsterdam-1	Amsterdam, Netherlands	AMS	OC1	1
Saudi Arabia West (Jeddah)	me-jeddah-1	Jeddah, Saudi Arabia	JED	OC1	1
Singapore (Singapore)	ap-singapore-1	Singapore,Singapore	SIN	OC1	1
South Africa Central (Johannesburg)	af-johannesburg-1	Johannesburg, South Africa	JNB	OC1	1
South Korea Central (Seoul)	ap-seoul-1	Seoul, South Korea	ICN	OC1	1

Table 2-1 (Cont.) Full Stack Disaster Recovery Regions

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
South Korea North (Chuncheon)	ap-chuncheon-1	Chuncheon, South Korea	YNY	OC1	1
Spain Central (Madrid)	eu-madrid-1	Madrid, Spain	MAD	OC1	1
Sweden Central (Stockholm)	eu-stockholm-1	Stockholm, Sweden	ARN	OC1	1
Switzerland North (Zurich)	eu-zurich-1	Zurich, Switzerland	ZRH	OC1	1
UAE Central (Abu Dhabi)	me-abudhabi-1	Abu Dhabi, UAE	AUH	OC1	1
UAE East (Dubai)	me-dubai-1	Dubai, UAE	DXB	OC1	1
UK South (London)	uk-london-1	London, United Kingdom	LHR	OC1	3
UK West (Newport)	uk-cardiff-1	Newport, United Kingdom	CWL	OC1	1
US East (Ashburn)	us-ashburn-1	Ashburn, VA	IAD	OC1	3
US Midwest (Chicago)	us-chicago-1	Chicago, IL	ORD	OC1	3

Table 2-1 (Cont.) Full Stack Disaster Recovery Regions

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
US West (Phoenix)	us-phoenix-1	Phoenix, AZ	PHX	OC1	3
US West (San Jose)	us-sanjose-1	San Jose, CA	SJC	OC1	1

Related Topics

- Oracle Cloud Infrastructure Free Tier
- About Universal Credits

Preparing Object Storage Buckets for Operation Logs

Full Stack DR configurations use Object Storage to store Disaster Recovery (DR) operation logs.

Before you create any DR configurations, you must create Object Storage buckets to include in the DR configuration process.

Oracle recommends that you follow these guidelines when creating the Object Storage bucket:

- Use a separate dedicated bucket for each DR protection group.
- Use Standard storage tier, not **Archive**.
- Do not set up replication for this object store bucket.
- Do not use this bucket to write other data, reserve it exclusively for use for logs for one DR protection group.
- Ensure that the object store bucket is writable by the user running DR plan executions.
- Ensure that no retention policies are set on the object store bucket.

Related Topics

- Overview of Object Storage
- Create a Bucket for Storage

Preparing Compute Instances for Full Stack Disaster Recovery

If the disaster recovery topology contains any compute instances, use the following steps to prepare each compute instance for Full Stack DR:

1. Create a block storage volume group in the same Availability Domain (AD) as the compute instance.
2. Add the boot volume for the compute instance to the volume group created in Step 1.
3. Add all block volumes attached to the compute instance to the volume group created in Step 1.

4. If the DR workflow requires execution of user-defined (custom) scripts or commands on compute instances, set up IAM policies so you can run these commands using the Oracle Cloud Agent. Configure `sudo` access for a `root` administrator on these instances if you want to invoke the scripts or commands with a different user ID.

> **Note:**
>
> If a command requires administrator permissions, you must grant administrator permissions to the Compute Instance Run Command plugin to be able to run the command. See Running Commands with Administrator Privileges.

Related Topics

- Volume Groups
- Running Commands on an Instance

Preparing Block Storage for Full Stack Disaster Recovery

If the disaster recovery topology contains any block storage volumes attached to compute instances, or any stand-alone block storage volumes, use the steps in this topic to prepare the block storage volumes for Full Stack DR.

1. Create a block storage volume group in the same Availability Domain (AD) as the compute instance.
2. Add the block storage elements (boot volumes or block volumes) to the volume group you created in Step 1.
3. Identify the standby region (recovery region) in which you want to recover the block storage. This region can be the same as the primary region. However, Oracle recommends that you generally use cross-region DR configurations to protect against region-wide outages.
4. Configure volume group replication to the standby region for the volume group you created in Step 1. Alternatively, configure backups for the volume group to the standby region.

> **Note:**
>
> - Setting up block storage replication or backups is only required for DR topologies where compute instances move across regions. For an active or passive DR topology, you do not need to configure block storage replications or backups. However, you may require other mechanisms, such as `rsync`, to keep the primary and standby compute instance file systems in sync.
> - You do not need to configure both replication and backups for Full Stack DR, only one of them is required. Oracle recommends you to use replication over backups because replication provides a much better recovery point (the data in backups can be up to a day old). If both, replication and backups are configured, Full Stack DR chooses replicas over backups for restoring block storage.

Related Topics

- Volume Groups

- Cross Region Replication
- Overview of Block Volume Backups

Preparing Oracle Databases for Full Stack Disaster Recovery

If the disaster recovery topology contains any databases related to Oracle Base Database Service, Oracle Exadata Database Service on Dedicated Infrastructure, use the steps in this topic to prepare the databases for Disaster Recovery (DR).

1. Configure a Data Guard association for the primary database.
2. Ensure that the Data Guard peer (standby database) is in the standby region that you intend to use for disaster recovery.
3. Create a vault in the primary and standby regions.
4. Create a secret in each of the primary and standby vaults.
5. Store the database admin password in the secrets created in the primary and standby regions.

> **Note:**
>
> This secret containing the database password is required when creating a disaster recovery configuration for the database.

Related Topics

- Use Oracle Data Guard on a DB System
- Use Oracle Data Guard with Exadata Cloud Infrastructure
- OCI Vault
- Create a Vault
- Create a Vault Secret

Preparing Oracle Autonomous Database Serverless for Full Stack Disaster Recovery

If the disaster recovery topology contains any Oracle Autonomous Database Serverless, use the steps in this topic to prepare the databases for Full Stack DR.

1. Configure a cross-regional Autonomous Data Guard association for the primary database.
2. Ensure that the Autonomous Data Guard peer (standby database) is in the standby region that you intend to use for disaster recovery.

> **Note:**
>
> Full Stack DR does not support Autonomous Data Guard configurations where the peer (standby) autonomous database is in the same region.

Preparing Applications for Full Stack Disaster Recovery

If the disaster recovery topology contains any Oracle applications or custom applications, use the steps in this topic to prepare the applications for Full Stack DR.

1. Identify any scripts, commands, and associated parameters that you use to manage the application life cycle. For example, scripts or commands that start or stop the application, or validate application functionality.
2. Determine if each of these scripts or commands will reside locally on each application instance, or if they will reside in Object Storage.
3. Identify any specific user IDs that you want to use when invoking these scripts or commands.

   Note all the details you collect in these steps as you will need them when you customize Full Stack DR workflows for application recovery.

Preparing Other Full Stack Disaster Recovery Workflow Customizations

If your Disaster Recovery (DR) topology requires additional user-defined custom steps to be run as part of the DR workflow, then use the steps described in this topic to prepare such customizations.

1. Identify any scripts, commands, or Oracle Functions you will run in the user-defined custom steps.
2. Determine if each of the entities that you plan to use for customization are:
   - Scripts stored locally on one or more compute instances that are part of the disaster recovery topology.
   - Scripts stored in Object Storage.
   - Serverless code stored as an Oracle Function.
3. If the customization uses scripts or commands that are stored locally on a compute instance or in Object Storage, then identify any specific user IDs that you want to use to run these scripts or commands.

Related Topics

- Functions

Creating Default NFS Client Export Option for File System

You can add the IP addresses from the standby region and those are copied over during the switchover of the file system to the other region.

So, you can add the NFS export options for both the regions. This allows Full Stack DR to mount the file system on any new IP address of the compute instance in the standby region.

Setting Up Policies for Full Stack Disaster Recovery and Other Services

In addition to preparing various components for disaster recovery, you must set up IAM policies that allow Full Stack DR to manage these components during the disaster recovery process.

Related Topics

- Policies for Full Stack Disaster Recovery
  This chapter lists the Identity and Access Management (IAM) policies you must configure for Full Stack DR.

3 Manage Disaster Recovery Protection Groups

Learn how to view protection group information, modify protection group definitions, and delete protection groups.

- Overview of Disaster Recovery Protection Groups
  A Disaster Recovery (DR) Protection Group is a consistency grouping created for the purposes of disaster recovery. A DR Protection Group groups together all the components of a full stack application so that you can recover all the components together to restore the full stack application.
- Create Disaster Recovery Protection Groups
  Create a protection group for each region that you want to include in a disaster recovery plan.
- View Disaster Recovery Protection Groups
  The details page of a Disaster Recovery (DR) protection group displays details such as its Plans, Plan executions, Members, and associated Work requests.
- Associate a Disaster Recovery Protection Group
  Learn how to associate a Disaster Recovery (DR) Protection Group with a peer.
- Disassociate Disaster Recovery Protection Groups
  Remove the association between a primary protection group and its standby protection group by updating the details of either protection group.
- Modify Disaster Recovery Protection Groups
  Modify a Disaster Recovery (DR) protection group definition to rename the protection group, add or remove members, or move the protection group to a different compartment.
- Update Disaster Recovery Protection Group Role
  Learn how to update the role of a Disaster Recovery (DR) Protection Group.
- Delete a Disaster Recovery Protection Group
  Learn how to delete a Disaster Recovery (DR) Protection Group.
- Rename a Disaster Recovery Protection Group
  Learn how to rename a Disaster Recovery (DR) Protection Group.

Overview of Disaster Recovery Protection Groups

A Disaster Recovery (DR) Protection Group is a consistency grouping created for the purposes of disaster recovery. A DR Protection Group groups together all the components of a full stack application so that you can recover all the components together to restore the full stack application.

- Only OCI resources, such as compute instances, volume groups, and Oracle Databases can be added to a DR Protection Group. These resources are referred to as *members* of the DR Protection Group. A DR Protection Group is exclusively paired with only one other DR Protection Group to form a peer relationship.
- Typically, each of the DR Protection Groups in a peer relationship exists in separate regions so that you can use the two DR Protection Groups to perform cross-region disaster recovery. However, you can also perform disaster recovery across Availability Domains (ADs) within the same region.

- A DR Protection Group is a region-specific OCI resource and does not migrate across regions. However, depending on the DR configuration and topology, members contained in a DR Protection Group, such as compute instances, can migrate to other regions and become members of the peer DR Protection Group.
- The DR Protection Groups homepage displays the following information:
  - Table
  - Region map
- View Region Maps of Disaster Recovery Protection Groups
  A region map displays a graphical representation of the DR Protection Groups globally present in a compartment. A region map displays the following information about the DR Protection Groups:

Related Topics

- Lifecycle States of Full Stack Disaster Recovery Resources
  Lifecycle states of Full Stack DR resources like DR Protection Groups, DR Plans, and DR Plan Executions.

View Region Maps of Disaster Recovery Protection Groups

A region map displays a graphical representation of the DR Protection Groups globally present in a compartment. A region map displays the following information about the DR Protection Groups:

Figure 3-1 Overview of region map

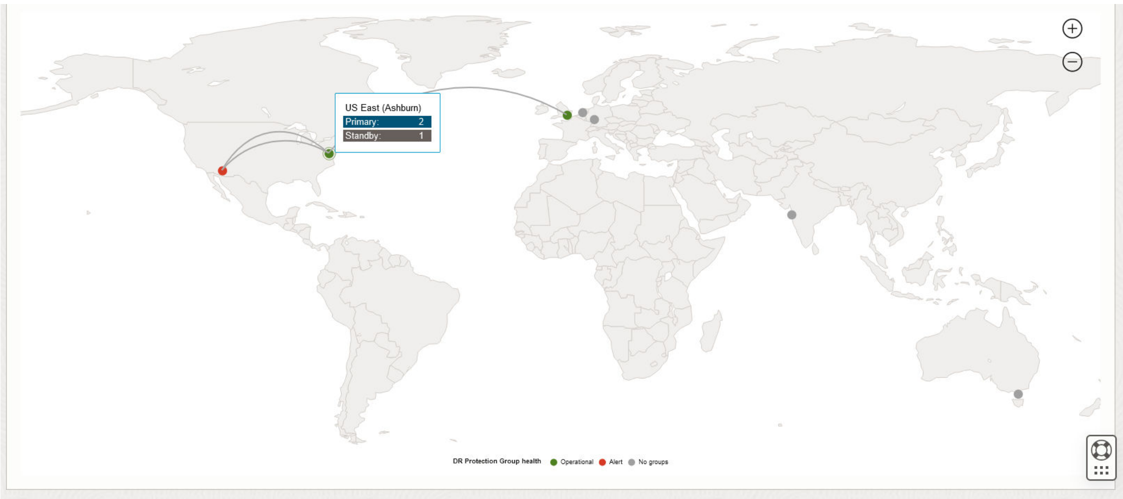

- A global view of all the subscribed regions in the tenancy, where each subscribed region is shown as a dot.
- Links connecting the current region with other regions indicate associations between groups in the current region with groups in other regions.
- Hover over a link (association) to view the names of the Primary and Standby groups in that association.
- Click on a link to visit the details page for the associated group in the current region.

- Hover over a region dot to see a summary of the number of Primary, Standby, and Not-configured groups available in that region.
- The color of a region dot indicates the health summary for all groups in that region. DR Protection Group health is shown as follows:
  - Operational (Green): Indicates that all the group(s) are working as expected.
  - Alert (Red): Indicates that the region has one or more group(s) in the **Failed** or **Needs attention** state.
  - No groups (Gray): Indicates that no group(s) are configured in the region.
- DR protection group details: Click on the region to view the DR protection group details page.
- Additionally, zoom in to view the following details as shown:

Figure 3-2 Additional information in a region map

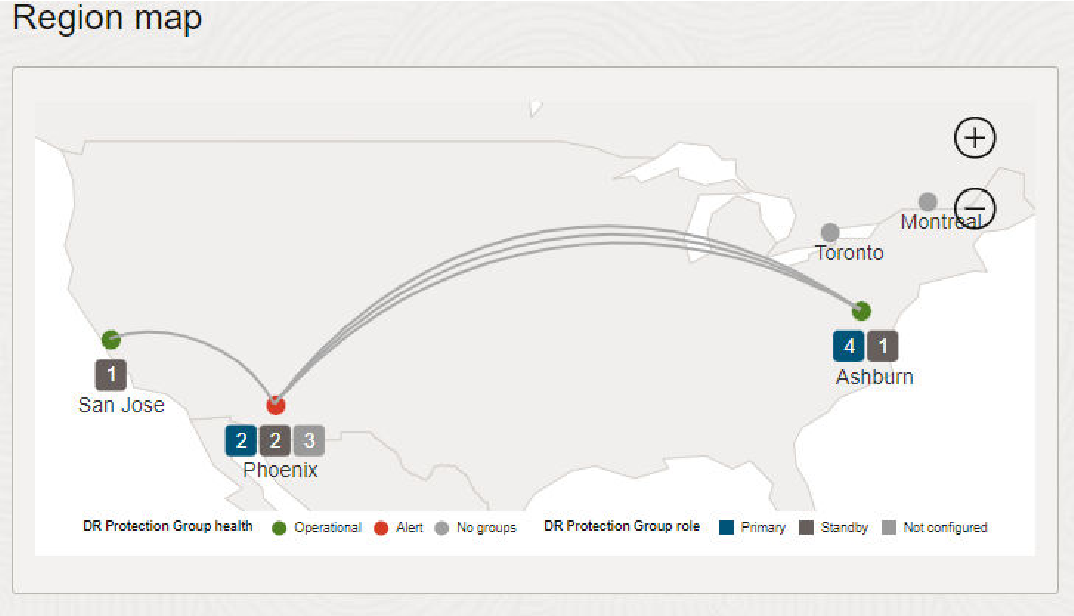

  - The number of DR Protection groups in each region based on the role of the DR Protection Group. DR Protection Group role is shown as follows:
    - Primary
    - Standby
    - Not configured

Create Disaster Recovery Protection Groups

Create a protection group for each region that you want to include in a disaster recovery plan.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR Protection Groups** to navigate to the home page.
3. Change the compartment to the compartment in which you want to create the DR Protection Group.

4. Click **Create DR protection group**.
5. Enter a name for the **DR protection group**.
6. Select the **Object Storage** bucket.

   This Object storage bucket will be used to store the DR Plans Execution logs created for this DR Protection Group.
7. **(Optional)**. From the list, select a role to assign to the **DR protection group**.
   a. If you have assigned a role, then select a peer **DR protection group**.
   b. Select the **peer region** and a peer **DR protection group** to associate with this DR Protection Group.

      DR will automatically assign the peer DR Protection Group the other role in the peer relationship.
   c. If you do not select a role and peer, you can select these after you create the DR Protection Group.
8. **(Optional)**. Add members to the DR Protection Group either at this point or later.
9. Click **Create** to create the DR Protection Group.

Related Topics

- Associate a Disaster Recovery Protection Group
  Learn how to associate a Disaster Recovery (DR) Protection Group with a peer.
- Add Members to a Disaster Recovery Protection Group
  Learn how to add members to a Disaster Recovery (DR) Protection Group.

View Disaster Recovery Protection Groups

The details page of a Disaster Recovery (DR) protection group displays details such as its Plans, Plan executions, Members, and associated Work requests.

To view information about a protection group:

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR protection groups** to navigate to the home page.
3. Change the compartment to the compartment in which you want to view DR Protection Groups.

   The DR Protection Groups in the compartment that you selected are listed.
4. Click one of the DR Protection Groups to navigate to the page for that DR Protection Group.

   This page provides detailed information on the DR Protection Group, including information about its role, its peer DR Protection Group (if applicable), and the log location where plan execution logs are stored.

Associate a Disaster Recovery Protection Group

Learn how to associate a Disaster Recovery (DR) Protection Group with a peer.

> **Note:**
>
> To associate a DR Protection Group with a peer, the former must have a role of **Not configured**. If a DR Protection Group is already associated with a peer, then you must first disassociate it from that peer before you can associate it with another peer.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR protection group** to navigate to the home page.
3. Change the compartment to the one that contains the DR Protection Group to be associated.

   The DR Protection Groups in the compartment that you selected are listed.
4. Click the **DR protection group** that you want to associate to navigate to the page for that DR Protection Group.
5. Click **Associate.**
6. Select a **Role** for the DR Protection Group.
7. Select a **Peer region** in which the peer DR Protection Group resides.
8. Select the **Peer DR protection group**.
9. Click **Associate** to associate the DR Protection Group with the selected peer.

Disassociate Disaster Recovery Protection Groups

Remove the association between a primary protection group and its standby protection group by updating the details of either protection group.

> **Note:**
>
> - If you disassociate a Disaster Recovery (DR) protection group from its peer DR protection group, all DR Plans, DR Plan Executions, and plan execution logs stored in Object Store for that DR Protection Group and its peer DR Protection Group will be deleted. These deletions are irreversible and cannot be undone. This plan deletion is irreversible and cannot be undone. You will have to recreate any required DR Plans.
> - To disassociate a DR Protection Group, it must already be associated with a peer and have a role of either **Primary** or **Standby**.

1. Open the navigation menu, click **Migration & Disaster Recovery**
2. Click **Disaster Recovery** or **DR protection groups** to navigate to the home page.
3. Change the compartment to the one that contains the DR Protection Group to be disassociated.

The DR Protection Groups in the compartment that you selected are listed.

4. Click the **DR protection group** that you want to disassociate to navigate to the page for that DR Protection Group.
5. Click **Disassociate** at the top of the page.
6. Accept the warning that all DR Plans will be deleted.
7. Click **Disassociate** to complete disassociating the DR Protection Group from its peer.

Modify Disaster Recovery Protection Groups

Modify a Disaster Recovery (DR) protection group definition to rename the protection group, add or remove members, or move the protection group to a different compartment.

> **Caution:**
>
> Adding members to a DR Protection Group or deleting members from a DR Protection Group triggers the deletion of all DR Plans for that respective DR Protection Group and its peer DR Protection Group. This plan deletion is irreversible and cannot be undone. You must recreate any required DR Plans.

- Add Members to a Disaster Recovery Protection Group
  Learn how to add members to a Disaster Recovery (DR) Protection Group.
- Delete Members from a Disaster Recovery Protection Group
  Learn how to delete members from a Disaster Recovery (DR) Protection Group.
- Edit Member Properties for a Disaster Recovery Protection Group
  Learn how to edit properties of members already added to a Disaster Recovery (DR) protection group.

Related Topics

- Member Properties Causing Plan Deletion Post Update

Add Members to a Disaster Recovery Protection Group

Learn how to add members to a Disaster Recovery (DR) Protection Group.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR Protection Groups** to navigate to the home page.
3. Change the compartment to the one that contains the DR Protection Group to which you want to add members.

   The DR Protection Groups in the compartment that you selected are listed.
4. Click the **DR Protection Group** to which you want to add members and navigate to the page for that DR Protection Group.
5. Click the **Member** link in the Resources panel to navigate to the Members page for the DR Protection Group.
6. Click **Add Member.**
7. Select the **Resource type** to add.

8. Select the warning that member addition or removal will delete all existing plans in the protection group.

9. Depending on the Resource (Member) type you select, refer to the following sections to add properties for that member.

- Add a Compute Instance to a Disaster Recovery Protection Group
  Learn how to add a compute instance to a Disaster Recovery (DR) Protection Group.
- Add a Volume Group to a Disaster Recovery Protection Group
  Learn how to add a volume group to a Disaster Recovery (DR) Protection Group.
- Add a Load Balancer or a Network Load Balancer to a Disaster Recovery Protection Group
  Learn how to add a load balancer or a network load balancer to a Disaster Recovery (DR) Protection Group.
- Add an Oracle Base Database Service or an Oracle Exadata Database Service to a Disaster Recovery Protection Group
  Learn how to add an Oracle Base Database Service or an Oracle Exadata Database Service to a Disaster Recovery (DR) Protection Group.
- Add a File System to a Disaster Recovery Protection Group
  Learn how to add a file system to a Disaster Recovery (DR) Protection Group.
- Add an Autonomous Database to a Disaster Recovery Protection Group
  Learn how to add an Autonomous Database to a Disaster Recovery (DR) Protection Group.

Add a Compute Instance to a Disaster Recovery Protection Group

Learn how to add a compute instance to a Disaster Recovery (DR) Protection Group.

> **Note:**
>
> Currently, you can only edit the existing members of the **Compute** resource type.

- Add a Non-Moving Instance to a Disaster Recovery Protection Group
  Learn how to add a non-moving Instance to a Disaster Recovery (DR) Protection Group.
- Add a Moving Instance to a Disaster Recovery Protection Group
  Learn how to add a moving instance to a Disaster Recovery (DR) Protection Group.

Related Topics

- Dedicated Virtual Machine Hosts
- Network Security Groups

Add a Non-Moving Instance to a Disaster Recovery Protection Group

Learn how to add a non-moving Instance to a Disaster Recovery (DR) Protection Group.

> **Note:**
>
> Currently, you can only edit the existing members of the **Compute** resource type.

1. From the **Resource Type** menu, Select **Compute**.

2. Accept the warning that indicates that all DR Plans will be deleted.
3. Select the instance to add from the list of instances.
4. Select or deselect **Move instance on switchover or failover** based on your requirement.

Non-Moving Instance

A non-moving instance is not moved from a DR Protection Group to its peer DR Protection Group during DR operations. You use non-moving instances in Active-Passive DR topologies where instances which comprise the application stack are pre-deployed in both regions and application software components. You start or stop these instances during DR operations in order to transition the service from one region to another.

The instances always stay in the region where you deploy them and do not transition from one DR Protection Group to another.

When you add an instance, a non-moving instance is added to the DR Protection Group. No additional member properties (configuration information) are required.

Provide the following configuration information:

In the **Show advanced options**, for a new non moving instance, the following options are available:

- In **Settings**, you can select **Start and stop instance on failover/switchover**.
- In **Block volumes** enter the following information:
  - **Block volume in** *compartment*: Select a compute instance.

    > **Note:**
    >
    > If the compute instance is a member of primary DR Protection Group, then select the block volume which is attached to that compute instance. If the compute instance is a member of the standby DR Protection Group, then select the block volume from the primary region.

  - **(Optional) Volume attachment reference instance in *compartment*:** Volume attachment reference is the peer compute instance used to get the attachment details.

    > **Note:**
    >
    > The value of the reference compute instance should always be from the peer region irrespective of the region (primary or standby).

  - **(Optional) Mount point**: The physical mount point used for mounting and unmounting.
  - Select the **+ Another block volume mapping** to add other block volumes.
- In **File systems** enter the following information:
  - **Export path**: Export path for the file system. Example: </fs-export-path>
  - **Mount point**: Physical mount point for the file system. Example: </mnt/yourmountpoint>
  - **Target in** *compartment*: Select the mount target which is used to unmount on primary DRPG member and is then used to mount on the standby DRPG Member.

  - Select the **+ Another export mapping** to add other file system.
- Click **Add** to add the compute instance to the DR Protection Group.

Related Topics

- Block Volume

Add a Moving Instance to a Disaster Recovery Protection Group

Learn how to add a moving instance to a Disaster Recovery (DR) Protection Group.

> **Note:**
>
> Currently, you can only edit the existing members of the **Compute** resource type.

1. From the **Resource Type** menu, Select **Compute**.
2. Accept the warning that indicates that all DR Plans will be deleted.
3. Select the instance to add from the list of instances.
4. Select or deselect **Move instance on switchover or failover** based on your requirement.

Moving Instance

A moving instance is moved from a DR Protection Group to its peer DR Protection Group during DR operations.

Moving instances are typically used in Pilot Light DR topologies where instances which comprise the application stack are only deployed in the primary region. The instances are moved from the primary DR Protection Group to the standby DR Protection Group.

When you add a moving instance to a DR Protection Group, provide the following configuration information:

- Click **Add VNIC mapping**. For each of the instance's VNICs, provide a VNIC to destination subnet mapping that indicates the subnet in the destination (standby) region to which the VNIC will attach after the instance moves to the destination (standby) region. Provide the following details:
  - **VNIC**
  - **Destination subnet in *your compartment***
  - (Optional) **Destination primary private IP address**: Assign a value for the primary private IP address.

    > **Note:**
    >
    > If you provide the IP address in one of the moving compute instances, then ensure to add the IP address for the other compute instances as well.

> **Note:**
>
> For a moving instance, you can choose to reassign the same private IP address and hostname to an instance that is relocated during disaster recovery.

  - (Optional) **Destination primary private IP hostname label**: Assign a value for the primary private IP hostname label.
  - **Network Security Groups**: You can select the **Destination network security group** in your compartment and also add another network security groups.

- If the VNIC has a public IP address in the source subnet, the selected destination must support public IP addresses. Else, the VNIC mapping is treated as a configuration error. OCI Networking assigns an appropriate public IP address in the selected destination subnet. You cannot select public IP addresses.
- The rules for assigning primary and secondary private IP addresses of a VNIC in the destination region are:
  - If the CIDR block for the VNIC's source subnet matches the CIDR block for the VNIC's destination subnet, and the private IP address is available in the destination subnet, then the Full Stack DR assigns the same private IP address. Else, Full Stack DR assigns another available private IP address.
  - If the CIDR block for the VNIC's source subnet does not match the CIDR block for the VNIC's destination subnet, then Full Stack DR assigns another available private IP address in the destination subnet.

> **Note:**
>
> You must provide VNIC mapping information for all VNICs that you have configured for an instance. If you do not provide VNIC mapping information, you will receive errors when you create DR Plans for disaster recovery.

- The rules for assigning host names are:
  - If the host name assigned to a VNIC is available in the destination subnet, that same host name is assigned.
  - If the host name assigned to a VNIC is not available in the destination subnet, another host name is assigned.

> **Note:**
>
> Assigning host names only applies to the short host name. The long (fully qualified) host name is determined the by the destination subnet and its VCN.

  - **(Optional)** Provide one or more network security groups that you want to assign to the VNIC in the destination region.

In the **Show advanced options**, for a new movable instance, following options are available:

- In **Settings**, you can select **Retain fault domain**. When you check this box, the newly created compute instances are launched in the same fault domains as of the primary region. If you reserve any capacity on the standby region, ensure that the capacity reserved on the standby region is specific to that fault domain as the instance is launched in the same fault domain.

  - Example 1: If the primary region is on `faultdomain-1`, then the standby region must be in `faultdomain-1`. If you do not check this box, then there should be some capacity reservation on the standby region which is not specific to the fault domain. Consequently, the `First available` fault domain is used for capacity reservation.
  - Example 2: If there is some capacity reservation on the standby region for `faultdomain-2` and if you do not retain the fault domain, then the switchover fails during precheck.

- In **Destination**, select any of the following configurations:
  - **(Optional)**. Select **Destination compartment** to provide a compartment in the destination (standby) region to which you want to move the compute instance. If you do not provide a destination compartment, the instance moves to the same compartment in which it resides in the primary region.
  - **(Optional)**. **Default destination** This option is checked by default.
  - **(Optional)**. Select a pre-provisioned **Destination dedicated VM host** in the destination region where you want to launch the instance. If you do not select a dedicated VM host, then the instance will launch using standard OCI provisioning procedures for new instances.

> **Note:**
>
> It is not common to launch instances on dedicated VM hosts, unless those instances have specific hardware or capacity requirements.

  - **(Optional)**. Select a pre-provisioned **Destination capacity reservation** in the destination region where you want to launch the instance. If you do not select a capacity reservation, then the instance will launch using standard OCI provisioning procedures for new instances.
  - **(Optional)**. Select a pre-created **Destination capacity reservation in *compartment*** from the standby region. Use the `Compartment` list in the panel on the left to change the compartment. This ensures that the reserved capacity is used when the instance is created. If you do not select a reservation ID, then the instance will be created using on-demand capacity. Reserve the capacity on the standby region so that when you move the movable instance to the standby region, the pre-reserved capacity is used. See Capacity Reservations for more information.

> **Note:**
>
> The **Destination dedicated VM host** and **Destination capacity reservation** are mutually exclusive, only one can be provided for an instance member. The rules for assigning a destination capacity reservation ID are as follows:
>
> * The capacity reservation ID can be given only for moving instances.
> * The capacity reservation must be pre-created and must be in an active state.
> * The availability domain of the given capacity reservation must match the destination availability domain where you want to move the instance. Generally the availability domain is where you replicate the boot volume.
> * The destination capacity reservation can be included only if the destination dedicated VM host is excluded.
> * The provided capacity reservation must have a capacity configuration entry matching that of the compute instance's shape configuration (`instanceShape`, `memoryInGBs`, `ocpus`)

- In the **File systems** tab, enter the following information:
  - **Export path**: Export path for the file system. Example: `</fs-export-path>`
  - **Mount point**: Physical mount point for the file system. Example: `</mnt/yourmountpoint>`
  - **Unmount target in** *compartment*: Select the unmount target.
  - **Mount target in** *compartment*: Select the mount target.
  - Select the **+ Another export mapping** to add other file system.
- Click **Add** to add the compute instance to the DR Protection Group.

Related Topics

- Dedicated Virtual Machine Hosts
- Network Security Groups

Add a Volume Group to a Disaster Recovery Protection Group

Learn how to add a volume group to a Disaster Recovery (DR) Protection Group.

1. From the Resource Type menu, select **Volume Group**.
2. Select the volume group to add from the list of volume groups.
3. Accept the warning that all DR Plans will be deleted.
4. Click **Add** to add the volume group to the DR Protection Group.

Add a Load Balancer or a Network Load Balancer to a Disaster Recovery Protection Group

Learn how to add a load balancer or a network load balancer to a Disaster Recovery (DR) Protection Group.

Before adding a load balancer or a network load balancer, ensure to perform the following prerequisites:

1. Create a load balancer if you want to add a load balancer. Alternatively, Create a network load balancer if you want to add a network load balancer.
2. Update the load balancer.
3. Create a standby DR Protection Group.
4. Create a primary DR Protection Group.

Perform the following steps:

1. From the Resource Type menu, select **Load balancer**.
2. Select a load balancer or network load balancer to add from the list of load balancers.
3. Select the **Destination load balancer** in *your* compartment.
4. Select the **Source backend set**.
5. Select the **Destination backend set** .
6. Check the **Is backend set for non-moving instance** box if you do not want the load balancer to send any ingress traffic to the backend servers (Compute instances) in this backend set. You must select this checkbox for the backend set if all of its backend servers are non-moving compute instances.
   - For non-moving compute instances, if you check this box, Full Stack DR removes the primary backend set and updates the standby backend set from offline to online. Full Stack DR updates the primary backend set from online to offline. As a prerequisite, you must specify the IP address of the backend server you want to add to your standby backend set.
   - For moving compute instances, Full Stack DR removes all the backend servers from the primary backend set and adds them to the mapped backend set on the standby load balancers, to route the traffic to the standby backend set during DR Drill execution. Do not check this box for moving compute instances. Ensure that your standby backend sets are empty and primary backend sets are not empty. As a prerequisite, ensure that the IP addresses in the primary backend set point to the moving compute instances.

> **Note:**
>
> - You need to create separate backend sets for moving and non-moving compute instances and add them in different backend sets and ensure that they have separate backend sets.
> - The checkbox indicates that the backend servers added in the backend set are for non-moving compute instances.
> - For non-moving compute instances that are a part of a backend set, Full Stack DR updates and disables them to stop sending any ingress traffic and they stay in an offline mode. The backend servers are not removed from the set in this scenario.

7. Click the option to add the **+ Another backend set mapping** to add the mapping of the load balancer or a network load balancer.
8. Accept the warning that all DR Plans will be deleted.
9. Click **Add** to add the load balancer or a network load balancer to the DR Protection Group.

> **Note:**
>
> Ensure to add only those OCI resources such as compute instances to load balancers or network load balancers that are part of the primary DR Protection Group.

Related Topics

- Backend Servers for Load Balancers
- Adding a Load Balancer Backend Server

Add an Oracle Base Database Service or an Oracle Exadata Database Service to a Disaster Recovery Protection Group

Learn how to add an Oracle Base Database Service or an Oracle Exadata Database Service to a Disaster Recovery (DR) Protection Group.

> **Note:**
>
> You can only add databases of type Oracle Base Database Service deployed on bare metal or virtual machines or Oracle Exadata Database Service deployed on dedicated infrastructure as members of a DR Protection Group.

Add an Oracle Base Database Service

1. From the **Resource type** menu, select **Database**.
2. From the **Database type** list, select **Oracle Base Database**.
3. Select the **Database system** which hosts the database to add.
4. Select the **Database home** for the database to add.

5. Select the **Database** to add.
6. Select a **Database password secret** which contains the database password.
7. Accept the warning that all DR Plans will be deleted.
8. Click **Add** to add the database to the DR Protection Group.

Add an Oracle Exadata Database Service

1. From the **Resource type** menu, select **Database**.
2. From the **Database type** list, select **Oracle Exadata on Oracle Public Cloud**.
3. Select the **VM cluster** which hosts the Exadata database to add.
4. Select the **Database home** for the database to add.
5. Select the **Database** to add.
6. Select a **Database password secret** which contains the database password.
7. Accept the warning that all DR Plans will be deleted.
8. Click **Add** to add the database to the DR Protection Group.

> **Note:**
>
> Database passwords stored as Secrets (in the Vault Service) are only retrieved and used when performing database DR operations such as switchover or failover. They are not stored locally anywhere within the Full Stack Disaster Recovery Service. To setup a Vault with a Secret, see Preparing Oracle Databases for Full Stack Disaster Recovery.

Add a File System to a Disaster Recovery Protection Group

Learn how to add a file system to a Disaster Recovery (DR) Protection Group.

Before adding a file system, ensure to perform the following prerequisites:

1. Create a file system.
2. Create a replication.
3. Set up the security rules.
4. Create a standby DR Protection Group.
5. Create a primary DR Protection Group.

Perform the following steps to add a file system to a Disaster Recovery (DR) Protection Group:

1. From the Resource Type menu, select **File system**.
2. Select the file system to add from the list of file systems.
3. Select the **Destination availability domain**.

> **Note:**
>
> Full Stack DR checks whether the replication is created in the selected **Destination availability domain**. If no replication is available, then the plan creation fails. As a prerequisite, ensure to setup replication before selecting the **Destination availability domain** in the standby region.

4. Select the **Source export path**.
5. Select the **Destination mount target in** *your* compartment.

   Specify the export path that should be attached to the mount target in the standby region.
6. Click the option to add the **+ Another export mapping** to add the mapping of the file systems.
7. Accept the warning that all DR Plans will be deleted.
8. Click **Add** to add the file system to the DR Protection Group.

> **Note:**
>
> Do not include file systems mounted to compute instances that are members of different DR Protection Groups. Any compute instance in completely different DR protection groups for other application stacks will lose access to the file systems during a DR operation such as a switchover.

Add an Autonomous Database to a Disaster Recovery Protection Group

Learn how to add an Autonomous Database to a Disaster Recovery (DR) Protection Group.

1. From the **Resource type** menu, select **Autonomous database**.

> **Note:**
>
> You can add only databases of type Oracle Autonomous Database deployed on shared infrastructure as members of type Database to a DR Protection Group. DR does not support databases of type Oracle Autonomous Database deployed on dedicated infrastructure.

2. Select the **Autonomous database** you want to add.
3. Accept the warning that all DR Plans will be deleted.
4. Click **Add** to add the Autonomous database to the DR Protection Group.

Delete Members from a Disaster Recovery Protection Group

Learn how to delete members from a Disaster Recovery (DR) Protection Group.

1. Navigate to the resource page for the DR Protection Group.
2. In the **Resources** panel on the left, click the **Members** link to navigate to the Members section for the DR Protection Group.

Remove a Single Member

1. Click the three-dots Action menu selector on the right of a member and select **Remove**.
2. Accept the warning that all DR Plans will be deleted.
3. Click **Remove** to remove the member from the DR Protection Group.

Remove multiple members

1. Select the members that you want to delete by clicking the box to the left of each member.
2. After all members are selected, click **Remove members** at the top of the table.
3. Accept the warning that all DR Plans will be deleted.
4. Click **Remove** to remove all the selected members from the DR Protection Group.

Edit Member Properties for a Disaster Recovery Protection Group

Learn how to edit properties of members already added to a Disaster Recovery (DR) protection group.

> **Note:**
>
> Unlike adding members to a DR protection group or deleting members from a DR protection group, editing the properties of an existing member does not delete DR plans created for that DR protection group.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery Service** or **DR Protection Groups** to navigate to the home page.
3. Change the compartment to the one which contains the DR Protection Group to which you want to add members.
4. Full Stack DR lists the DR Protection Groups in the compartment that you selected.
5. Click the DR Protection Group for which you want to delete members and navigate to the page for that DR Protection Group.
6. Click the **Members** link in the **Resources** panel to navigate to the **Members** page for the DR Protection Group
7. Click the 3-dot **Action** menu for the member whose properties you want to edit
8. Edit the member properties and click **Save Changes** to save the edited properties

Related Topics

- Add Members to a Disaster Recovery Protection Group
  Learn how to add members to a Disaster Recovery (DR) Protection Group.

Update Disaster Recovery Protection Group Role

Learn how to update the role of a Disaster Recovery (DR) Protection Group.

1. Navigate to the resource page for the DR Protection Group.

2. In the **DR Protection Group Information** pane, identify the current **Role** of the DR Protection Group and click the **Edit** link next to the role.
3. Click **Change to primary** or **Change to standby** to change the role of the DR Protection Group.
4. When you change the role of a DR Protection Group, the role of its peer is also automatically updated. DR plans stored at the primary group move into an **Inactive** state, and DR plans stored at the standby group move into an **Active** state.

> **Note:**
>
> To update the role of a DR Protection Group, you must associate it with another peer DR Protection Group. You cannot update the role of a DR Protection Group with a role of **Not Configured** (not associated with a peer).

Delete a Disaster Recovery Protection Group

Learn how to delete a Disaster Recovery (DR) Protection Group.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR protection groups** to navigate to the home page.
3. Change the compartment to the compartment that contains the DR Protection Group to which you want to add members.

   The DR Protection Groups in the selected compartment are listed.
4. Click the **DR protection group** you want to delete.
5. Select **Delete** from the **More Actions** menu.
6. Click **Delete** to confirm deletion of the DR Protection Group.

   The DR Protection Group moves to a **Deleted** state and is removed after some time. Once deleted, the DR Protection Group becomes unusable.

> **Note:**
>
> To delete a DR Protection Group, you must assign it the role of **Not configured** (not associated with a peer). You must disassociate a DR Protection Group that is associated with a peer from that peer before you delete it.

Rename a Disaster Recovery Protection Group

Learn how to rename a Disaster Recovery (DR) Protection Group.

1. Open the navigation menu, click **Migration & Disaster Recovery**.
2. Click **Disaster Recovery** or **DR protection groups** to navigate to the home page.
3. Change the compartment to the compartment that contains the DR Protection Group to which you want to add members.

   The DR Protection Groups in the selected compartment are listed.
4. Click the **DR protection group** you want to rename.

5. Click **Rename**.
6. Enter a new **Name** for the DR Protection Group.
7. Click **Save** to apply the changes.

4 Manage Disaster Recovery Plans

Full Stack DR automates the generation of Disaster Recovery (DR) plans that are used to perform disaster recovery operations. DR plans can be switchover plan or failover plans.

Overview of Disaster Recovery Plans

A Disaster Recovery (DR) plan is an automated DR workflow (a DR runbook) created by Full Stack DR to perform disaster recovery for all the resources in the primary DR protection group.

The DR plan consists of a sequence of steps that define how all the application stack components in a primary DR protection group in one region or availability domain (AD) are transitioned to its peer standby DR protection group in another region or AD.

These are some of the key features of DR plans:

1. Full Stack DR automatically creates DR plans after performing an intelligent analysis (introspection) of the contents of the primary and standby DR protection groups.
2. A DR plan consists of a sequence of plan groups, and each of these plan groups consists of plan steps.
3. When creating a DR plan, the optimal sequence of plan groups and steps within those groups are automatically determined by Full Stack DR.
4. When a DR plan is executed, plan groups in the DR plan execute sequentially, and plan steps in each group execute in parallel.
5. Plan groups and steps generated by Full Stack DR are called **Built-in**, whereas plan groups and steps added by you are called **User-defined**.

6. DR plans are highly flexible and you can customize them. For more details about customizing a DR plan, refer Modify a Disaster Recovery Plan
   - You can add your own user-defined groups and steps to a DR plan after it is created by Full Stack DR.
   - You can reorder the sequence of **Built-in** and **User-defined** groups within a DR plan.
   - You can customize the execution behavior of **Built-in** and **User-defined** groups and steps.
7. Full Stack DR can create DR plans only for a DR protection group that has a standby role. To create DR plans at a primary DR protection group, you must first transition the primary DR protection group to a standby role by executing a DR plan to perform a DR transition.
8. A pair of associated DR protection groups can have multiple DR plans created for performing DR transitions between the two DR protection groups.
9. DR plans can exist at both protection groups in an associated pair, but only DR plans at the standby DR protection group are in an **Active** state and available for modification or execution. DR plans at the primary DR protection group are held in an **Inactive** state and cannot be modified or executed until that DR protection group assumes a standby role.
10. After you perform a DR transition and the roles of the DR protection groups in the pair are reversed, DR plans at the new standby DR protection group become Active and the DR plans at the new primary DR protection group become Inactive.

A typical sample workflow for creating and executing DR plans works as follows:

1. Create an associated pair of primary and standby DR protection groups and add application stack members and other resources to the DR protection groups. You can name the primary DR protection group as DRPG-IAD, and the standby DR protection as DRPG-PHX.
2. Create one or more DR plans at DRPG-PHX because it is the current standby. Customize these DR plans if required.
3. Execute any one of these DR plans at DRPG-PHX to perform a DR transition of application stack from DRPG-IAD to DRPG-PHX. After the DR plan execution is complete, DRPG-IAD is now the new standby and DRPG-PHX is the new primary.
4. Create another set of DR plans at DRPG-IAD because it now has a standby role. Customize these DR plans if required.
5. Execute any one of these plans at DRPG-IAD to perform a reverse DR transition from DRPG-PHX back to DRPG-IAD.
6. Now that there are DR plans created and stored at both DRPG-IAD and DRPG-PHX. Use the appropriate plans stored at the current standby DR protection group to perform DR transitions back and forth between these two DR protection groups at any time.

Related Topics

- Lifecycle States of Full Stack Disaster Recovery Resources
  Lifecycle states of Full Stack DR resources like DR Protection Groups, DR Plans, and DR Plan Executions.

Types of Disaster Recovery Plans

Full Stack DR provides the ability to create the following different types of Disaster Recovery (DR) plans.

- **Start drill** - A DR drill is often performed to check if the standby stack can be brought up successfully. Start drill plan generates the plan to perform the DR drill without interrupting the production environment. It creates a replica of your production stack in the standby DR protection group. This replica of the production stack is similar to the stack that is brought up in the standby DR protection group when you perform a switchover or failover.
- **Stop drill** - A type of DR plan that stops the DR Drill. This removes the replica of your production stack created earlier by **Start drill**.

> **Note:**
>
> A replica drill stack must be present in the standby DR protection group when you run the **Stop drill** plan.

- **Switchover (planned)** – A type of DR plan that performs a planned transition of services from the primary DR protection group to the standby DR protection group. Switchover plans are used to perform an orderly transition by shutting down the application stack in the primary region and then bringing it up in the standby region. Therefore, a switchover plan requires that application stack components and other required OCI services be available in both regions. Switchover plans are typically used for the purposes of planned site maintenance, software patching, DR testing, and validation.
- **Failover (unplanned)** – A type of DR plan that performs an unplanned transition of services to the standby DR protection group. Failover plans usually perform an immediate transition by bringing up the application stack in the standby region, without attempting to shutdown service in the primary region. Therefore, a failover plan only requires that OCI services be available in the standby region. Failover plans are generally used to perform DR transitions when an outage or disaster affects the primary region.
  When you run a **Start drill** plan:
  - The primary and standby DR protection groups transition to a **DrillInProgress** state.
  - When the DR protection groups are in a **DrillInProgress** state, if you attempt to execute a Switchover, Failover or Start Drill plan (or an associated precheck), you will get an error message indicating that this is a disallowed operation and that you must first execute a **Stop Drill** plan.
  - You cannot add or delete Members to/from DR protection groups when a drill is in progress.
  - Ensure to update the DR protection group pair to an **Active** state by executing a **Stop Drill** plan, before executing a switchover or failover DR plan.
  - The roles of the two DR protection groups do not change when a drill is in progress.
  - When the primary and standby DR protection groups states are **DrillInProgress**, you can execute only the **Stop Drill** plans.
  - The roles of the two DR protection groups remain as **Primary** and **Standby**. The lifecycle states change from **Active** to **Inactive** and lifecycle substate of **DrillInProgress** is set when a drill is started. The lifecycle states revert to **Active** when the drill is stopped.

Create a Disaster Recovery Plan

You can create Disaster Recovery (DR) plans only in a DR protection group that is associated and has a standby role.

1. Navigate to the **Resource** page for the DR protection group.

2. Click the **Plans** link in the **Resources** panel to navigate to the **Plans** section for the DR protection group.
3. Click **Create plan** to open the plan creation dialog.
4. Enter a **Name** for the new DR plan.
5. Select the **Plan type** for the new DR plan.
6. Click **Create** to finish creating the DR plan.

View a Disaster Recovery Plan

You can view a Disaster Recovery (DR) plan in the **Plans** section to check details of an existing plan.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to navigate to the **Plans** section for the DR protection group
3. Click the name of the plan you want to view.

   The plan details with a sequential list of plan groups are displayed.
4. Click the **>** symbol next to each plan group's name to expand the group and see the steps in the plan group.

Modify a Disaster Recovery Plan

You can modify a Disaster Recovery (DR) plan to rename a DR plan, edit plan group attributes, and edit attributes of steps in a plan group.

Full Stack DR creates out-of-the-box DR plans by performing intelligent introspection on the members in the primary and standby DR protection groups. These out-of-the-box DR plans contain groups and steps that are called built-in. Once a built-in DR plan is generated, you can modify the plan to add your own groups and steps called user-defined groups and steps. In addition to adding user-defined groups and steps to a plan, you can also modify the execution attributes of built-in groups and steps.

> **Note:**
>
> You cannot delete a built-in group and steps, but you can disable them.

> **Note:**
>
> When using the `UpdateDrPlan` API to update an existing DR Plan, the payload must contain the entire list of groups and steps in the DR Plan, including any updates you are making to the DR Plan.

- Rename a Disaster Recovery Plan
  You can rename a Disaster Recovery (DR) plan to assign a new name to an existing plan.
- Edit Attributes of a Plan Group
  You can edit attributes of a Disaster Recovery (DR) plan to rename the plan group, enable or disable steps, and change timeouts.

- Edit Attributes of Steps in a Plan Group
  You can edit attributes of steps in a Disaster Recovery (DR) plan group to change the step name, error mode, timeout, and enable or disable the step.
- Add a Pause Group in a Plan Group
  You can add a pause group in the list of plan groups to pause the plan executions.
- Change the Order of Execution of Plan Groups
  Modify the order in which you run Disaster Recovery (DR) plan groups.
- Add User-Defined Plan Groups and Steps
  You can add user-defined groups and steps to a Disaster Recovery (DR) plan after Full Stack DR creates an initial DR plan with built-in groups and steps.
- Remove User-Defined Plan Groups
  You can remove a user-defined plan group in a Disaster Recovery (DR) plan.
- Remove User-Defined Steps From a Plan Group
  You can remove user-defined steps from a plan group in a Disaster Recovery (DR) plan.

Rename a Disaster Recovery Plan

You can rename a Disaster Recovery (DR) plan to assign a new name to an existing plan.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the 3-dot **Action** menu for the DR plan that you want to rename.
4. Select **Rename** from the dropdown menu to open the DR plan rename dialog.
5. Enter a new name for the selected plan.
6. Click **Save** to save the changes.

An alternate method for renaming a DR plan is:

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to rename.
4. Click **Rename** to open the DR plan rename dialog.
5. Type in a new name for the selected plan.
6. Click **Save** to save the changes.

Edit Attributes of a Plan Group

You can edit attributes of a Disaster Recovery (DR) plan to rename the plan group, enable or disable steps, and change timeouts.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to view.

   The plan details with a sequential list of plan groups are displayed.

4. Click the 3-dot **Action** menu for the plan group for which you want to edit the attributes.
5. Select an attribute from the drop-down list to edit. You can edit the following attributes:
   - **Rename** – Change the name of the Plan Group
   - **Enable all steps** – Enable execution of all the steps in the Plan Group. This option is grayed out if all the steps are already enabled.
   - **Disable all steps** – Disable execution of all the steps in the Plan Group. The option is grayed out if all the steps are already disabled.
   - **Change timeouts** – Provide a new timeout value for all the steps in the Plan Group.
   - **Change error modes** – Provide a new error mode for all the steps in the Plan Group.
6. Click **Save** in each dialog, to save the changes.

Edit Attributes of Steps in a Plan Group

You can edit attributes of steps in a Disaster Recovery (DR) plan group to change the step name, error mode, timeout, and enable or disable the step.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to view.

   The plan details with a sequential list of plan groups are displayed.
4. Click the **>** symbol next to each plan group's name to expand the group and show the steps in the group.
5. Click the 3-dot **Action** menu for the step for which you want to change the attributes and select **Edit** to open the Step edit dialog.
6. You can change the following attributes in the Step edit dialog:
   - **Step name** – Change the name of the step.
   - **Error mode** – The error mode for step execution.
     - **Stop on error mode** – Indicates that DR plan execution should stop if step execution fails.
     - **Continue on error mode** – Indicates that DR plan execution should continue even if the step execution fails.
   - **Timeout in seconds** – The total duration allowed for the step execution before the step is considered to have timed out. For built-in steps, the timeout value cannot be less than 600 seconds.
   - **Enable step** – Indicates whether the step is enabled for execution. Unchecking the box causes the DR plan execution to skip this step.
7. Click **Update** to close the step edit dialog and save the changes.

Add a Pause Group in a Plan Group

You can add a pause group in the list of plan groups to pause the plan executions.

1. Navigate to the **Resource** page for the DR protection group.

2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to view.

   The plan details with a sequential list of plan groups are displayed.
4. Click the **Actions** dropdown menu and select **Add pause**.
5. In the **Add pause** page, enter the following details:
   a. Add the **Group name**.
   b. Select **Enable pause** to enable the pause action. By default, this option is selected.
   c. Click **Add after** or **Add before** to add a group after or before a particular group.

      > **Note:**
      >
      > The plan fails if you try to add a step before the **Built In Prechecks** group. Also, you cannot add a pause group after the last group.

   d. Select the **Group** after or before which you need to add a pause group.
   e. Click **Add**

   The pause group is added in the list of plan groups. You can enable, disable, and also edit the other attributes of the pause group.

To resume a pause group, see Resume a Disaster Recovery Plan Execution.

Change the Order of Execution of Plan Groups

Modify the order in which you run Disaster Recovery (DR) plan groups.

By default, plan groups are run in the order in which they are listed in the plan groups section of the switchover plan or failover plan.

> **Note:**
>
> Oracle recommends that you do not reorder the sequence of built-in plan groups. Built-in plan groups are ordered using an optimal sequence and changing this order can cause undesirable side-effects and even failures when the DR plan is executed. Oracle recommends that you restrict reordering to changing the sequence of user-defined groups.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to view.

   The plan details with a sequential list of plan groups are displayed.
4. Click the **Actions** dropdown menu and select **Reorder groups**.
5. Use the up and down arrows to move the plan groups to their desired location.
6. Click **Save changes** to save the new order of the plan groups.

Add User-Defined Plan Groups and Steps

You can add user-defined groups and steps to a Disaster Recovery (DR) plan after Full Stack DR creates an initial DR plan with built-in groups and steps.

> **Note:**
>
> You can run user-defined steps only using Oracle Cloud Agent's **Run Command** feature on instances that are present in either the primary or standby DR protection group. You can also run user-defined steps using Oracle Functions.

> **Note:**
>
> For a failover plan, when you add a function to a step, ensure to select a function from the standby region.

You can add a new user-defined group and a step to a plan, or you can add a new step to an existing user-defined group.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan to which you want to add groups and steps.

   The plan details with a sequential list of plan groups are displayed.
4. Click the **Add group** button to add a new user-defined group with a step, or click the three-dots **Action** menu for an existing user-defined group and select **Add step** to add a step to the user-defined group.
5. Provide a name for the new group (if adding a new group).
6. Click **Add after** or **Add before** to add a group after or before a particular group.

   > **Note:**
   >
   > Ensure to add a group after the **Built In Prechecks** group. The plan fails if you try to add a step before the **Built In Prechecks** group.
7. Select a **Group** from the list of groups.
8. Provide a name for the new step.
9. Select **Enable step** to enable the step for execution. If you deselect the box, DR plan execution will skip this step.
10. Select either of the following error modes for the new step:
    a. **Stop on error** – Indicates that DR plan execution should stop if step execution fails.
    b. **Continue on error** – Indicates that DR plan execution should continue even if the step execution fails.

11. Provide a timeout value for the step. This is the total duration allowed for step execution before the step is considered timed out.
12. Select a region which contains the instance on which this step will execute. This only applies to steps where an object store or local script executes on an instance. This region must be region where the instance is currently present, and not the region where the step will execute.

    The following sections describe the three different choices for configuring the user-defined step to execute user-provided script or function:

    - Run object storage script
    - Run local script
    - Invoke function

- Configure a Step to Execute an Object Storage Script
  You can configure a user-defined step to execute a script that resides in Object Storage. The script type and format must comply with all the script type and format restrictions.
- Configure a Step to Execute a Local Script
  You can configure a user-defined step to execute a script that resides on the instance. The script type and format must comply with all the script type and format restrictions.
- Configure a Step to Execute Oracle Functions
  You can invoke a user-defined step as Oracle Functions. You can use this method if you require a serverless execution mode instead of deploying an instance for executing user-defined scripts or code.

Configure a Step to Execute an Object Storage Script

You can configure a user-defined step to execute a script that resides in Object Storage. The script type and format must comply with all the script type and format restrictions.

For a list of script type and format restrictions, see Limitations and Considerations.
Scripts that exit with a non-zero code are considered to have failed and will result in the user-defined step failing execution.

> **Note:**
>
> A script stored in object storage is executed on the instance using the default `ocarun` user ID. You can not execute object storage scripts using a different user ID. Furthermore, you can not provide any additional arguments for object storage scripts. To execute scripts using a different user ID or to provide additional arguments to the script, see Configure a Step to Execute a Local Script.

1. Select the **Region** in which the instance currently resides.

   > **Note:**
   >
   > When selecting the region for the instance, ensure that the instance is currently located in the selected region. Even if the step runs after the instance relocates to another region, the selected region must match the current region of the instance.

2. Select the **Run object storage script** option.

3. Select the **Target instance**. This is the instance on which you want to execute the script.

   > **Note:**
   >
   > If the **Target instance** is in a private subnet, ensure to set up a NAT Gateway. See Setting Up a NAT Gateway for more information.

4. Select the **Object storage bucket** that contains the **Script** (the object storage object).
5. Select the script that you want to execute.
6. Click **Add Step** to finish adding the user-defined group and step.

Configure a Step to Execute a Local Script

You can configure a user-defined step to execute a script that resides on the instance. The script type and format must comply with all the script type and format restrictions.

For a list of script type and format restrictions, see Limitations and Considerations.

> **Note:**
>
> If you want a user-defined local script step to execute on a movable compute instance after it has moved to the standby region:
>
> 1. Configure the user-defined group and the step for the script to execute on the primary compute instance.
> 2. Edit the DR plan and reorder the groups to move this user-defined group and step to a location in the DR plan that is after the group where the movable compute instance is launched in the standby region.

Scripts that exit with a non-zero code are considered failed and result in the user-defined step failing execution.

1. Select the region in which the instance currently resides.

   > **Note:**
   >
   > When selecting the region for the instance, you must ensure that the instance is currently located in the selected region. Even if the step executes after the instance relocates to another region, the selected region must match the current region of the instance.

2. Select the option for **Run local script**.
3. Select the **Target instance**. This is the instance on which the script resides and will be executed.

> **Note:**
>
> If the **Target instance** is in a private subnet, ensure to set up a NAT Gateway. See Setting Up a NAT Gateway for more information.

4. Provide the **Script parameter** for execution. This should be the full path of the script including all parameters required for execution. For example:

   ```
   /home/opc/scripts/myscript.sh arg1 arg2 arg3
   ```

5. Optionally, provide a **Run as user** to execute the script using a user ID that is different from the default user ID `ocarun`. For example, provide `opc` as the **Run as user** to execute the script as the `opc` user.

   > **Note:**
   >
   > The **Run as user** option is not supported on a Windows instance. However, the **Run as user** option is supported on a Linux instance.

6. Click **Add Step** to finish adding the user-defined group and step.

Configure a Step to Execute Oracle Functions

You can invoke a user-defined step as Oracle Functions. You can use this method if you require a serverless execution mode instead of deploying an instance for executing user-defined scripts or code.

For more details on deploying applications and functions, refer Overview of Functions.

1. Select the region in which Oracle Functions resides.
2. Select the button for **Invoke function**.
3. Select the target application in Oracle Functions.
4. Select the target function which is part of the target application.
5. Optionally, provide a request-body for the function using the **Payload** field. For more information about invoking the functions, refer Invoking Functions.
6. Click **Add Step** to finish adding the user-defined group and step.

Remove User-Defined Plan Groups

You can remove a user-defined plan group in a Disaster Recovery (DR) plan.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan that you want to remove.

   The plan details with a sequential list of plan groups are displayed.
4. Click the three-dots **Action** menu for the plan group that you want to remove and select **Remove**.
5. Click **Remove** to confirm removal of the plan group from the plan.

Remove User-Defined Steps From a Plan Group

You can remove user-defined steps from a plan group in a Disaster Recovery (DR) plan.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the name of the plan from which you want to remove user-defined steps.

   The plan details with a sequential list of plan groups are displayed.
4. Click on the **>** symbol next to each plan group's name to expand the group and display the steps in the group.
5. Click the three-dots **Action** menu for the step that you want to remove and select **Remove**.
6. Click the **Remove** button to confirm removal of the step from the DR plan group.

Delete a Disaster Recovery Plan

Learn how to delete a Disaster Recovery (DR) Plan.

1. Navigate to the **Resource** page for the DR protection group.
2. Click the **Plans** link in the **Resources** panel to list the DR plans for the DR protection group.
3. Click the 3-dot **Action** menu for the DR plan that you want to delete.
4. Click **Delete** to confirm deletion of the DR plan.

   The DR plan moves to a **Deleted** state and is removed after some time. Once deleted, the DR plan becomes unusable.

5 Manage Disaster Recovery Plan Executions

A Disaster Recovery (DR) plan execution is created when any DR plan is executed or a precheck for the DR plan is executed. Disaster recovery operations include switchover and failover.

- Overview of Disaster Recovery Plan Executions
  A Disaster Recovery (DR) plan execution contains all the groups and steps that are contained in the DR plan, along with their execution attributes, such as Error Mode, Timeout, and Enabled/Disabled.
- Create Disaster Recovery Plan Executions
  You can manually create a Disaster Recovery (DR) plan execution to start executing a DR plan.
- Prechecks for Disaster Recovery Plan Executions
  Run Precheck performs a comprehensive validation of all the steps in a Disaster Recovery (DR) plan and the members associated with the steps.
- Monitor Disaster Recovery Plan Executions
  Monitor the Disaster Recovery (DR) plan executions to check progress of the plan execution and view logs to check for errors.
- Pause a Disaster Recovery Plan Execution
  You can pause only the Disaster Recovery (DR) plan executions that are in progress. You cannot pause prechecks.
- Resume a Disaster Recovery Plan Execution
  You can resume only the Disaster Recovery (DR) plan executions that are paused. You can not pause or resume prechecks.
- Cancel a Disaster Recovery Plan Execution
  You can cancel a Disaster Recovery (DR) plan execution to stop executing the plan. You cannot resume or retry a canceled DR plan execution.
- Retry a Group in Disaster Recovery Plan Execution
  Retry a group execution is possible when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its steps fails to execute.
- Retry a Step in Disaster Recovery Plan Execution
  Retry a step is possible when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its steps fails to execute.
- Skip a Step in Disaster Recovery Plan Execution
  You can skip a step only when a Disaster Recovery (DR) plan execution is in a **Failed** state. A DR plan execution enters a **Failed** state when one of its steps fails to execute.
- Skip a Group in Disaster Recovery Plan Execution
  You can skip a group only when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its groups or steps fails to execute.

- Delete Disaster Recovery Plan Executions
  You can delete a Disaster Recovery (DR) plan execution that is in a terminal state, which is either **Succeeded** or **Canceled.**

Overview of Disaster Recovery Plan Executions

A Disaster Recovery (DR) plan execution contains all the groups and steps that are contained in the DR plan, along with their execution attributes, such as Error Mode, Timeout, and Enabled/Disabled.

A DR plan execution sequentially executes all the groups in a DR plan in the order in which they are arranged. While the groups execute sequentially, all the steps within each group are executed in parallel. The execution attributes for each group and step are applied during the execution of that group or step.

A precheck performs a comprehensive validation of all the steps in a DR plan and the members associated with those steps. A precheck is a completely passive check and does not alter any part of the DR topology or configuration. The goal of the precheck is to ensure that the DR plan as it is presently constructed does not have any inconsistencies or misconfigured elements which can cause the DR plan execution to fail.

A precheck can be executed independently by itself, known as the stand-alone precheck, or as an immediate precursor to the DR plan execution, known as the built-in precheck. Running stand-alone precheck on a regular basis is highly recommended as it ensures DR readiness. A successful precheck maximizes the probability that the DR plan execution will succeed when launched. Running a built-in precheck in conjunction with a plan execution allows you to validate the plan elements before it begins execution.

> **Note:**
>
> Only one precheck or DR plan execution can be in progress at a time for any given pair of DR protection groups. When a precheck or a DR plan execution is in progress, the DR protection group enters in **Updating** state and cannot be modified until the operation has completed.

Related Topics

- Lifecycle States of Full Stack Disaster Recovery Resources
  Lifecycle states of Full Stack DR resources like DR Protection Groups, DR Plans, and DR Plan Executions.

Create Disaster Recovery Plan Executions

You can manually create a Disaster Recovery (DR) plan execution to start executing a DR plan.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plans** link in the **Resources** panel to navigate to the **Plans** page for the DR protection group.
3. Click the **Execute DR plan** button to launch the **Execute DR plan** dialog.
4. Select **Enable prechecks** to enable the execution of prechecks before you execute the actual switchover or failover plan.

The **Ignore warnings** selection box is reserved for future use.

5. Optionally, enter a name for the DR plan execution.
6. Click **Execute DR plan** to start plan execution.

> **Note:**
>
> When new resources are created during the switchover or failover (DrPlanExecution) operations, the new resources will retain the name of the original resources. However, the new resources will have an additional Free-form Tag added to that name.

Prechecks for Disaster Recovery Plan Executions

Run Precheck performs a comprehensive validation of all the steps in a Disaster Recovery (DR) plan and the members associated with the steps.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plans** link in the **Resources** panel to navigate to the plans page for the DR protection group.
3. Click **Run Prechecks** to launch the **Run Prechecks** dialog.

   The **Ignore Warnings** selection box is reserved for future use.
4. Optionally, provide a name for the precheck execution.
5. Click **Run Prechecks** to start running the precheck.

Monitor Disaster Recovery Plan Executions

Monitor the Disaster Recovery (DR) plan executions to check progress of the plan execution and view logs to check for errors.

- Monitor Progress of Disaster Recovery Plan Executions
  You can monitor a Disaster Recovery (DR) plan execution to see progress of a precheck, plan status, or progress of a DR plan execution.
- View Logs for Disaster Recovery Plan Executions
  Download and view the logs for a Disaster Recovery (DR) plan execution to check execution status and errors.

Monitor Progress of Disaster Recovery Plan Executions

You can monitor a Disaster Recovery (DR) plan execution to see progress of a precheck, plan status, or progress of a DR plan execution.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The DR plan execution or precheck, which is currently executing, displays at the top of the list and has a state of **In Progress**.

4. Click the DR plan execution or precheck that is currently executing to go to the page for the DR plan execution or precheck.

   A list of all plan groups in the DR plan execution or precheck is displayed.

5. The displayed plan groups can be in one of four states:
   - **Queued** – The plan group has not yet begun execution.
   - **In Progress** – The plan group is currently executing.
   - **Succeeded** – The plan group has finished execution and all the steps succeeded.
   - **Failed** – The plan group has finished execution but one or more steps failed.

6. Click the **>** symbol next to any plan group to expand the group and display the details for the steps in the group.

View Logs for Disaster Recovery Plan Executions

Download and view the logs for a Disaster Recovery (DR) plan execution to check execution status and errors.

1. Navigate to the **Resource** page for the DR plan execution as described in Monitor Progress of Disaster Recovery Plan Executions .

   The list of all plan groups in the DR plan execution or precheck is displayed.

2. Click the **>** symbol next to any plan group to expand the group and display the details for the steps in that group.

3. Click the three-dots **Action** menu to the right of any step and select **View log** from the menu.

   The detailed execution log for that step is displayed.

4. Click the **Download log** button to download and save the log file locally.

> **Note:**
>
> You can download the logs without viewing them by clicking the three-dots Action menu for the step, and then selecting **Download log**.

Only 1 KB of step log data is displayed. If the step log size exceeds this limit, only the last 1 KB of step log content is displayed. To view the entire log, download the log to a local device.

Pause a Disaster Recovery Plan Execution

You can pause only the Disaster Recovery (DR) plan executions that are in progress. You cannot pause prechecks.

1. Navigate to the **Resource** page for the standby DR protection group.

2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.

3. The DR plan execution or precheck that is currently executing is displayed at the top of the list and has a state of **In Progress**.

4. Click the DR plan execution or precheck that is currently executing to go to the page for the DR plan execution or precheck.
5. Click the **Pause** button to pause execution of the DR plan execution.

   The DR plan execution moves to a paused state after the currently executing plan group has finished execution. The DR protection group remains in an updating state while the DR plan execution is paused.

Resume a Disaster Recovery Plan Execution

You can resume only the Disaster Recovery (DR) plan executions that are paused. You can not pause or resume prechecks.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The DR plan execution that is currently paused is displayed at the top of the list and has a state of **Paused**.
4. Click the DR plan execution that is currently paused to go to the page for the DR plan execution.
5. Click the **Resume** button to resume execution of the DR plan execution.

   The DR plan execution moves to an In Progress state and the next queued plan group resumes execution. The DR protection group remains in an Updating state after the DR plan execution has resumed.

Cancel a Disaster Recovery Plan Execution

You can cancel a Disaster Recovery (DR) plan execution to stop executing the plan. You cannot resume or retry a canceled DR plan execution.

> **Note:**
>
> Canceling in-progress plan executions can leave the components in the application stack in an inconsistent state. You may need to manually recover individual components in the application stack after you cancel a plan execution.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The DR plan execution that is currently executing or paused is displayed at the top of the list.
4. Click the DR plan execution or precheck that is currently executing or paused to go to the page for the DR plan execution or precheck.
5. Click the **Cancel** button to cancel execution of the DR plan execution.

The DR plan execution moves to a **Canceled** state after the currently executing plan group finishes execution. The DR protection group moves to a **Needs Attention** state after the DR plan execution is canceled.

A DR protection group in a **Needs Attention** state indicates that user intervention is required to resolve any underlying inconsistencies in the DR topology that may exist because DR plan execution was canceled. The DR Protection Group can be moved back to an **Active** state using the following steps:

1. Navigate to the Resource page for the standby DR protection group.
2. Click the **More Actions** menu and select **Reset State** to reset the DR Protection Group's state back to **Active**.

> **Note:**
>
> This state reset must be performed separately for each DR Protection Group in a pair of associated DR Protection Groups.

Retry a Group in Disaster Recovery Plan Execution

Retry a group execution is possible when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its steps fails to execute.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The most recent DR plan execution or precheck, which has failed, is displayed at the top of the list and has a state of **Failed**.
4. Click the failed DR plan execution or precheck to go to the page for the DR plan execution or precheck.

   A list of all plan groups in the DR plan execution or precheck is displayed.
5. Identify the failed plan group, which has a state of **Failed**.
6. Click the 3-dot **Action** menu or the failed plan group and select **Retry** to retry execution of all the failed steps in the plan group.

   The plan execution resumes and all the failed steps in that group are retried.

Retry a Step in Disaster Recovery Plan Execution

Retry a step is possible when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its steps fails to execute.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.

3. The most recent DR plan execution or precheck, which has failed, is displayed at the top of the list and has a state of **Failed**.
4. Click the failed DR plan execution or precheck to go to the page for the DR plan execution or precheck.

   A list of all plan groups in the DR plan execution or precheck is displayed.
5. Identify the failed plan group, which has a state of **Failed**.
6. Click the **>** symbol next to any plan group to expand the group and display the details for the steps in the group.
7. Identify the failed step, which has a state of **Failed**.
8. Click the 3-dot **Action** menu to the right of the failed step and select **Retry** to retry execution of the failed step.

   The plan execution resumes, starting with a retry of the failed step.

Skip a Step in Disaster Recovery Plan Execution

You can skip a step only when a Disaster Recovery (DR) plan execution is in a **Failed** state. A DR plan execution enters a **Failed** state when one of its steps fails to execute.

> **Note:**
>
> Failed steps in a precheck cannot be skipped because a failed precheck step does not stop execution of the precheck. To retry a failed step in a precheck, the entire precheck must be re-run.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The most recent DR plan execution or precheck, which has failed, is displayed at the top of the list and has a state of **Failed**.
4. Click the failed DR plan execution or precheck to go to the page for the DR plan execution or precheck.

   A list of all plan groups in the DR plan execution or precheck is displayed.
5. Identify the failed plan group, which has a state of **Failed**.
6. Click the **>** symbol next to any plan group to expand the group and display the details for the steps in the group.
7. Identify the failed step, which has a state of **Failed**.
8. Click the 3-dot **Action** menu to the right of the failed step and select **Skip** to skip execution of the failed step.

   The plan execution resumes, skipping the failed step.

Skip a Group in Disaster Recovery Plan Execution

You can skip a group only when a Disaster Recovery (DR) plan execution or precheck is in a **Failed** state. A DR plan execution or precheck enters a **Failed** state when one of its groups or steps fails to execute.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. The most recent DR plan execution or precheck, which has failed, is displayed at the top of the list and has a state of **Failed**.
4. Click the failed DR plan execution or precheck to go to the page for the DR plan execution or precheck.

   A list of all plan groups in the DR plan execution or precheck is displayed.
5. Identify the failed plan group, which has a state of **Failed**.
6. Identify the failed plan group, which has a state of **Failed**.
7. Click the 3-dot **Action** menu for the failed group and select **Skip** to skip execution of all the steps in the failed group.

   The plan execution resumes with the group following the failed group.

Delete Disaster Recovery Plan Executions

You can delete a Disaster Recovery (DR) plan execution that is in a terminal state, which is either **Succeeded** or **Canceled.**

To force a failed DR plan execution to move into a terminal state, see Cancel a Disaster Recovery Plan Execution.

1. Navigate to the **Resource** page for the standby DR protection group.
2. Click the **Plan Executions** link in the **Resources** panel to navigate to the **Plan Executions** page for the DR protection group.

   A list of all DR plan executions and prechecks for the protection group is displayed.
3. Click the 3-dot menu to the right of the DR plan execution or precheck that you want to delete.
4. Select **Delete** to delete the DR plan execution or precheck.
5. Click **Delete** to confirm deletion of the DR plan execution or precheck and all the logs associated with it.

   The DR plan execution or precheck moves to a **Deleted** state and is removed after some time. Once deleted, the DR plan execution or precheck becomes unusable.

> **Note:**
>
> You can also delete a DR Plan Execution by navigating to the page for that execution, clicking the **More Actions** menu, and selecting **Delete**.

6
Policies for Full Stack Disaster Recovery

This chapter lists the Identity and Access Management (IAM) policies you must configure for Full Stack DR.

- About IAM Policies for Disaster Recovery
  You can define Identity and Access Management (IAM) policies for Disaster Recovery (DR) either for an individual resource type or for the entire family.
- Resource Types and Permissions
  Lists different permissions available for defining Identity and Access Management (IAM) policies for each resource type.
- Permissions for REST API Operations
  Lists the REST API operations available in Full Stack Disaster Recovery (Full Stack DR) and the permissions required for each Disaster Recovery (DR) operation.
- Details of Verb and Resource-Type Combinations
  The level of access is cumulative as you go from `inspect`, `read`, `use`, `manage`.
- Examples of IAM Policy Definitions for Disaster Recovery
  You can create policy statements to allow a group of users to administer Disaster Recovery (DR) operations, create DR configurations, and execute prechecks.

About IAM Policies for Disaster Recovery

You can define Identity and Access Management (IAM) policies for Disaster Recovery (DR) either for an individual resource type or for the entire family.

Refer to the How Policies Work guide to know more about general concepts of the IAM policies before defining the IAM policies for DR.

You can define IAM policies in either of the following two ways:

- Define policies for an individual resource type
- Define policies for the entire family

The following table lists the family and resource types available for defining IAM policies:

Table 6-1 Family and Resource Types

Family Type	Individual Resource Types
`disaster-recovery-family`	• `disaster-recovery-protection-groups` • `disaster-recovery-plans` • `disaster-recovery-plan-prechecks` • `disaster-recovery-plan-executions` • `disaster-recovery-workrequests`

Resource Types and Permissions

Lists different permissions available for defining Identity and Access Management (IAM) policies for each resource type.

Table 6-2 Permissions for Resource Types

Resource Type	Permissions
`disaster-recovery-protection-groups`	• `DISASTER_RECOVERY_PROTECTION_GROUP_INSPECT` • `DISASTER_RECOVERY_PROTECTION_GROUP_READ` • `DISASTER_RECOVERY_PROTECTION_GROUP_CREATE` • `DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE` • `DISASTER_RECOVERY_PROTECTION_GROUP_DELETE`
`disaster-recovery-plans`	• `DISASTER_RECOVERY_PLAN_INSPECT` • `DISASTER_RECOVERY_PLAN_READ` • `DISASTER_RECOVERY_PLAN_CREATE` • `DISASTER_RECOVERY_PLAN_UPDATE` • `DISASTER_RECOVERY_PLAN_DELETE`
`disaster-recovery-plan-prechecks`	• `DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT` • `DISASTER_RECOVERY_PLAN_EXECUTION_READ` • `DISASTER_RECOVERY_PLAN_PRECHECK_CREATE` • `DISASTER_RECOVERY_PLAN_PRECHECK_UPDATE` • `DISASTER_RECOVERY_PLAN_PRECHECK_DELETE`
`disaster-recovery-plan-executions`	• `DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT` • `DISASTER_RECOVERY_PLAN_EXECUTION_READ` • `DISASTER_RECOVERY_PLAN_EXECUTION_CREATE` • `DISASTER_RECOVERY_PLAN_EXECUTION_UPDATE` • `DISASTER_RECOVERY_PLAN_EXECUTION_DELETE`
`disaster-recovery-workrequests`	• `DISASTER_RECOVERY_WORKREQUEST_INSPECT` • `DISASTER_RECOVERY_WORKREQUEST_READ` • `DISASTER_RECOVERY_WORKREQUEST_DELETE`

Table 6-2 (Cont.) Permissions for Resource Types

Resource Type	Permissions
disaster-recovery-family	• DISASTER_RECOVERY_PROTECTION_GROUP_INSPECT • DISASTER_RECOVERY_PROTECTION_GROUP_READ • DISASTER_RECOVERY_PROTECTION_GROUP_CREATE • DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE • DISASTER_RECOVERY_PROTECTION_GROUP_DELETE • DISASTER_RECOVERY_PLAN_INSPECT • DISASTER_RECOVERY_PLAN_READ • DISASTER_RECOVERY_PLAN_CREATE • DISASTER_RECOVERY_PLAN_UPDATE • DISASTER_RECOVERY_PLAN_DELETE • DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT • DISASTER_RECOVERY_PLAN_EXECUTION_READ • DISASTER_RECOVERY_PLAN_PRECHECK_CREATE • DISASTER_RECOVERY_PLAN_PRECHECK_UPDATE • DISASTER_RECOVERY_PLAN_PRECHECK_DELETE • DISASTER_RECOVERY_PLAN_EXECUTION_CREATE • DISASTER_RECOVERY_PLAN_EXECUTION_UPDATE • DISASTER_RECOVERY_PLAN_EXECUTION_DELETE • DISASTER_RECOVERY_WORKREQUEST_INSPECT • DISASTER_RECOVERY_WORKREQUEST_READ • DISASTER_RECOVERY_WORKREQUEST_DELETE

Permissions for REST API Operations

Lists the REST API operations available in Full Stack Disaster Recovery (Full Stack DR) and the permissions required for each Disaster Recovery (DR) operation.

Refer to Permissions in Oracle Cloud Infrastructure Documentation for more information about permissions.

Table 6-3 Permissions Required for REST API Operations

REST API Operation	Permissions (| = or, & = and)	Description
CancelDrPlanExecution	DISASTER_RECOVERY_PLAN_EXECUTION_CREATE	Cancel the DR Plan Execution identified by drPlanExecutionId.
AssociateDrProtectionGroup	DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE	Create an association between the DR Protection Group identified by drProtectionGroupId and another DR Protection Group in a different region.
CancelWorkRequest	DISASTER_RECOVERY_WORKREQUEST_DELETE	Cancel the work request identified by workRequestId.
CreateDrPlan	DISASTER_RECOVERY_PLAN_CREATE	Creates a new DR Plan of the specified DR Plan type.

Table 6-3 (Cont.) Permissions Required for REST API Operations

REST API Operation	Permissions (| = or, & = and)	Description
`CreateDrPlanExecution`	`DISASTER_RECOVERY_PLAN_EXECUTION_CREATE | DISASTER_RECOVERY_PLAN_PRECHECK_CREATE`	Execute a DR Plan for a DR Protection Group.
`CreateDrProtectionGroup`	`DISASTER_RECOVERY_PROTECTION_GROUP_CREATE`	Create a new DR protection group.
`ChangeDrProtectionGroupCompartment`	`DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE`	Move the DR Protection Group identified by `drProtectionGroupId` to a different compartment.
`DeleteDrPlan`	`DISASTER_RECOVERY_PLAN_DELETE`	Delete a DR plan identified by `drPlanId`.
`DeleteDrPlanExecution`	`DISASTER_RECOVERY_PLAN_EXECUTION_DELETE | DISASTER_RECOVERY_PLAN_PRECHECK_DELETE`	Delete the DR Plan Execution identified by `drPlanExecutionId`.
`DeleteDrProtectionGroup`	`DISASTER_RECOVERY_PROTECTION_GROUP_DELETE`	Delete the DR Protection Group identified by `drProtectionGroupId`.
`DisassociateDrProtectionGroup`	`DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE`	Delete the association between the DR Protection Group identified by `drProtectionGroupId`. and its peer DR Protection Group.
`GetDrPlan`	`DISASTER_RECOVERY_PLAN_READ`	Get details for the DR Plan identified by `drPlanId`.
`GetDrPlanExecution`	`DISASTER_RECOVERY_PLAN_EXECUTION_READ`	Get details of a DR plan execution identified by `drPlanExecutionId`.
`GetDrProtectionGroup`	`DISASTER_RECOVERY_PROTECTION_GROUP_READ`	Get details of a DR protection group identified by `drProtectionGroupId`.
`GetWorkRequest`	`DISASTER_RECOVERY_WORKREQUEST_READ`	Get the status of the work request identified by `workRequestId`.
`IgnoreDrPlanExecution`	`DISASTER_RECOVERY_PLAN_EXECUTION_CREATE`	Ignore failed group or step in DR Plan Execution identified by `drPlanExecutionId` and resume execution.
`ListDrPlanExecutions`	`DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT`	Get a summary list of all DR Plan Executions for a DR Protection Group.
`ListDrPlans`	`DISASTER_RECOVERY_PLAN_INSPECT`	Get a summary list of all DR Plans for a DR Protection Group.
`ListDrProtectionGroups`	`DISASTER_RECOVERY_PROTECTION_GROUP_INSPECT`	Get a summary list of all DR Protection Groups in a compartment.
`ListWorkRequestErrors`	`DISASTER_RECOVERY_WORKREQUEST_INSPECT`	Lists work request errors for the work request identified by `workRequestId`.

Table 6-3 (Cont.) Permissions Required for REST API Operations

REST API Operation	Permissions (| = or, & = and)	Description
ListWorkRequestLogs	DISASTER_RECOVERY_WORKREQUEST_INSPECT	Returns a (paginated) list of logs for the work request identified by workRequestId.
ListWorkRequests	DISASTER_RECOVERY_WORKREQUEST_INSPECT	List work requests in a compartment.
PauseDrPlanExecution	DISASTER_RECOVERY_PLAN_EXECUTION_CREATE	Pause the DR plan execution identified by drPlanExecutionId.
ResumeDrPlanExecution	DISASTER_RECOVERY_PLAN_EXECUTION_CREATE	Resume the DR plan execution identified by drPlanExecutionId.
RetryDrPlanExecution	DISASTER_RECOVERY_PLAN_EXECUTION_CREATE	Retry failed group or step in DR Plan Execution identified by drPlanExecutionId and resume execution.
UpdateDrPlan	DISASTER_RECOVERY_PLAN_UPDATE	Update the DR Plan identified by drPlanId.
UpdateDrPlanExecution	DISASTER_RECOVERY_PLAN_EXECUTION_UPDATE | DISASTER_RECOVERY_PLAN_PRECHECK_UPDATE	Update the DR Plan Execution identified by drPlanExecutionId.
UpdateDrProtectionGroup	DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE	Update the DR Protection Group identified by drProtectionGroupId.
UpdateDrProtectionGroupRole	DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE	Update the role of the DR Protection Group identified by drProtectionGroupId.

Details of Verb and Resource-Type Combinations

The level of access is cumulative as you go from inspect, read, use, manage.

A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas **no extra** indicates no incremental access. For example, the read verb for the disaster-recovery-protection-groups resource-type includes the same permissions and API operations as the inspect verb. The DISASTER_RECOVERY_PROTECTION_GROUP_READ provides permissions and enables the GetDrProtectionGroup API operation. The use verb covers UpdateDrProtectionGroup API operations with additional permissions required compared to read. Lastly, manage requires more permissions compared to use and enables additional operations.

Table 6-4 Disaster Recovery (DR) Protection Group Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
inspect	DISASTER_RECOVERY_PROTECTION_GROUP_INSPECT	ListDrProtectionGroups
read	inspect +DISASTER_RECOVERY_PROTECTION_GROUP_READ	GetDrProtectionGroup

Table 6-4 (Cont.) Disaster Recovery (DR) Protection Group Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
use	read +DISASTER_RECOVERY_PROTECTION_GROUP_UPDATE	UpdateDrProtectionGroup
inspect	use +DISASTER_RECOVERY_PROTECTION_GROUP_CREATE DISASTER_RECOVERY_PROTECTION_GROUP_DELETE	• CreateDrProtectionGroup • DeleteDrProtectionGroup

Table 6-5 Disaster Recovery Plan Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
inspect	DISASTER_RECOVERY_WORKREQUEST_INSPECT	ListDrPlans
read	inspect +DISASTER_RECOVERY_PLAN_READ	GetDrPlan
use	read +DISASTER_RECOVERY_PLAN_UPDATE	UpdateDrPlan
inspect	use +DISASTER_RECOVERY_PLAN_CREATE DISASTER_RECOVERY_PLAN_DELETE	• CreateDrPlan • DeleteDrPlan

Table 6-6 Disaster Recovery Plan Execution Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
inspect	DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT	ListDrPlanExecutions
read	inspect +DISASTER_RECOVERY_PLAN_EXECUTION_READ	GetDrPlanExecution
use	read +DISASTER_RECOVERY_PLAN_EXECUTION_UPDATE	UpdateDrPlanExecution (for DR plan executions)
inspect	use +DISASTER_RECOVERY_PLAN_EXECUTION_CREATE	• CreateDrPlanExecution (for DR plan executions) • PauseDrPlanExecution (for DR plan executions) • ResumeDrPlanExecution (for DR plan executions) • RetryDrPlanExecution (for DR plan executions)
manage	use +DISASTER_RECOVERY_PLAN_EXECUTION_DELETE	DeleteDrPlanExecution (for DR plan executions)

Table 6-7 Precheck Execution Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
inspect	DISASTER_RECOVERY_PLAN_EXECUTION_INSPECT	ListDrPlanExecutions
read	inspect +DISASTER_RECOVERY_PLAN_EXECUTION_READ	GetDrPlanExecution
use	read +DISASTER_RECOVERY_PLAN_PRECHECK_UPDATE	UpdateDrPlanExecution (for precheck executions)

Table 6-7 (Cont.) Precheck Execution Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
`manage`	`use` `+DISASTER_RECOVERY_PLAN_PRECHECK_CREATE`	`CreateDrPlanExecution` (for precheck executions)
`manage`	`use` `+DISASTER_RECOVERY_PLAN_PRECHECK_DELETE`	`DeleteDrPlanExecution` (for precheck executions)

Table 6-8 Work Request Verbs and Permissions

Access Level	Permissions (| = or, & = and)	REST APIs Covered
`inspect`	`DISASTER_RECOVERY_WORK_REQUEST_INSPECT`	• `ListWorkRequests` • `ListWorkRequestLogs` • `ListWorkRequestErrors`
`read`	`inspect` `+DISASTER_RECOVERY_WORKREQUEST_READ`	`GetWorkRequest`
`manage`	`read` `+DISASTER_RECOVERY_WORKREQUEST_DELETE`	`CancelWorkRequest`

Examples of IAM Policy Definitions for Disaster Recovery

You can create policy statements to allow a group of users to administer Disaster Recovery (DR) operations, create DR configurations, and execute prechecks.

The following example allows a group of users to administer all aspects of DR operations in the entire tenancy.

```
Allow group DRUberAdmins to manage disaster-recovery-family in tenancy
```

These policy statements allow the group `DRUberAdmins` to be superusers for all disaster recovery operations.

The following example allows a group of users to create DR configurations and execute prechecks.

```
Allow group DRMonitors to manage disaster-recovery-protection-groups in
compartment ApplicationERP
Allow group DRMonitors to manage disaster-recovery-plans in compartment
ApplicationERP
Allow group DRMonitors to manage disaster-recovery-prechecks in compartment
ApplicationERP
```

These policy statements allow the group `DRMonitors` to create DR configurations and plans, and also execute prechecks but not actually execute DR operations. This ability is limited to just the `ApplicationERP` compartment.

The following example allows a group of users to create DR configurations in a specific compartment.

```
Allow group DRConfig to manage disaster-recovery-protection-groups in
compartment ApplicationERP
Allow group DRConfig to manage disaster-recovery-plans in compartment
ApplicationERP
```

These policy statements allow the group `DRConfig` to create DR configurations and plans only but not execute any DR operations. This ability is limited to one compartment.

7 Policies for Other Services Managed by Full Stack Disaster Recovery

Full Stack DR service implements DR workflows by managing other OCI resources that are part of the application stack.

To enable DR service to manage these OCI resources, you must configure policy-based access to allow access to these resources. Refer to How Policies Work to learn about general concepts of the IAM policies before defining the IAM policies for DR.

- Policies for Compute Service
  Shows how to allow Disaster Recovery (DR) to manage compute instances that are part of the application stack.
- Policies for Oracle Cloud Agent
  Shows how to allow Disaster Recovery (DR) to execute scripts using Oracle Cloud Agent on compute instances that are part of the application stack.
- Policies for Block Volume Service
  Shows how to allow Disaster Recovery (DR) to manage block storage volumes and volume groups that are part of the application stack.
- Policies for Networking Service
  Shows how to allow Disaster Recovery (DR) to manage networking components for compute instances that are part of the application stack.
- Policies for Functions Service
  Shows how to allow Disaster Recovery (DR) to invoke Oracle Functions as part of a user-defined step in a DR plan.
- Policies for Oracle Exadata Database Service and Oracle Base Database Service
  Shows how to allow Disaster Recovery (DR) to manage DR for Oracle Exadata Database Service and Oracle Base Database Service databases that are part of the application stack.
- Policies for Oracle Autonomous Database Service
  Shows how to allow Disaster Recovery (DR) to manage Oracle Autonomous Database Service databases that are part of the application stack.
- Policies for Object Storage Service
  Shows how to allow Disaster Recovery (DR) to manage Object Storage buckets and objects. This access is required to write logs to Object Storage during DR plan executions.
- Policies for Tags Service
  Shows how to allow Disaster Recovery (DR) to use Tag namespaces. This access is required when launching compute instances as part of DR plan executions.
- Policies for Vault Service
  Shows how to allow Disaster Recovery (DR) to use the Vault service. This access is required when reading Exadata and Enterprise database passwords required for DR plan executions.
- Policies for File Systems
  Shows how to allow Disaster Recovery (DR) to manage file systems that are part of the application stack.

- Policies for Load Balancers
  Shows how to allow Disaster Recovery (DR) to manage load balancers that are part of the application stack.
- Policies for Network Load Balancers
  Shows how to allow Disaster Recovery (DR) to manage network load balancers that are part of the application stack.
- Policies for Compartments
  Shows how to allow Disaster Recovery (DR) to manage compartments (or the tenancy) that contain resources that are part of the application stack.

Policies for Compute Service

Shows how to allow Disaster Recovery (DR) to manage compute instances that are part of the application stack.

```
Allow group DrAdmins to manage instance-family in compartment compartment_name
```

To know more about the Identity and Access Management (IAM) policies for compute instances, refer Details for the Core Services.

> **Note:**
>
> If a movable instance in a compute instance has a block volumes attached with the iSCSI attachment type, ensure to configure the OCA policies. Configure `sudo` access for a `root` administrator. See Policies for Oracle Cloud Agent.

Policies for Oracle Cloud Agent

Shows how to allow Disaster Recovery (DR) to execute scripts using Oracle Cloud Agent on compute instances that are part of the application stack.

```
Allow group group_name to manage instance-agent-command-family in compartment
compartment_name
```

In addition to granting policy access for Disaster Recovery to execute workflow scripts, you may need to configure additional Identity and Access Management (IAM) policies for Oracle Cloud Agent to allow access to the Run Command feature. For additional details on the IAM policy configurations, refer Required IAM Policy.

To execute scripts or commands with Linux administrator (`sudo`) privileges, you may need additional configuration for a Linux instance. For more details about administrator privileges, refer Running Commands with Administrator Privileges.

Policies for Block Volume Service

Shows how to allow Disaster Recovery (DR) to manage block storage volumes and volume groups that are part of the application stack.

```
Allow group group_name to manage volume-family in compartment compartment_name
```

For more details on Identity and Access Management (IAM) policies for block volume storage, refer Details for the Core Services.

Policies for Networking Service

Shows how to allow Disaster Recovery (DR) to manage networking components for compute instances that are part of the application stack.

```
Allow group group_name to read virtual-network-family in compartment
compartment_name
Allow group group_name to use subnets in compartment compartment_name
Allow group group_name to use vnics in compartment compartment_name
Allow group group_name to use network-security-groups in compartment
compartment_name
Allow group group_name to use private-ips in compartment compartment_name
```

For more details about Identity and Access Management (IAM) policies for networking, refer Details for the Core Services.

Policies for Functions Service

Shows how to allow Disaster Recovery (DR) to invoke Oracle Functions as part of a user-defined step in a DR plan.

```
Allow group group_name to read fn-app in compartment compartment_name
Allow group group_name to read fn-function in compartment compartment_name
Allow group group_name to use fn-invocation in compartment compartment_name
```

For more details about Identity and Access Management (IAM) policies for Oracle Functions, refer Details for Functions.

Policies for Oracle Exadata Database Service and Oracle Base Database Service

Shows how to allow Disaster Recovery (DR) to manage DR for Oracle Exadata Database Service and Oracle Base Database Service databases that are part of the application stack.

```
Allow group group_name to manage database-family in compartment
compartment_name
```

A more restrictive policy that allows DR to only perform switchover and failover operations on databases is similar to the following:

```
Allow group group_name to update databases in compartment compartment_name
```

For additional details about Identity and Access Management (IAM) policies for Oracle Exadata Database Service and Oracle Base Database Service, refer Details for the Database Service.

Policies for Oracle Autonomous Database Service

Shows how to allow Disaster Recovery (DR) to manage Oracle Autonomous Database Service databases that are part of the application stack.

```
Allow group group_name to manage autonomous-database-family in compartment
compartment_name
```

A more restrictive policy that allows DR to only perform switchover and failover operations on autonomous databases is similar to the following:

```
Allow group group_name to update autonomous-databases in compartment
compartment_name
```

For additional about the Identity and Access Management (IAM) policies for Oracle Autonomous Database, refer Details for the Database Service.

Policies for Object Storage Service

Shows how to allow Disaster Recovery (DR) to manage Object Storage buckets and objects. This access is required to write logs to Object Storage during DR plan executions.

```
Allow group group_name to manage buckets in compartment compartment_name
Allow group group_name to manage objects in compartment compartment_name
```

For additional about the Identity and Access Management (IAM) policies for Object Storage, refer Details for Object Storage, Archive Storage, and Data Transfer.

Policies for Tags Service

Shows how to allow Disaster Recovery (DR) to use Tag namespaces. This access is required when launching compute instances as part of DR plan executions.

```
Allow group group_name to use tag-namespaces in tenancy tenancy_name
```

For additional about the Identity and Access Management (IAM) policies for Tags, refer Required IAM Policy.

Policies for Vault Service

Shows how to allow Disaster Recovery (DR) to use the Vault service. This access is required when reading Exadata and Enterprise database passwords required for DR plan executions.

```
Allow group group_name read vaults in compartment compartment_name
Allow group group_name read secret-family in compartment compartment_name
```

For additional about the Identity and Access Management (IAM) policies for Vault, refer Details for the Vault Service.

Policies for File Systems

Shows how to allow Disaster Recovery (DR) to manage file systems that are part of the application stack.

```
Allow group <group_name> to manage file-family in compartment <compartment_name
or compartment_ocid>
```

Policies for Load Balancers

Shows how to allow Disaster Recovery (DR) to manage load balancers that are part of the application stack.

```
Allow group <group_name> to manage load-balancers in compartment
<compartment_name or compartment_ocid>
```

Policies for Network Load Balancers

Shows how to allow Disaster Recovery (DR) to manage network load balancers that are part of the application stack.

```
Allow group <group_name> to manage network-load-balancers in compartment
<compartment_name or compartment_ocid>
```

Policies for Compartments

Shows how to allow Disaster Recovery (DR) to manage compartments (or the tenancy) that contain resources that are part of the application stack.

To allow Disaster Recovery (DR) to access and manage resource in a specific compartment.

To allow Disaster Recovery (DR) to access and manage resource in the entire tenancy.

```
Allow group <group_name> to read all-resources in compartment <compartment_name
or compartment_ocid>
```

```
Allow group <group_name> to read all-resources in tenancy
```

8 Troubleshooting

This section contains troubleshooting related information.

- Launching DR Operations During a Home Region Outage
  Issue: During an outage of the home IAM region, access to Full Stack DR service can be interrupted when using the OCI Console.
- Resetting DR Configuration After a Failover
  Issue: A failover plan execution does not clean up resources in the primary DR protection group or configure DR member properties to prepare for a DR transition in the reverse direction.
- Updating Backend Server IP Addresses in OCI Load Balancer Backend Sets after DR operation
  Issue: After performing switchover, failover, or start drill DR operations, there is a slight possibility that the load balancer may not pass traffic to newly launched resources, such as compute instances that were added as backend servers in the load balancer backend set on the new primary region.

Launching DR Operations During a Home Region Outage

Issue: During an outage of the home IAM region, access to Full Stack DR service can be interrupted when using the OCI Console.

Solution: This can prevent the execution of DR Plans using the OCI Console. In this situation, you can still use the REST API or any of the other API surfaces (CLI, Terraform, Python, and so on) to initiate a DR Plan execution and recovery.

The following example shows how you can execute a DR plan using the OCI CLI:

```
oci disaster-recovery dr-plan-execution create --plan-id
ocid1.drplan.oc1.phx.exampleocid --execution-options file:///path/to/execution-
options.json

where 'execution-options.json' is a file containing the execution options as shown below:

Content of the 'execution-options.json' file for executing a failover:
{
    "arePrechecksEnabled": true,
    "areWarningsIgnored": true,
    "planExecutionType": "FAILOVER"
}

Content of the 'execution-options.json' file for executing a failover precheck:
{
    "areWarningsIgnored": true,
    "planExecutionType": "FAILOVER_PRECHECK"
}
```

Resetting DR Configuration After a Failover

Issue: A failover plan execution does not clean up resources in the primary DR protection group or configure DR member properties to prepare for a DR transition in the reverse direction.

Solution: After a failover plan execution has succeeded, you must manually perform this cleanup and reverse the configuration of properties to prepare the DR configuration for switchover or failover in the reverse direction when needed.

> **Note:**
>
> Performing the manual cleanup step listed here will delete all the existing DR plans.

Perform the following cleanup and reconfiguration tasks manually at the end of a failover plan execution:

1. Delete all the volume groups from the primary DR protection group.
2. Delete all the file systems from the primary DR protection group.
3. Detach and delete all the block volumes attached to moving instances in the primary DR protection group.
4. Delete all the moving instances from the primary DR protection group. Ensure that you also delete the attached boot volumes when deleting instances.
5. Perform a Data Guard reinstate operation on all the Oracle Exadata Database Service and the Oracle Base Database Service databases in the old primary DR protection group. Before performing the reinstate operation verify using the OCI Database console that the new primary database has a **Primary** role, and the old primary database has a **Disabled Standby** role. See Perform Database Switchover and Failover and Reinstate a Database for additional details.
6. If you had configured any non-moving instances in the primary DR protection group to detach block volumes, then detach these block volumes and delete them.
7. If you had configured any non-moving instances in the primary DR protection group to unmount file systems, then unmount these file systems.
8. If you had configured any non-movable instances in the primary DR protection group to unmount file systems, then remove the existing **Mount target** details and re-configure the **Mount target** properties in the **File system** tab for those members.
9. If you had configured any non-movable instances in the primary DR protection group to detach block volumes, then remove the existing attach/detach properties for block volumes in the **Block volumes** tab for those members and re-configure them using block volumes in the new primary DR protection group.

Updating Backend Server IP Addresses in OCI Load Balancer Backend Sets after DR operation

Issue: After performing switchover, failover, or start drill DR operations, there is a slight possibility that the load balancer may not pass traffic to newly launched resources, such as compute instances that were added as backend servers in the load balancer backend set on the new primary region.

Solution: To verify whether the load balancers added to the DR protection group are functioning correctly, check each load balancer and ensure that their health checks are successful. If you find that the health checks are failing, you must update the IP addresses of the respective load balancer's backend servers in the backend set.

A
Reference

This section contains reference materials.

- Events
  Disaster Recovery (DR) resources that emit events.
- Lifecycle States of Full Stack Disaster Recovery Resources
  Lifecycle states of Full Stack DR resources like DR Protection Groups, DR Plans, and DR Plan Executions.
- Prechecks Performed by Full Stack Disaster Recovery
  Full Stack Disaster Recovery performs prechecks for resources such as DR Protection Groups, DR Plans, and DR Plan Executions.
- How to Migrate from an Existing COMPUTE_INSTANCE to a New Instance Type?
  The migration process from the legacy Compute Instance to a new Movable Compute instance or Non-Movable Compute instance is done by removing and adding back the compute instance member to the DR Protection Group.
- Member Properties Causing Plan Deletion Post Update
  Following is the list of member properties which can cause plan deletion after the update:

Events

Disaster Recovery (DR) resources that emit events.

- Table A-1
- Table A-2
- Table A-3

Disaster Recovery Protection Group Event Types

Table A-1 DR Protection Group Event Types

Friendly Name	Event Type
Create DR Protection Group Begin	`com.oraclecloud.disasterrecovery.createdrprotectiongroup`
Create DR Protection Group End	`com.oraclecloud.disasterrecovery.createdrprotectiongroup.end`
Update DR Protection Group Begin	`com.oraclecloud.disasterrecovery.updatedrprotectiongroup.begin`
Update DR Protection Group End	`com.oraclecloud.disasterrecovery.updatedrprotectiongroup.end`

Table A-1 (Cont.) DR Protection Group Event Types

Friendly Name	Event Type
Associate DR Protection Group Begin	`com.oraclecloud.disasterrecovery.associatedrprotectiongroup.begin`
Associate DR Protection Group End	`com.oraclecloud.disasterrecovery.associatedrprotectiongroup.end`
Disassociate DR Protection Group Begin	`com.oraclecloud.disasterrecovery.disassociatedrprotectiongroup.begin`
Disassociate DR Protection Group End	`com.oraclecloud.disasterrecovery.disassociatedrprotectiongroup.end`
Change DR Protection Group Compartment Begin	`com.oraclecloud.disasterrecovery.changedrprotectiongroupcompartment.begin`
Change DR Protection Group Compartment End	`com.oraclecloud.disasterrecovery.changedrprotectiongroupcompartment.end`
Update DR Protection Group Role Begin	`com.oraclecloud.disasterrecovery.updatedrprotectiongrouprole.begin`
Update DR Protection Group Role End	`com.oraclecloud.disasterrecovery.updatedrprotectiongrouprole.end`
Delete DR Protection Group Begin	`com.oraclecloud.disasterrecovery.deletedrprotectiongroup.begin`
Delete DR Protection Group End	`com.oraclecloud.disasterrecovery.deletedrprotectiongroup.end`

Disaster Recovery Protection Group Event Example

This is a reference event for Create DR Protection Group Begin.

Example A-1 DR Protection Group Event

```
{
   "eventType":
"com.oraclecloud.disasterrecovery.createdrprotectiongroup.begin",
   "cloudEventsVersion": "0.1",
   "eventTypeVersion": "2.0",
   "source": "DisasterRecovery",
   "eventTime": "2022-10-12T21:19:24Z",
   "contentType": "application/json",
   "data": {
      "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
      "compartmentName": "example_name",
      "resourceName": "my_drprotectiongroup",
      "resourceId": "ocid1.drprotectiongroup.oc1..<unique_ID>",
      "availabilityDomain": "<availability_domain>",
   }
   "eventID": "<unique_ID>",
   "extensions": {
      "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
   }
}
```

Disaster Recovery Plan Event Types

Table A-2 DR Plan Event Types

Friendly Name	Event Type
Create DR Plan Begin	`com.oraclecloud.disasterrecovery.createdrplan`
Create DR Plan End	`com.oraclecloud.disasterrecovery.createdrplan.end`
Delete DR Plan Begin	`com.oraclecloud.disasterrecovery.deletedrplan`
Delete DR Plan End	`com.oraclecloud.disasterrecovery.deletedrplan.end`
Update DR Plan Begin	`com.oraclecloud.disasterrecovery.updatedrplan.begin`
Update DR Plan End	`com.oraclecloud.disasterrecovery.updatedrplan.end`

Disaster Recovery Plan Event Example

This is a reference event for Create DR Plan Begin.

Example A-2 DR Plan Event

```
{
   "eventType": "com.oraclecloud.disasterrecovery.createdrplan",
   "cloudEventsVersion": "0.1",
   "eventTypeVersion": "2.0",
   "source": "DisasterRecovery",
   "eventTime": "2022-10-12T21:19:24Z",
   "contentType": "application/json",
   "data": {
      "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
      "compartmentName": "example_name",
      "resourceName": "my_drplan",
      "resourceId": "ocid1.drplan.oc1..<unique_ID>",
      "availabilityDomain": "<availability_domain>",
   }
   "eventID": "<unique_ID>",
   "extensions": {
        "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
    }
}
```

Disaster Recovery Plan Execution Event Types

Table A-3 DR Plan Execution Event Types

Friendly Name	Event Type
Create Switchover DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createswitchoverdrplanexecution`
Create Switchover DR Plan Execution End	`com.oraclecloud.disasterrecovery.createswitchoverdrplanexecution.end`
Create Switchover PreCheck DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createswitchoverprecheckdrplanexecution`
Create Switchover PreCheck DR Plan Execution End	`com.oraclecloud.disasterrecovery.createswitchoverprecheckdrplanexecution.end`
Create Failover DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createfailoverdrplanexecution`
Create Failover DR Plan Execution End	`com.oraclecloud.disasterrecovery.createfailoverdrplanexecution.end`
Create Failover PreCheck DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createfailoverprecheckdrplanexecution`
Create Failover PreCheck DR Plan Execution End	`com.oraclecloud.disasterrecovery.createfailoverprecheckdrplanexecution.end`

Table A-3 (Cont.) DR Plan Execution Event Types

Friendly Name	Event Type
Create Startdrill DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createstartdrilldrplanexecution`
Create Startdrill DR Plan Execution End	`com.oraclecloud.disasterrecovery.createsstartdrilldrplanexecution.end`
Create Startdrill PreCheck DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createstartdrillprecheckdrplanexecution`
Create Startdrill PreCheck DR Plan Execution End	`com.oraclecloud.disasterrecovery.createstartdrillprecheckdrplanexecution.end`
Create Stopdrill DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createstopdrilldrplanexecution`
Create Stopdrill DR Plan Execution End	`com.oraclecloud.disasterrecovery.createsstopdrilldrplanexecution.end`
Create Stopdrill PreCheck DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.createstopdrillprecheckdrplanexecution`
Create Stopdrill PreCheck DR Plan Execution End	`com.oraclecloud.disasterrecovery.createstopdrillprecheckdrplanexecution.end`
Update DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.updatedrplanexecution.begin`

Table A-3 (Cont.) DR Plan Execution Event Types

Friendly Name	Event Type
Update DR Plan Execution End	`com.oraclecloud.disasterrecovery.updatedrplanexecution.end`
Cancel DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.canceldrplanexecution.begin`
Cancel DR Plan Execution End	`com.oraclecloud.disasterrecovery.canceldrplanexecution.end`
Pause DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.pausedrplanexecution.begin`
Pause DR Plan Execution End	`com.oraclecloud.disasterrecovery.pausedrplanexecution.end`
Resume DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.resumedrplanexecution.begin`
Resume DR Plan Execution End	`com.oraclecloud.disasterrecovery.resumedrplanexecution.end`
Retry DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.retrydrplanexecution.begin`
Retry DR Plan Execution End	`com.oraclecloud.disasterrecovery.retrydrplanexecution.end`
Ignore DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.ignoredrplanexecution.begin`
Ignore DR Plan Execution End	`com.oraclecloud.disasterrecovery.ignoredrplanexecution.end`
Delete DR Plan Execution Begin	`com.oraclecloud.disasterrecovery.deletedrplanexecution.begin`

Table A-3   (Cont.) DR Plan Execution Event Types

Friendly Name	Event Type
Delete DR Plan Execution End	`com.oraclecloud.disasterrecovery.deletedrplanexecution.end`

Disaster Recovery Plan Execution Event Example

This is a reference event for Create DR Plan Execution Begin:

Example A-3   DR Plan Execution Event

```
{
   "eventType":
"com.oraclecloud.disasterrecovery.createswitchoverdrplanexecution",
   "cloudEventsVersion": "0.1",
   "eventTypeVersion": "2.0",
   "source": "DisasterRecovery",
   "eventTime": "2022-10-12T21:19:24Z",
   "contentType": "application/json",
   "data": {
      "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
      "compartmentName": "example_name",
      "resourceName": "my_drplanexecution",
      "resourceId": "ocid1.drplanexecution.oc1..<unique_ID>",
      "availabilityDomain": "<availability_domain>",
   }
   "eventID": "<unique_ID>",
   "extensions": {
      "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
   }
}
```

Lifecycle States of Full Stack Disaster Recovery Resources

Lifecycle states of Full Stack DR resources like DR Protection Groups, DR Plans, and DR Plan Executions.

- Table A-4
- Table A-6
- Table A-7

Lifecycle States for DR Protection Groups

Table A-4   Lifecycle States for DR Protection Groups

Lifecycle State	Description
Creating	The DR Protection Group is in the process of being created. When in this state, you cannot modify or delete a group. Wait for the group to reach the Active or Failed state before you attempt to modify or delete it.
Active	The DR Protection Group is available for use. When the group is in this state, you can modify or delete it. The group is also available for you to create DR Plans or run DR Plan Executions.
Updating	The DR Protection Group is being updated and not available for modification. The group enters this state during the following situations: • The DR Protection Group is being modified. This can include addition or deletion of members. • A DR Plan Execution for the DR Protection Group is in-progress or has failed, and is waiting for user intervention. The DR Protection Group usually moves back to a Active state after the modification or plan execution is complete, or it may move to a Needs attention state. If a DR Protection Group is stuck in the Updating state, ensure that plan executions for the DR Protection Group are not in a Failed or Paused state. The DR Protection Group exits the Updating state after the plan execution succeeds or is canceled.
Needs attention	The DR Protection Group requires user attention and intervention because one of the following conditions is true: • During DR Plan Execution, contact with the peer DR Protection Group is unsuccessful (contact with the peer region is lost). • During a role change at the end of DR Plan Execution, updating the lifecycle states of DR Plans was unsuccessful. You can move the DR Protection Group back to the Active state by going to the resource page for the DR Protection Group, clicking the **More Actions** drop-down, and selecting **Reset State.** Additionally, you can update the roles of the two paired DR Protection Groups to correctly set their roles to primary and standby.
Deleting	The DR Protection Group is being deleted and cannot be modified.
Deleted	The DR Protection Group is deleted and cannot be modified. Deleted resources are removed after some time.
Failed	The DR Protection Group failed during creation, association, or modification. To remove the DR Protection Group from this state reattempt the association or modification. You can also delete a DR Protection Group in this state.
Inactive	When you enable the **Start drill** plan and the drill is in progress for a DR Protection Group, then the state is updated to Inactive. **Note:** No updates are allowed when the DR Protection Groups are in an Inactive state.

Lifecycle Sub-states for DR Protection Groups

Table A-5 Lifecycle Sub-states for DR Protection Groups

Lifecycle Sub-state	Description
Drill in progress	The Drill in progress lifecycle sub-state would be set after you run the **Start drill** plan and when the **Start drill** execution is completed.

Lifecycle States for DR Plans

Table A-6 Lifecycle States for DR Plans

Lifecycle State	Description
Creating	The DR Plan is in the process of being created. You cannot modify or delete a DR plan in this state. Wait for the DR Plan to reach the Active or Failed state before you try to modify or delete it.
Updating	The DR Plan is being updated and you cannot modify or delete it. The DR Plan enters this state when you are modifying it.
Active	The DR Plan is available for modification or for launching plan executions. This is the state in which plans exist at the standby DR Protection Group after they are created and available for use. When the primary DR Protection Group transitions to a standby role, all the Inactive plans stored at the DR protection group will move to an Active state.
Inactive	The DR Plan is not available for modification or for launching plan executions. In this state plans exist at the primary DR Protection Group after you have created them. When the standby DR Protection Group transitions to a primary role, all the Active plans stored at the DR Protection Group move to an Inactive state.
Deleting	The DR Plan is being deleted and you cannot modify it.
Deleted	The DR Plan has been deleted and you cannot modify it. Full Stack DR removes deleted resources after some time.
Failed	The DR Plan has failed due to an internal error and you cannot use it in its current state. You can delete Failed DR Plans.
Needs attention	The DR Plan requires user attention and intervention because a plan lifecycle state update was unsuccessful. To resolve this issue, create a new DR Plan for the DR Protection Group or execute a precheck for any existing plan which is in a Needs attention state.

Lifecycle States for DR Plan Execution

Table A-7 Lifecycle States for DR Plan Execution

Lifecycle State	Description
Accepted	The DR Plan Execution is accepted but has not started.
In progress	The DR Plan Execution has started and is in progress.
Canceling	The DR Plan is being canceled because the user attempted to cancel it (abandon the execution).
Canceled	The DR Plan Execution is canceled. Canceled plan executions reach a terminal state and you cannot modify or resume a it. However, you can delete canceled plan executions.

Table A-7 (Cont.) Lifecycle States for DR Plan Execution

Lifecycle State	Description
Succeeded	The DR Plan Execution is successful.
Failed	The DR Plan Execution failed. • You can retry or skip the failed step or group of a failed DR Plan Execution. • You must cancel a failed DR Plan Execution to bring the parent DR Protection Group out of the Updating state. • You cannot launch a new DR Plan Execution when another failed DR Plan Execution exists in an uncanceled state. • You must first cancel a failed DR Plan Execution before you can delete it.
Deleting	The DR Plan Execution is being deleted and you cannot resume, restart, or modify it.
Deleted	The DR Plan Execution is deleted and you cannot be resume, restart, or modify it. Deleted resources are removed after some time.
Pausing	The DR Plan Execution is being paused. You cannot modify, resume, or cancel a DR Plan Execution when it is being paused.
Paused	The DR Plan Execution is paused. You can update, resume, or cancel a DR Plan Execution when it is paused. You cannot delete a paused DR Plan Execution.
Resuming	The DR Plan Execution is resuming. You cannot modify, pause, cancel or delete a DR Plan Execution when it is resuming. The DR Plan Execution can move back to an In progress or Failed state after it is resumed.

Prechecks Performed by Full Stack Disaster Recovery

Full Stack Disaster Recovery performs prechecks for resources such as DR Protection Groups, DR Plans, and DR Plan Executions.

Prechecks for Compute Instance

Full Stack DR first performs the following storage prechecks for each VM present in the primary DRPG. Full Stack DR verifies that:

- Volume group replication is configured or backup is configured with a backup policy and cross-region copy is enabled.
- A volume group replica or at least one volume group backup exists in the standby region. Multiple backups can also exist as Full Stack DR uses the latest volume group backup.
- All the boot and block volumes of the VMs of the members in a DRPG are added to the volume group.
- Volume group contains only the boot and block volumes attached to the VM of the members in a DRPG.
- Whether the user is trying to add moving compute instances to a standby DR Protection Group, which is not allowed.

Prechecks for Mount File System on Compute Instance:

Full Stack DR first performs the following prechecks for the mount file system on compute instance:

- Movable Compute Instance:

- Create DR Plan Validations:
  - Validation for mount details for File System:
    - Validates that the mount details property is present on instance property (fileSystemOperationDetails).
    - Validates that the mount target of mount details matches with standby region.
    - Validates that the combination of mount point and export is unique (avoid multiple mounting on the same mount point).
    - Validates that the mount target of mount details is in an active state.
  - Validates that the instance operating system is not WINDOWS.
- Create DR Plan Execution Pre-check Validations:
  - Validation for mount details for File System:
    - Validates that the mount details property is present on instance property (fileSystemOperationDetails).
    - Validates that the mount target of mount details matches with standby region.
    - Validates that the combination of mount point and export is unique (avoid multiple mounting on same mount point).
    - Validates that the mount target of mount details is in an active state.
  - Validates that the compute instance and mount target of mount details are having correct TCP/UDP protocol enable. See Configuring VCN Security Rules for File Storage
  - In the case of start drill or failover:

    > **Note:**
    >
    > For a switchover, this check is performed in the unmount pre-check step. However you do not need to check for the stop drill as there is no mount operation.

    - Validates that the compute instance have the Compute Instance Run Command plugin enabled.
    - Validates that the compute instance has a root-access on ocarun user. For information on how to get the root access on compute instance, see Running Commands on an Instance.
    - Validates that the compute instance has nfs-client installed. For information on how to install nfs-client on compute, see Mounting File Systems From UNIX-Style Instances.

- Non-Movable Compute Instance:
  - Create DR Plan Validations in Standby DR Protection Group
    - Validation for mount details for File System:
      - Validates that the mount target property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of mount details matches with standby region.
      - Validates that the mount target of mount details is in active state.

        - Validates that the combination of mount point and export is unique (avoid multiple mounting on the same mount point).
  - Validates that the instance operating system is not WINDOWS.
  - Create DR Plan Execution Pre-check Validations in Standby DR Protection Group
    - Validation for mount details for File System:
      - Validates that the mount target property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of mount details matches with standby region.
      - Validates that the mount target of mount details is in active state.
      - Validates that the combination of mount point and export is unique (avoid multiple mounting on the same mount point).
    - Validates that the compute instance and mount details are having the correct TCP/UDP protocol enabled.
    - Validates that the compute instance has Compute Instance Run Command plugin enabled.
    - Validates that the instance operating system is not WINDOWS.
    - Validates that the compute instance has root-access on ocarun user. For information on how to get the root access on compute instance, see Running Commands on an Instance.
    - Validates that the compute instance has nfs-client installed. For information on how to install nfs-client on compute, see Mounting File Systems From UNIX-Style Instances.

Prechecks for Unmount File System on Compute Instance:

Full Stack DR first performs the following prechecks for the unmount file system on compute instance:

- Movable Compute Instance
  - Create DR Plan Validations
    - Validation for unmount details for File System:
      - Validates that the unmount details property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of unmount details matches with primary region.
      - Validates that the mount target of unmount details is in active state.
      - Validates that the export path is present on the mount target of unmount details.
    - Validates that the instance operating system is not WINDOWS.
  - Create DR Plan Execution Pre-check Validations
    - Validation for unmount details for File System:
      - Validates that the unmount details property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of unmount details matches with primary region.

        - Validates that the mount target of unmount details is in active state.
        - Validates that the export path is present on the mount target of unmount details.
      - Validates that the compute instance and mount target of unmount details are having the correct TCP/UDP protocol enabled.
      - Validates that the compute instance has the Compute Instance Run Command plugin enabled.
      - Validates that the mount point is present on the compute instance.
      - Validates that the instance operating system is not WINDOWS.
      - Validates that the compute instance has root-access on ocarun user. For information on how to get the root access on compute instance, see Running Commands on an Instance.
      - Validates that the compute instance has nfs-client installed. For information on how to install nfs-client on compute, see Mounting File Systems From UNIX-Style Instances.
- Non Movable Compute Instance
  - Create DR Plan Validations for Primary DR Protection Group
    - Validation for unmount details for File System:
      - Validates that the mount target property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of unmount details matches with primary region.
      - Validates that the mount target of unmount details is in active state.
      - Validates that the export path is present on the mount target of unmount details.
    - Validates that the instance operating system is not WINDOWS.
  - Create DR Plan Execution Pre-check Validations for Primary DR Protection Group
    - Validation for unmount details for File System:
      - Validates that the mount target property is present on instance property (fileSystemOperationDetails).
      - Validates that the mount target of unmount details matches with primary region.
      - Validates that the mount target of unmount details is in active state.
      - Validates that the export path is present on the mount target of unmount details.
    - Validates that the compute instance and unmount details are having the correct TCP/UDP protocol enabled.
    - Validates that the mount point is present on the compute instance.
    - Validates that the compute instance have the Compute Instance Run Command plugin enabled.
    - Validates that the instance operating system is not WINDOWS.

  - Validates that the compute instance has root-access on ocarun user. For information on how to get the root access on compute instance, see Running Commands on an Instance.
  - Validates that the compute instance has nfs-client installed. For information on how to install nfs-client on compute, see Mounting File Systems From UNIX-Style Instances.

Prechecks for Volume Groups (Block Storage)

Full Stack DR first performs the following prechecks for all volume groups added to the primary DRPG. Full Stack DR verifies that:

- The volume group is in an `Available` state.
- The volume group has either replication or backups configured in the standby region. If both are configured, Full Stack DR uses replicas and ignore backups.
- For intra-region DR that any destination (standby) region replicas are not in the same availability domain (AD).
- The replica in the standby region is in an `Available` state, or if backups are used, that at least one backup exists and is `Available`.
- The list of volumes in the source volume group match the list of volumes in the standby region replica or backup.

Prechecks for Block volume for Non-Movable compute instances

Full Stack DR first performs the following prechecks for the block volume for non-movable compute instances:

If the compute instance is added as member to a DRPG with the role **Primary**, then perform the following validations for each block volume ID provided in the new member property list:

- The block volume ID should be a valid OCID of a block volume.
- The block volume should not have duplicates in the member properties of the same compute instance.
- Block volume should be already attached to the compute instance.
- The block volume should be a part of some volume group member of the DRPG.
- If a **Volume attachment reference instance ID** is provided in the attachment details, then that instance should be a member of the standby DR Protection Group and the block volume ID should be added in its member properties.
- If the **Volume attachment reference instance ID** is not provided in the attachment details, then only one compute instance in the standby DRPG should have a member property defined with this block volume ID.
- The mount points that are defined should be unique.

If the compute instance is added as member to a DRPG with the role **Standby**, then perform the following validations for each block volume ID provided in the new member property list:

- The block volume ID should be a valid OCID of a block volume.
- The block volume should not have duplicates in the member properties of the same compute instance.
- The block volume should be from the region of the primary DRPG.
- The block volume should be a part of some volume group member of the primary DRPG.

- The volume group's destination/target AD (where the backup or replica will be activated) should match the AD of this standby compute instance.
- If a **Volume attachment reference instance ID** is provided in the attachment details, then that peer instance should be a member of primary DRPG and the block volume should be attached to it.
- If the **Volume attachment reference instance ID** is not provided in the attachment details, then only one compute instance in the primary DRPG should have the block volume attached to it.
- The mount points that you define should be unique.
- No two block volumes should be configured to attach using a same device path.
- If the attachment uses device paths, then the device paths must not be in use.
- If a block volume is configured to be attached to more than one compute instance, then the attachment must have a shareable access.

Prechecks for Database (Oracle Base Database Service and Oracle Exadata Cloud Service) Instance

Full Stack DR performs the following prechecks if a database member (Oracle BaseDatabase Service, Oracle Exadata Database Service on Dedicated Infrastructure) is a part of the DRPG. Full Stack DR verifies that:

- Database member properties are not empty or null and password secret vault location is a part of the database member properties.
- You are able to access the secret vault in which the database password is stored database and peer database is in an `Available` state.
- Database and peer Database have Data Guard enabled and they are Data Guard peers of each other.
- Database and peer Database have the correct Data Guard roles.
- Database and peer Database are a part of the two associated DR protection groups that are a part of the configuration. Primary database is a part of the primary DR protection group and standby database is a part of the standby DR protection group.

Prechecks for Autonomous Database Instance

Full Stack DR performs the following prechecks if an Autonomous database member is part of the DRPG. Full Stack DR verifies that:

- Autonomous database member properties are not empty or null.
- The primary Autonomous database does not have an empty standby database list.
- The primary Autonomous database does not have more than one standby databases configured.
- The standby Autonomous database is not in the same region as the primary database region and is not a local peer.
- The Autonomous database and the peer Autonomous database are a part of the two associated DR protection groups that are a part of the configuration.
- Remote Data Guard is configured.
- Remote peer database belongs to the remote DRPG.
- The primary database lifecycle state is `AVAILABLE`

For switchover prechecks, Full Stack DR performs the following additional validations on the standby database:

- Verifies that remote peer standby is in the correct (`STANDBY`) state.
- Verifies that remote peer standby has only one peer configured which is the primary database.

How to Migrate from an Existing COMPUTE_INSTANCE to a New Instance Type?

The migration process from the legacy Compute Instance to a new Movable Compute instance or Non-Movable Compute instance is done by removing and adding back the compute instance member to the DR Protection Group.

> **Note:**
>
> All the existing plans will be deleted because this process counts as deleting and adding members to a DR Protection Group.

Regenerate the DR plans, including re-adding any user-defined steps and customizations, after migrating to one of the new instance types. By migrating to new compute instance types, you can use new features such as capacity reservation, block volume operations, FSS mount/unmount and so on.

- Migrating Using the Console
  Refer to the following steps to perform migration using the Console.
- Migrating Using the REST API
  Refer to the following steps to perform migration using the REST API.

Migrating Using the Console

Refer to the following steps to perform migration using the Console.

1. In the console, navigate to **DR Protection Groups** and select the DR protection group from which you need to migrate the members.
2. In the **Resources** section, select **Members**.
3. Select the check box for the compute instance members and select the legacy compute instances which needs to be migrated.

   > **Note:**
   >
   > Ensure to mark only the compute instances marked as legacy that need migration.

4. Click **Remove Members**.
5. Accept the warning about the plan deletion and click **Remove**.

   Now the selected compute instance members will be deleted from the DRPG.

6. Now add back the compute instances to the DR Protection Group by clicking **Add Members** in members section of the DR Protection Group.
7. Select **Compute** in the resource type and accept the warning about the plan deletion.
8. Select whether the compute instance is a Movable or Non-Movable instance.
9. Provide any additional member properties and configuration information for this compute instance.
10. Click **Add** to add the new compute instance member to the DR Protection Group.
11. Repeat the same steps to replace all the legacy compute instance members in the primary and standby DR Protection Groups.

> **Note:**
>
> You can generate the plans and add the user defined steps to the plans once the compute instance members are migrated.

Migrating Using the REST API

Refer to the following steps to perform migration using the REST API.

1. GET DR protection group details. Endpoint:`GET /drProtectionGroups/{drProtectionGroupId}`

```
Response:

{
  "id": "ocid1.drprotectiongroup.oc1.iad.xxxxxxxx",
  "compartmentId": "ocid1.compartment.oc1..xxxxxxx",
  "displayName": "IAD-DRPG-1234",
  "role": "PRIMARY",
  "peerId": "ocid1.drprotectiongroup.oc1.phx.xxxxxxx",
  "peerRegion": "us-phoenix-1",
  "logLocation": {
    "namespace": "xyz1234",
    "bucket": "bucket-1234",
    "object": null
  },
  "members": [
    {
      "memberId": "ocid1.instance.oc1. iad.xxxxxxxx1",
      "memberType": "COMPUTE_INSTANCE",
      "isMovable": true,
      "vnicMapping": [
        {
          "sourceVnicId": "ocid1.vnic.oc1.iad.xxxxxxxx",
          "destinationSubnetId": "ocid1.subnet.oc1.phx.xxxxxx"
        }
      ]
    },
    {
      "memberId": "ocid1.instance.oc1.iad.xxxxxxxx2",
      "memberType": "COMPUTE_INSTANCE",
      "isMovable": false
    },
    {
      "memberId": "ocid1.volumegroup.oc1.iad.xxxxxxx",
      "memberType": "VOLUME_GROUP"
    }
  ],
  "timeCreated": "2023-10-17T00:14:28.860Z",
  "timeUpdated": "2023-10-17T00:14:44.544Z",
  "lifecycleState": "ACTIVE",
  "lifeCycleDetails": null,
  "lifecycleSubState": null,
  "freeformTags": null,
  "definedTags": null,
  "systemTags": null
}
```

2. Update the DRPG deleting legacy compute members with `memberType`: `COMPUTE_INSTANCE` by invoking an update DRPG API (PUT). Use the response of the GET request to build the update request payload by removing the legacy members. `Endpoint :PUT / drProtectionGroups/{drProtectionGroupId}`

```
Request:
{
  "members": [
    {
      "memberType": "VOLUME_GROUP",
      "memberId": "ocid1.volumegroup.oc1.iad.xxxxxxx"
    }
  ]
}
```

3. Update the DRPG adding new compute members replacing the legacy compute members having `isMovable`: `false` with `memberType`: `COMPUTE_INSTANCE_NON_MOVABLE` and which are having `isMovable`: `true` with `memberType`:`COMPUTE_INSTANCE_MOVABLE`.

> **Note:**
>
> The `vnicMapping` attribute in the legacy compute member object is changed to `vnicMappings` in new object.

```
Endpoint :PUT /drProtectionGroups/{drProtectionGroupId}

Request:
{
  "members": [
    {
      "memberId": "ocid1.instance.oc1. iad.xxxxxxxx1",
      "memberType": "COMPUTE_INSTANCE_MOVABLE",
      "vnicMappings": [
        {
          "sourceVnicId": "ocid1.vnic.oc1.iad.xxxxxxxx",
          "destinationSubnetId": "ocid1.subnet.oc1.phx.xxxxxx"
        }
      ]
    },
    {
      "memberId": "ocid1.instance.oc1.iad.xxxxxxxx2",
      "memberType": "COMPUTE_INSTANCE_NON_MOVABLE"
    },
    {
      "memberId": "ocid1.volumegroup.oc1.iad.xxxxxxx",
      "memberType": "VOLUME_GROUP"
    }
  ]
}
```

Member Properties Causing Plan Deletion Post Update

Following is the list of member properties which can cause plan deletion after the update:

- Any Resource type: Adding or deleting members from the DRPG.
- Any Compute instance: Changing the **Compute instance type** from moving instance to non-moving instance and vice versa.

  > **Note:**
  >
  > This option is disabled from the UI.

- Moving instance:
  - Adding, deleting or updating **Destination capacity reservation** in the **Destinations** tab.
  - Adding, deleting or updating File systems' export mappings in the **File Systems** tab.
- Non-moving instance:
  - Enabling/disabling Start and stop instance on failover/switchover in the **Settings** tab.
  - Adding, deleting or updating block volume in block volume mappings will delete the plans in the **Block Volumes** tab.
  - Updating **Volume attachment reference instance** or **mount point** in existing block volume mapping will not delete the plans in the **Block Volumes** tab.
- Adding, deleting or updating File systems' export mappings in the **File Systems** tab.