Disaster Recovery Framework Concepts for Oracle® Solaris Cluster 4.4

Part No: E69311
August 2018

Disaster Recovery Framework Concepts for Oracle Solaris Cluster 4.4

Part No: E69311

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.

Référence: E69311

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` ou le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` si vous êtes malentendant.

Contents

Figures

Using This Documentation

- **Overview** – Describes how Oracle Solaris Cluster disaster recovery framework adds to Oracle Solaris Cluster.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster and the disaster recovery framework, and expertise with the volume manager software that is used with the disaster recovery framework.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Documentation and resources for this product and related products are available at `http://docs.oracle.com/cd/E69294_01`.

Feedback

Provide feedback about this documentation at `http://www.oracle.com/goto/docfeedback`.

♦ ♦ ♦ CHAPTER 1

Introduction to the Disaster Recovery Framework

The Oracle Solaris Cluster disaster recovery framework (formerly called Geographic Edition) is a layered extension of the Oracle Solaris Cluster software. The disaster recovery framework protects applications from unexpected disruptions by using multiple clusters that are separated by long distances and by using a redundant infrastructure that replicates data between these clusters. The clusters can be global clusters, zone clusters, or a combination of both.

An Oracle Solaris Cluster configuration running the disaster recovery framework enables applications to tolerate disasters by managing replication of data and migration of services to a geographically separated secondary cluster. A configuration running the disaster recovery framework can also be used to manage a remote Oracle database that uses Oracle Data Guard replication to a secondary site, even when the database system is not running Oracle Solaris Cluster software.

A disaster such as an earthquake, a fire, or a storm might disable the cluster at the primary site. If a disaster occurs, the disaster recovery framework can continue to provide services by using the following levels of redundancy:

- A secondary cluster
- Duplicated application configuration on the secondary cluster
- Replicated data on the secondary cluster

This chapter gives a high-level overview of the disaster recovery framework. It contains the following sections:

- "Business Continuity" on page 12
- "Making Applications Highly Available With the Disaster Recovery Framework" on page 12
- "Recovering From a Disaster" on page 13
- "Key Features of the Disaster Recovery Framework" on page 14
- "Administration and Configuration Tools" on page 15

Business Continuity

Business continuity is a wide-ranging subject that analyzes all aspects of how a business maintains service to its customers when faced with an unexpected disaster. When creating a business continuity plan, companies must trade off the cost of the additional required infrastructure such as hardware, software, telecommunications, and buildings, against the risks, such as the costs of a prolonged outage. As a result, systems that are critical to the business and those for which there is a legal requirement are the top priorities.

For a service to be available, all the constituent pieces must also be available. The key issue is rapid recovery of individual service elements. Outages can occur from hardware failure such as component or power failures, or from software failures such as operating system panics and application crashes. Network connectivity failures can also affect service availability. Most of these failures can be masked through component redundancy or by having a standby server ready to take over the workload. The disaster recovery framework is a building block for disaster tolerance which provides a framework that enables data services to be moved between a primary cluster and a geographically separated secondary cluster in a controlled fashion.

Making Applications Highly Available With the Disaster Recovery Framework

The disaster recovery framework provides a suite of tools to manage and configure geographically separated clusters with a migration of services between sites. The disaster recovery framework can manage availability across multiple physical locations through robust security, application service migration, and data replication to tolerate disaster across an enterprise system.

The disaster recovery framework enables an improved combination of performance, cost, and separation of data recovery points. This combination contrasts with campus or metro clustering, which consists of a single cluster with widely separated nodes. The disaster recovery framework provides the management and configuration tools for geographically separated clusters.

A configuration that is running the disaster recovery framework consists of a set of clusters that are geographically distributed. The primary cluster provides application services. The secondary cluster in the set is an alternative site that can take over the primary cluster services if a disaster occurs. The disaster recovery framework manages configuration, data replication, and heartbeat monitoring between the two clusters and enables data to be decentralized across multiple disaster recovery sites.

Recovering From a Disaster

Disaster tolerance is the ability of a system to restore an application on a secondary cluster when the primary cluster fails. Disaster tolerance is based on data replication and failover. The disaster recovery framework enables disaster tolerance by redundantly deploying the following:

- Highly available clusters that are geographically separated
- Data replication at either the host or the storage level
- Backups and restoration and data vaulting

Data replication is the process of continuously copying data from the primary cluster to the secondary cluster. Through data replication, the secondary cluster has a recent copy of the data on the primary cluster. The secondary cluster can be geographically separated from the primary cluster.

The disaster recovery framework supports two types of migration of services: a switchover and a takeover.

- A switchover is a planned migration of services from the primary cluster to the secondary cluster. During a switchover, the primary cluster is connected to the secondary cluster and coordinates the migration of services with the secondary cluster. This coordination enables the data replication to complete and ensures that services can be transferred from the primary cluster to the secondary cluster without loss or corruption of data.
- A takeover is an emergency migration of services from the primary cluster to the secondary cluster. A system administrator can initiate a takeover to recover from a disaster. Unlike a switchover, the primary cluster is not connected to the secondary cluster during a takeover. Therefore, the primary cluster cannot coordinate with the secondary cluster to migrate the services. Because of this lack of coordination, the risk of data loss and data corruption in a takeover is higher than it is with a switchover. The disaster recovery framework uses dedicated recovery procedures during a takeover to minimize data loss and data corruption.

These operations intentionally require manual initiation, rather than occur automatically like failover between cluster nodes. Business continuity covers all aspects of a company's response to a disaster, not only information technology (IT) but also staff availability and welfare, phones, buildings, and so forth. A good business continuity plan will include all these things and will outline the actions to be taken. When a disaster occurs, it can be extremely difficult to obtain accurate information about what is happening. Having one part of the infrastructure attempting an automatic recovery while other areas are still trying to work out what is happening can often make matters worse.

General best practice is to have a designated Business Continuity Manager involved in disaster recovery decisions, to review status and decide on appropriate action. Once an action is decided upon, it must then be performed correctly, preferably in an automated, tested way.

This is the basis of the disaster recovery framework takeover operation. For example, if a brief power outage has crashed systems at one site, switching to a remote site might not be the correct response. If the remote site is in another time zone, where staff are not on duty, such a takeover will require that staff be paged, and potentially all communications services redirected. After the outage is corrected, the process must be reversed. It might, in the circumstances, be much more effective to simply restart the primary site. Having the IT infrastructure take over automatically while the situation is being evaluated will not help recovery.

Key Features of the Disaster Recovery Framework

The disaster recovery framework provides the following features:

- Failure detection of multiple clusters that are geographically separated
- Configurable heartbeat monitoring between clusters
- Application resource switchover from one cluster to another cluster
- Remote management of partner clusters through a browser interface and a command-line interface (CLI)
- Management of Oracle Data Guard that is running on remote systems which are not running Oracle Solaris Cluster software.
- Data replication between geographically separated clusters
- Secure administration interfaces through rights profiles
- Secure Sockets Layer (SSL) authentication and encryption for communication between nodes or clusters
- Configurable IPsec security for data replication between clusters and for heartbeat communication between clusters
- Ability to automatically run a script when a heartbeat-loss notification is issued
- Ability to start, stop, switch over, or take over predetermined sets of protection groups in a single operation

The disaster recovery framework provides tools for managing data replication between geographically separated clusters. The software supports the following data replication products:

- MySQL
- Oracle Data Guard, for use only with HA for Oracle Database or Oracle Real Application Clusters (Oracle RAC), with the Oracle database running locally or accessed remotely by using HA for Oracle External Proxy
- Oracle Solaris ZFS snapshots
- Oracle ZFS Storage Appliance

- Script-based plug-ins feature of the disaster recovery framework, to integrate user-written replication modules

Administration and Configuration Tools

You can configure, control, and monitor partnerships, heartbeats, and protection groups either through the Oracle Solaris Cluster Manager browser interface or through the command-line interface (CLI). The disaster recovery framework CLI contains a set of dedicated commands.

For more information about Oracle Solaris Cluster Manager, see Chapter 12, “Using the Oracle Solaris Cluster Manager Browser Interface” in *Administering an Oracle Solaris Cluster 4.4 Configuration*. For more information about the disaster recovery framework CLI, see *Disaster Recovery Framework Reference for Oracle Solaris Cluster 4.4*.

♦ ♦ ♦ CHAPTER 2

Key Concepts for Disaster Recovery Framework

This chapter describes the key concepts for using the disaster recovery framework. These concepts can help you understand the relationships between the disaster recovery framework components.

This chapter contains the following sections:

- "Data Replication" on page 17
- "Cluster Partnerships" on page 21
- "Protection Groups" on page 22
- "Multigroups and Sites" on page 25
- "Heartbeat Monitoring and Plug-ins" on page 25

Data Replication

This section provides the following information:

- "Overview of Data Replication" on page 17
- "Oracle Data Guard Software" on page 18
- "Oracle Solaris ZFS Snapshots Feature" on page 19
- "Oracle ZFS Storage Appliance Software" on page 19
- "Disaster Recovery Framework Script-Based Plug-Ins" on page 19
- "Resource Groups and Data Replication Components" on page 20

Overview of Data Replication

Data replication enables controlled migration of production services from a primary cluster to a secondary cluster either in the event of a disaster or as part of a planned procedure. Data is continuously replicated from the primary cluster to the secondary cluster either synchronously or asynchronously, or a combination of both, depending on the recover point objectives of the application services that are supported by the clusters.

The disaster recovery framework supports the following software for data replication:

- MySQL software
- Oracle Data Guard software, where the Oracle database is running on a local Oracle Solaris Cluster node or on a remote system that is not running Oracle Solaris Cluster software
- Oracle Solaris ZFS snapshots feature
- Oracle ZFS Storage Appliance
- Disaster recovery framework script-based plug-ins

Oracle Solaris ZFS snapshots is a host-based data replication facility which replicates data at the file system or logical volume level within the operating system. Oracle ZFS Storage Appliance software uses a storage-based data replication facility which replicates data at the storage system level and provides a transparent service to applications. Oracle Data Guard software and MySQL software are application-based data replication facilities that maintain one or more standby databases as synchronized replicas of a production database.

MySQL Software

MySQL database software has a built-in replication component that is an application-based replication facility. MySQL replication keeps one or more standby databases in sync with the active master. disaster recovery framework supports the use of MySQL replication between two databases. Only asynchronous replication is available.

You add a plug-in for MySQL replication to a protection group. This plug-in monitors the status of the replication and controls the replication direction. For more information about MySQL software, see the product documentation.

Oracle Data Guard Software

Oracle Data Guard software is an application-based replication facility that maintains one or more standby databases as synchronized replicas of a production database. disaster recovery framework supports the use of Oracle Data Guard replication where the Oracle database is running locally on an Oracle Solaris Cluster node or remotely on a separate system that is not running Oracle Solaris Cluster software. When the Oracle database runs locally, the HA for Oracle Database or Support for Oracle RAC data service must be used. When the database runs remotely, the HA for Oracle External Proxy data service must be configured on the local cluster for the remote database.

You add an Oracle Data Guard broker configuration that is controlled by the Oracle Data Guard software to a protection group. The disaster recovery framework creates a shadow server proxy resource group for each Oracle Data Guard broker configuration. The shadow server proxy resource group "shadows" the database resource group to manage and monitor the Oracle databases that are under the control of disaster recovery framework. For more information about Oracle Data Guard software, see the product documentation.

Oracle Solaris ZFS Snapshots Feature

Oracle Solaris ZFS snapshots is a host-based replication facility that enables ZFS dataset snapshots to be transmitted from the primary cluster to the secondary cluster. Oracle Solaris ZFS has this built-in feature to take consistent snapshots of a ZFS dataset, which can be transmitted and replayed on another system that hosts a copy of the ZFS dataset. This feature of ZFS provides a method of replicating ZFS datasets within a disaster recovery (DR) environment and enables the primary cluster to transmit ZFS snapshots of a ZFS dataset to a secondary cluster. For more information about Oracle Solaris ZFS snapshots, see Chapter 8, “Working With Oracle Solaris ZFS Snapshots and Clones” in *Managing ZFS File Systems in Oracle Solaris 11.4*.

Oracle ZFS Storage Appliance Software

Oracle ZFS Storage 7000 appliances support snapshot-based replication of projects and shares from a source appliance to any number of target appliances manually, on a schedule or continuously. The replication includes both data and metadata. Remote replication, or just replication, is a general-purpose feature optimized for disaster recovery, data distribution, disk-to-disk backup and data migration. For more information about Oracle ZFS Storage 7000 appliances, see the product documentation.

Disaster Recovery Framework Script-Based Plug-Ins

The disaster recovery framework provides a generic interface module referred to as *script-based plug-ins*. With this module, you can configure custom scripts for additional replication technologies to use in a disaster recovery framework configuration. For more information about disaster recovery framework script-based plug-ins, see Chapter 13, “Script-Based Plug-Ins” in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

Resource Groups and Data Replication Components

This section provides the following information:

- "Overview of Resource Groups and Data Replication Components" on page 20
- "Replication Resource Groups" on page 20
- "Data Replication Components" on page 20

Overview of Resource Groups and Data Replication Components

Resource groups and data replication components enable the disaster recovery framework to manage data replication and takeover between clusters. For more information about configuring data replication, see the guide for the data replication product you are using:

- Chapter 1, "Replicating Data With MySQL" in *Oracle Solaris Cluster Data Replication Guide for MySQL*
- Chapter 1, "Replicating Data With Oracle Data Guard Software" in *Oracle Solaris Cluster Data Replication Guide for Oracle Data Guard*
- Chapter 2, "Administering Oracle ZFS Storage Appliance Protection Groups" in *Oracle Solaris Cluster Remote Replication Guide for Oracle ZFS Storage Appliance*
- Chapter 1, "Setting Up Oracle Solaris ZFS Snapshot Protection Groups" in *Oracle Solaris Cluster Data Replication Guide for ZFS Snapshots*

Replication Resource Groups

The disaster recovery framework extends disaster recovery framework resource management features to integrate the data replication products. When you configure a protection group, the disaster recovery framework creates replication resource groups for monitoring and controlling data replication.

Data Replication Components

A data replication component is the replication object that disaster recovery framework monitors and manages. The data replication component is added to a disaster recovery framework protection group, which can then start, stop, switch over, or take over the replication component. The disaster recovery framework coordinates these actions between the clusters and the applications being protected.

Types of data replication components include Oracle Data Guard broker configurationsand Oracle ZFS Storage Appliance replicated projects. For information about configuring data replication components in disaster recovery framework configurations, see the guide for the data replication product you are using.

Cluster Partnerships

A partnership establishes heartbeat monitoring between two clusters that are running the disaster recovery framework. Clusters in a partnership exchange heartbeats to monitor each other's presence and health. You can configure a partnership between only two clusters, and only one partnership can be defined between the clusters.

If the partnership uses either application-based replication, such as Oracle Data Guard, or no replication, a member of the partnership can be either a global cluster or a zone cluster. Such a partnership can consist of two clusters of the same type or consist of one global cluster and one zone cluster.

If the partnership uses storage-based replication, only global clusters can be members of the partnership.

The two clusters must have an Internet connection with each other. A partnership establishes heartbeats between the clusters.

The disaster recovery framework uses the IP interconnect between partner clusters for management as well as heartbeats. For additional security when a public network is in use, use `IPsec` to secure the IP interconnect.

The disaster recovery framework enables you to specify a command to execute when a heartbeat-loss notification is issued. This command is executed with root permissions. You can also specify a list of email addresses to notify when a heartbeat-loss notification has been issued.

The following figure illustrates a partnership between two clusters.

FIGURE 1 Partnerships Between Clusters

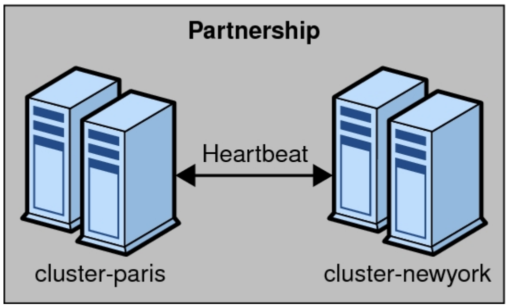

A single cluster can participate in more than one partnership with other clusters, but two clusters cannot be in more than one partnership with each other.

Protection Groups

This section provides the following information:

- "Overview of Protection Groups" on page 22
- "Relationship Between Partnerships and Protection Groups" on page 23
- "Protection Group Status" on page 24
- "Application Resource Groups" on page 24

Overview of Protection Groups

Protection groups enable a set of clusters to tolerate and recover from disaster by managing the resource groups for services. Protection groups can exist only in a partnership. You must create a partnership before you can create a protection group for that partnership. One partner cluster is the primary cluster of the protection group, and the other partner cluster is the secondary cluster. A protection group contains application resource groups and properties for managing data replication for those application resource groups. You must duplicate the application resource group configuration on partner clusters. The configuration for a protection group is identical on partner clusters, so partner clusters must have the application resource groups of the protection group defined in their configuration. The disaster recovery framework propagates protection group configurations between partners.

You can specify a data replication type in the protection group to indicate the mechanism that is used for data replication between partner clusters. A protection group supports only one data replication type. A protection group can manage one or more application resource groups. When a service is protected from disaster by data replication, the protection group also contains replication resource groups. Protection groups link an application in a resource group with the application data that should be replicated. This linkage and replication enable the application to fail over seamlessly from one cluster to another cluster.

Relationship Between Partnerships and Protection Groups

Clusters in a protection group must be defined as partners. Protection groups require a partnership that defines the clusters that can host the protection group. A cluster can be defined in more than one protection group, and the cluster can have a different role in each protection group. For example, the primary cluster of one protection group can also be the secondary cluster of another protection group. A partnership can have any number of protection groups.

The following figure illustrates two clusters that are defined in one cluster partnership and two protection groups.

FIGURE 2 Example Configuration of Two Clusters in Protection Groups

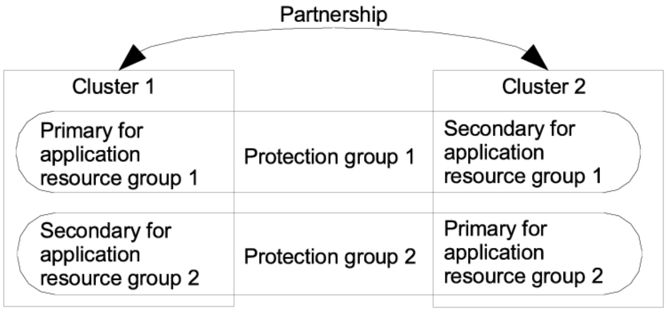
Partnership
Cluster 1
Cluster 2
Primary for application resource group 1
Protection group 1
Secondary for application resource group 1
Secondary for application resource group 2
Protection group 2
Primary for application resource group 2

The following figure illustrates three clusters that are defined in two cluster partnerships and two protection groups.

FIGURE 3 Example Configuration of Three Clusters in Protection Groups

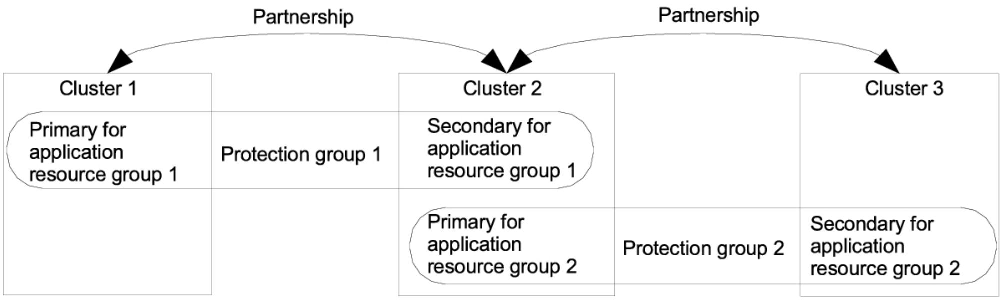
Partnership
Partnership
Cluster 1
Cluster 2
Cluster 3
Primary for application resource group 1
Protection group 1
Secondary for application resource group 1
Primary for application resource group 2
Protection group 2
Secondary for application resource group 2

Protection Group Status

The disaster recovery framework monitors the status of a protection group on each cluster. The framework then combines the local status of each cluster into a global view of the protection group status. The global status reflects the overall status of the protection group.

You can view the protection group status from the Oracle Solaris Cluster Manager browser interface or through the CLI.

For more information about the status of protection groups, see the *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

Application Resource Groups

To be highly available, an application must be managed as a resource in an application resource group. You can configure an application resource group for a takeover application or a scalable application. You must also configure application resources and application resource groups on both the primary cluster and the secondary cluster. The data that the application resource accesses must be replicated on the secondary cluster.

The replication for the data volumes that an application resource accesses must be in the same protection group as the applications.

Support for data replication might limit how you configure application resource groups. These requirements and limitations vary with the data replication type you choose. For more information about these requirements, see the *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

Multigroups and Sites

A *multigroup* is a set of protection groups that is preconfigured for migration, as a group, between predetermined sets of clusters, called *sites*. A switchover or takeover of all protection groups in a multigroup is performed in a single operation.

A protection group in a multigroup can be configured in a *dependency chain*, which defines which protection groups have a dependency on another protection group. A multigroup can contain multiple dependency chains. The dependency chain configuration determines in which order the protection groups must be taken offline and brought online, based their roles in the dependency chain. This ensures that interdependent services can be stopped and restarted in the correct order. For example, you might want an application that depends on a database to be stopped before the database is shut down, and after switchover the database be restarted before the application is restarted.

Heartbeat Monitoring and Plug-ins

The disaster recovery framework uses heartbeats to monitor the state between partner clusters. Heartbeats are sent over the public network to detect cluster failures at geographically separated sites. Heartbeat monitoring is part of a partnership configuration. For example, a cluster failure occurs when all nodes of a cluster shut down. Heartbeats could also be lost when a cluster loses access to the public network and no communication occurs between the partner clusters.

The heartbeat monitor uses plug-in modules to query the heartbeat status of its partners. The disaster recovery framework offers default plug-ins for monitoring through a TCP/UDP connection.

The disaster recovery framework enables you to specify a command to execute when a heartbeat-loss notification is issued. You can also specify a list of email addresses to notify when a heartbeat-loss notification has been issued.

You can use customized plug-ins to provide a data path over alternate communication links, such as email, HTTP, satellite, and microwave towers.

♦ ♦ ♦ CHAPTER 3

Disaster Recovery Framework Architecture

The disaster recovery framework enables a group of clusters to be managed and viewed as a single, large system. This chapter presents a high-level overview of the disaster recovery framework architecture, which you can use in preparation for installing, configuring, and administering the disaster recovery framework.

This chapter contains the following topics:

- "Disaster Recovery Framework Software Environment" on page 27
- "Disaster Recovery Framework Hardware Environment" on page 29
- "Data Replication Configuration" on page 29
- "Geographically Distributed Cluster Topology" on page 32
- "Three-Data-Center (3DC) Topologies" on page 35
- "Multigroups and Cluster Sites" on page 35

Disaster Recovery Framework Software Environment

The disaster recovery framework provides tools for managing geographically separated clusters. The disaster recovery framework also maintains availability of services between clusters by utilizing disaster recovery framework resource management features.

The following software components form a disaster recovery framework cluster:

- Oracle Solaris software
- Oracle Solaris Cluster software
- Disaster recovery framework software
- Application data service agents
- Data replication software
- Solaris Volume Manager software

The following figure illustrates how components of the disaster recovery configuration interrelate.

FIGURE 4 Overview of Disaster Recovery Framework Software Architecture

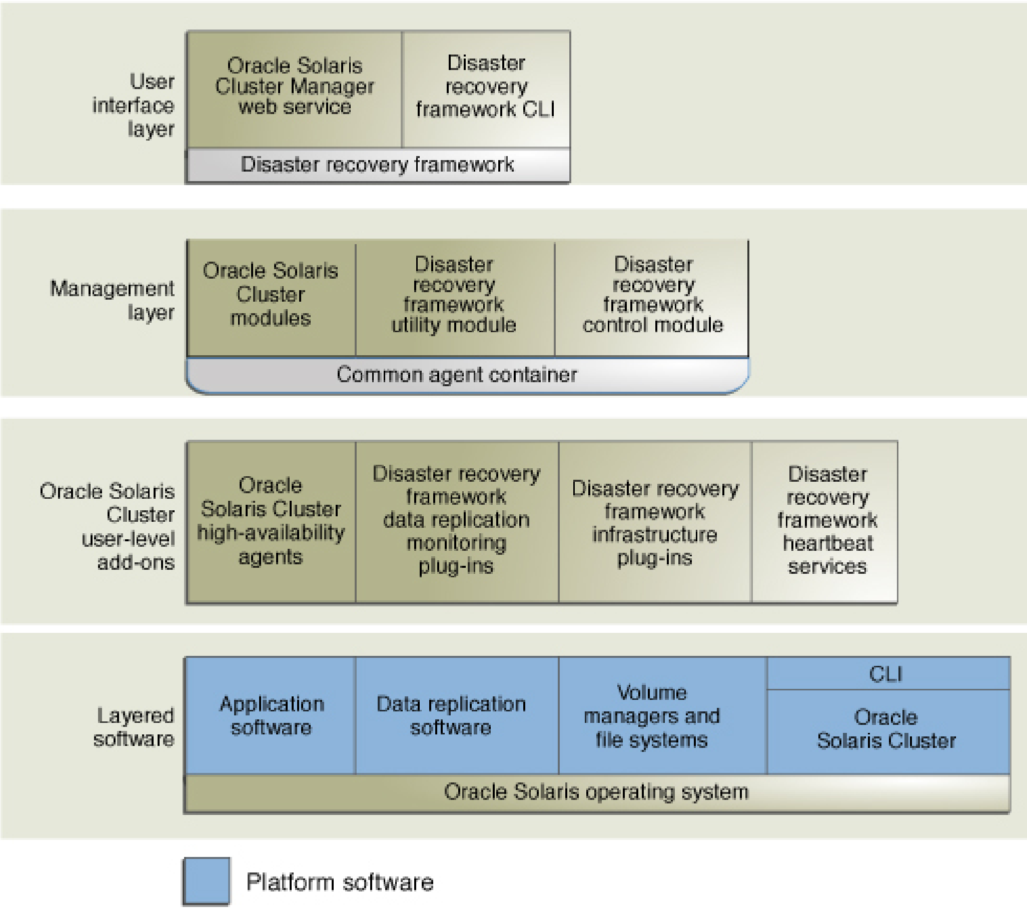
User interface layer
Oracle Solaris Cluster Manager web service
Disaster recovery framework CLI
Disaster recovery framework
Management layer
Oracle Solaris Cluster modules
Disaster recovery framework utility module
Disaster recovery framework control module
Common agent container
Oracle Solaris Cluster user-level add-ons
Oracle Solaris Cluster high-availability agents
Disaster recovery framework data replication monitoring plug-ins
Disaster recovery framework infrastructure plug-ins
Disaster recovery framework heartbeat services
Layered software
Application software
Data replication software
Volume managers and file systems
CLI
Oracle Solaris Cluster
Oracle Solaris operating system
Platform software

The clusters in a disaster recovery framework configuration can run different versions of Oracle Solaris Cluster software, as long as the versions are no more than one consecutive release different. For example, the clusters in one disaster recovery framework configuration could run either version 4.2 or 4.3, or run either version 4.2 or 4.1. But you cannot have clusters running either version 4.1 or 4.3 in the same disaster recovery framework configuration.

You can install and remove the disaster recovery framework software independently of the underlying Oracle Solaris Cluster installation. The installation and the uninstallation processes do not require an additional node reboot or cluster downtime.

Disaster Recovery Framework Hardware Environment

The disaster recovery framework hardware configuration is the basis for a disaster recovery framework cluster.

The following additional hardware components form a disaster recovery framework cluster:

- Disaster recovery framework hardware installations, with attached data storage
- Internet connections for inter-cluster management communication between the Oracle Solaris Cluster installations
- Internet connections for inter-cluster heartbeats
- Connections for data replication
- Connections for custom heartbeats

The disaster recovery framework hardware environment supports the following topologies:

- N+1, where multiple clusters that are located at multiple sites are communicating with a single backup cluster
- Cluster pair, where both clusters are online and providing services

Figure 6, "Data Replication From a Two-Node Cluster to a Two-Node Cluster," on page 32 illustrates a high-level view of the disaster recovery framework hardware architecture.

Data Replication Configuration

The disaster recovery framework does not limit the distance between partner clusters. Partner clusters require data replication connections to support the protection groups that are hosted by the partnership. Partner clusters must be compatibly configured to support data replication between the clusters.

The disaster recovery framework supports replication from a single-node cluster to a single-node cluster, from a multinode cluster to a single-node cluster, and from a multinode cluster to a multinode cluster.

While you can use a single-node cluster at both the primary and backup sites, a single-node cluster offers no internal redundancy. To ensure no single point of failure, you must have a minimum of two nodes in a cluster at the primary site. Use a single-node cluster at the secondary site as a cost effective backup solution if the secondary site is used only for backup purposes. Do not use single-node cluster to run mission critical applications.

Primary clusters and secondary clusters can have any configuration that is supported by the disaster recovery framework if the data replication characteristics of these clusters are compatible. The level of compatibility varies with each data replication product.

The following requirements determine the data replication connection:

- The distance between the partner clusters
- The amount of data that is being replicated
- Data replication configuration parameters

The disaster recovery framework enables you to balance between data consistency and the cost of the connection, where data consistency is the acceptable amount of data loss.

The following figure illustrates a data replication configuration from a two-node cluster to a single-node cluster.

FIGURE 5 Data Replication From a Two-Node Cluster to a Single-Node Cluster

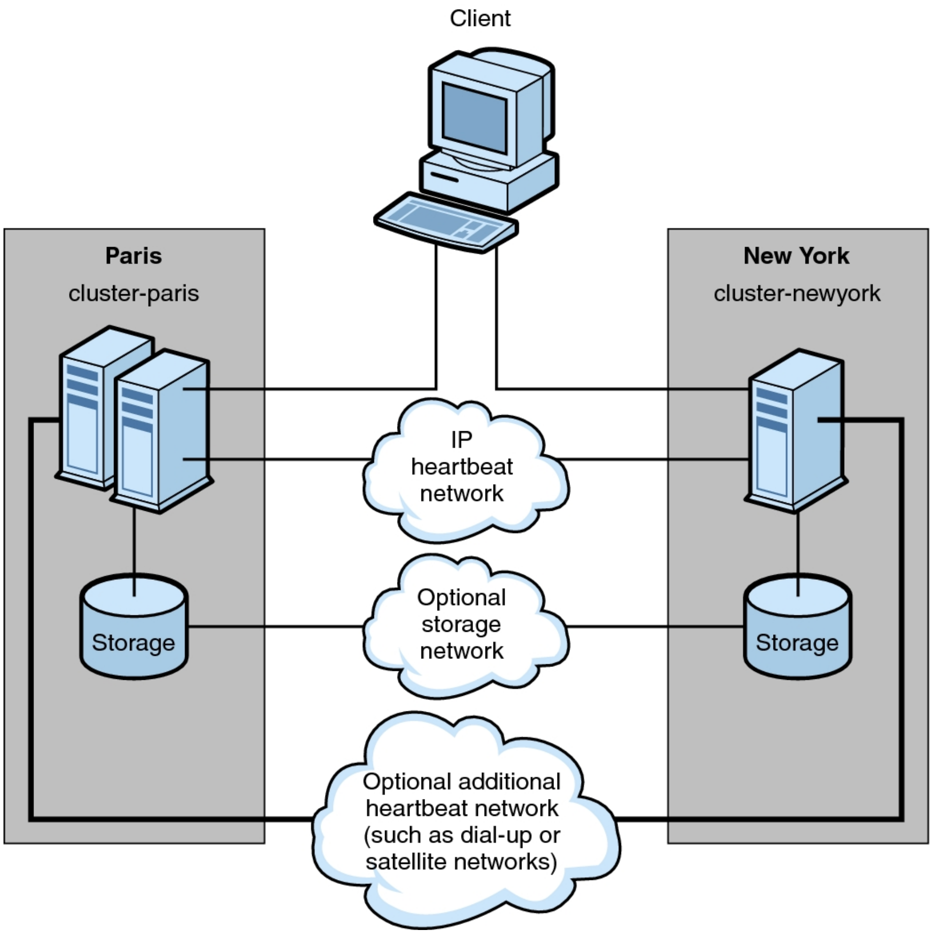

The following figure illustrates a data replication configuration from a two-node cluster to a two-node cluster.

FIGURE 6 Data Replication From a Two-Node Cluster to a Two-Node Cluster

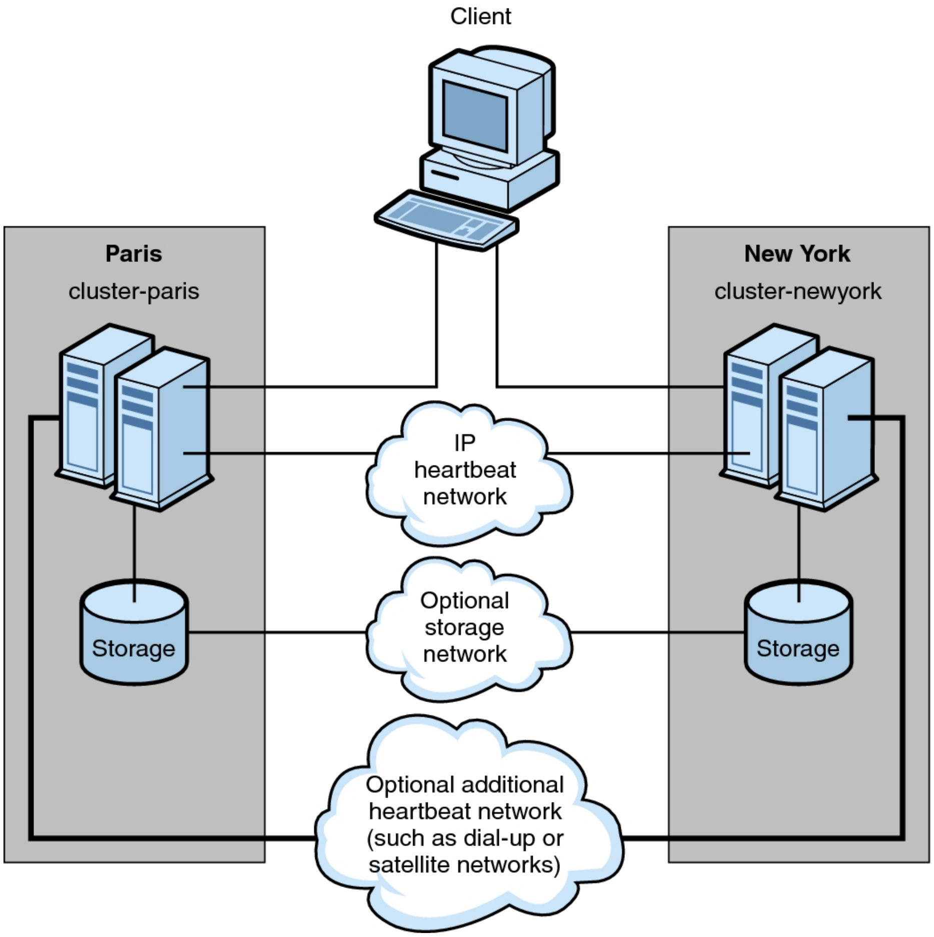

Geographically Distributed Cluster Topology

A partnership establishes communication and a heartbeat between clusters. One cluster can participate in several partnerships. A protection group establishes data replication between partner clusters. You can configure several protection groups for a partnership, with each protection group replicating different data between the partner clusters.

The following figure illustrates a geographically distributed topology that demonstrates intercluster relationships.

FIGURE 7 Geographically Distributed Topology

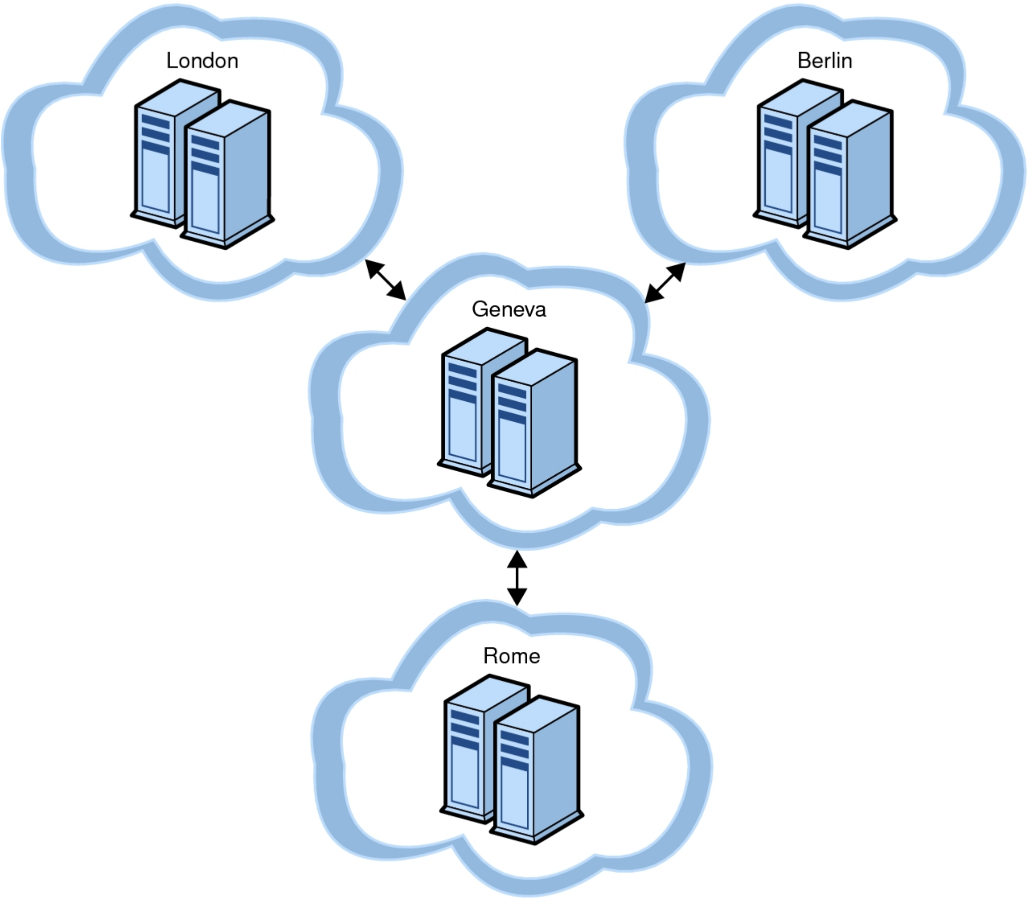

The disaster recovery framework enables you to configure multiple partnerships for a cluster with heartbeats between partner clusters. For example, the Geneva-London-Rome-Berlin topology contains a central cluster in Geneva that forms three separate partnerships with clusters in London, Rome, and Berlin. The partnerships require two-way Internet connections between the following cluster pairs: London and Geneva, Rome and Geneva, and Berlin and Geneva. These partnerships enable the cluster in Geneva to detect failures of the clusters in London, Berlin, and Rome by exchanging heartbeats.

Each partnership has a protection group so that the primary clusters in London, Rome, and Berlin can replicate data to the cluster in Geneva as a secondary cluster.

The following figure illustrates a geographically distributed topology that demonstrates intercluster relationships.

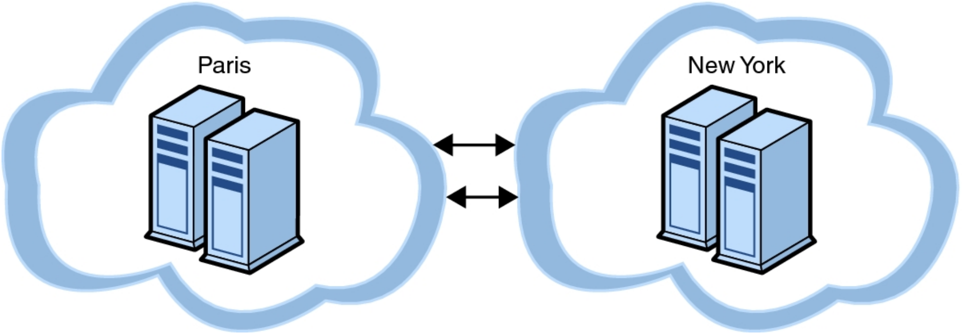

The Paris-New York topology has two clusters that form a partnership with two protection groups. Each cluster is the primary cluster for one protection group, and the secondary cluster for the other protection group. The partnership requires a two-way Internet connection between the two clusters for intercluster management and heartbeats. The two clusters must have a data replication link to support data replication for two protection groups.

In the Geneva-London-Rome-Berlin topology, the cluster in Geneva could become the primary cluster for any of the three protection groups. However, the cluster in Geneva must have adequate provisioning to run all the services that are offered by the application resource groups.

For example, if the cluster in Rome is shut down for maintenance, the cluster in Geneva could be the new primary cluster by using a controlled switchover for the Rome-Geneva protection group. As the new primary cluster for the Rome-Geneva protection group, the cluster in Geneva would host the services that are provided by the application resource groups of the Rome-Geneva protection group. The cluster in Geneva would simultaneously serve as the secondary cluster for the clusters in London and Berlin.

Similarly, in the Paris-New York topology, either cluster could take over as the primary cluster to both protection groups if the partner cluster were unexpectedly lost.

Three-Data-Center (3DC) Topologies

You can use a campus cluster for the primary cluster and a secondary cluster for disaster recovery, which creates a three-site or three–data–center (3DC) configuration of primary, backup, and disaster recovery sites.

There are two ways to perform 3DC replication:

- Use volume-manager mirroring or ZFS pool-level mirroring within the campus cluster. Use application-based replication, such as Oracle Data Guard, MySQL, or ZFS snapshots software replication, to the disaster recovery site.
- Use synchronous storage-based replication within the campus cluster . Use application-based replication, such as Oracle Data Guard, MySQL, or ZFS snapshots software to replicate the application from the campus cluster to a disaster recovery site.

You cannot create a daisy chain of disaster recovery framework pairs (for example, London-Paris-Rome) for the same managed application.

Multigroups and Cluster Sites

The disaster recovery framework enables you to configure predetermined sets of protection groups and target clusters on which to perform start, stop, switchover. or takeover operations in a single operation.

The following figure illustrates a group of clusters that are defined in sites and a multigroup defined for protection groups between two of the sites.

FIGURE 8 Example Configuration of a Multigroup

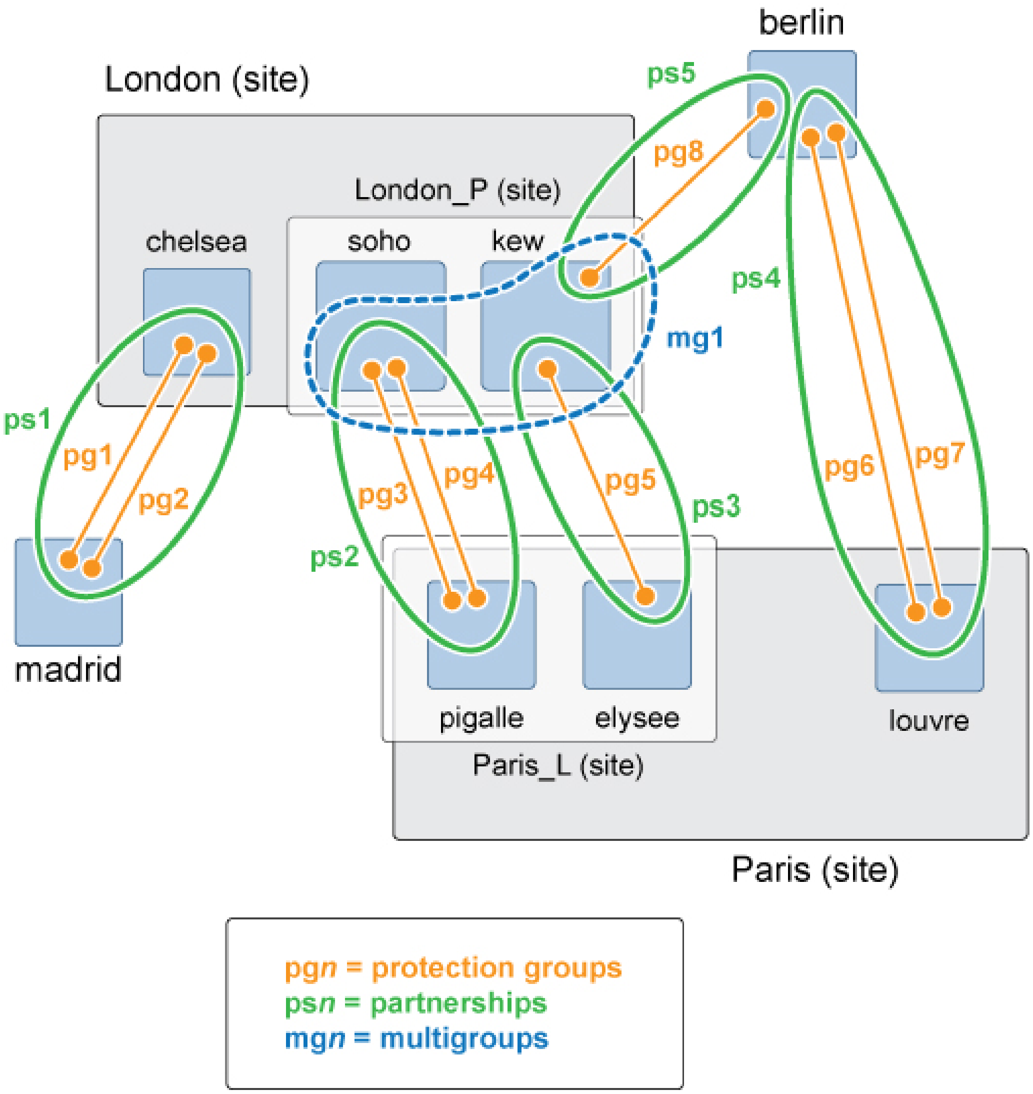

Glossary

active–active clusters Two clusters where each cluster is both the primary cluster for some services and the secondary cluster for other services.

application resource An application that is managed as a resource to its increased availability.

application resource group An Oracle Solaris Cluster resource group that is configured by the user to make an application highly available in an Oracle Solaris Cluster configuration. An application resource group can be configured into a protection group to make it disaster tolerant and highly available.

campus cluster A cluster that supports data replication between geographically separate nodes within *one* cluster. The maximum distance between nodes is limited.

data replication The copying of data from data storage devices in a primary cluster to data storage devices in a secondary cluster. Through data replication, the secondary cluster has a recent copy of the data on the primary cluster. The primary and secondary clusters can be geographically separated.

In a campus cluster, the two data storage devices are on the same cluster. In a geographically separated cluster that runs the disaster recovery framework, the two data storage devices are on different clusters.

data replication component A data replication component is the replication object that the Geographic Edition framework monitors and manages.

data replication resource An Oracle Solaris Cluster resource that monitors the state and status of data replication.

disaster tolerance The ability of a system to restore an application on a secondary cluster when the primary cluster fails. Disaster tolerance is based on data replication and failover.

disconnected partnership An error scenario in which two clusters in a protection group act as the primary cluster. In a disconnected partnership, the system administrator must execute a takeover, making one cluster the primary cluster and the other cluster the secondary cluster.

heartbeat A signal that is emitted from a cluster and detected by its partner cluster. Heartbeats enable a cluster to monitor the presence and failure of its partner cluster.

inactive cluster A primary cluster where an application is not running and data is not being replicated to a secondary cluster. Alternatively, a secondary cluster where data is not being replicated from the primary cluster.

multigroup A set of protection groups that is preconfigured for migration as a group between predetermined sets of clusters, called sites.

partnership A relationship between two geographically separated clusters that are installed with the Oracle Solaris Cluster software and disaster recovery framework software. These two clusters monitor each other's health by exchanging heartbeats.

primary cluster A cluster that is in a cluster partnership, that hosts the application resources, and that holds the primary copy of replicated data. Protection groups define whether a cluster is primary or secondary. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.

protection group An entity that manages application resource groups for services that are protected from disaster. Clusters in a protection group must be defined as partners. A cluster can have different roles in different protection groups. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.

A protection group has the following characteristics:

- A set of resource groups and resources for services that are protected from disaster
- Device group entries
- A primary cluster that hosts the protection group
- A secondary cluster that is able to host the protection group
- A data replication service

replication resource group A resource group that contains data replication resources.

resource An Oracle Solaris Cluster resource.

resource group An Oracle Solaris Cluster resource group. A resource group can be an application resource group or a replication resource group.

scalable application An application that runs on several nodes of one cluster to create a single, logical service. If a node that is running a scalable application fails, failover does not occur. The application continues to run on the other nodes of the cluster.

secondary cluster A cluster in a cluster partnership that is capable of hosting a protection group. The secondary cluster receives mirrored data from the primary cluster. If the primary cluster fails, the secondary cluster can become the new primary cluster.

A secondary cluster can be associated with a protection group. If a primary cluster fails, the protection group is migrated to a secondary cluster. Protection groups define whether a cluster is primary or secondary. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.

secondary node A node that is in a cluster, but does not host the application services. If the primary node fails, the secondary node becomes the new primary node.

site A location that houses one or more clusters that run the disaster recovery framework. To participate in the disaster recovery environment, a cluster must have a partner cluster on a geographically separated site.

standby cluster A cluster with a minimal configuration that acts as a secondary cluster. A standby cluster can take over from the primary cluster in an emergency situation, but supports a reduced service only. A standby cluster is a low-cost alternative to a secondary cluster.

switchover The planned migration of services from the primary cluster to the secondary cluster.

Unlike a takeover, the primary cluster is connected to the secondary cluster during a switchover. During a switchover, the primary cluster is connected to the secondary cluster and coordinates the migration of services with the secondary cluster. This coordination enables the data replication to be completed and ensures that services can be transferred from the primary cluster to the secondary cluster with minimal loss or corruption of data.

takeover The emergency migration of services from the primary cluster to the secondary cluster. A system administrator can initiate a takeover to recover from a disaster scenario.

Unlike a switchover, the primary cluster is not connected to the secondary cluster during a takeover. Therefore, the primary cluster cannot coordinate with the secondary cluster to migrate the services. Because of this lack of coordination, the risk of data loss and data corruption is higher than with a switchover. During a takeover, dedicated recovery procedures are used to minimize data loss and data corruption.

Index