Oracle® Fail Safe

Concepts and Administration Guide

Release 3.4.2 for Microsoft Windows

E14973-01

October 2009

ORACLE®

Primary Author: Ashmita Bose

Contributing Author: Janelle Simmons

Contributor: Paul Mead

Contents

3   Designing an Oracle Fail Safe Solution

4   Management for High Availability

5   The FSCMD Command-Line Interface

6   Troubleshooting Tools

7 Configuring Single-Instance Databases for High Availability and Disaster Tolerance

8   Configuring Generic Services for High Availability

9   Configuring Oracle Management Agent for High Availability

10   Configuring Oracle Application Server Components for High Availability

A   Network Configuration Considerations

B   Contacting Oracle Support Services

Glossary

Index

List of Figures

List of Tables

x

Preface

This guide describes how to use Oracle Fail Safe running on a Microsoft cluster system to configure the following for high availability:

- Oracle single-instance databases
- Applications installed as Windows generic services
- Oracle Management Agent
- Oracle Application Server components

Audience

This guide is intended for anyone who is interested in how Oracle Fail Safe minimizes downtime for software components running on a Microsoft cluster.

Readers should be familiar with Failover Clusters, Oracle Net networking, and the applications for which they want to provide high availability.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at http://www.oracle.com/accessibility/.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at `http://www.fcc.gov/cgb/consumerfacts/trs.html`, and a list of phone numbers is available at `http://www.fcc.gov/cgb/dro/trsphonebk.html`.

Related Documents

Refer to the following documentation for more information about Oracle Fail Safe:

- For more information about updates to the software, access to online documentation, and other release-specific information, see *Oracle Fail Safe Release Notes*.

- For installation, deinstallation, and upgrade instructions, see *Oracle Fail Safe Installation Guide*.

- For online assistance, Oracle Fail Safe Manager provides help topics online. To access the online help topics, click **Help** on the menu bar in Oracle Fail Safe Manager.

- For more information about Oracle Call Interface, see *Oracle Call Interface Programmer's Guide*.

- For more information about ODBC, see Microsoft ODBC documentation.

For more information about other related products, see the documentation for those products.

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.
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

Introduction to Oracle Fail Safe

Increasingly, businesses expect products and services to be available 24 hours a day, 365 days a year. While no solution can ensure 100% **availability,** Oracle Fail Safe minimizes the **downtime** of Oracle databases and other applications running on Microsoft clusters and configured with **Microsoft Cluster Services (MSCS)**.

This chapter discusses the following topics:

- What Is Oracle Fail Safe?

- Benefits of Oracle Fail Safe

- A Typical Oracle Fail Safe Configuration

- Deploying Oracle Fail Safe Solutions

1.1 What Is Oracle Fail Safe?

Oracle Fail Safe is a user-friendly software that works with Microsoft Cluster Server (MSCS) to provide highly available business solutions on Microsoft clusters. A **cluster** is a configuration of two or more Microsoft Windows systems that makes them appear to network users as a single, highly available system. Each system in a cluster is referred to as a **cluster node**.

Oracle Fail Safe works with MSCS cluster software to provide high availability for applications and single-instance databases running on a cluster. When a cluster node fails, the cluster software moves its workload to the surviving node based on parameters that you configure using Oracle Fail Safe. This operation is called a **failover**.

With Oracle Fail Safe, you can reduce downtime for single-instance Oracle databases and almost any application that can be configured as a Microsoft Windows service.

Oracle Fail Safe consists of Oracle Services for MSCS and Oracle Fail Safe Manager:

- Oracle Services for MSCS works with the MSCS software to configure fast, automatic failover during planned and unplanned outages for resources configured for high availability. These **resources** can be the Oracle database, or other Microsoft Windows services (also the software and hardware upon which these items depend). Also, Oracle Services for MSCS can attempt to restart a failed software resource so that a failover from one cluster node to another may not be required.

> **Note:** Oracle Services for MSCS was referred to as Oracle Fail Safe Server in previous releases of Oracle Fail Safe.

- Oracle Fail Safe Manager provides a user-friendly interface and wizards that help you to configure and manage cluster resources, and troubleshooting tools that help you to diagnose problems.

Together, these components enable rapid deployment of highly available database, application, and Internet business solutions.

1.2 Benefits of Oracle Fail Safe

Oracle Fail Safe provides the key benefits discussed in the following sections:

- Highly Available Resources and Applications
- Ease of Use
- Ease of Integration with Applications

1.2.1 Highly Available Resources and Applications

Oracle Fail Safe works with MSCS to configure both hardware and software resources for high availability. Once configured, the multiple nodes in the cluster appear to end users and clients as a single *virtual server*; end users and **client applications** connect to a single, fixed network address, called a **virtual address**, without requiring any knowledge of the underlying cluster. If one node in the cluster becomes unavailable, then MSCS moves the workload of the failed node (and client requests) to another node.

For example, the left side of Figure 1–1 shows a two-node cluster configuration where both nodes are available and actively processing transactions. On the surface, this configuration may seem no different from setting up two independent servers, except that the storage subsystem is configured so that the disks are connected physically to both nodes by a **shared storage interconnect**. Although both nodes are physically connected to the same disks, MSCS ensures that each disk can be owned and accessed by only one node at a time.

The right side of Figure 1–1 shows how, when hardware or software becomes unavailable on one node, its workload automatically moves (*fails over*) to the surviving node and is restarted, without administrator intervention. During the failover, ownership of the cluster disks is released from the failed server (Node A) and acquired by the surviving server (Node B). If a single-instance Oracle database was running on Node A, then Oracle Fail Safe restarts the database instance on Node B. Clients then can access the database through Node B using the same virtual address that they used to access the database when it was hosted by Node A.

Figure 1–1   Failover with Oracle Fail Safe in a Microsoft Cluster

1.2.2 Ease of Use

Because of the numerous hardware and software components involved, configuring software and all of its dependent components (for example, disks, IP addresses, network) to work in a cluster can be a complex process. In contrast, Oracle Fail Safe is designed to be easy to install, administer, and use and simplifies configuration of software in a cluster.

Installation: Using Oracle Universal Installer, you can install Oracle Fail Safe either interactively or in silent mode. With the **silent mode** installation method, you install software by supplying input to Oracle Universal Installer with a response file. Also, you can perform **rolling upgrade**s of both the operating system and application software. Rolling upgrades minimize downtime by allowing one cluster node to continue hosting the cluster workload while the other system is being upgraded. See *Oracle Fail Safe Installation Guide* for more information.

Administration and Use: Oracle Fail Safe Manager provides a user-friendly interface to set up, configure, and manage applications and databases on the cluster. Oracle Fail Safe Manager provides wizards that automate the configuration process and ensure that the configuration is replicated consistently across cluster nodes.

Oracle Fail Safe Manager includes:

- A tree view of objects that displays multiple views of the same data to help you find information efficiently

- Wizards that automate and simplify resource configuration, and drag-and-drop capabilities that help you quickly perform routine system maintenance, such as moving resources across nodes to balance the workload

- An integrated family of verification tools that automatically diagnose and fix common configuration problems both before and after configuration

- Online documentation, including a tutorial, help, and manuals available in HTML and PDF formats

■ A command-line interface (FSCMD) for managing the cluster through batch programs or scripts

Figure 1–2 shows an Oracle Fail Safe Manager window. The left pane displays a tree view showing multiple views (and the current state) of clusters and cluster resources. The right pane displays a property page that lists all groups on the cluster that have been selected from the tree view and the current state of those groups. Depending on the object chosen from the tree view, the display in the right pane changes. When you select a particular cluster, node, group, or resource, the property sheet for that cluster, node, group, or resource is displayed.

Figure 1–2   Oracle Fail Safe Manager
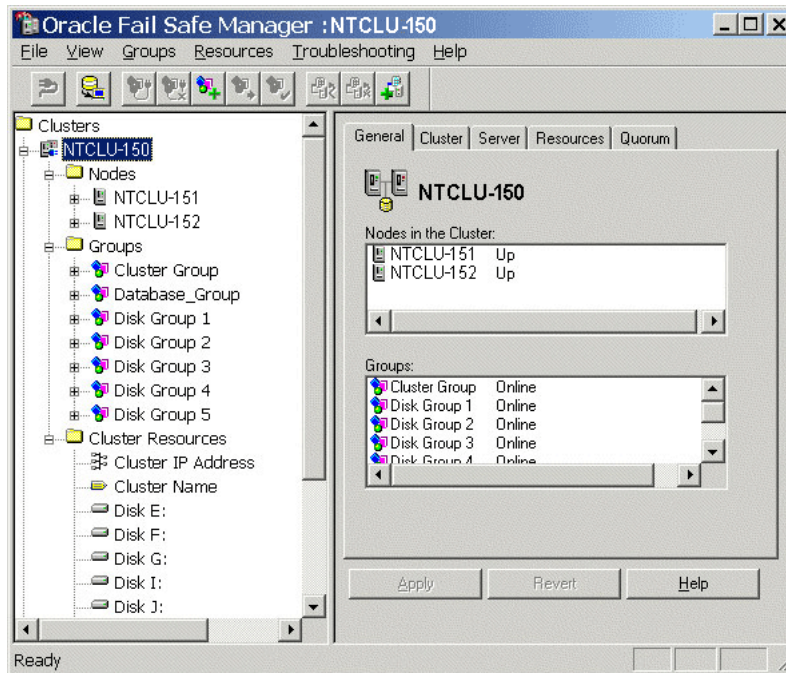
Figure 1–3 shows the Oracle Fail Safe menus and the items within each menu.

Figure 1–3   Oracle Fail Safe Manager Menus and Contents
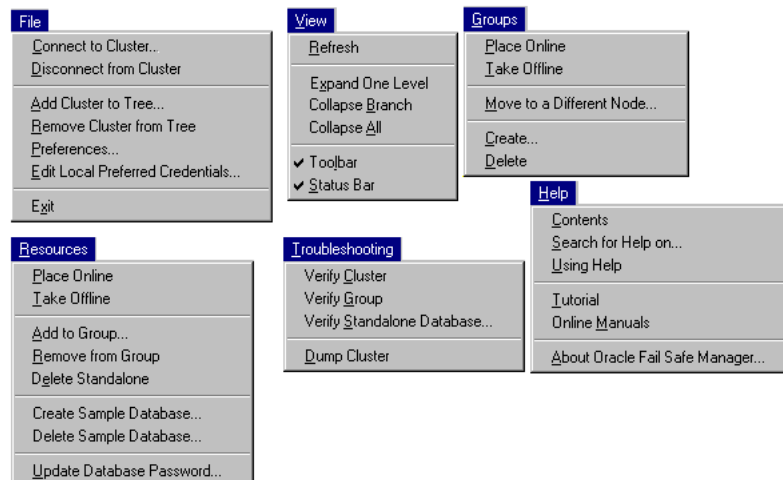
1.2.3 Product Accessibility

Oracle Fail Safe has two user interfaces: the FSCMD Command-Line Interface and the Oracle Fail Safe Manager GUI. However, the Oracle Fail Safe Manager GUI is used more widely. The Oracle Fail Safe Manager GUI presents two panels: a navigation tree in the left panel and the right panel displaying its corresponding selection. Most of the time the right panel displays a tabbed set of property pages for the selection made in the navigation tree. At other times it displays a list of objects. Wizard pages are displayed when the user selects an action that requires multiple steps, such as adding a resource to a group.

The Oracle Fail Safe Manager uses specific keyboard sequences for navigation, that can be used instead of a pointer device. Oracle Fail Safe can be easily accessed in the following ways:

- In the main user interface window use CTRL-T to navigate between the tree view and the multi-page (tabbed) property page.

- Branches in the tree view are expanded using keyboard sequence ALT+V X and closed by using ALT+V B.

1.2.4 Ease of Integration with Applications

If you want to configure an existing application to access databases or other applications configured with Oracle Fail Safe, then few or no changes are required. Because applications always access cluster resources at the same virtual address, applications treat failover as a quick node restart.

After a failover occurs, database clients or users must reconnect and replay any transactions that were left undone (such as database transactions that were rolled back during instance recovery). Applications developed with OCI (including ODBC clients that use the Oracle ODBC driver) can take advantage of automatic reconnection after failover. See Section 7.9 for more information.

1.3 A Typical Oracle Fail Safe Configuration

Oracle Fail Safe solutions can be deployed on any Microsoft Windows cluster certified by Microsoft for configuration with MSCS.

Most clusters are configured similarly, differing only in choice of storage interconnect (SCSI, Fibre Channel, or SAN) and in the way applications are deployed across the cluster nodes.

A typical cluster configuration includes the following hardware and software:

- Hardware
  - Microsoft cluster nodes, each with one or more local (private) disks where executable application files are installed.
  - Private (heartbeat) interconnect between the nodes for intracluster communications.
  - Public interconnect (Internet, Intranet, or both) to the local area network (LAN) or wide area network (WAN).
  - NTFS formatted disks on the **shared storage interconnect** (SCSI, Fibre Channel, or SAN). All **data file**s, log files, and other files that must fail over from one node to another are located on these cluster disks.

> **Note:** See the documentation for your cluster hardware for information about using redundant hardware, such as RAID, to further ensure high availability.

- Additional **redundant components** (UPS, network cards, disk controllers, and so on).

■ Software (installed on each node)

- Microsoft Windows

- Oracle Services for MSCS

- Oracle Fail Safe Manager (installed on one or more cluster nodes, one or more client workstations, or both)

- One or more of the following resources that are highly available, such as:

  * Oracle single-instance databases

  * Oracle Management Agent

  * Oracle applications or third-party applications that can be configured as Windows generic services

See *Oracle Fail Safe Release Notes* for information about the supported releases of these components.

Figure 1–4 shows the hardware and software components in a two-node cluster configured with Oracle Fail Safe. Note that the executable application files are installed on a private disk on each cluster node and the application data and log files reside on a shared cluster disk.

Figure 1–4   Hardware and Software Components Configured with Oracle Fail Safe
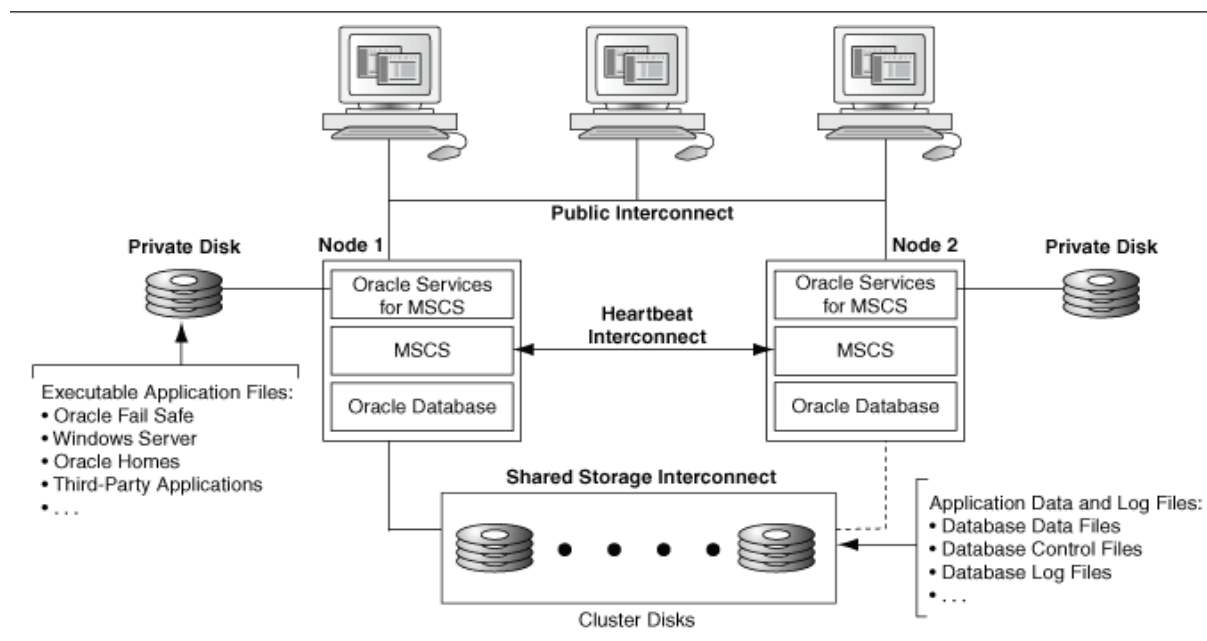
1.4 Deploying Oracle Fail Safe Solutions

Oracle Fail Safe works with MSCS to configure resources running on a cluster, to provide fast failover, and to minimize downtime during planned (system upgrades) and unplanned (hardware or software failure) outages.

Clusters provide high availability by managing:

- Unplanned group failover

  Clusters manage **unplanned group failover**s (failure of hardware or software components) in a way that is transparent to users. When one node on the cluster becomes unavailable, another node temporarily serves both its own workload and the workload from the failed node. When a resource fails and cannot be restarted on the current node, another node takes ownership of that resource (and any other resources upon which it depends) and attempts to restart it.

- Planned failover

  Clusters manage **planned group failover**s (those which you intentionally start, such as when you upgrade software on the cluster). You can fail over the resources to another node, perform a software or hardware upgrade, and then return the resources to the original node. (This is called failing back the resources.) Then, perform the same upgrade process on the other nodes in the cluster.

Oracle Fail Safe also ensures efficient use of resources in the cluster environment by managing the following:

- Independent workloads

  The cluster nodes can serve separate workloads. For example, one node can host an Oracle database, and the others can host applications.

- Load balancing

  You can balance resources across the cluster nodes. For example, a database can be moved from a node that is heavily loaded to one that has spare capacity.

Oracle Fail Safe has a variety of deployment options to satisfy a wide range of failover requirements. Chapter 3 explains how to configure an Oracle Fail Safe solution for your business needs, including active/passive solutions and active/active solutions.

2

Cluster Concepts

Oracle Fail Safe high-availability solutions use Microsoft cluster hardware and
Microsoft Cluster Server (MSCS) software.

- A Microsoft **cluster** is a configuration of two or more independent computing
  systems (called nodes) that are connected to the same disk subsystem.

- **Microsoft Cluster Services (MSCS)** software, included with Microsoft Windows
  software, enables you to configure, monitor, and control applications and
  hardware components (called resources) that are deployed on a Windows cluster.

To take advantage of the high-availability options that Oracle Fail Safe offers, you
must understand MSCS concepts.

This chapter discusses the following topics:

- Cluster Technology
- Resources, Groups, and High Availability
- Groups, Virtual Addresses, and Virtual Servers
- Allocating IP Addresses for Virtual Addresses
- Cluster Group and Cluster Alias
- Failover
- Failback

2.1 Cluster Technology

The Windows systems that are members of a cluster are called **cluster node**s. The
cluster nodes are joined together through a public shared storage interconnect as well
as a private internode network connection.

The internode network connection, sometimes referred to as a heartbeat connection,
allows one node to detect the availability of another node. Typically, a **private
interconnect** (that is distinct from the public network connection used for user and
client application access) is used for this communication. If one node fails, then the
cluster software immediately fails over the workload of the unavailable node to an
available node, and remounts on the available node any cluster resources that were
owned by the failed node. Clients continue to access cluster resources without any
changes.

Figure 2–1 shows the network connections in a two-node Microsoft cluster
configuration.

Figure 2–1   Microsoft Cluster System
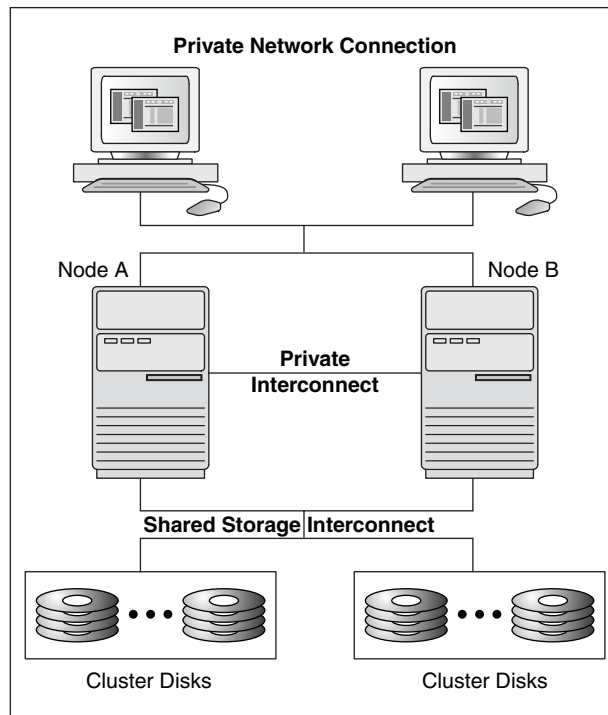
2.1.1 How Clusters Provide High Availability

Until cluster technology became available, **reliability** for PC systems was attained by hardware redundancy such as RAID and mirrored drives, and dual power supplies. Although disk redundancy is important in creating a highly available system, this method alone cannot ensure the availability of your system and its applications.

By connecting servers in a Windows cluster with MSCS software, you provide server redundancy, with each server (node) having exclusive access to a subset of the cluster disks during normal operations. A cluster is far more effective than independent standalone systems, because each node can perform useful work, yet still is able to take over the workload and disk resources of a failed cluster node.

By design, a cluster provides high availability by managing component failures and supporting the addition and subtraction of components in a way that is transparent to users. Additional benefits include providing services such as failure detection, recovery, and the ability to manage the cluster nodes as a single system.

> **Note:**   See your hardware documentation for information about using redundant hardware, such as RAID technology, to increase high availability.

2.1.2 System-Level Configuration

There are different ways to set up and use a cluster configuration. Oracle Fail Safe supports the following configurations:

- Active/passive configurations
- Active/active configurations

See Chapter 3 for information about these configurations.

2.1.3 Disk-Level Configuration

When an MSCS cluster is recovering from a **failure**, a surviving node gains access to the failed node's disk data through a shared-nothing configuration.

In a **shared-nothing configuration,** all nodes are cabled physically to the same disks, but only one node can access a given disk at a time. Even though all nodes are physically connected to the disks, only the node that owns the disks can access them.

Figure 2–2 shows that if a node in a two-node cluster becomes unavailable, then the other cluster node can assume ownership of the disks and application workloads that were owned by the failed node and continue processing operations for both nodes.

Figure 2–2   Shared-Nothing Configuration
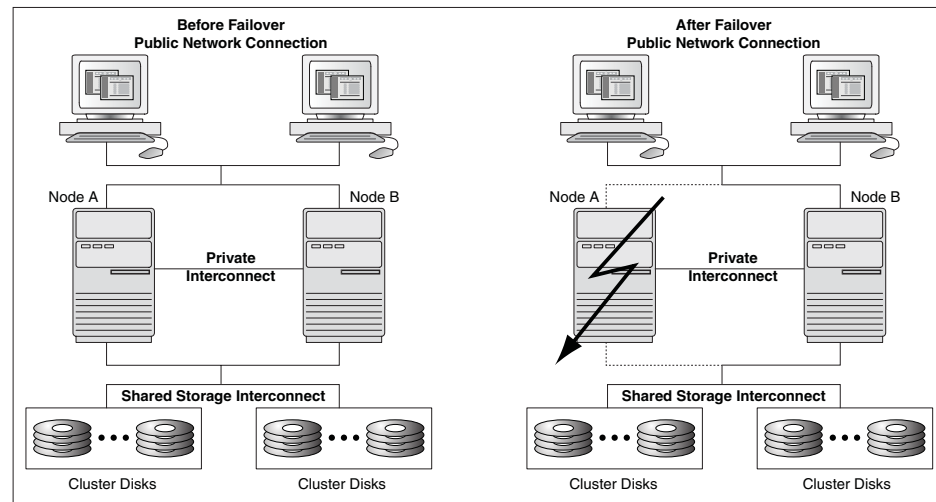
2.2 Resources, Groups, and High Availability

When a server node becomes unavailable, its cluster *resources* (for example, disks, Oracle databases and applications, and IP addresses) that are configured for high availability are moved to an available node in units called *groups*. The following sections describe resources and groups, and how they are configured for high availability.

2.2.1 Resources

A **cluster resource** is any physical or logical component that is available to a computing system and has the following characteristics:

■   It can be brought online and taken offline.

■   It can be managed in a cluster.

■   It can be hosted by only one node in a cluster at a given time, but can be potentially owned by another cluster node. (For example, a resource is owned by a given node. After a failover, that resource is owned by another cluster node. However, at any given time only one of the cluster nodes can access the resource.)

2.2.2 Group

A **group** is a logical collection of cluster resources that forms a minimal unit of failover. A group is sometimes referred to as a service or application. During a

failover, the group of resources is moved to another cluster node. A group is owned by only one cluster node at a time. All resources required for a given workload (database, disks, and other applications) must reside in the same group.

For example, a service or application created to configure an Oracle database for high availability by using Oracle Fail Safe may include the following resources:

- All disks used by the Oracle database

- An Oracle database **instance**

- One or more virtual addresses, each one consisting of:

    – An IP address

    – A **network name**

- An Oracle Net network **listener** that listens for connection requests to databases in the group

- An Oracle Management Agent that manages communications between Oracle Enterprise Manager and the databases in the service or application

Note that when you add a resource to a group, the disks it uses are also included in the group. For this reason, if two resources use the same disk, then they cannot be placed in different groups. If both resources are to be fail-safe, then both must be placed in the same group.

Oracle Fail Safe helps you to create groups and add the resources needed to run applications. For step-by-step instructions on creating a group, see *Oracle Fail Safe Tutorial*.

2.2.3 Resource Dependencies

Figure 2–3 shows a group created to make a Sales database highly available. When you add a resource to a group, Oracle Fail Safe Manager automatically adds the other resources upon which the resource you added depends; these relationships are called **resource dependencies**. For example, when you add a single-instance database to a group, Oracle Fail Safe adds the shared-nothing disks used by the database **instance** and configures Oracle Net files to work with each group. Oracle Fail Safe also tests the ability of each group to fail over on each node.

Figure 2–3   Designing a Group
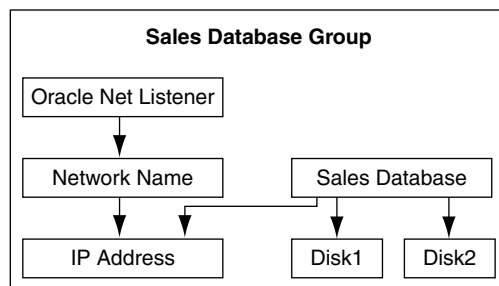
Each node in the cluster can own one or more groups. Each group is composed of an independent set of related resources. The dependencies among resources in a group define the order in which the cluster software brings the resources online and offline. For example, a failure causes the Oracle application or database (and Oracle Net listener) to be brought offline first, followed by the physical disks, network name, and IP address. On the **failover node,** the order is reversed; MSCS brings the IP address

online first, then the network name, then the physical disks, and finally the Oracle database and Oracle Net listener or application.

2.2.4 Resource Types

Each resource type (such as a generic service, physical disk, Oracle database, and so on) is associated with a resource dynamic-link library (DLL) and is managed in the cluster environment by using this resource DLL. There are standard MSCS resource DLLs as well as custom Oracle resource DLLs. The same resource DLL may support several different resource types.

MSCS provides resource DLLs for the resource types that it supports, such as IP addresses, physical disks, generic services, and many others. (A **generic service** resource is a Windows service that is supported by a resource DLL provided in MSCS.)

Oracle Fail Safe uses many of the MSCS resource DLLs to monitor resource types for which Oracle Fail Safe provides custom support, such as generic services.

Oracle provides a custom DLL for the Oracle database resource type. MSCS uses the Oracle resource DLL to manage the Oracle database resources (bring online and take offline) and to monitor the resources for availability.

Oracle Fail Safe provides the following resource DLL files to enable MSCS to communicate with and monitor Oracle database resources:

- FsResOdbs.dll provides functions that enable MSCS to bring an Oracle database and its listener online or offline and check its status through Is Alive polling.

- FsResOdbsEx.dll provides a resource administration extension DLL file that is used by the MSCS Cluster Administrator to display the properties of the Oracle database resource.

For example, when you use Oracle Fail Safe Manager to add an Oracle database to a group, Oracle Fail Safe creates the database resource and an Oracle listener resource.

Because Oracle Fail Safe has more information than MSCS about Oracle cluster resources, Oracle recommends that you use Oracle Fail Safe Manager (or the FSCMD command) to configure and administer Oracle databases and applications.

> **See Also:** *Oracle Fail Safe Installation Guide* for complete information about the custom resource DLLs provided by Oracle Fail Safe, and the MSCS documentation set for information about standard resource types and resource DLLs

2.3 Groups, Virtual Addresses, and Virtual Servers

A **virtual address** is a network address at which resources in a group can be accessed, regardless of the cluster node hosting those resources. A virtual address provides a constant node-independent network location that lets clients to easily access resources without the need to know which physical cluster node is hosting those resources.

Because groups move from an unavailable node to an available node during a failure, a client cannot connect to an application that uses an address that is identified with only one node. You identify a virtual address for a group in Oracle Fail Safe Manager by adding a unique network name and IP address to a group.

Figure 2–4 shows the wizard page in Oracle Fail Safe Manager that helps you add one or more virtual addresses to a group. For step-by-step instructions on adding a virtual address to a group, see *Oracle Fail Safe Tutorial*.

Figure 2–4   Add Resource to Group - Virtual Address Wizard Page

Once you add a virtual address to a group, the group becomes a **virtual server**. Although at least one virtual address is required for each group for client access, you can assign multiple virtual addresses to a group. You may assign multiple virtual addresses to provide increased bandwidth or to segment security for the resources in a group.

Each group appears to users and **client applications** as a highly available virtual server, independent of the physical identity of one particular node. To access the resources in a group, clients always connect to the virtual address of the group. To the client, the virtual server is the interface to the cluster resources and looks like a physical node.

Figure 2–5 shows a two-node cluster with a group configured on each node. Clients access these groups through Virtual Servers A and B. By accessing the cluster resources through the virtual address of a group, as opposed to the physical address of an individual node, you ensure successful remote connection regardless of which cluster node is hosting the group.

Figure 2–5 Accessing Cluster Resources Through a Virtual Server
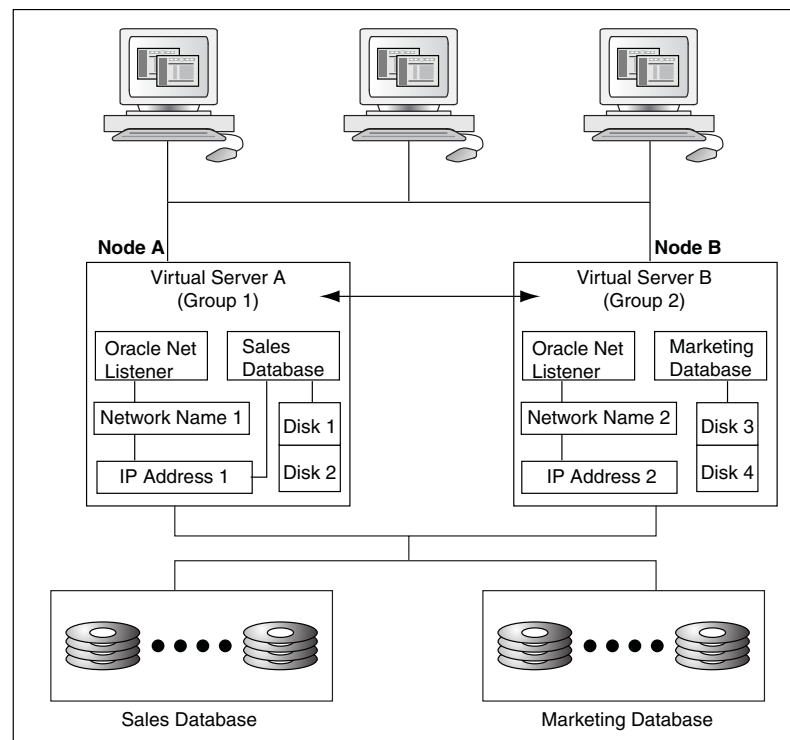
2.4 Allocating IP Addresses for Virtual Addresses

When you set up a cluster, allocate at least the following number of IP addresses:

- One IP address for each cluster node

- One IP address for the cluster alias (described in Section 2.5)

- One IP address for each group

For example, the configuration in Figure 2–5 requires five IP addresses: one for each of the two cluster nodes, one for the cluster alias, and one for each of the two groups.

> **Note:**   You can specify multiple virtual addresses for a group. See Section 4.7 for details.

> **See Also:**   *Oracle Fail Safe Installation Guide* for more information about allocating IP addresses for your Oracle Fail Safe environment

2.5 Cluster Group and Cluster Alias

The **cluster alias** is a node-independent network name that identifies a cluster and is used for cluster-related system management. MSCS creates a group called the Cluster Group, and the cluster alias is the virtual address of this group. Oracle Services for MSCS is a resource in the Cluster Group, making it highly available and ensuring that Oracle Services for MSCS is always available to coordinate Oracle Fail Safe processing on all cluster nodes.

In an Oracle Fail Safe environment, the cluster alias is used only for system management. Oracle Fail Safe Manager interacts with the cluster components and MSCS using the cluster alias.

To populate the tree view in Oracle Fail Safe Manager, specify the cluster alias as shown in Figure 2–6. The cluster alias is not the same as the computer name of any node in the cluster. By specifying the cluster alias when you add the cluster to the tree view, you ensure that when Oracle Fail Safe Manager connects to that cluster it uses the virtual server where Oracle Services for MSCS is running. The cluster alias is always in the Cluster Group (the same group as Oracle Services for MSCS).

> **See Also:** *Oracle Fail Safe Tutorial* for step-by-step instructions on populating the Oracle Fail Safe Manager tree view and connecting to a cluster

Figure 2–6   Cluster Alias in Add Cluster to Tree Dialog Box
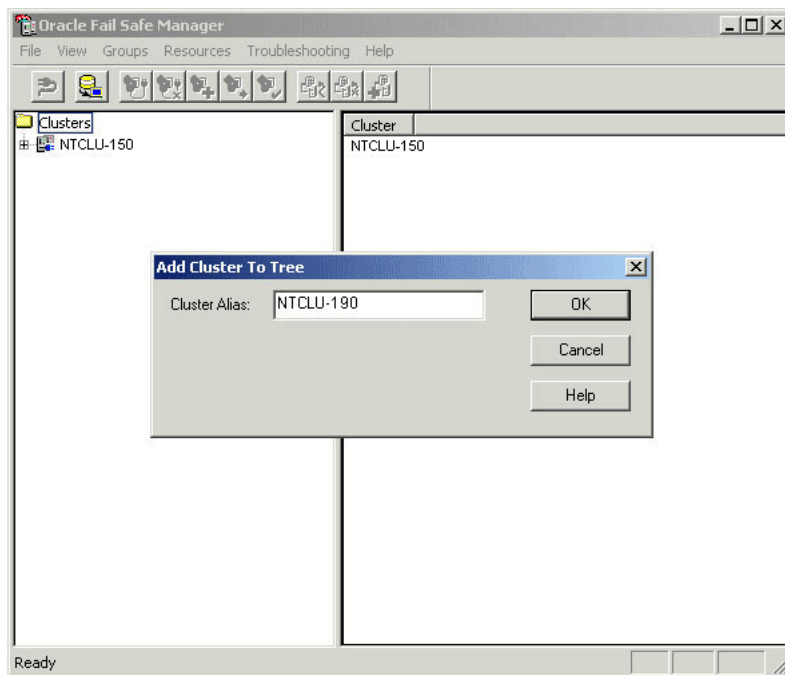
Client applications do not use the cluster alias when communicating with a cluster resource. Rather, clients use one of the virtual addresses of the group that contains that resource.

2.6 Failover

The process of taking a group offline on one node and bringing it back online on another node is called **failover**. After a failover occurs, resources in the group are accessible as long as one of the cluster nodes that is configured to run those resources is available. MSCS continually monitors the state of the cluster nodes and the resources in the cluster.

A failover can be unplanned or planned:

- An unplanned failover occurs automatically when the cluster software detects a node or resource failure.

- A planned failover is a manual operation that you use when you must perform such functions as load balancing or software upgrades.

The following sections describe these types of failover in more detail.

2.6.1 Unplanned Failover

There are two types of **unplanned group failover**s, which can occur due to one of the following:

- Failure of a resource configured for high availability
- Failure or unavailability of a cluster node

2.6.1.1 Unplanned Failover Due to a Resource Failure

An unplanned failover due to a resource failure is detected and performed as follows:

1. The cluster software detects that a resource has failed.

   To detect a resource failure, the cluster software periodically queries the resource (through the resource DLL) to see if it is up and running. See Section 2.6.4 for more information.

2. The cluster software implements the **resource restart policy**. The restart policy states whether or not the cluster software must attempt to restart the resource on the current node, and if so, how many attempts within a given time period must be made to restart it. For example, the resource restart policy may specify that Oracle Fail Safe must attempt to restart the resource three times in 900 seconds.

   If the resource is restarted, then the cluster software resumes monitoring the software (Step 1) and failover is avoided.

3. If the resource is not, or cannot be, restarted on the current node, then the cluster software applies the **resource failover policy**.

   The resource failover policy determines whether or not the resource failure must result in a group failover. If the resource failover policy states that the group must not fail over, then the resource is left in the failed state and failover does not occur.

   Figure 2–10 shows the property page on which you can view or modify the resource restart and failover policies.

   If the resource failover policy states that the group must fail over if a resource is not (or cannot be) restarted, then the group fails over to another node. The node to which the group fails over is determined by which nodes are running, the resource's possible owner nodes list, and the group's preferred owner nodes list. See Section 2.6.7 for more information about the resource possible owner nodes list, and see Section 2.6.10 for more information about the group preferred owner nodes list.

4. Once a group has failed over, the group failover policy is applied. The **group failover policy** specifies the number of times during a given time period that the cluster software must allow the group to fail over before that group is taken offline. The group failover policy lets you prevent a group from repeatedly failing over. See Section 2.6.8 for more information about the group failover policy.

5. The **failback policy** determines if the resources and the group to which they belong are returned to a given node if that node is taken offline (either due to a failure or an intentional restart) and then placed back online. See Section 2.7 for information about failback.

In Figure 2–7, Virtual Server A is failing over to Node B due to a failure of one of the resources in Group 1.

Figure 2–7   Resource Failover

2.6.1.2 Unplanned Failover Due to Node Failure or Unavailability

An unplanned failover that occurs because a cluster node becomes unavailable is performed as described in the following list:

1. The cluster software detects that a cluster node is no longer available.

   To detect node failure or unavailability, the cluster software periodically queries the nodes in the cluster (using the private interconnect).

2. The groups on the failed or unavailable node fail over to one or more other nodes as determined by the available nodes in the cluster, each group's preferred owner nodes list, and the possible owner nodes list of the resources in each group. See Section 2.6.7 for more information about the resource possible owner nodes list, and see Section 2.6.10 for more information about the group preferred owner nodes list.

3. Once a group has failed over, the group failover policy is applied. The **group failover policy** specifies the number of times during a given time period that the cluster software must allow the group to fail over before that group is taken offline. See Section 2.6.8 for more information about the group failover policy.

4. The **failback policy** determines if the resources and the groups to which they belong are moved to a node when it becomes available once more. See Section 2.7 for information about failback.

Figure 2–8 shows Group 1 failing over when Node A fails. Client applications (connected to the failed server) must reconnect to the server after failover occurs. If the application is performing updates to an Oracle database and uncommitted database transactions are in progress when a failure occurs, the transactions are rolled back.

> **Note:**   Steps 3 and 4 in this section are the same as steps 4 and 5 in
> Section 2.6.1.1. Once a failover begins, the process is the same,
> regardless of whether the failover was caused by a failed resource or a
> failed node.

Figure 2–8   Node Failover

2.6.2 Planned Group Failover

A **planned group failover** is the process of intentionally taking client applications and
cluster resources offline on one node and bringing them back online on another node.
This lets administrators perform routine maintenance tasks (such as hardware and
software upgrades) on one cluster node while users continue to work on another node.
Besides performing maintenance tasks, you can perform a planned failover to balance
the load across the nodes in the cluster. In other words, you can use a planned failover
to move a group from one node to another. In fact, to implement a planned failover,
you perform a move group operation in Oracle Fail Safe Manager (see the online help
in Oracle Fail Safe Manager for instructions).

During a planned failover, Oracle Services for MSCS works with MSCS to efficiently
move the group from one node to another. Client connections are lost and clients must
manually reconnect at the virtual server address of the application, unless you have
configured transparent application failover (see Section 7.9 for information about
transparent application failover). Then, you can take your time performing the
upgrade, because Oracle Fail Safe lets clients work uninterrupted on another cluster
node while the original node is offline. (If a group contains an Oracle database, then
the database is checkpointed prior to any planned failover to ensure rapid database
recovery on the new node.)

2.6.3 Group and Resource Policies That Affect Failover

Values for the various resource and group failover policies are set to default values
when you create a group or add a resource to a group using Oracle Fail Safe Manager.

However, you can reset the values in these policies with the Group Failover property page, the Group Failback property page, and the Resource Policies property page. You can set values for the group failback policy at group creation time or later, using the Group Failback property page.

Figure 2–9 shows the page for setting group failover policies. To access this page, select the group of interest in the Oracle Fail Safe Manager tree view and then click the Failover tab.

Figure 2–10 shows the page for setting resource policies. To access this page, select the resource of interest in the Oracle Fail Safe Manager tree view and then click the Policies tab.

Figure 2–9   Group Failover Property Page
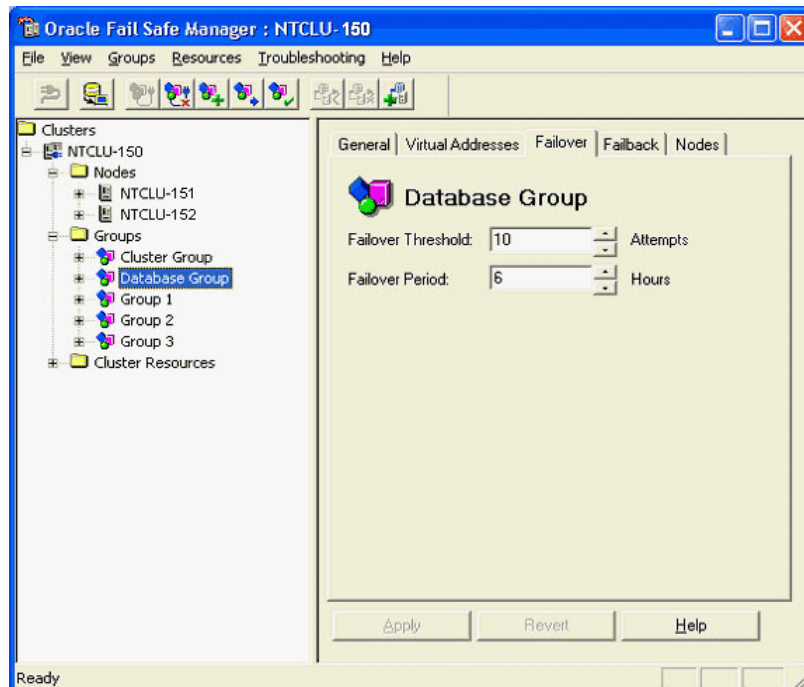
Figure 2–10   Resource Policies Property Page
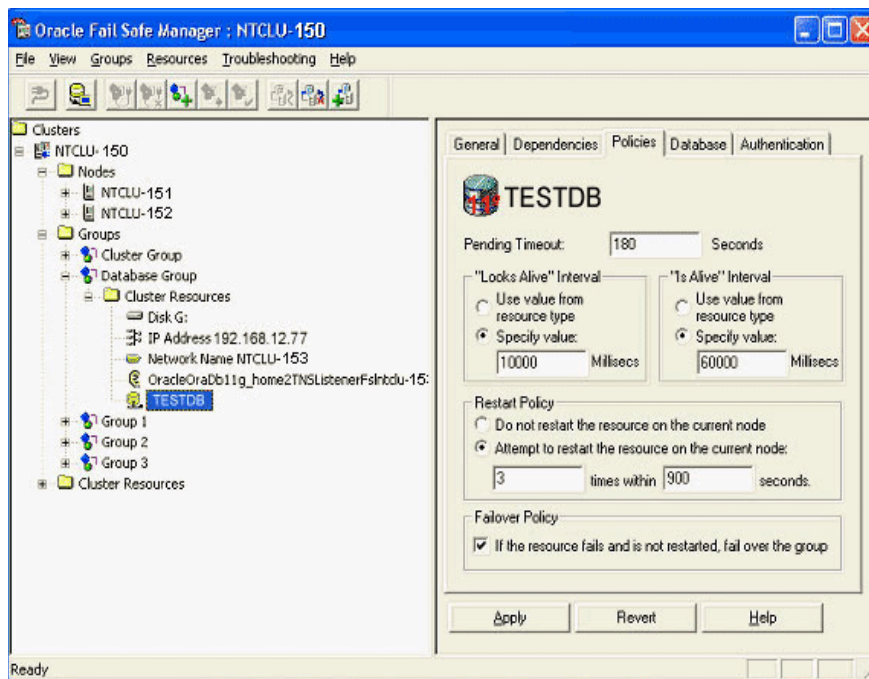
2.6.4 How a Resource Failure Is Detected

All resources that have been configured for high availability are monitored for their status by the cluster software. Resource failure is detected based on three values:

- Pending timeout value

  The pending timeout value specifies how long the cluster software must wait for a resource in a pending state to come online (or offline) before considering that resource to have failed. By default, this value is 180 seconds.

- Is Alive interval

  The Is Alive interval specifies how frequently the cluster software must check the state of the resource. You can use either the default value for the resource type or specify a number (in milliseconds). This check is more thorough, but uses more system resources than the check performed during a Looks Alive interval.

- Looks Alive interval

  The Looks Alive interval specifies how frequently the cluster software must check the registered state of the resource to determine if the resource appears to be active. You can use either the default value for the resource type or specify a number (in milliseconds). This check is less thorough, but also uses fewer system resources, than the check performed during an Is Alive interval.

2.6.5 Resource Restart Policy

Once it is determined that a resource has failed, the cluster software applies the restart policy for the resource. The resource restart policy provides two options, as shown in Figure 2–10:

- The cluster software must not attempt to restart the resource on the current node. Instead, it must immediately apply the resource failover policy.

■ The cluster software must attempt to restart the resource on the current node a specified number of times within a given time period. If the resource cannot be restarted, then the cluster software must apply the resource failover policy.

2.6.6 Resource Failover Policy

The resource failover policy determines whether or not the group containing the resource must fail over if the resource is not (or cannot be) restarted on the current node. If the policy states that the group containing the failed resource must not fail over, then the resource is left in the failed state on the current node. (The group may eventually fail over anyway; if another resource in the group has a policy that states that the group containing the failed resource must fail over, then it will.) If the policy states that the group containing the failed resource must fail over, then the group containing the failed resource fails over to another cluster node as determined by the group preferred owner nodes list. See Section 2.6.10 and Section 2.7.1 for a description of the preferred owner nodes list.

2.6.7 Resource Possible Owner Nodes List

The **possible owner nodes list** consists of all nodes on which a given resource is permitted to run. A node on which a resource is permitted to run must satisfy the following conditions:

■ The DLL for the given resource must be installed on the node.

■ You must not specify that the node must be excluded from the possible owner nodes list.

In addition, although it is not a requirement, you must ensure that all resources that are permitted to run on a given node are also configured to run on that node. Otherwise, a group containing that resource may fail over to the node, but be unable to restart the resource. A resource is configured to run on a possible owner node when you do either of the following:

■ Add the resource to a group that currently includes that node as a possible owner node for the group.

   See the chapter that describes configuring your particular resource for high availability for information about adding that particular resource type to a group:

   – Configuring Single-Instance Databases for High Availability and Disaster Tolerance

   – Configuring Generic Services for High Availability

   – Configuring Oracle Management Agent for High Availability

   – Configuring Oracle Application Server Components for High Availability

■ Run the `Verify Group` command.

   If the node becomes a **possible owner node** for the group after you have added the resource to the group, then the `Verify Group` command prompts you to configure the group on that node. The `Verify Group` command ensures that all possible owner nodes for a group are configured for the group. See Section 6.1.2 for information about the `Verify Group` command.

As mentioned previously, you can specify that a node must be excluded from the possible owner nodes list. For example, suppose you have a four-node cluster and each node has the Oracle database and the Oracle Fail Safe database resource DLLs installed. You have the choice of specifying that all four nodes are possible owner

nodes for the resource. However, suppose Node 3 does not have sufficient memory to run both the database instance and the rest of its workload. You may decide to remove Node 3 from the possible owner nodes list for the database resource.

You specify the possible owner nodes list for a resource when you add it to a group. You can adjust the possible owner nodes list for a resource that has been made highly available using one of the following property pages:

- The General property page for the resource

  The General property page for the resource does not show you how modifications to the possible owner nodes list of the resource affects the group to which the resource belongs. If you use this property page to modify the possible owner nodes list of a resource, then ensure you do not inadvertently create a situation where none of the resources in the group have common nodes in their possible owner nodes lists.

- The Nodes property page for the group containing the resource

  The Nodes property page presents the possible owner nodes list for the group. However, the possible owner nodes list is not actually an attribute of a group. Oracle Fail Safe determines which nodes to present in the possible owner nodes list for a group by finding the intersection of the possible owner nodes list of each resource in the group. Using this property page, you can see if removing one of the possible owner nodes results in no nodes being a possible owner node for a group. Figure 2–11 is an example of the Nodes property page. Note that if you make a change to the possible owner nodes list for a group, then this change is applied to all resources in the group, except disk resources.

In a two-node cluster, the possible owner nodes list for every resource usually includes both nodes. To provide failover capabilities, at least two cluster nodes must be possible owner nodes for a resource.

> **Note:** Assume you add a new node to the cluster and Oracle Fail Safe or MSCS DLLs (or both) are installed on that node. That node becomes a possible owner for resources supported by the installed DLLs. If resources have not yet been configured for high availability on that node, then a group can fail over to that node and be unable to restart the resources on that node.
>
> However, if you run the `Verify Group` command, then Oracle Fail Safe checks that the resources in the specified group are configured to run on each node that is a possible owner for the group. If it finds a possible owner node where the resources in the group are not configured to run, then Oracle Fail Safe configures them for you.
>
> Therefore, Oracle strongly recommends you run the `Verify Group` command for each group for which the new node is listed as a possible owner. Section 6.1.2 describes the `Verify Group` command.

Figure 2–11   Nodes Property Page
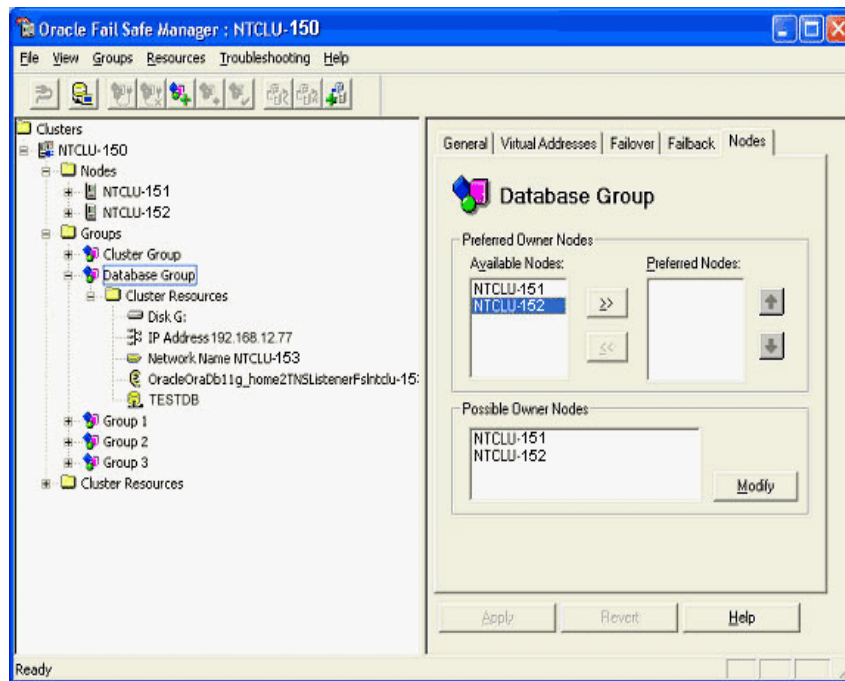
2.6.8 Group Failover Policy

If the resource failover policy states that the group containing the resource must fail over if the resource cannot be restarted on the current node, then the group fails over and the group failover policy is applied. Similarly, if a node becomes unavailable, then the groups on that node fail over and the group failover policy is applied.

The group failover policy specifies the number of times during a given time period that the cluster software must allow the group to fail over before that group is taken offline. The failover policy provides a means to prevent a group from failing over repeatedly.

The group failover policy consists of a failover threshold and a failover period:

- Failover threshold

  The **failover threshold** specifies the maximum number of times failover can occur (during the failover period) before the cluster software stops attempting to fail over the group.

- Failover period

  The **failover period** is the time during which the cluster software counts the number of times a failover occurs. If the frequency of failover is greater than that specified for the failover threshold during the period specified for the failover period, then the cluster software stops attempting to fail over the group.

For example, if the failover threshold is 3 and the failover period is 5, then the cluster software allows the group to fail over 3 times within 5 hours before discontinuing failovers for that group.

When the first failover occurs, a timer to measure the failover period is set to 0 and a counter to measure the number of failovers is set to 1. The timer is not reset to 0 when the failover period expires. Instead, the timer is reset to 0 when the first failover occurs after the failover period has expired.

For example, assume again that the failover period is 5 hours and the failover threshold is 3. As shown in Figure 2–12, when the first group failover occurs at point A, the timer is set to 0. Assume a second group failover occurs 4.5 hours later at point B, and the third group failover occurs at point C. Because the failover period has been exceeded when the third group failover occurs (at point C), group failovers are allowed to continue, the timer is reset to 0, and the failover counter is reset to 1.

Assume that another failover occurs at point D (after 7 total hours have elapsed since point A, and 2.5 hours have elapsed since point B). You may expect that failovers will be discontinued. The failovers at points B, C, and D have occurred within a 5-hour timeframe. However, because the timer for measuring the failover period was reset to 0 at point C, the failover threshold has not been exceeded, and the cluster software allows the group to fail over.

Assume that another failover occurs at point E. When a problem that ordinarily results in a failover occurs at point F, the cluster software does not fail over the group. Three failovers have occurred during the 5-hour period that has passed since the timer was reset to 0 at point C. The cluster software leaves the group on the current node in a failed state.

Figure 2–12   Failover Threshold and Failover Period Timeline

2.6.9 Effect of Resource Restart Policy and Group Failover Policy on Failover

Both the resource restart policy and the failover policy of the group containing the resource affect the failover behavior of a group.

For example, suppose the Northeast database is in a group called Customers, and you specify the following:

- On the Policies property page for the Northeast database:
    - Attempt to restart the database on the current node 3 times within 600 seconds (10 minutes)
    - If the resource fails and cannot be restarted, fail over the group
- On the Failover property page for the Customers group:
    - The failover threshold for the group containing the resource is 20
    - The failover period for the group containing the resource is 1 hour

    Assume a database failure occurs. Oracle Fail Safe attempts to restart the database instance on the current node. The attempt to restart the database instance fails three times within a 10-minute period. Therefore, the Customers group fails over to another node.

    On that node, Oracle Fail Safe attempts to restart the database instance and fails three times within a 10-minute period, so the Customers group fails over again. Oracle Fail Safe continues its attempts to restart the database instance and the

Customers group continues to fail over until the database instance restarts or the group has failed over 20 times within a 1-hour period. If the database instance cannot be restarted, and the group fails over fewer than 20 times within a 1-hour time period, then the Customers group fails over repeatedly. In such a case, consider reducing the failover threshold to eliminate the likelihood of repeated failovers.

2.6.10 Group Failover and the Preferred Owner Nodes List

When you create a group, you can create a preferred owner nodes list for both group failover and failback. (When the cluster contains only two nodes, you specify this list for failback only.) You create an ordered list of nodes to indicate the preference you have for each node in the list to own that group.

For example, in a four-node cluster, you may specify the following preferred owner nodes list for a group containing a database:

- Node 1
- Node 4
- Node 3

This indicates that when all four nodes are running, you prefer for the group to run on Node 1. If Node 1 is not available, then your second choice is for the group to run on Node 4. If neither Node 1 nor Node 4 is available, then your next choice is for the group to run on Node 3. You have omitted Node 2 from the preferred owner nodes list. However, if it is the only choice available to the cluster software (because Node 1, Node 4, and Node 3 have all failed), then the group fails over to Node 2. (This happens even if Node 2 is not a possible owner for all resources in the group. In such a case, the group fails over, but remains in a failed state.)

When a failover occurs, the cluster software uses the preferred owner nodes list to determine the node to which it must fail over the group. The cluster software fails over the group to the top-most node in the list that is up and running and is a possible owner node for the group. Section 2.6.11 describes in more detail how the cluster software determines the node to which a group fails over.

See Section 2.7.1 for information about how the group preferred owner nodes list affects failback.

2.6.11 Determining the Failover Node for a Group

The node to which a group fails over is determined by the following three lists:

- List of available cluster nodes

  The list of available cluster nodes consists of all nodes that are running when a group failover occurs. For example, suppose you have a four-node cluster. If one node is down when a group fails over, then the list of available cluster nodes is reduced to three.

- List of possible owner nodes for each resource in the group (See Section 2.6.7.)

- List of preferred owner nodes for the group containing the resources (See Section 2.6.10.)

The cluster software determines the nodes to which your group can possibly fail over by finding the intersection of the available cluster nodes and the common set of possible owners of all resources in the group. For example, assume you have a

four-node cluster and a group on Node 3 called Test_Group. You have specified the possible owners for the resources in Test_Group as shown in Table 2–1.

Table 2–1   Example of Possible Owners for Resources in Group Test_Group

Possible Owners for Resource 1	Possible Owners for Resource 2	Possible Owners for Resource 3
Node 1 - Yes	Node 1 - Yes	Node 1 - Yes
Node 2 - Yes	Node 2 - No	Node 2 - Yes
Node 3 - Yes	Node 3 - Yes	Node 3 - Yes
Node 4 -Yes	Node 4 - Yes	Node 4 - Yes

By reviewing Table 2–1, you see that the intersection of possible owners for all three resources is:

- Node 1

- Node 3

- Node 4

Assume that Node 3 (where Test_Group currently resides) fails. The available nodes list is now:

- Node 1

- Node 4

To determine the nodes to which Test_Group can fail over, the cluster software finds the intersection of the possible owner nodes list for all resources in the group and the available nodes list. In this example, the intersection of these two lists is Node 1 and Node 4.

To determine the node to which it must fail over Test_Group, the cluster software uses the preferred owner nodes list of the group. Assume that you have set the preferred owner nodes list for Test_Group to be:

- Node 3

- Node 4

- Node 1

Because Node 3 has failed, the cluster software fails over Test_Group to Node 4.   If both Node 3 and Node 4 are not available, then the cluster software fails over Test_Group to Node 1. If Nodes 1, 3, and 4 are not available, then the group fails over to Node 2. However, because Node 2 is not a possible owner for all of the resources in Test_Group, the group remains in a failed state on Node 2.

2.7 Failback

A **failback** is a process of automatically returning a group of cluster resources to a preferred owner node from the failover node after a preferred owner node returns to operational status. A **preferred owner node** is a node on which you want a group to reside when possible (when that node is available).

You can set a **failback policy** that specifies when (and if) groups must fail back to a preferred owner node from the failover node. For example, you can set a group to fail back immediately or between specific hours that you choose. Or, you can set the failback policy so that a group does not fail back, but continues to run on the node

where it currently resides. Figure 2–13 shows the property page for setting the failback policy for a group.

Figure 2–13   Group Failback Policy Property Page
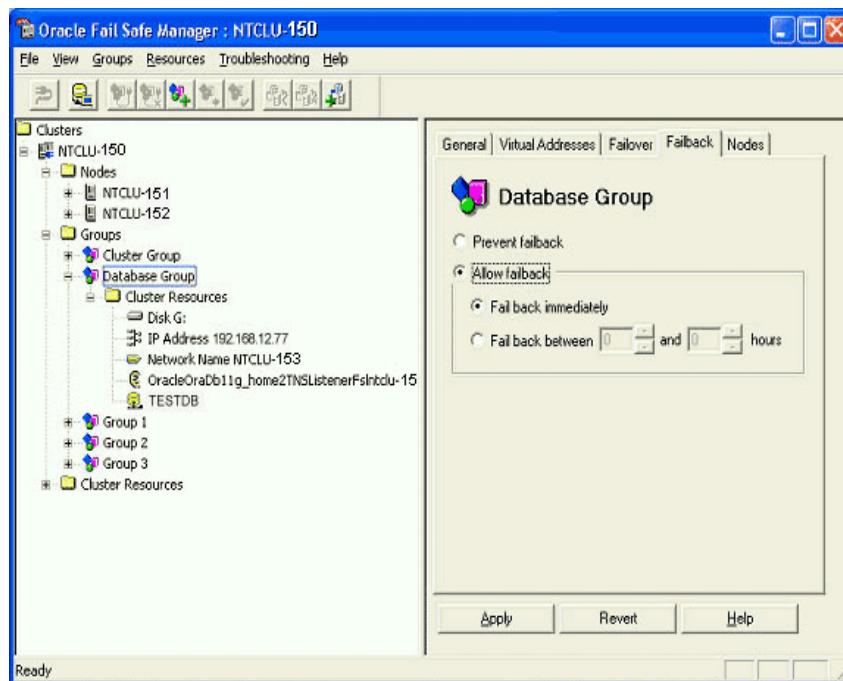
2.7.1 Group Failback and the Preferred Owner Nodes List

When you create a group on a cluster, you can create a preferred owner nodes list for group failover and failback. When the cluster contains two nodes, you specify this list for failback only. You create an ordered list of nodes that indicates the nodes where you prefer a group to run. When a previously unavailable node comes back online, the cluster software reads the preferred owner nodes list for each group on the cluster to determine whether or not the node that just came online is a preferred owner node for any group on the cluster. If the node that just came online is higher on the preferred owner nodes list than the node on which the group currently resides, then the group is failed back to the node that just came back online.

For example, in a four-node cluster, you may specify the following preferred owner nodes list for the group called My_Group:

- Node 1

- Node 4

- Node 3

Assume that My_Group has failed over to, and is currently running on, Node 4 because Node 1 had been taken offline. Now Node 1 is back online. The cluster software reads the preferred owner nodes list for My_Group (and all other groups on the cluster); it finds that the preferred owner node for My_Group is Node 1. It fails back My_Group to Node 1, if failback is enabled.

If My_Group is currently running on Node 3 (because both Node 1 and Node 4 are not available) and Node 4 comes back online, then My_Group fails back to Node 4 if failback is enabled. Later, when Node 1 becomes available, My_Group fails back once more, this time to Node 1. When you specify a preferred owner nodes list, be careful

not to create a situation where failback happens frequently and unnecessarily. For most applications, two nodes in the preferred owner nodes list is sufficient.

A scenario with unexpected results is exhibited when a group has been manually moved to a node. Assume all nodes are available and `My_Group` is currently running on Node 3 (because you moved it there with a move group operation). If Node 4 is restarted, then `My_Group` fails back to Node 4, even though Node 1 (the highest node in the preferred owner node list of `My_Group`) is also running.

When a node comes back online, the cluster software checks to see if the node that just came back online is higher on the preferred owner nodes list than the node where each group currently resides. If so, all such groups are moved to the node that just came back online.

See Section 2.6.10 for information about how the group preferred owner nodes list affects failover.

2.7.2 Client Reconnection After Failover

Node failures affect only those users and applications:

- That are directly connected to applications hosted by the failed node

- Whose transactions were being handled when the node failed

Typically, users and applications connected to the failed node lose the connection and must reconnect to the failover node (through the node-independent virtual address) to continue processing. With a database, any transactions that were in progress and uncommitted at the time of the failure are rolled back. Client applications that are configured for transparent application failover experience a brief interruption in service; to the client applications, it appears that the node was quickly restarted. The service is automatically restarted on the failover node—without operator intervention.

See Section 7.9 for information about transparent application failover.

3

Designing an Oracle Fail Safe Solution

Oracle Fail Safe provides a number of configuration options to satisfy your architecture or failover requirements.

This chapter discusses the following topics:

- Customizing Your Configuration
- Integrating Clients and Applications

3.1 Customizing Your Configuration

You can deploy highly available solutions using the following configurations:

- Active/Passive Configuration
- Active/Active Configuration

These configurations differ in the way work is allocated among the cluster nodes, but share the following features:

- One or more Oracle homes are created on a private disk (usually the system disk) on each node.
- All Oracle product executable files are installed in the Oracle homes on each node.
- All **data files,** configuration files, log files, html files, and so on that are required by the application being made highly available are placed on cluster disks, so that they can be accessed by each cluster node.

The Oracle Services for MSCS software automatically runs as needed on one or more cluster nodes to ensure proper configuration and failover.

Figure 1–4 shows the software and hardware components in a cluster configured with Oracle Fail Safe.

3.1.1 Active/Passive Configuration

In an **active/passive configuration,** one or more nodes host the entire cluster workload, but one node remains idle (as a standby server), ready to take over processing in case a node running an application fails. This solution ensures that the performance for the fail-safe workload is the same before and after failover.

Figure 3–1 shows a two-node configuration with Oracle Services for MSCS, Oracle Application Server and an Oracle database running on Node 1, and with Node 2 as a standby node. Currently, nothing is running on Node 2. Node 2 takes over the workload of Node 1 in the event of a failover.

Figure 3–1   Active/Passive (Standby) Two-Node Configuration
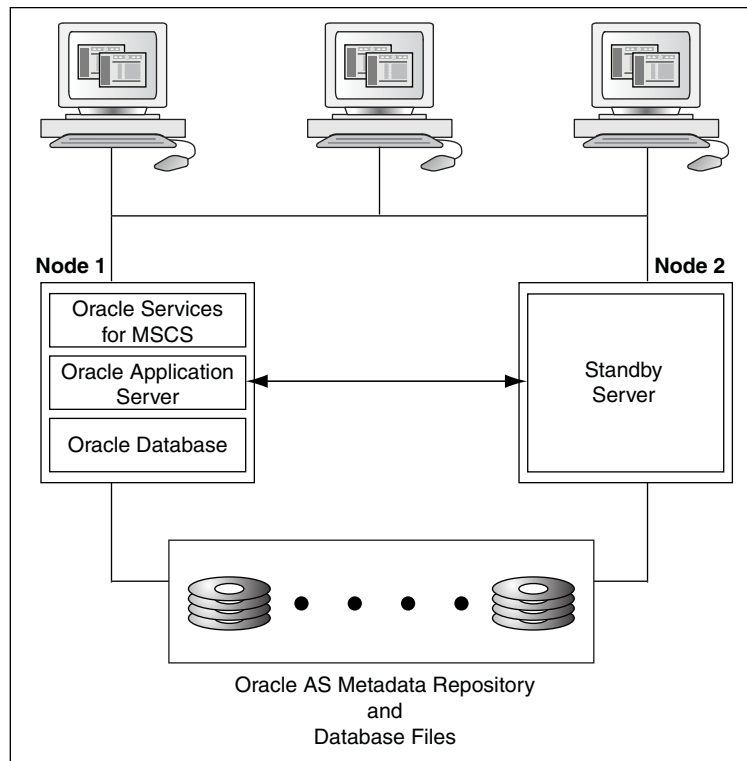
Figure 3–2 shows a four-node configuration with Oracle Services for MSCS and an Oracle database running on Node 1, an Oracle Application Server and an Oracle database running on Node 2, and an Oracle Application Server and an Oracle database running on Node 3. Node 4 is the **standby node**. Currently, nothing is running on Node 4. In the event of a failover, Node 4 takes over the failover workload.

Figure 3–2   Active/Passive (Standby) Four-Node Configuration
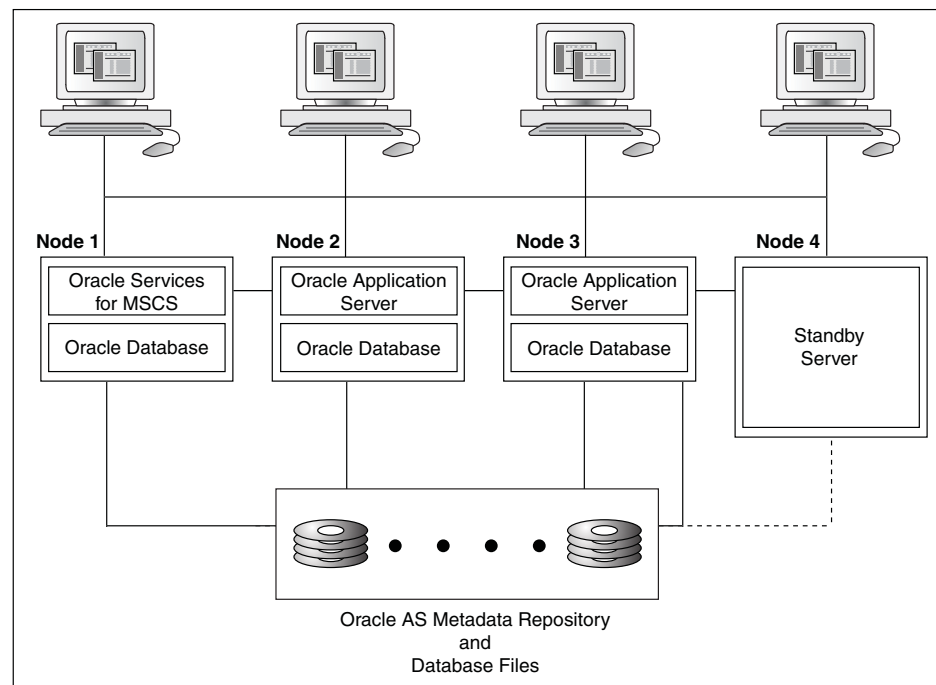
An active/passive configuration is the fastest failover configuration, because the passive standby node has no workload of its own.

3.1.2 Active/Active Configuration

In an active/active configuration each node shares the application processing tasks, and also backs up other nodes in the event of a failure. If one node fails, then another node runs its own applications and services as well as those that fail over from the failed node. The active/active configuration is more cost-effective than the active/passive configuration. This configuration provides a flexible architecture that allows you divide the workload to best meet your business needs.

Figure 3–3 shows a two-node active/active configuration with an Oracle database running on both cluster nodes. In addition, an Oracle Application Server and a generic service are running on Node 1, and Oracle Services for MSCS and an Oracle Application Server are running on Node 2. In Figure 3–3, an Oracle database is used for marketing on Node 1, and for sales on Node 2. The cluster disks owned by Node 1 store the marketing files, and the cluster disks owned by Node 2 store the sales files.

Figure 3–3   Active/Active Configuration
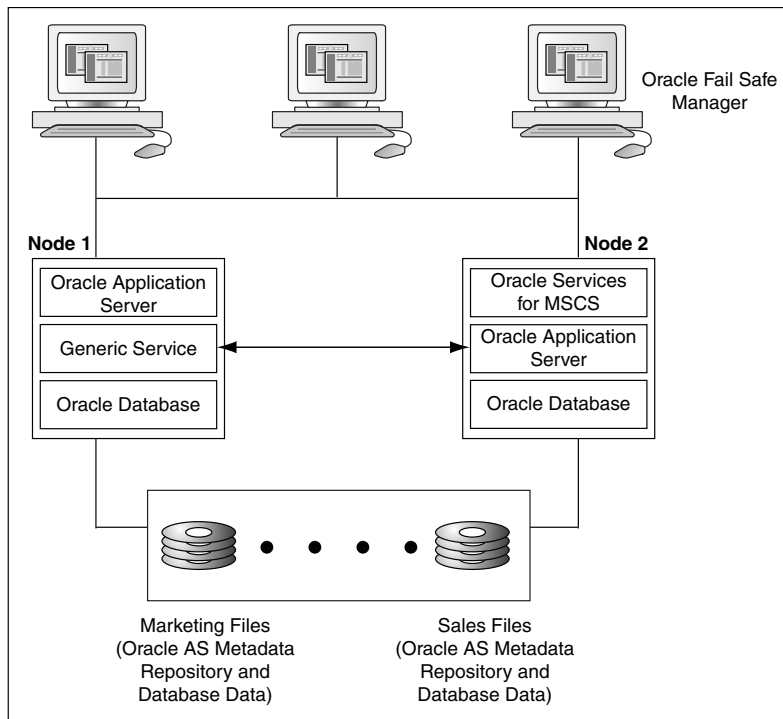
In the **active/active configuration,** all nodes actively process applications during normal operations. This configuration provides better performance (higher throughout) when all nodes are operating, but slower failover and possibly reduced performance when a node fails. Also, the client connections are distributed over all nodes.

Balancing workload means making trade-offs concerning the size of the normal workload on each system. If all systems run at nearly full capacity, then few resources are available to handle the load of another system in an outage, and client systems experience significantly slower response during and after a failover. If you have the resources to quickly repair or replace a failed system, then the temporary period during which one cluster node serves both workloads will be small; a short period of slow response can be tolerated better than a long one. In fact, some businesses actually prefer having applications run more slowly than usual than to have a period of downtime.

Alternatively, running all systems slightly under 75% to 50% capacity (depending on the number of nodes in the cluster) ensures that clients do not experience loss of response time after a failover, but the equivalent of an entire system can remain idle under normal conditions, much like an active/passive configuration.

Oracle Fail Safe can be configured to avoid some of the performance problems with this type of configuration. For example, you can:

■   Enable failover only for your **mission-critical application**s

■   Use different database parameter files on each node so that fewer system resources are used after a failover

■   Configure each component (Oracle database, Oracle Application Server, and so on) into a separate group with its own failover and failback policies

   This is possible because Oracle Fail Safe enables you to configure each cluster node to host several virtual servers.

- Combine the scripting support of Oracle Fail Safe (using the `FSCMD` command described in Chapter 5) with a system monitoring tool (such as Oracle Enterprise Manager) to automate the movement of groups for load-balancing purposes.

Although the nodes do not need to be physically identical, you must select servers with enough power, memory, disk host adapters, and disk drives to support an adequate level of service if a failover occurs at a busy time of the day.

3.2 Integrating Clients and Applications

To operate in an Oracle Fail Safe environment, **client applications** do not require any special programming or changes. Client applications that work with an Oracle resource on a single node continues to function correctly in an Oracle Fail Safe environment without recoding, recompiling, or relinking. This is because clients can use the virtual server to access the application.

Chapter 7, and Chapter 8 contain a section specific to how you can integrate clients and applications. Chapter 7 describes how to make your clients and applications transparently fail over when a database fails over to another node in the cluster.

4

Management for High Availability

The unique advantage offered by Oracle Fail Safe is its ability to help you easily configure resources in a Windows cluster environment. This chapter discusses the following topics:

- What Does It Mean to Configure Failover?
- How Does Oracle Fail Safe Use the Wizard Input?
- Managing Cluster Security
- Discovering Standalone Resources
- Renaming Resources
- Using Oracle Fail Safe in a Multiple Oracle Homes Environment
- Configurations Using Multiple Virtual Addresses
- Adding a Node to an Existing Cluster

For step-by-step procedures to configure standalone resources into groups, and for information about managing those resources once they are in groups, refer to Chapters 7, 8, 9, and 10 in this manual and to *Oracle Fail Safe Tutorial* and online help.

4.1 What Does It Mean to Configure Failover?

Using Oracle Fail Safe Manager wizards, you can easily configure failover automatically and with minimal work by a network manager. Oracle Fail Safe Manager helps to configure resources into groups so that when one node in a cluster fails, another cluster node immediately takes over the resources in the failed node's groups.

The wizards minimize the risk of introducing configuration problems during implementation and also reduce the level of expertise required to configure resources for high availability. Most policies that you set with the wizards can be modified later with Oracle Fail Safe Manager.

The following list summarizes the basic tasks to perform to implement failover for resources. Except for the first task, you must perform all of these tasks using Oracle Fail Safe Manager.

1. Ensure that the products that you want to configure with Oracle Fail Safe are properly installed. (This is described in the *Oracle Fail Safe Installation Guide*.)

2. Start Oracle Fail Safe Manager.

3. Verify the cluster.

4. Create a group.

5. Add one or more virtual addresses to the group.

6. If you are adding a standalone Oracle database, then use the `Verify Standalone Database` tool to verify the database.

7. Add resources to the group.

8. Verify the group.

9. Update any Oracle Net file (such as the `tnsnames.ora` file) on client systems.

> **Note:** Depending on the type of resource you are configuring, there may be additional steps or considerations.

Refer to the tutorial and online help in Oracle Fail Safe Manager for step-by-step guidance on using the Oracle Fail Safe Manager wizards.

4.2 How Does Oracle Fail Safe Use the Wizard Input?

Once the wizard collects all the required information, Oracle Fail Safe Manager interacts with Oracle Services for MSCS (which in turn interacts with MSCS) to facilitate a high-availability environment.

Based on the information that you provide with the wizards, Oracle Fail Safe derives any additional information it requires to configure the environment.

Most resources are configured by Oracle Fail Safe using a similar series of steps. Oracle Fail Safe performs the following specific steps to configure a highly available Oracle database:

1. Configures access to the database using a virtual address:

   a. Configures Oracle Net to use the virtual address or addresses associated with the database on all nodes listed in the possible owner nodes list for the database. (On a two-node cluster, this is both cluster nodes. On clusters that consist of more than two nodes, you are asked to specify the possible owner nodes for a resource as a step in the Add Resource to Group Wizard.)

   b. Duplicates the network configuration information on all nodes in the possible owner nodes list.

2. Configures the database to:

   a. Verify that all data files used by the database resource are on cluster disks and are not currently used by applications in other groups. If the cluster disks are in another group, but not used by applications in that group, then Oracle Fail Safe moves the disks into the same group with the database resource.

   b. Create the failback policy for the database resources based on choices you made in the wizard.

   c. Populate the group with these resources:

      * Each disk resource used by the cluster group

      * Oracle database

      * Oracle Net listener

3. Performs the following steps on each of the possible owner nodes for the group to which the database has been added, one at a time:

   a. Creates an Oracle instance with the same name on the node.

      b. Verifies that the node can bring the database online and offline by failing it over to the node to ensure that the failover policy works.

4. Shuts down Oracle database after testing failover on all nodes in the possible owner nodes list. If the preferred owner node list is empty, then the group remains on the last node to which it was failed over as part of the configuration process.

By performing these steps, Oracle Fail Safe ensures that the resource is correctly configured and capable of failing over and failing back to all possible owner nodes of the group to which it has been added.

Figure 4–1 shows a two-node active/active cluster configuration in which each node hosts a group with a database.

Figure 4–1   Virtual Servers and Addressing in an Oracle Fail Safe Environment
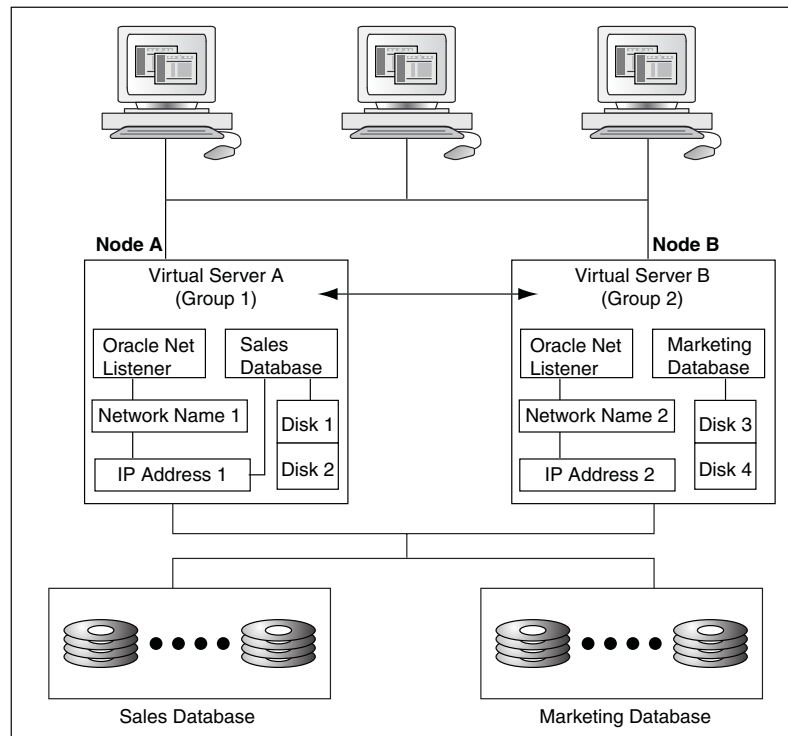
The virtual servers (A and B) and their network addresses are known by all clients and cluster nodes. The `listener.ora` file on each cluster node and the `tnsnames.ora` file on each client workstation contain the network name and address information for each virtual server.

For failover to work properly, the **host name** (virtual address), database instance, SID entry, and protocol information in each `tnsnames.ora` and `listener.ora` file must match on each server node that is a possible owner of the resources in the group and the client system.

For example, during normal operations, Virtual Server A is active on Node A. Node B is the failover node for Virtual Server A. The cluster disks are connected to both nodes so that resources can run on either node in the cluster, but service for the resources in each group is provided by only one cluster node at a time.

If a system failure occurs on Node A, then Group 1 becomes active on Node B using the same virtual address and port number as it had on Node A. Node B takes over the workload from Node A transparently to clients, which continue to access Group 1

using Virtual Server A and Group 2 using Virtual Server B. Clients continue to access the resources in a group using the same virtual server name and address, without considering the physical node that is serving the group.

4.3 Managing Cluster Security

To accomplish administrative tasks associated with Oracle Fail Safe, you need the appropriate privileges to manage Oracle resources and applications and to perform operations through Oracle Fail Safe Manager.

Table 4–1 provides a quick reference for the privileges required for the services you use in an Oracle Fail Safe environment. For more information, refer to the sections listed in the last column.

Table 4–1   Permissions and Privileges

Service	Required Privileges	Reference
Oracle Services for MSCS	Domain user account that has Administrator privileges on all cluster nodes	Section 4.3.1
Oracle Fail Safe Manager	Domain user account that has Administrator privileges on all cluster nodes	Section 4.3.2
Oracle database	Database administrator account with SYSDBA privileges	Section 7.5
Generic services	By default, a generic service runs under the local system account. If you specify that the generic service must run under a user account, then it must have the "Log on as a service" privilege.	Section 8.4

4.3.1 Oracle Services for MSCS

Oracle Fail Safe accesses database resources from two different Windows services: the Cluster Service service and the OracleMSCSServices service. The Cluster Service service implements the database resource DLL functions, that is, the common resource functions that start and stop the database resource, and determine if the database resource is functioning properly by issuing simple database queries against the database ("Is Alive" polling). The OracleMSCSServices service issues database queries that are related to configuration of database resources, such as verifying that all database files are on shared cluster disks, the database starts and stops successfully, and so on.

Each of these services executes in the context of the Log On As user specified for the particular service. The OracleMSCSServices service executes under the account provided to the Oracle Services for MSCS Security Setup tool during the installation of Oracle Fail Safe. Prior to Windows Server 2008, the Cluster Service service executed under the cluster account specified when the cluster was configured. In Windows Server 2008 and later the Cluster Service service executes as user Local System.

All database connections must be properly authenticated, so Oracle Fail Safe must execute from a context that is authorized to connect to a database. If operating system authentication is being used to access a database (the database parameter REMOTE_ LOGIN_PASSWORDFILE is set to *NONE*) then Oracle Database authenticates the access from the Windows service using the account name for that service. For the OracleMSCSServices service, that means that authentication is done using the Log On As account specified for the OracleMSCSServices service. For the Cluster Service service, on installations that are using a Windows Server version that is older than 2008, the cluster account is used. In Windows Server 2008 and later, the Oracle Fail

Safe database resource DLL impersonates the OracleMSCSServices account when connecting to the database. Thus in Windows Server 2008 and later, even though the Cluster Service service is executing as Local System, database access authentication is done using the OracleMSCSServices account.

Prior to Windows Server 2008 it was possible for Oracle Fail Safe to access databases from two different user accounts: the one specified for the Cluster Service service and the OracleMSCSServices service. On systems using Windows Server 2008 and later, when using operating system authentication, Oracle Fail Safe only attempts to authenticate database access using the account specified for the OracleMSCSServices service. See Section 7.3.3.4, "Database Authentication" for more information regarding database authentication.

4.3.1.1 Account Updates Using the Oracle Fail Safe Security Setup Tool

Oracle Fail Safe provides a Security Setup tool that you can use to update the information for the account under which Oracle Services for MSCS runs. The Oracle Services for MSCS Security Setup tool is installed when you install Oracle Services for MSCS.

On a cluster node, you can access the Oracle Services for MSCS Security Setup tool from the Windows taskbar. To do so, from the Windows **Start** menu, select **Programs** (or **All Programs**), then *Oracle_Home*, and finally, **Oracle Services for MSCS Security Setup.**

Note: Be sure that you use the Oracle Services for MSCS Security Setup tool to update the security information on all cluster nodes, and that you use the same account on all cluster nodes.

Figure 4–2 shows the setup for user account Administrator in the domain EXAMPLE\Admin.

Figure 4–2 Windows User Account Settings for the Oracle Services for MSCS
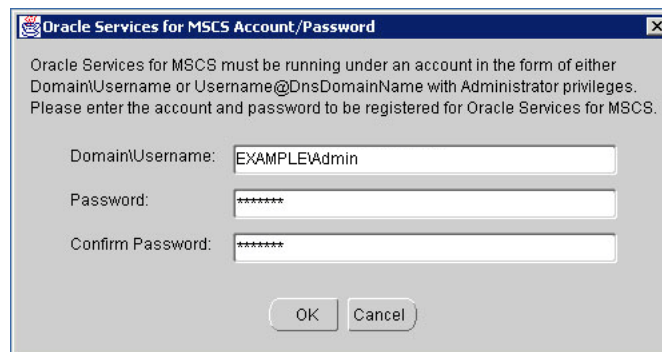
4.3.2 Oracle Fail Safe Manager

The account you use to log in to Oracle Fail Safe Manager must be a domain user account (not a local account) that has Administrator privileges on all cluster nodes.

4.4 Discovering Standalone Resources

Oracle Services for MSCS automatically discovers (locates) and displays **standalone resource**s in the Oracle Fail Safe Manager tree view when you select the Standalone

Resources folder from the tree view. Chapter 7, Chapter 9, and Chapter 10 contain information about how Oracle Fail Safe discovers each type of component that you can configure for high availability with Oracle Fail Safe.

4.5 Renaming Resources

Once a resource is added to a group, you must not change the resource name. If the resource name must be changed, then use Oracle Fail Safe Manager to remove the resource from the group and then, add it back to the group using the new name.

4.6 Using Oracle Fail Safe in a Multiple Oracle Homes Environment

Oracle Fail Safe supports the multiple Oracle homes feature. The following list describes the requirements for using Oracle Fail Safe in a multiple Oracle homes environment:

- Install Oracle Services for MSCS in any one Oracle home on all cluster nodes. Only one version of Oracle Services for MSCS can be installed and running on a node.

- Use the latest release of Oracle Fail Safe Manager to manage multiple clusters. See *Oracle Fail Safe Release Notes* for information about the compatibility of various versions on Oracle Fail Safe Manager and the Oracle Fail Safe server component.

> **Note:** You can install multiple versions of Oracle Fail Safe Manager on a system, but each version must be installed in a different Oracle home, and the latest release of Oracle Fail Safe Manager must be installed last.

- Each resource to be configured for high availability must be installed in the same Oracle home on all cluster nodes that are possible owners. The Verify Cluster operation validates this symmetry. See Section 6.1.1 for information about the Verify Cluster operation.

- All databases and listeners in a group must come from the same Oracle home.

  When you add a database to a group, an Oracle Net listener resource is added to the group also. Optionally, you can add an Oracle Management Agent resource to the group. See Section 9.2 for more information.

  The listener is created in the same Oracle home where the database resides.

4.7 Configurations Using Multiple Virtual Addresses

Before any resources, other than generic services, can be added to a group using Oracle Fail Safe Manager, one or more virtual addresses must be added to the group. Client applications connect to the resources in a group using one of the virtual addresses in the group.

You can add up to 32 virtual addresses to a group, prior to adding resources, by invoking the Add Resource to Group Wizard. In Oracle Fail Safe Manager, on the **Resources** menu, select **Add to Group.**

Note the following restrictions:

- At least one virtual address must be added to a group before you can add another resource to the group. Only generic services can be added to a group that does not already contain a virtual address.

- If the group contains one or more Oracle databases, then:

    – All virtual addresses that you plan to configure with one or more databases in a group must be added to the group before you can add any databases to the group.

    – All databases in a group must use the same set of virtual addresses that you specify for the first database that you add to the group. (The set of virtual addresses can contain as few as one address.)

    See Section 7.3.3.2 for more information about configuring multiple virtual addresses with Oracle databases.

When you add a virtual address to a group, the group is accessible by clients at the same network address, regardless of which cluster node is hosting the cluster.

Multiple virtual addresses in a group provide flexible configuration options. For example, users can access a database over the public network while you perform a database backup operation over the private network. Or you can allocate different virtual addresses on different network segments to control security, with administrators accessing the database on one segment, while users access the database on another segment.

When you add more than one virtual address to a group, Oracle Fail Safe Manager asks you to specify the address that clients can use to access the resources in that group. If you add more than one resource to a group (for example, a database and an Oracle Application Server), then you can dedicate one virtual address for users to access the database directly and another for users to access the Oracle Application Server. Alternatively, if there are many database users, then you can have some users access the database using one virtual address and the others use the other virtual address, to balance the network traffic.

See the online help in Oracle Fail Safe Manager for information about adding a virtual address to a group.

4.8 Adding a Node to an Existing Cluster

Instructions for installing the software to add a new node to an existing cluster are described in the *Oracle Fail Safe Installation Guide*. Once that task is completed, there is one final step. You must run the `Verify Group` command on each group on the cluster for which the new node is a possible owner.

Assume you add a new node to the cluster and install Oracle Fail Safe on that node along with the DLLs for the resources you intend to run on that node. The new node becomes a possible owner for these resources. If these resources have not yet been configured to run on the new node, when the group or groups containing them fail over to that node, then these resources cannot be restarted on that new node.

However, if you run the `Verify Group` command, then Oracle Fail Safe checks that the resources in the verified groups are configured to run on each node that is a possible owner for the group. If it finds a possible owner node where the resources in the group are not configured to run, then Oracle Fail Safe configures them for you.

Therefore, Oracle strongly recommends that you run the `Verify Group` operation for each group for which the new node is listed as a possible owner. Section 6.1.2 describes the `Verify Group` operation. You can also verify groups using the `FSCMD` command, as described in Chapter 5.

5

The FSCMD Command-Line Interface

Oracle Fail Safe provides a command-line interface as an alternative to using Oracle Fail Safe Manager for managing resources in a cluster. For example, using the `FSCMD` command at the command prompt, you can take Oracle resources offline or bring them online. The `FSCMD` command is useful if you want to manage Oracle resources from batch programs or scripts.

FSCMD

The FSCMD command enables you to perform many operations that you can perform with Oracle Fail Safe Manager. To run the FSCMD command, specify the full path, beginning with the location where Oracle Fail Safe Manager is installed: *Oracle_ Home*\fs\fsmgr\bin\fscmd.exe.

If you do not use this path, then the Windows operating system is not able to locate the FSCMD command.

Format

To use the FSCMD command, open a command prompt window and type the FSCMD command line using the following syntax:

FSCMD *action name* /CLUSTER=cluster-name [*qualifier*]

> **Note:**   For clarity, the full path for FSCMD is omitted from the syntax and examples in this chapter.

Description

Oracle Fail Safe supplies the FSCMD command, which you can use in scripts to do many functions that you can do with Oracle Fail Safe Manager. For example, you can use the FSCMD command to take cluster resources offline before you perform nightly backups, and then place the resources online again when the backups are complete.

You can run the FSCMD command on any system where Oracle Fail Safe Manager is installed. (The FSCMD software is a component of Oracle Fail Safe Manager.)

When you run the FSCMD command, you can specify the name of a log file to capture the results of long-running operations.

Command Parameters

The following are some command parameters.

action
Specifies the action that can be applied to the group, resource, or cluster. Use one of the actions described in the following table:

Action	Description
DISABLEISALIVE	Disables Is Alive and Looks Alive polling for the named database instance until one of the following occurs:
	■ You explicitly enable Is Alive polling with the FSCMD ENABLEISALIVE command.
	■ The instance is placed online.
	The Verify Group command displays a warning if it finds that Is Alive polling is disabled for an instance, but it does not enable Is Alive polling again.
	When Is Alive polling is disabled, the resource DLL writes a warning event to the Microsoft Windows event log to indicate that database polling is disabled.

Action	Description
DUMPCLUSTER	Prints cluster configuration information. Output from this operation is written to the file name that you specify with the /LOGFILE qualifier.
ENABLEISALIVE	Enables Is Alive and Looks Alive polling for the named database instance if they were disabled with the FSCMD DISABLEISALIVE command.
MOVEGROUP	Moves the group of resources to the node that you specify with the /NODE command qualifier. Output from this operation is written to the file name that you specify with the /LOGFILE qualifier.
ONLINEGROUP	Places the group online.
ONLINERESOURCE	Places the resource online.
OFFLINEGROUP	Takes the group offline.
OFFLINERESOURCE	Takes the resource offline. For Oracle database resources, this operation requires the /OFFLINE command qualifier.
VERIFYGROUP	Verifies a single group configured by Oracle Fail Safe. Output from this operation is written to the file name that you specify with the /LOGFILE qualifier.
VERIFYALLGROUPS	Verifies all groups configured by Oracle Fail Safe. Output from this operation is written to the file name that you specify with the /LOGFILE qualifier.
VERIFYCLUSTER	Verifies the cluster configuration. Output from this operation is written to the file name that you specify with the /LOGFILE qualifier.

The action parameter must be the first argument to the FSCMD command.

name
The name of the resource or group on which you want the FSCMD command to take action. For example, PERSONNEL.world is a valid name for a single-instance database resource.

This parameter must follow the action parameter. The name parameter is required for all FSCMD command actions except for the DUMPCLUSTER, VERIFYALLGROUPS, and VERIFYCLUSTER actions.

Command Qualifiers

The following are some command qualifiers.

/CLUSTER=cluster-name
Specifies the name of the cluster on which the FSCMD command will run.

This qualifier is required.

/LOGFILE=file-name
Specifies the location of the log file that a long-running operation creates when the DUMPCLUSTER, MOVEGROUP, VERIFYCLUSTER, VERIFYGROUP, or VERIFYALLGROUPS action is performed. If you do not specify the LOGFILE qualifier, then the log file is written to the current output device, which is typically the system console.

This qualifier is optional.

/NODE=*node-name*

Use this qualifier only with the MOVEGROUP action to specify the name of the node where you want to move a group.

This qualifier is required when you specify the MOVEGROUP action.

/OFFLINE=*offline-option*

Use this qualifier only when you specify the OFFLINERESOURCE action to take an Oracle resource offline. For Oracle databases, if you do not supply one of the offline-option modes from the following table, then the resource is taken offline in the immediate mode (the default):

Mode	Description
Abort	Shuts down a database instantaneously by aborting the database instance. If possible, use the abort mode of shutdown *only* in the following situations: ■ The database or one of its applications is functioning irregularly *and* neither the immediate nor the normal mode of shutdown works. ■ You must shut down the database instantaneously (for example, you know a power shutdown is going to occur in one minute). Database recovery procedures are performed when the database is restarted.
Immediate	Shuts down a database immediately by terminating SQL statements in progress, rolling back uncommitted transactions, and disconnecting users. The immediate mode is the default mode to take a resource offline. If possible, choose the immediate mode *only* in the following situations: ■ A power shutdown is going to occur soon. ■ The database or one of its applications is functioning irregularly. ■ You are preparing to perform a database backup operation.
Normal	Shuts down a database by: ■ Disallowing new connections after the database shutdown command is issued ■ Waiting for all connected users to disconnect before actually shutting down the database
Transactional	Shuts down the database only after all of the current transactions have completed.

/DOMAIN=*domain-name*

Specifies the domain in which the user account (specified with the /USER qualifier) is located.

/PWD=*password*

Specifies the password for the account that is specified with the /USER qualifier.

/USER=*username*

Specifies the user name for a domain account that has Administrator privileges on all cluster nodes.

Usage Notes

The following are some FSCMD command usage notes:

■ The *action* parameter must be the first argument to the FSCMD command.

■ The parser accepts either a slash (/) or a hyphen (-) as the delimiter between command qualifiers.

- The command parameters and qualifiers are not case-sensitive; you can use either uppercase or lowercase letters, or a combination of both cases.

- A group or resource name that contains one or more spaces must be within quotation marks; for example, "Sales Group."

- Authentication information (domain, username or password) for the FSCMD utility is optional. If no authentication information is supplied, then FSCMD attempts to connect to the server using its default local authentication information. If any authentication is specified on the command line (domain, username or password) but one of the items is missing, then FSCMD prompts for the missing information. For example, when an FSCMD command is issued, if the /domain and /username switches are present but the /password switch is not provided then FSCMD displays a prompt asking for the user to supply the password; the password does not echo to the console.

- You must enter a space between qualifiers in the command line. The following examples show correct and incorrect spacing in the command line:

  Correct usage:

  ```
  FSCMD onlineresource salesdb.world /CLUSTER=ntclu-160 /USER=smith
  /PASSWORD=smithpwd  /DOMAIN=newengland
  ```

  Incorrect usage:

  ```
  FSCMD onlineresource
  salesdb.world/CLUSTER=ntclu-160/USER=smith/PWD=smithpwd/DOMAIN=newengland
  Invalid number of parameters.
  ```

- Use the Windows Event Viewer to display events that have been reported during the processing of FSCMD commands.

Command Examples

The following are some command examples.

Example 1

The following command places online an Oracle database named `salesdb.world`:

```
FSCMD onlineresource salesdb.world /CLUSTER=ntclu-160 /USER=smith /PWD=smithpwd
/DOMAIN=newengland
```

Example 2

The following command verifies all groups on the cluster NTCLU-160 and writes verify operation output to the log file `C:\temp\fsverify.log`.

```
FSCMD verifyallgroups /LOGFILE=c:\temp\fsverify.log /CLUSTER=ntclu-160
 /USER=smith /PWD=smithpwd /DOMAIN=ORANT
```

Example 3

The following command takes the Oracle database offline immediately:

```
FSCMD offlineresource salesdb.world /CLUSTER=NTCLU-160 /USER=smith
 /PWD=smithpwd /DOMAIN=ORANT /OFFLINE=immediate
```

Example 4

The following command takes a group called Disk Group 1 offline:

```
FSCMD offlinegroup "Disk Group 1" /CLUSTER=ntclu-160 /USER=smith /PWD=smithpwd
/DOMAIN=ORANT
```

Script Example

The following script performs an online backup of a database called `db.world`. This example assumes the database is contained in a group called FS Group1 running on a cluster called NTCLU-140, on node NTCLU-141. (To perform a closed backup of the database, you would use the `FSCMD offlineresource` command to take the database resource offline in Step 2, and place it online in Step 6 using the `FSCMD onlineresource` command, instead of disabling and enabling Is Alive polling.)

```
REM    This script shows an example of performing an online backup operation
REM    on an Oracle Fail Safe database.
REM
REM    1. Move the group FS Group1 that contains the database to the node on
REM    which the backup operation will run. Alternatively, you can create file
REM    share resources for each cluster disk to let the backup software
REM    access the drives through a virtual server address regardless of which
REM    cluster node currently owns them.

fscmd movegroup "FS Group1" /node=NTCLU-141 /cluster=NTCLU-140 /USER=smith
 /PWD=smithpwd /DOMAIN=ORANT

REM    2. Disable Is Alive polling for the database resource. This step allows
REM       you to keep the database online during the backup operation, but
REM       prevents Oracle Fail Safe from attempting to fail over the database
REM       during the online backup operation.

fscmd disableisalive db.world /cluster=NTCLU-140 /USER=smith /PWD=smithpwd
/DOMAIN=ORANT

REM    3. Mark the beginning of the online tablespace backup operation for each
REM       database tablespace. In this example there are two tablespaces, users
REM       and indx. The database must be in ARCHIVELOG mode to back up the
REM       tablespaces. (If you use Recovery Manager (RMAN) for the online
REM       backup operation, you are not required to mark the beginning
REM       and end of the tablespace backup operations, nor would you copy the
REM       database files, as described in steps 3, 4, and 5.)

SQL> ALTER TABLESPACE users BEGIN BACKUP;
SQL> ALTER TABLESPACE indx BEGIN BACKUP;

REM    4. When you are certain the previous operations have completed, begin the
REM       backup operation. As an example, the following lines
REM       copy files using the copy function of the operating system.

copy e:\ofsdb\ofs1\data\*.ora e:\backup\data
copy e:\ofsdb\ofs1\log\*.ora e:\backup\log
copy e:\ofsdb\ofs1\param\*.ora e:\backup\param

REM    5. Indicate the end of the online backup operation of each tablespace.

SQL> ALTER TABLESPACE users END BACKUP;
SQL> ALTER TABLESPACE indx END BACKUP;

REM    6. Reenable Is Alive polling for the database resource.

fscmd enableisalive db.world /cluster=ntclu-140 /USER=smith /PWD=smithpwd
/DOMAIN=ORANT
```

```
REM   The backup operation is complete.
```

6

Troubleshooting Tools

This chapter provides general information about the troubleshooting tools provided with Oracle Fail Safe Manager. The following topics are discussed in this chapter:

- Verify Operations
- Dump Cluster
- Verify Security Parameters
- Finding Additional Troubleshooting Information

Note that Oracle Fail Safe provides a centralized message facility. When you perform an action that results in an error, the system locates the message associated with the error and displays it. You can find more information about these messages in the *Oracle Fail Safe Error Messages* manual.

6.1 Verify Operations

Oracle Fail Safe provides a family of tools to help you verify cluster components and the cluster environment to validate the status of nodes, groups, and resources. If a discrepancy or a problem is found, then the verify operation takes the appropriate action to fix any potential or actual problems.

Figure 6–1 shows the verify commands in the Troubleshooting menu.

Figure 6–1   Troubleshooting Menu and Verify Commands

Table 6–1 describes the verify commands and provides references for more information.

Table 6–1   Verify Commands for Troubleshooting

Tool	Description	Reference
Verify Cluster	Validates the Oracle Fail Safe installation, the Oracle product installation (including Oracle homes and product version numbers), cluster network configuration, and cluster resource DLL registration.	Section 6.1.1

Table 6–1   (Cont.) (Cont.) Verify Commands for Troubleshooting

Tool	Description	Reference
Verify Group	Validates that the group resources and their dependencies are configured correctly.	Section 6.1.2
Verify Standalone Database	Validates the standalone database instance and removes any old configuration information that may remain on another node.	Section 6.1.3

You can use the verify commands at any time to validate your cluster, group, or standalone database. If problems are found during verification, then Oracle Fail Safe prompts you to fix them or returns an error message that further describes the problem.

If errors are returned when you run one of the verify commands, then fix the errors and then rerun the verify command. Repeat this process until the verify operation runs without errors.

6.1.1 Verify Cluster

The Verify Cluster operation validates the installation and network configuration of the cluster. You can perform a cluster verification at any time. From the Oracle Fail Safe Manager menu bar, select **Troubleshooting**, then select **Verify Cluster.**

The first time you connect to a cluster after installing or upgrading the Oracle Fail Safe software, you are prompted to run Verify Cluster. You can run Verify Cluster at any time, however, you must run it whenever the cluster configuration changes. The Verify Cluster operation verifies that:

- Each Oracle home name into which Oracle software is installed is the same on all cluster nodes

  If, for example, OFS is the Oracle home name for the Oracle Fail Safe software on one cluster node, then OFS must be the Oracle home name on all nodes in the cluster where Oracle Fail Safe is installed. Similarly, if OfsDb is the Oracle home name for the Oracle database software on one cluster node, then it must be the Oracle home name on all nodes in the cluster where the Oracle database software is installed.

- The Oracle Services for MSCS release is identical on all nodes

- The resource providers (components) are configured identically on at least two of the nodes that are possible owners for each resource

- The Host Name/IP Address mappings resolve consistently across all nodes in the cluster

  If there is a problem with inconsistent mapping, then the Verify Cluster command returns errors indicating that the order of network adapters may be incorrect. See Appendix A for details.

Verify Cluster also registers Oracle resource DLLs with Microsoft Cluster Server (MSCS).

Figure 6–2 shows the output from a typical Verify Cluster operation.

Figure 6–2   Clusterwide Operation Window for Verify Cluster
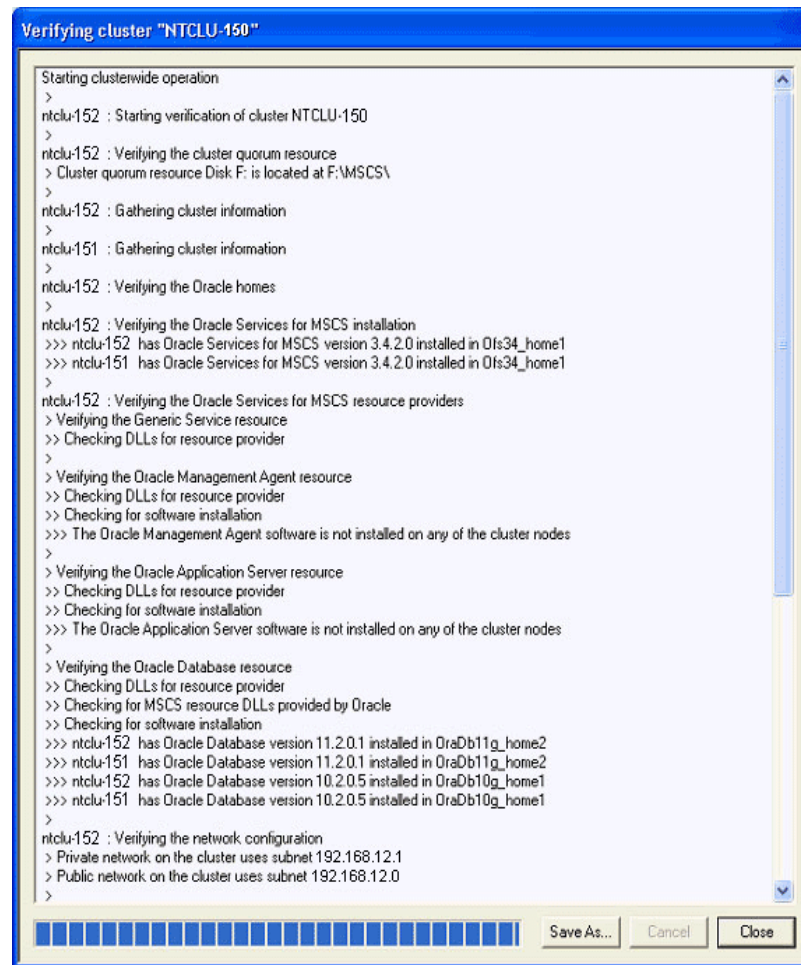
If you run the `Verify Cluster` operation and it does not complete successfully, then it may indicate one or more of the following problems:

- A problem exists in the configuration of the hardware, network, or the MSCS software.

- A problem exists in the symmetry of the Oracle homes and versions.

- A problem exists with the Oracle Fail Safe installation (for example, with the symmetry of the resource providers).

If the operation completes successfully, but you are having problems with Oracle Fail Safe, then the problem is based in the Oracle Fail Safe configuration.

6.1.2 Verify Group

The `Verify Group` operation does the following to ensure that a group performs correctly:

- Checks all resources in a group and confirms that they have been configured correctly on all nodes that are possible owners for the group.

- Updates the dependencies among resources in the group.

- After prompting you, repairs a group that is misconfigured.

You can run the `Verify Group` operation at any time. However, you *must* run it when any of the following occurs:

- A group or resource in a group does not come online.

- Failover or failback do not perform as you expected.

- You add a node to the cluster.

To verify a group select the group from the Oracle Fail Safe Manager tree view and then from the Oracle Fail Safe Manager menu bar, select **Troubleshooting**, then **Verify Group.**

Or, you can run a `Verify Group` operation using the `FSCMD` command `VERIFYGROUP` (see Chapter 5). The `FSCMD` command also provides a `VERIFYALLGROUPS` command that lets you verify all groups configured by Oracle Fail Safe on a given cluster. You can run the `VERIFYGROUP` and `VERIFYALLGROUPS` commands in scripts as batch jobs.

You can watch the progress of the `Verify Group` operation and view the status of the individual resources in the group as Oracle Fail Safe verifies the group.

Figure 6–3 shows the output from a `Verify Group` operation.

Figure 6–3   Clusterwide Operation Window for Verify Group
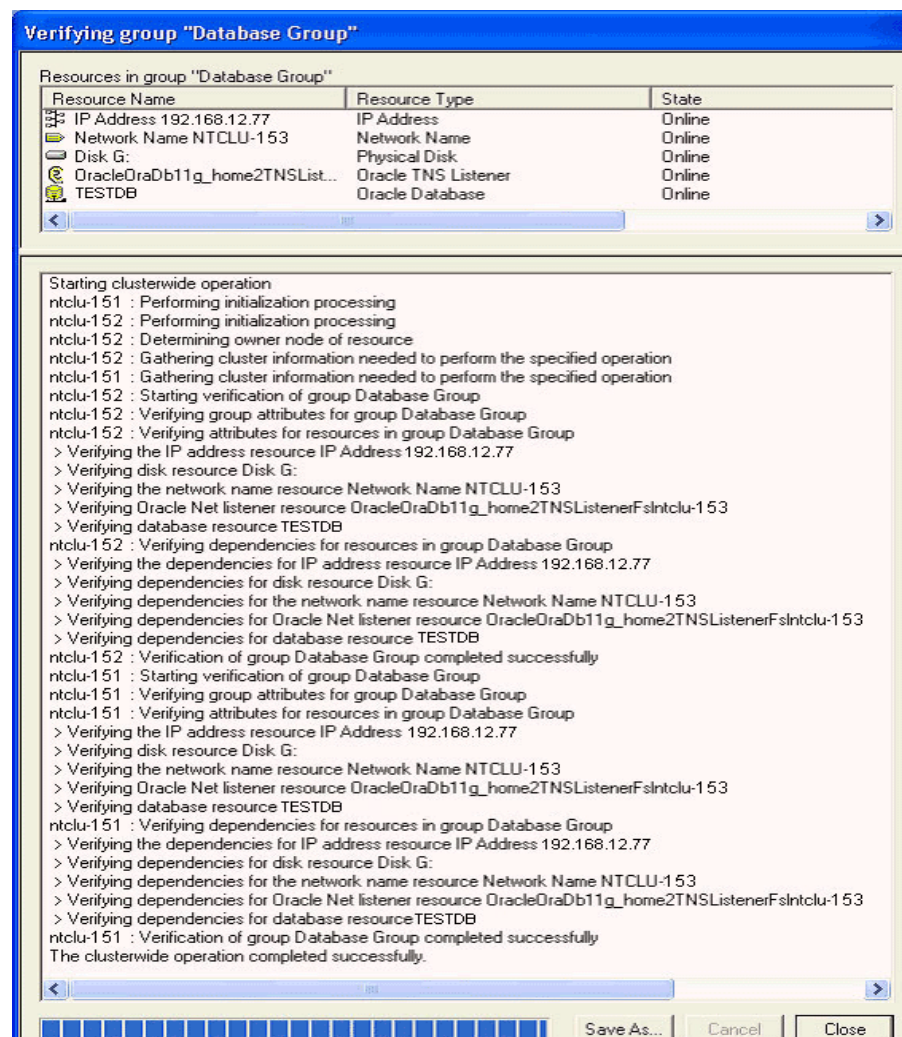
6.1.3 Verify Standalone Database

You can validate a standalone database at any time by using the `Verify Standalone Database` operation. To run the `Verify Standalone Database` command, select the database from the Oracle Fail Safe Manager tree view, and then from the Oracle Fail Safe Manager menu bar, select **Troubleshooting**, then select **Verify Standalone Database.**

The `Verify Standalone Database` operation performs validation checks to ensure that the standalone database is configured correctly on the node where it resides and to remove any references to the database that may exist on other cluster nodes. (References to the database may exist on other cluster nodes if the database was once added to a group and then later removed.) This ensures that the database can be made highly available using Oracle Fail Safe.

Oracle recommends that you use the `Verify Standalone Database` command on a standalone database before you add it to a group. You can also use it whenever you have trouble accessing a standalone database. However, note that Oracle Fail Safe stops and restarts the database during the verify operation.

For example, you may perform a verification:

- If a failure occurs when you try to add a database to a group.

- If you used an administrator tool other than Oracle Fail Safe Manager to perform an operation on the database and the database now is inaccessible.

- If you removed or deinstalled the MSCS software from the cluster nodes without first removing the Oracle Fail Safe software (for example, during a software upgrade). This is described in more detail in the *Oracle Fail Safe Installation Guide*.

Figure 6–4 shows the Verify Standalone Database dialog box in which you enter valid database information and account information for a standalone database.

Figure 6–4   Verify Standalone Database Dialog Box

To use the Verify Standalone Database dialog box, you must specify:

- The service name of the standalone database, in the Service Name field

- The instance name of the standalone database, in the Instance Name field

- The database name of the standalone database, in the Database Name field

- The parameter file disk, path name, and file name for the initialization parameter file for the standalone database, in the Parameter File field

- The account that Oracle Fail Safe must use to attach to the database, in the Account area.

Oracle Fail Safe uses this information to:

- Fix clusterwide problems with Oracle Net

- Check that the standalone database is on a cluster disk

- Ensure that Oracle Fail Safe can attach to the database

If a standalone database is open and you run a `Verify Standalone Database` operation, then the operation does not restart the database.

If a standalone database is not open or if the database is stopped, then Oracle Fail Safe asks your permission to stop and restart the database instance. Subsequently, Oracle Fail Safe opens the database for access.

Figure 6–5 shows the output from a typical `Verify Standalone Database` operation in a Clusterwide Operation window.

Figure 6–5   Clusterwide Operation Window for Verify Standalone Database
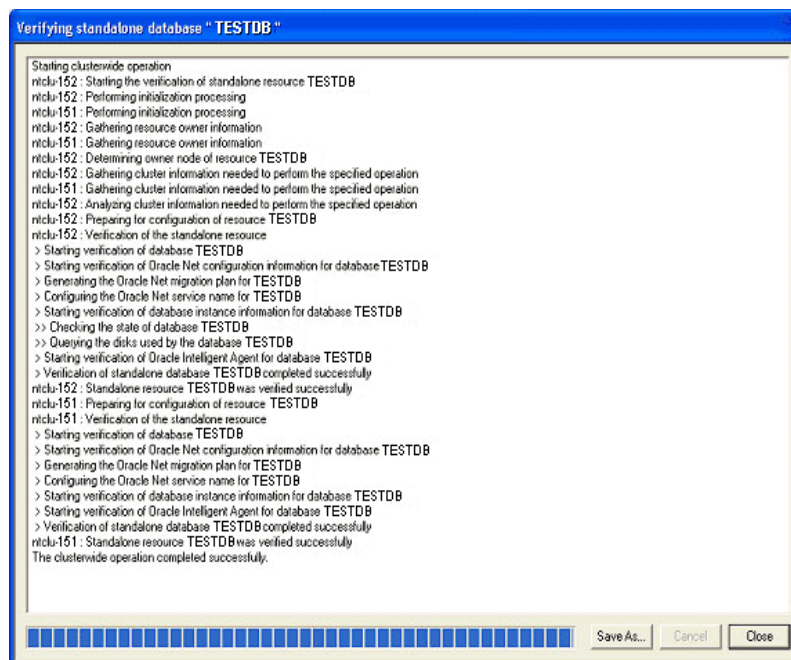
If any problems are found during verification, then the `Verify Standalone Database` operation prompts you before it attempts to fix them. For example, imagine that you try to add a database to a group, but the operation fails because of an Oracle Net problem. You can run the `Verify Standalone Database` command to fix the network problem and subsequently add the database to a group.

6.2 Dump Cluster

Oracle Fail Safe provides the Dump Cluster command to display Oracle Fail Safe Manager cluster data in a window. Run this command periodically (and save the output) to maintain a record of changes made to the cluster over time, or run it at the request of customer support so as to provide a snapshot of the cluster environment.

Data displayed when you run the Dump Cluster command includes:

- Information related to the operating system (including the location of the quorum disk)
- Public and private network information
- Resources registered with the cluster
- Group failover and failback policies

You can optionally save the Dump Cluster data to a file by clicking **Save As**.

To run the Dump Cluster command, select the cluster from the Oracle Fail Safe Manager tree view, and then from the Oracle Fail Safe Manager menu bar, select **Troubleshooting**, and then select **Dump Cluster.**

Figure 6–6 shows the portion of the Dump Cluster command output that provides information about the NTCLU-150 cluster and some of its resources.

Figure 6–6   Dump Cluster Clusterwide Operation
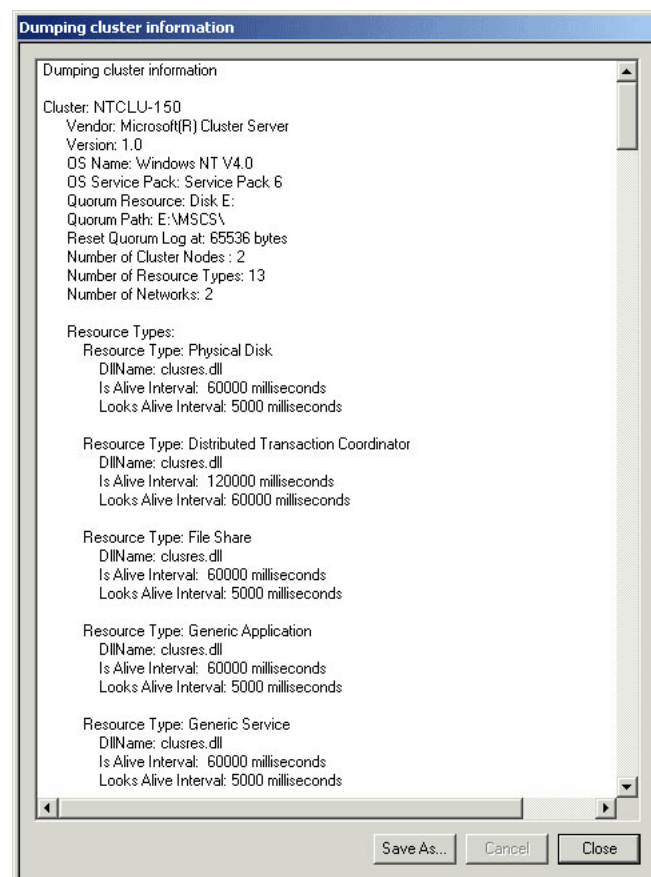
6.3 Verify Security Parameters

Oracle Fail Safe provides the `fssvr` command qualifier, `/GETSECURITY`, which displays security information about the system where the command is run. Run the `fssvr` command qualifier, `/GETSECURITY` on each cluster node to help diagnose FS-1075n errors (where n is a value between 0 and 7, inclusive).

The command and its associated output must be similar to the following:

```
fssvr /getsecurity

Looking up user account information for OracleMSCSServices.
The user account must be a domain user acount with local Administrator
privileges.  The user account must also have the 'Log on as batch job'
privilege.

    User account specified for OracleMSCSServices is NEDCDOMAIN\cluadmin
    User account specified has local Administrator privileges
    User account has the 'Log on as batch job' privilege

Looking up user account information for Cluster Service. The user account
must be a domain user account with local Administrator privileges. The user
account must also have the 'Log on as batch job' privilege.

    User account specified for Cluster Service is NEDCDOMAIN\cluadmin
    User account specified has local Administrator privileges
    User account has the 'Log on as batch job' privilege

Checking to see if DCOM is enabled.  DCOM must be enabled.
    DCOM is enabled.
```

6.4 Finding Additional Troubleshooting Information

This chapter describes how to use the Oracle Fail Safe Manager family of troubleshooting tools. Additional information is available as follows:

■ Information about troubleshooting a specific component can be found in Chapters 7 through 9, each of which describes how to configure a particular component for high availability.

■ Information about troubleshooting network configuration problems is described in Appendix A.

■ Because Oracle Fail Safe is layered upon Microsoft Cluster Server software, you may need to refer to the MSCS documentation to troubleshoot problems with the cluster service, interconnect, and hardware configuration.

■ If you are unable to start Oracle Fail Safe, then start the Windows Event Viewer and look at the application log. Oracle Services for MSCS usually logs an event identifying the problem.

7

Configuring Single-Instance Databases for High Availability and Disaster Tolerance

Oracle Fail Safe provides high availability for single-instance Oracle databases (both Oracle Database Standard and Enterprise Editions) running on Windows clusters configured with MSCS.

By making a single-instance Oracle database highly available, you ensure that even when a cluster node is shut down or fails, applications that access that database suffers only a momentary loss of connection with the database while the database is restarted on another cluster node. Applications can automatically reconnect to the database after such a failover event occurs using transparent application failover, resulting in a failover that is not apparent to users.

This chapter discusses the following topics:

- Discovering Standalone Single-Instance Databases
- Oracle Net Configuration for Standalone Single-Instance Databases
- Adding Single-Instance Oracle Databases to a Group
- Oracle Net Listener Resource Creation and Configuration
- Security Requirements for Single-Instance Databases
- Optimizations for Single-Instance Database Recovery
- Performing Administrative Tasks on a Single-Instance Fail-Safe Database
- Database Homes
- Configuring Transparent Application Failover (TAF)
- Handling Errors and Troubleshooting Problems with Databases

7.1 Discovering Standalone Single-Instance Databases

Oracle Services for MSCS discovers standalone single-instance databases (those that are not in a group) to display them in the Oracle Fail Safe Manager tree view. Standalone single-instance databases created with the `Create Sample Database` command (sample databases) and standalone single-instance databases created using other methods are discovered in different ways, as follows:

- Sample standalone databases

  Sample standalone single-instance databases are discovered by looking in the Windows registry under the key:

  `HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\FailSafe\SampleDB`

■   Other standalone single-instance databases

Other standalone single-instance databases (those not created with `Create Sample Database` command) are discovered by parsing the `tnsnames.ora` file on each cluster node and looking for valid **net service name** entries. Standalone single-instance databases are discovered when the following conditions are true:

–   The standalone single-instance database has an instance on a cluster node.

–   The standalone single-instance database has a valid entry in the `tnsnames.ora` file, including:

*   A `SID` that matches the database instance on the node or a `SERVICE_NAME` that matches the service name in the database parameter file.

*   A **host name** or IP address (Oracle Fail Safe does not support the use of alias network names.)

If no valid net service name entry can be found for an instance, then the instance name is used to represent the database instance in Oracle Fail Safe Manager.

To view the standalone single-instance databases (and other standalone resources) in the Oracle Fail Safe Manager tree view, expand each node folder, then expand the Standalone Resources folder.

7.2 Oracle Net Configuration for Standalone Single-Instance Databases

The following sections briefly summarize the Oracle Net configuration for standalone single-instance databases.

7.2.1 Updating the Oracle Net Configuration for a Database Created Using DBCA

If you use the Database Configuration Assistant (DBCA) to create a standalone single-instance database, then DBCA adds information about the new database in the Oracle Net configuration, as follows:

1. DBCA adds the `SID_DESC` parameter for the database to the default listener in the default `listener.ora` file. This `SID_DESC` parameter consists of the database `SID` name.

2. DBCA adds the net service name entry for the database to the `tnsnames.ora` file in the database's Oracle home only. This entry consists of the `SERVICE_NAME` parameter.

After DBCA configures the Oracle Net information, Oracle Fail Safe displays the new single-instance database in the Oracle Fail Safe Manager tree view (under Standalone Resources for the cluster node on which the database was created).

If there are multiple Oracle homes on the cluster node where the database is created, then run the `Verify Standalone Database` command in Oracle Fail Safe Manager on the new database. Oracle Fail Safe checks the `tnsnames.ora` files in all Oracle homes. When it detects that a `tnsnames.ora` file does not contain the net service name entry for the database, it asks if you want to update that file. If you select Yes, then Oracle Fail Safe adds a net service name entry for the new database.

If the default domain name values (for example, values of `names.default_domain` parameters in the `sqlnet.ora` files) are different across Oracle homes, then the net service name entry is not accessible from some Oracle homes. To resolve this problem, edit the `tnsnames.ora` file in each Oracle home and append the default domain name of the respective Oracle home to the net service name entry.

7.2.2 Listener Must Use IP Address, Not Host Name

If the system **host name** is used in the definition of a listener, then this listener listens on all IP addresses on that node, not just the IP address associated with the host name. This causes conflicts later when you add a database to a group and listeners configured with Oracle Fail Safe are defined to listen on a virtual address.

To avoid this problem, the listener must use the node IP address for its host entry instead of the host name. When you add a single-instance database to a group, if Oracle Fail Safe finds a listener using a host entry, then you are asked if you want Oracle Fail Safe to modify this entry to use an IP address. If you do not, then the Add Resource to Group operation does not continue.

The following is an example of an *invalid* entry in an Oracle Fail Safe environment:

```
LISTENER =
   ....
   (ADDRESS=
       (PROTOCOL=TCP)
       (HOST=NTCLU-152)
       (PORT=1521)
   )
```

The following is an example of a *valid* entry in an Oracle Fail Safe environment:

```
LISTENER =
   ....
   (ADDRESS=
       (PROTOCOL=TCP)
       (HOST=192.0.2.254)
       (PORT=1521)
   )
```

7.2.3 SID List Entries and Upgrades to Oracle Database Software

When you upgrade Oracle database software, you must also upgrade the listener. To maintain your entries in the listener SID list from an earlier version, move them to the version of the listener to which you are upgrading, as follows:

1. Locate the listener.ora files that correspond to the database software release from which you will upgrade.

2. Copy the SID_DESC entries from the SID list of the listener you are upgrading from and add them to the SID list of the listener that corresponds to the Oracle database software release to which you are upgrading. For example, the updated SID list may appear similar to the following:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME= c:\oracle)
      (PROGRAM = extproc)
    )
    (SID_DESC=
      (SID_NAME=ORCL)
    )
    (SID_DESC=
      (SID_NAME=OFS2)
    )
  )
```

3. Stop the listener of the database from which you are upgrading and change the startup state of the listener to manual.

4. Start the listener of the database to which you have upgraded.

   If you do not use the listener that corresponds to the release of the database software you are using, then you will run into problems when you perform the following Oracle Fail Safe operations:

   – Add database resource to group

   – Verify group

   – Remove database resource from group

7.2.4 Configuring Oracle Net on Nodes with Multiple Listeners

When Oracle Fail Safe searches for a standalone database listener, it looks at the listeners defined in all Oracle homes. When multiple Oracle servers are installed (in one or more Oracle homes), it is possible to have multiple listeners. Even when there are multiple listeners, Oracle Fail Safe can find the standalone database listener.

After you install additional Oracle homes on your system, you must decide how many listeners to use and which listeners to use on the system. Then:

1. Remove the definitions of listeners that you do not use from the listener.ora files. This ensures that Oracle Fail Safe does not find unnecessary listeners.

2. Ensure that no two listeners listen on the same address of the SID. (An Oracle listener does not start if another active listener is listening on the same address or SID.)

3. Start the listeners to be used. Ensure that these listeners are set to start automatically in the Windows Services startup list.

The state of a listener defined in a listener.ora file affects the result when Oracle Fail Safe searches for the listener of a standalone database. The following list shows the order in which Oracle Fail Safe looks for the listeners:

1. Listener is started.

2. Listener is stopped.

3. Listener has no Windows service defined.

For example, suppose you have a database with two listeners in two different Oracle homes, and assume that the listener in Home 1 is stopped and the listener in Home 2 is started. When Oracle Fail Safe looks through the two homes on the system, it finds the listener in Home 2 because in the first pass, Oracle Fail Safe looks only for listeners that are started. As soon as a started listener is found, Oracle Fail Safe stops looking. If a listener that is started is not found, Oracle Fail Safe looks through the list for listeners that are stopped, and so on.

> **Note:** Ensure that the listeners of standalone single-instance databases are in the intended state, stopped or started, before you run any Oracle Fail Safe operations.

7.2.5 Shared Server Configuration and a Standalone Database

When a database is configured for high availability, Oracle Fail Safe makes adjustments to the default listener. This affects the Oracle Net configuration for all

databases, including standalone databases. Therefore, all standalone databases in an Oracle Fail Safe environment require some adjustments to the Oracle Net configuration if any database in the cluster has been made highly available.

If the shared server configuration for standalone single-instance databases relies on the default listener, then no listener parameters are specified in the database parameter file. (The default listener is a listener that listens on the **host name** of the node, the default port number, and TCP protocol.) In this case, the configuration will no longer work after Oracle Fail Safe has changed the default listener to use an IP address in place of the host name.

Resolve this problem by doing the following:

1. Add the LOCAL_LISTENER parameter to the database initialization parameter file. The LOCAL_LISTENER parameter specifies a network name that resolves to an address of the Oracle Net default listener.

   Locate the database initialization parameter file of the database and add the LOCAL_LISTENER parameter to the file.

   ```
   LOCAL_LISTENER = network-name
   ```

2. Determine the address of the Oracle Net default listener.

   Find the definition of the default listener in the listener.ora file of the database home. In the definition, identify the first address that uses the TCP protocol.

   For example, assume that the default listener is defined as follows:

   ```
   LISTENER =
     (DESCRIPTION_LIST=
       (DESCRIPTION=
         (ADDRESS_LIST=
           (ADDRESS=
             (PROTOCOL=TCP)
             (HOST=192.0.2.1)
             (PORT=1521)
           )
         )
       )
     )
   ```

   Then the first address is:

   ```
   (ADDRESS_LIST=
           (ADDRESS=
             (PROTOCOL=TCP)
             (HOST=192.0.2.1)
             (PORT=1521)
   ```

3. Create a network-name entry in the tnsnames.ora file.

   In the tnsnames.ora file, create an entry for the network-name using the address found in Step 2.

   In this example, the entry is as follows:

   ```
   network-name= (ADDRESS=
                     (PROTOCOL=TCP)
                     (HOST=192.0.2.1)
                     (PORT=1521)
                     )
   ```

This change will take effect when the database is restarted.

7.3 Adding Single-Instance Oracle Databases to a Group

To configure a single-instance Oracle database for high availability, you add it to a group that currently contains at least one virtual address. Oracle Fail Safe adds all other resources that the single-instance Oracle database requires. Typically, the group includes the following resources:

- One or more virtual addresses, each of which consists of an IP address and network name

- The Oracle database instance

- All disks used by the Oracle database

- An Oracle Net network listener that listens on the virtual address (or addresses) of the group for connection requests to the databases in the group

- An Oracle Management Agent configured to use a group's virtual addresses (if Oracle Enterprise Manager is used to manage the database)

7.3.1 Before You Get Started

Before you add a single-instance database to a group, note the following:

- All files used by the single-instance database must be on the shared cluster disks, except the database initialization parameter file, which can be placed on a private disk or on a shared cluster disk. See Section 7.3.3.3 for more information about the placement of the initialization parameter file.

- Resources must belong to one group only. Therefore, if two single-instance databases share the same disk drives, then both databases must be in the same group.

- In a failover, the data in a temporary table does not fail over. Operations that involve the use of temporary tables and tablespaces (such as sorts and hash joins) re-create any needed temporary objects when restarted on the failover node. However, you must review applications that rely on the existence of specific data in temporary tables to be sure they function as expected.

  Refer to the Temporary Tables discussion in the *Oracle Database Concepts* manual for more information about temporary tables.

- The group must contain at least one virtual address.

- Database service names must be unique across the cluster.

  To ensure that the properties of standalone and cluster resources are discovered and displayed correctly by Oracle Fail Safe Manager and Oracle Enterprise Manager, each resource must have a unique name within the cluster. It may be necessary to specify names that are different from default values or to change the default names of resources.

7.3.2 Configuration Steps

Table 7–1 provides a quick reference to the tasks needed to configure a single-instance Oracle database for high availability. For detailed instructions about a particular task, see the online help and tutorial. From the Oracle Fail Safe Manager menu bar, select

Help, then select **Search for Help on** or **Help,** and then select **Tutorial** for step-by-step instructions.

Table 7–1   Steps for Configuring Databases

Step	Procedure	Oracle Fail Safe Manager Procedure
1	Ensure that the Oracle database software is installed on a private disk on each node in the cluster that you intend to be a possible owner for the Oracle database.	See the Oracle database documentation for installation information.
2	Create a group and add one or more virtual addresses.	On the **Groups** menu, select **Create** to open the Create Group Wizard. The wizard helps you to set up failover and failback policies and automatically opens the Add Resource to Group Wizard to let you add a virtual address to the group. To add additional virtual addresses to the group select **Resources**, then select **Add to Group**.
3	Create a sample database, if desired.	On the **Resources** menu, select **Create Sample Database** to create a sample standalone single-instance database on which you can try out the features of Oracle Fail Safe before using them on a production database. Do not use the sample database for production work.
4	Verify the standalone database.	On the **Troubleshooting** menu, select **Verify Standalone Database** to validate the database and Oracle Net configuration for the database. This command ensures that Oracle Fail Safe can attach to the database, and confirms that the standalone database is located on a cluster disk.
5	Add the Oracle database to the group.	On the **Resources** menu, select **Add to Group,** and then select **Oracle Database** to open the Add Resource to Group Wizard. The wizard helps you configure the single-instance Oracle database for high availability.
6	Modify the tnsnames.ora file on each client system.	Configure clients (modify the tnsnames.ora file on each client system using a network configuration tool) to recognize the virtual server. See Section 7.4 for more information.

7.3.3 Configuration Data for Oracle Databases

Oracle Fail Safe Manager provides the Add Resource to Group Wizard to assist you in configuring a single-instance Oracle database for high availability. The pages presented in the wizard vary, depending on the number of virtual addresses currently in the group, and the number of nodes in the cluster.

Typically, each group has one virtual address, but more complex configurations may have multiple virtual address. To perform a typical configuration using the Add Resource to Group Wizard, you need the following data:

- Identity of the single-instance Oracle database, including the service name, instance name, database name, and specification for the database initialization parameter file

- The database SYSDBA (usually, SYS) account and password

If you add a database to a group that currently contains multiple virtual address, then you are also asked to specify the virtual address or addresses that you want your clients to use to access the database.

The following sections describe in detail the configuration requirements for single-instance databases.

7.3.3.1 Choose Nodes

If you are adding a database to a group and the cluster consists of more than two nodes, then you are asked to specify the nodes, which must be possible owners for the database by specifying a list of selected nodes, as shown in Figure 7–1. To specify that a particular node must not be a possible owner for the database, select the node from the Selected Nodes list and click the left arrow.

Section 2.6.7 describes in detail the concept of the possible owner nodes list.

Figure 7–1   Choose Nodes Wizard Page When All Nodes Are Available
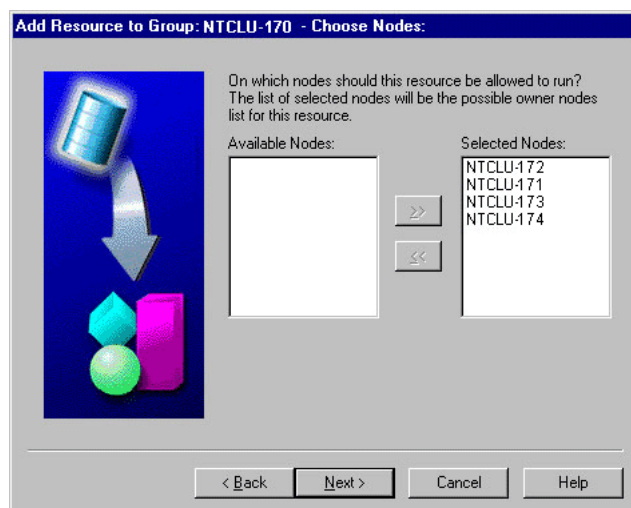
If you are adding a single-instance database to a group and the cluster consists of two or more nodes, but one or more nodes are unavailable, then you are also asked to specify which nodes must be possible owners for the database. In this case, the wizard page displays which nodes are unavailable and why, as shown in Figure 7–2.

Figure 7–2   Choose Nodes Wizard Page When Any Node Is Unavailable
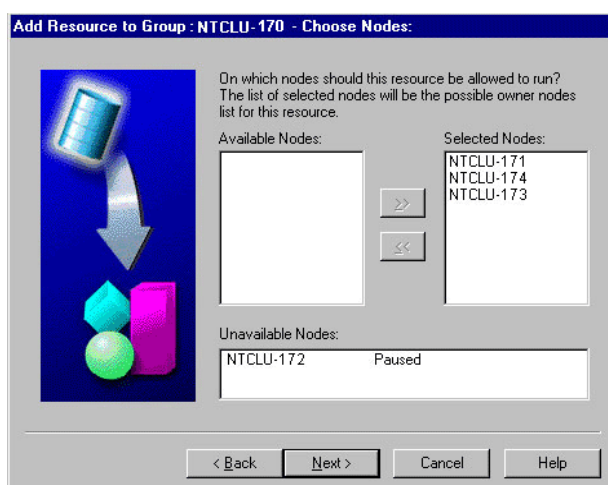
7.3.3.2 Virtual Address

If the group to which you are adding a single-instance database contains multiple virtual address, then the Add Resource to Group Wizard asks you which of the virtual addresses in the group you want clients to use when they access the database or

databases in the group, as shown in Figure 7–3. This page is not displayed if the group to which you are adding a database contains only one virtual address.

Figure 7–3   Database Virtual Address Wizard Page

Oracle Fail Safe includes support for multiple virtual addresses in a group. All databases in a group must use the same virtual addresses, and the virtual addresses must be added to the group before you add the databases to the group. The sequence for building a group is as follows:

1. Create a group.

2. Add one or more virtual addresses to the group.

3. Add one or more single-instance databases to the group.

For example, if a group contains a database that is using two virtual addresses and you add a second database to the group, then the second database must use the same virtual addresses as the first database that was configured into the group. Oracle Fail Safe Manager checks to ensure that the same virtual addresses are used for all single-instance databases that you add to a group.

See Section 4.7 for information about configuring a resource in a group with multiple virtual addresses.

7.3.3.3 Database Identity

The Add Resource to Group Wizard requests database identity information to uniquely identify the single-instance database that is being configured for high availability, as shown in Figure 7–4.

Figure 7–4   Database Identity Wizard Page
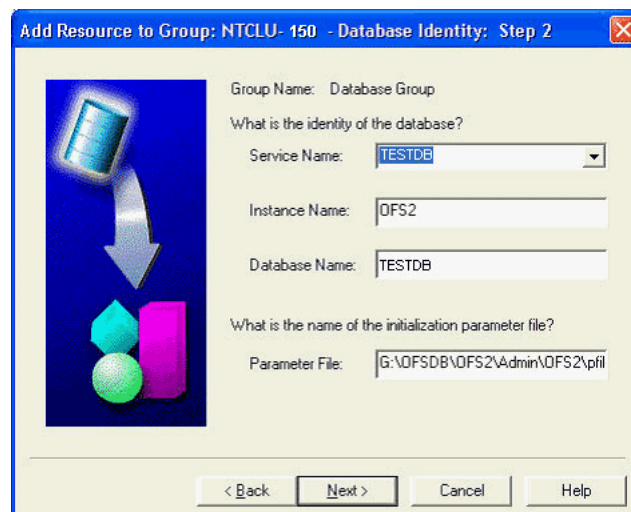
Oracle Fail Safe uses this data to configure the database into the cluster (for example, to update the `tnsnames.ora` file). It also passes the data that you supply to MSCS, where it is registered for use when the database is brought online, taken offline, or when Is Alive polling is performed. Oracle Fail Safe requests the following information:

- Service name

  This is the **net service name**. This is the name that appears in the Oracle Fail Safe Manager tree view and the MSCS tree view. This is the name that client applications specify in a connection request.

  If you do not specify a domain name in the Oracle Net service name, then Oracle Fail Safe chooses a domain name to append to the net service name, as described in Section 7.4.3.1.

- Instance name

  This is the name of the database instance, also referred to as a `SID`.

- Database name

  This is the `db_name` parameter used to identify the database in the initialization parameter file. It is the name that was used when the database was created (for example, in the `SQL CREATE DATABASE` statement).

- Name and location of the initialization parameter file

  When an Oracle database starts, it uses the initialization parameter file to specify the name of the database, the amount of memory to allocate, the names of control files, and various limits and other system parameters.

  In most cases, you place the parameter file on a cluster disk so that it can be accessed regardless of which cluster node is currently hosting the database. However, you can place a copy of the initialization parameter file on each node's private disk, if you ensure that the file exists at the same location on all cluster nodes that are configured to run a database. You may decide to place the parameter file on each node's private disk to set different parameters for the database, depending on which node is hosting it. This can be useful if some nodes have less memory or processing capabilities than others. See Section 7.3.3.3.2 for special considerations for using this method on a database that was created with Database Configuration Assistant.

> **Note:** If needed, you can move the initialization parameter file after a database has been configured for high availability. See the Oracle Fail Safe Manager Help for information about how this is performed.

7.3.3.3.1 Parameter File and Databases That Use an SPFILE Oracle Fail Safe requires that a text initialization parameter file (`PFILE`) be specified in the Parameter File field. To use a binary server parameter file (`SPFILE`) with databases configured for high availability, specify the location of the `SPFILE` from within the `PFILE` using the `SPFILE=`*`SPFILE-location`* parameter. For example, the contents of the PFILE may include the following parameters:

```
SPFILE=I:\Oracle_Home\admin\oradb\pfile\spfileoradb.ora
REMOTE_LOGIN_PASSWORDFILE=none
```

(If you specify an `SPFILE` in the `PFILE` that Oracle Fail Safe uses, then be careful if and when you export the `SPFILE`. If you use a `CREATE PFILE FROM SPFILE` command without including file specifications, then you overwrite the `PFILE` that Oracle Fail Safe is using. Therefore, be sure to specify a unique file name for the `PFILE` to which the `SPFILE` is exported. See *Oracle Database Administrator's Guide* for detailed information about server parameter files.)

7.3.3.3.2 Parameter File and Databases Created with DBCA When you use Database Configuration Assistant to create a database, a text initialization parameter file (`init.ora`) is created in the *`ORACLE_HOME`*`\database` directory. This file contains the `IFILE` parameter. The `IFILE` parameter value is the file specification of another text initialization parameter file, which contains the initialization parameters of the database.

The file specification of the initialization parameter file `IFILE` parameter is stored in the Windows registry. By default, Oracle Fail Safe displays the file specification of the initialization parameter file that contains the `IFILE` parameter (except for databases created with the `Create Sample Database` command) in the Parameter File field. Databases created with the Create Sample Database command displays the initialization parameter file that the `IFILE` parameter specifies in the Parameter File field.

Oracle recommends that in a cluster environment, you enter the file that the `IFILE` parameter specifies in the Parameter File field. This provides the most direct route to the database initialization parameters. In addition, if you choose not to do this and you also choose to keep the parameter file on a private disk, then you must remember to copy both the initialization parameter file containing the `IFILE` parameter and the file that the `IFILE` parameter specifies to the private disk of each cluster node.

7.3.3.4 Database Authentication

The Authentication page is presented if the account under which Oracle Services for MSCS was installed is not in one of the following Windows operating system groups: the `ORA_DBA` group or the `ORA_`*`SID`*`_DBA` group associated with the database. When the account under which Oracle Services for MSCS was installed is in the `ORA_DBA` group or the `ORA_`*`SID`*`_DBA` group, it can use operating system authentication to access the database. If the account is not a member of the `ORA_DBA` group or the `ORA_`*`SID`*`_DBA` group, then it must use the `SYS` account to access the database.

This page lets you specify whether Oracle Services for MSCS should use operating system authentication or the SYS account to access the database and its instances, as shown in Figure 7–5.

> **Note:** The Database Authentication page is not presented if the account under which Oracle Services for MSCS was installed, is a member of a group that lets it access the database using operating system authentication.

To specify operating system authentication, select **Use operating system authentication**. To specify the database SYS account, select **Use SYS account**, then select **SYS**, and then enter and confirm the password for the account.

Figure 7–5 Database Authentication Wizard Page

If you select the "Use operating system authentication" option, then Oracle Services for MSCS opens the Confirm Add to DBA Group window, as shown in Figure 7–6.

The text in Figure 7–6 indicates that Oracle Services for MSCS was installed under the EXAMPLE\Admin account. If you click **Yes**, then the EXAMPLE\Admin account gets added to the Windows operating system group ORA_OFS2_DBA.

Figure 7–6 Confirm Add to DBA Group Window

7.3.3.5 Database Password

If Oracle Services for MSCS detects that the standalone database has a password file associated with it, then the Add Resource to Group Database Wizard asks if you want Oracle Services for MSCS to create the password file on all nodes that will be the possible owner nodes for the fail-safe database, as shown in Figure 7–7.

Oracle recommends that you select the "Yes, create the password file" option. A password file is often required when you perform remote operations. For example, Recovery Manager (RMAN) requires a password file when connecting to the target database over a nonsecure Oracle Net connection.

If you select the "No, do not create the password file" option, then all users must access the database using operating system authentication, and users will not be able to perform remote database administration operations.

> **Note:** If Oracle Services for MSCS does not detect password files for the database that you are adding to a group, then it does not present the Database Password page.

Figure 7–7 Database Password Wizard Page
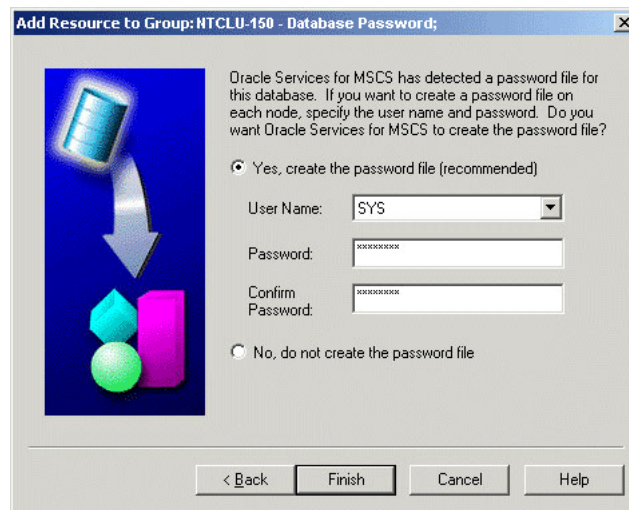
Oracle Services for MSCS makes the following adjustments to the database initialization parameter file (if needed), depending on whether you choose to have Oracle Services for MSCS create the password file on all cluster nodes that are possible owner nodes for the database:

- Yes, create the password file

  Sets the REMOTE_LOGIN_PASSWORDFILE parameter to EXCLUSIVE.

- No, do not create the password file

  Sets the REMOTE_LOGIN_PASSWORDFILE parameter to NONE.

> **Note:** Oracle Services for MSCS does not support setting the Windows registry DBA_AUTHORIZATION parameter to the value of BYPASS.

If you want to change the password for the SYS account after the database has been added to a group, then you must also update the password through Oracle Fail Safe Manager. See Section 7.5.2 for information about how to update the password for this account after the database has been added to a group.

7.4 Oracle Net Listener Resource Creation and Configuration

When you add a single-instance database to a group, Oracle Fail Safe creates and configures the Oracle Net listener resource and the database resource in the group. When it brings the database online, Oracle Fail Safe first ensures that the listener resource is online. If Oracle Fail Safe is not able to connect to the database through the listener during Is Alive polling of the database resource, then it uses the **bequeath protocol** adapter to connect to the database. Oracle fail Safe also logs an event to notify users of the failure. If you see this event, then check the listener resource. The failure of the listener resource does not affect the existing connections to the database, but prevents new users from connecting to the database. The listener resource has its own restart policy, so it will be restarted automatically by the cluster in the event of failure.

Oracle Fail Safe creates a dependency between the database and the IP address associated with the listener but not on the listener itself. This dependency is created to avoid a situation in which clients would stop responding when an IP address was taken offline before the database.

7.4.1 Using Shared Sockets in Dedicated Server Mode

You can set the `USE_SHARED_SOCKET` parameter to true to enable the use of shared sockets. If this parameter is set to true, then the network listener passes the socket descriptor for client connections to the database thread. As a result, the client does not need to establish a new connection to the database thread and database connection time improves. Also, all database connections share the port number used by the network listener, which can be useful if you are setting up third-party proxy servers.

This parameter only works in dedicated server mode in a TCP/IP environment.

7.4.2 Client Connections to Highly Available Single-Instance Databases

Network objects (including databases) are identified by a network address. For a connection between a client and a database to be made, the network address in the `tnsnames.ora` file on the client and the network address in the `listener.ora` file on the server must match. In other words, a client uses a network address to send a connection request to a particular network object location, and the recipient listens for requests on this address and grants a connection based on its address information matching its client information.

When you add a single-instance database to a group, Oracle Fail Safe creates a listener for the group in the same Oracle home where the database resides. When Oracle Fail Safe configures the virtual address information, it updates the `tnsnames.ora` files in all Oracle homes on cluster nodes that are possible owners for the database, and on the client system from which you run Oracle Fail Safe Manager. This enables Oracle Fail Safe to access the database instance using the updated configuration.

Section 7.4.3 describes how Oracle Fail Safe creates an entry in the `listener.ora` file and updates the `tnsnames.ora` file after you add a database to a group so that clients can connect to the database, regardless of which cluster node is hosting the database.

> **Note:** Oracle Fail Safe does not support the use of the `TNS_ADMIN` Windows environment variable or registry parameter. Oracle Fail Safe retrieves and updates Oracle Net files in the `Oracle_Home\network\admin` directory; it ignores the `TNS_ADMIN` Windows environment variable or registry parameter if either is specified.

7.4.3 Updated Oracle Net Configuration After Adding a Database to a Group

When you add a single-instance database to a group, Oracle Fail Safe changes the Oracle Net configuration for the database in the `tnsnames.ora` file, the `listener.ora` file, and the `sqlnet.ora` file as described in the following sections.

7.4.3.1 Updates That Oracle Fail Safe Makes to the tnsnames.ora File

When you add a single-instance database to a group, Oracle Fail Safe updates the net service name entry in the `tnsnames.ora` file for the database to use the virtual addresses of the group. If there are multiple Oracle homes on the node, then all `tnsnames.ora` files are updated. In addition:

- Oracle Fail Safe may make adjustments to the net service name, and the `SERVICE_NAME` parameter or the `SID` parameter.

  When you create a sample single-instance database or add a single-instance database to a group, if you do not specify a domain name in the Oracle Net service name, then Oracle Fail Safe chooses a domain name to append to the net service name as follows:

  - Oracle Fail Safe looks for the default domain name in the Oracle home of the latest database version on the node. If found, this default domain name is appended to the net service name. For example, assuming Oracle Database 10*g* is the latest database version on the node, if you specify `MyDB` as the Oracle Net service name, and the default domain name in the Oracle Database 10*g* home is `example.com`, then the net service name will become `MyDB.example.com`.

  - If there is no default domain name in the Oracle home of the latest database version on the node, then Oracle Fail Safe appends nothing to the net service name. For example, if you specify `MyDB`, then the net service name will also be `MyDB`.

  In either case, when the net service name is decided, Oracle Fail Safe checks the `tnsnames.ora` file for an existing entry with the same name and updates the `tnsnames.ora` file, as follows:

  - If an entry for the net service name exists, then Oracle Fail Safe updates the entry so that it connects to the virtual host. Oracle Fail Safe does not change the `CONNECT_DATA` parameter, including the `SERVICE_NAME`.

  - If an entry does not exist, then Oracle Fail Safe writes a new entry to the `tnsnames.ora` file. In this case, the `SID` is used in the `CONNECT_DATA` parameter. (Users can still connect to a database using a net service name that refers to the `SID`, rather than the service name.)

- If you define an archive log destination as a service name, as shown in the following example, then Oracle Fail Safe will not automatically update the `tnsnames.ora` file on all cluster nodes. You must manually edit or add the service name entry to the `tnsnames.ora` file on each cluster node.

```
log_archive_dest_2='SERVICE=standby OPTIONAL REOPEN=120'
```

■ On server nodes, addresses that use virtual **host name**s configured only for internal cluster communications will be placed at the top of the address list. Therefore, applications running on the server nodes, including the database resource monitor, attempts to use addresses that use the private network interconnect among the cluster nodes before using addresses that use the **public interconnect** for connections to the database resource.

On client nodes, addresses that use virtual host names configured only for internal cluster communications will not be included in the address list.

> **Note:** Oracle Fail Safe updates the tnsnames.ora files on all cluster nodes (that are possible owners for the single-instance database) and on the client node where you run Oracle Fail Safe Manager. If you must enable remote clients (that are not running Oracle Fail Safe) to process work against a single-instance Oracle database through a cluster node, you must edit the tnsnames.ora files to update the host name with the virtual address information. Edit each client's local tnsnames.ora file using a network configuration tool.

7.4.3.2 Updates That Oracle Fail Safe Makes to the listener.ora File

When you add a single-instance database to a group, Oracle Fail Safe makes the following changes to the listener.ora file:

1. Removes the SID_DESC parameter from the standalone database listener

2. Creates a new Oracle Fail Safe listener that is configured to listen on the virtual address associated with the single-instance database

3. Adds the SID_DESC parameter to the new Oracle Fail Safe listener

4. Stops and restarts the standalone database listener to accept the changes that have been made

5. Starts the new Oracle Fail Safe listener

7.4.3.3 Updates That Oracle Fail Safe Makes to the sqlnet.ora File

When you add a single-instance database to a group, if operating system authentication has been chosen for the database, then Oracle Fail Safe adds the SQLNET.AUTHENTICATION_SERVICES=(NTS) parameter to the sqlnet.ora file (assuming the parameter is not set).

7.4.4 Using External Procedures with Databases Configured for High Availability

Oracle Fail Safe configures the address of an external procedure in the Oracle Net listener definition for a group (the listener.ora file). As the first database is added to a group, Oracle Fail Safe determines whether or not external procedures are configured in the original Oracle Net listener serving the database. If they are configured, then Oracle Fail Safe creates both an IPC listener address and a SID descriptor (SID_DESC) in the listener.ora file for the group.

Oracle Fail Safe creates the key of the IPC address by concatenating the prefix EXTPROC to the network name of the first virtual address in the group. Oracle Fail Safe uses the alphabetic order of the virtual address network names in the group to determine which network name is first. Therefore, if the network name of the first

virtual address is `ntclu45`, then the key value is `EXTPROCntclu45` and the IPC address entry will be as follows:

```
(ADDRESS=
   (PROTOCOL=IPC)
   (KEY=EXTPROCntclu45)
)
```

Oracle Fail Safe copies the `SID` descriptor information from the original Oracle Net listener for the database to the Oracle Net listener for the group. For example, if the original Oracle Net listener for the database has the following `SID` descriptor, then the Oracle Net listener for the group will have the same descriptor:

```
(SID_DESC=
  (SID_NAME=PLSExtProc)
  (ORACLE_HOME=C:\app\oracle\product\11.2.0\db_1)
  (PROGRAM=extproc)
)
```

When a PL/SQL or SQL application calls an external procedure, the application specifies the `EXTPROC_CONNECTION_DATA` net service name entry in the `tnsnames.ora` file. This entry contains the information needed by the listener to start a process for the external procedure. The IPC address of the Oracle Net listener in a group is added to the `EXTPROC_CONNECTION_DATA` net service name entry in all `tnsnames.ora` files on the cluster.

7.4.5 Support for Databases Using Shared Servers

The following sections describe how Oracle Fail Safe supports single-instance databases that use a shared server configuration.

> **Note:**   When you set up a database to use a shared servers configuration, you must ensure that Oracle Fail Safe can continue to use a dedicated server connection for its internal operations. You do this by specifying the `(SERVER=DEDICATED)` parameter in the connect data portion of the net service name entry for the database in the `tnsnames.ora` file on each cluster server node. (By default, if shared servers are used and no `SERVER` parameter is specified, then the listener establishes a connection using shared servers.)

7.4.5.1 Shared Servers for Databases

To use a shared server configuration, you may be required to make modifications to the database parameter file.

You can specify listener information in either the `LOCAL_LISTENER` or the `DISPATCHERS` parameter for a shared server configuration.

If the shared server configuration uses the `LOCAL_LISTENER` parameter to specify full listener information (full listener information specifies both host and port values), then Oracle Fail Safe automatically updates the database parameter file for the shared server configuration during the Add Resource to Group operation.

The single-instance database runs in shared server mode after you add it to a group. You must not make any further changes to the database parameter file.

The following example shows a shared server configuration that will be updated automatically by Oracle Fail Safe:

```
dispatchers = "(PROTOCOL=TCP)(DISPATCHERS=1)"
local_listener = "(ADDRESS=(PROTOCOL=TCP)(HOST=124.7.56.1)(PORT=1521))"
```

After you add the database to a group, Oracle Fail Safe updates the LOCAL_LISTENER parameter to use the listener information for the group.

However, if the shared servers configuration uses the DISPATCHERS parameter to specify full listener information, then you must remove the host and port values from the DISPATCHERS parameter. Oracle Fail Safe always writes the LOCAL_LISTENER parameter to the database parameter file.

When you remove a database from a group using Oracle Fail Safe Manager, it deletes the LOCAL_LISTENER parameter from the database initialization file. You must add the parameter back into the database initialization file by following the instructions in Section 7.2.5.

7.5 Security Requirements for Single-Instance Databases

To manage a single-instance Oracle database, use a database administrator account that has SYSDBA privileges. This lets you administer Oracle databases from a remote client.

When you create a single-instance **sample database** or add a single-instance database to a group, Oracle Fail Safe must use operating system authentication or the SYS user account to access the database. Use an authentication password file and set the initialization parameter, REMOTE_LOGIN_PASSWORDFILE, in the database initialization parameter file (init*database-name*.ora) to either SHARED or EXCLUSIVE if users access the database using the SYS account. Set the REMOTE_LOGIN_PASSWORDFILE to NONE if users only access the database using operating system authentication.

> **Note:** Oracle Fail Safe does not support setting the Windows registry DBA_AUTHORIZATION parameter to the value of BYPASS.

Refer to *Oracle Database Administrator's Guide* for more information about database administrator authentication and the REMOTE_LOGIN_PASSWORDFILE parameter.

7.5.1 Synchronizing Password Files on Cluster Nodes

Database password files are stored on private disks. Changes that you make to the password file on one cluster node are not automatically applied to the corresponding file on the other cluster nodes.

Therefore, if you add an account to the password file on one cluster node, then you must add that account to the password file on the other cluster nodes that are configured to run the database instance. If there are accounts in addition to SYS stored in a password file, then you must grant SYSOPER and SYSDBA privileges for the additional accounts on the other cluster nodes for a single-instance fail-safe database.

If you add a single-instance database to a group with the Oracle Fail Safe Manager Add Resource to Group Wizard, then Oracle Services for MSCS creates a database instance on the other nodes that are configured to run the database and uses the default value for the maximum number of users in the password file. The password file on the node where the instance is created contains only the password for the SYS account that you supply in the Add Resource to Group Wizard.

On the other nodes configured to run the database instance, perform the following steps to synchronize the password files on the other cluster nodes:

1. If the number of accounts in the password file exceeds the default maximum, then create a new password file. Otherwise, skip to Step 2.

    To create a new password file, refer to instructions about creating password files in the Administrator's Guide for your Oracle database release.

2. Move the group containing the single-instance database to another node configured to run the database instance.

3. Grant privileges to accounts other than SYS on the node to which you move the database.

4. Repeat Step 2 and 3 for each node in the cluster configured to run the database.

    Now the local copies of the password file are identical on all nodes configured to run the database instance.

7.5.2 Changing the SYSDBA Account Password

If you want to change the SYSDBA (SYS) account password for one or more databases after those databases have been added to groups, then use Oracle Fail Safe Manager to make this change so that Oracle Services for MSCS use the correct password to access the database when it performs operations such as Is Alive polling.

You can use the Update Database Password Wizard, or the Authentication tab in the fail-safe database property sheets to update the SYSDBA password, as follows:

- To change the SYSDBA password for several databases, Oracle recommends that you use the Update Database Password Wizard. Access the Update Database Password Wizard in Oracle Fail Safe Manager by choosing **Resources** from the menu bar, then **Update Database Password**.

- To change the SYSDBA password for a single database, Oracle recommends that you use the database Authentication tab. From the tree view, choose the fail-safe database for which you want to change the password, then select the **Authentication** tab.

If the REMOTE_LOGIN_PASSWORDFILE parameter is set to SHARED the password update is queued until the database is restarted. That is, the actual password change is done the next time the database is taken offline and then placed back online. A message is displayed to inform the user that the password cannot be changed at that point and time and the password change will occur later when the database is restarted.

If operating system authentication is enabled (without the use of a password file), then you cannot change the SYSDBA account information using Oracle Fail Safe Manager unless you first remove the database from the group, then add it to the group again with the Add Resource to Group Wizard. When you add the database to the group again, select **Use SYS Account** on the Database Authentication page of the wizard.

See the online help for the following additional information:

- Detailed instructions on using the Update Database Password Wizard and the Authentication property page.

- Information about using these tools when the SYSDBA password has been changed with a tool other than Oracle Fail Safe Manager.

7.5.3 Upgrading a Fail-Safe Database with the Oracle Database Upgrade Assistant

This section describes how to use the Oracle Database Upgrade Assistant to upgrade a single-instance fail-safe database from one release to another or to move a single-instance Oracle database from one Oracle home to another.

For each single-instance database you upgrade or move to a new home, perform the following steps:

1. Remove the single-instance database from the group. In the Oracle Fail Safe Manager tree view, select the database, then from the menu bar select **Resources** followed by **Remove from Group.**

2. Run the Oracle Database Upgrade Assistant from the Oracle home to which you are moving or upgrading your single-instance database.

3. Be prepared to provide the location of the database parameter file for the single-instance database you are upgrading. During a database upgrade, the database parameter file is converted. If the database parameter file is on a cluster disk, then your parameter file is appropriately located for Oracle Fail Safe to make the conversion. If the database parameter file is located on a private disk, then the Oracle Database Upgrade Assistant only converts the local copy. In this case, you must edit the copy on the other cluster nodes and make the appropriate changes.

4. The Oracle Database Upgrade Assistant also asks you to specify the location of the converted database files. Either leave the data files in their current location, or specify a cluster disk that is currently accessible by the local node. If you choose the latter, then ensure the cluster disk is not being used by another group.

5. When all databases in the group have been upgraded or moved to a new home with the Oracle Database Upgrade Assistant, use Oracle Fail Safe Manager to put the databases back into the group and then place the databases online, as follows:

   a. On the **Resources** menu, select **Add to Group.**

   b. Follow the steps through the Add Resource to Group Wizard.

All databases in the group being moved must be from the same Oracle home. If one database in a group is moved with the Oracle Database Upgrade Assistant to a new Oracle home, then all databases in the group must be moved to the new Oracle home.

7.6 Optimizations for Single-Instance Database Recovery

Oracle databases configured with Oracle Fail Safe for high availability ensure fast failover and fast recovery during both unplanned and planned outages (such as software upgrades and scheduled maintenance). You can take advantage of Oracle fast-start and disaster-recovery features, control time spent during database recovery, and ensure continuous monitoring of databases configured with Oracle Fail Safe for high availability.

Oracle Fail Safe and Oracle database technology optimize the time it takes to shut down a database on one node and complete instance recovery on another node for both planned and unplanned failovers. The Oracle database checkpoint algorithms optimize the time it takes to perform instance recovery for planned and unplanned failovers.

When you use Oracle Fail Safe Manager (or FSCMD) to carry out a planned failover, Oracle Services for MSCS checkpoints the single-instance Oracle database before it is shut down. The single-instance database is started on the other node in a restricted mode so that instance recovery can be completed quickly and the database made

available to the database clients promptly. (If you use MSCS to carry out a planned failover, then it does not checkpoint the database before shutting it down.)

> **Note:** After you add a single-instance database to a group, use only Oracle Fail Safe Manager or the `FSCMD` command to place the Oracle database online and take the database offline. Otherwise, the database will not be checkpointed first. In addition, if you use a tool other than Oracle Fail Safe Manager, FSCMD, or MSCS to take a database offline, then Oracle Fail Safe considers it a failed resource and attempts to place it back online.

For unplanned failover, the instance recovery time is controlled by the database recovery processing. See the Oracle database documentation for details on fast-start recovery operations.

7.7 Performing Administrative Tasks on a Single-Instance Fail-Safe Database

You perform administrative tasks on a database configured for high availability as you would for any database, with one exception. You must use Oracle Fail Safe Manager or the `FSCMD` command-line interface (see Chapter 5) to take a database offline (and stop cluster monitoring of the database) during any operation that restricts access to the database or for which you want to temporarily disable the possibility of failover. This includes not only cold backup operations but also administrative operations that must be performed while users continue to access the database, or operations that could affect response times during the periodic Is Alive polling of the database by MSCS.

Use the following steps to perform administrative tasks on a database that is configured in a group with Oracle Fail Safe Manager:

1. Use Oracle Fail Safe Manager or the `FSCMD` command to take the database offline, shut down the database, and suspend monitoring of the database by the cluster. All users connected to the database will be disconnected.

2. Use a tool such as SQL*Plus to start the database and to perform your administrative tasks. While the database is started, users can access the database.

3. Once you have completed the administrative tasks, use a tool such as SQL*Plus to shut down the database.

4. Use Oracle Fail Safe Manager or the `FSCMD` command to place the database online again. The cluster will resume monitoring the database.

Chapter 5 provides an example of a script in which `FSCMD` commands are used to perform a backup operation.

If, during an administrative task, you perform an operation that changes the configuration of the database (such as adding a new tablespace and associated data file), then you must run the `Verify Group` operation. Adding a new data file can introduce a new disk dependency in the group. When you run the `Verify Group` operation, it checks to ensure that the disk is a cluster disk and that it does not belong to another group. If adding the new data file introduces a new disk dependency in the group, then the disk is added to the same group as the database and the information in the cluster registry is updated to ensure that the new disk correctly fails over with the database.

7.8 Database Homes

When Oracle Fail Safe accesses a database, it usually uses the same Oracle database home to access any database on the system. The database home that Oracle Fail Safe uses, is chosen when the server starts. Oracle Fail Safe scans all database homes to look for the home that has the highest software version. Initially it only looks at homes that have their \bin path included in the system PATH environment variable. If Oracle Fail Safe does not find any database homes in the system PATH then it scans all the database homes looking for the home with the highest version.

Note that since Oracle Fail Safe chooses a database home when it first starts, it will not be aware of any database homes that are installed after Oracle Fail Safe started. The Oracle Fail Safe server and resource monitors must be restarted before Oracle Fail Safe considers a new database home for use. After installing a newer version of Oracle database, the Cluster Service service should be restarted on all nodes so that Oracle Fail Safe can use the new database home.

If there are databases being managed by the cluster then there must be a database home installed on a local disk of each node in the cluster. If a database home is installed on a shared cluster disk then it's \bin directory should not be included in the system PATH environment variable.

There are two different Oracle Fail Safe components that may access a database:

1. The Oracle Fail Safe server (OracleMSCSServices)

2. The Oracle Fail Safe database resource DLL (FsResOdbs.dll)

The server will normally only access a database when configuring a database resource (add or delete resource) or during verify operations. During typical system operation the Oracle Fail Safe server does not access any databases.

The resource DLL is invoked by the Windows Cluster Service service when a database or listener resource is referenced by the cluster. For example, when a database is brought online, during IsAlive polling, when the resource fails over to another virtual node, and so on. On systems with multiple database homes, there may be a requirement to have each database on the system accessed, using the same database home software that is being used by the instance for that database. It is possible to configure a resource to run in a separate resource monitor process by selecting the "Run this resource in a separate resource monitor" check box on the resource properties page, displayed by the Oracle Fail Safe Manager. When that option is enabled, instead of always using the highest version database home on the system, the resource monitor process for the database uses the database home that is used to run the database instance when accessing the database. When referencing database listener resources, the resource DLL always uses the software from the \bin path used to start the listener service, regardless of the setting of the "Run this resource in a separate resource monitor" option.

See Section 4.3.1 for information regarding the user accounts used by Oracle Fail Safe components when accessing databases.

7.9 Configuring Transparent Application Failover (TAF)

For standalone single-instance databases, **transparent application failover** (TAF) instructs Oracle Net to reestablish a failed connection to a database by connecting to a different listener. This lets the user continue work using the new connection as if the original connection had never failed. The transparent application failover feature does not work the same way for a single-instance Oracle Fail Safe database as it does for a standalone single-instance database. For a Oracle Fail Safe database, a transparent

application failover instructs Oracle Net to reconnect to the *same* listener, which has moved to another cluster node due to a group failover.

For a standalone database, the term failover in the phrase "transparent application failover" refers to Oracle Net failing over a connection from one listener to another. For a Oracle Fail Safe database, the term failover in the phrase "transparent application failover" is a bit of a misnomer as the application does not fail over, but the listener to which it is connected fails over, and then a connection is reestablished.

These differences in implementation do not affect how you manage transparent application failover.

To take advantage of transparent application failover when connected to a database configured with Oracle Fail Safe, the **client application**s must:

- Connect through Oracle Net to an Oracle database

With transparent application failover, clients must not explicitly reconnect after a group fails over. The OCI connection handles reconnection and state recovery automatically for the client application. In fact, applications that are not actively updating the database at the time of a failure may not notice that failover is occurring.

Refer to the *Oracle Net Services Administrator's Guide* for complete information about transparent application failover.

7.10 Handling Errors and Troubleshooting Problems with Databases

The following sections describe how to specify a script to handle errors if they occur when Oracle Fail Safe attempts to bring a highly available single-instance database online and how to troubleshoot specific problems that you may encounter with single-instance Oracle databases configured for high availability. For general information about troubleshooting Oracle databases, see the Oracle database documentation.

7.10.1 Handling Errors That Occur When Bringing a Database Online

You can specify a script to handle errors that may occur when Oracle Fail Safe is attempting to place a single-instance database online. Oracle Fail Safe uses the same script for all single-instance fail-safe databases on the cluster.

To specify an error handling script:

1. Create a script to handle the error or errors.

2. Name the script `FsDbError.bat`.

3. Ensure that the script returns `0` if it succeeds and any nonzero integer if it fails.

4. Place the script in the following directory on each cluster node that is a possible owner for a database resource and ensure that the file owner has local Administrator privileges on that cluster node:

   `Oracle_Home\fs\fssvr\scripts`

If Oracle Fail Safe cannot bring a single-instance database online, then it spawns a process to run the script, then it passes the error code, the database name, the database `SID`, the TNS service name, and the database parameter file specification to the script and executes the script, as follows:

`FsDbError.bat error-code database-name SID TNS service name parameter-file-spec`

For example:

```
FsDbError.bat ORA-01113 OracleDB OracleDB OracleDB.WORLD
D:\Ora\admin\OracleDB\pfile\initOracleDB.ora
```

The process in which the script is running waits for the script to finish within the period of time specified as the Pending Timeout value for database resources. If the script does not finish within the pending timeout period, then the script is terminated.

Oracle Fail Safe logs an event to the Windows Event log to indicate whether the script succeeded, failed, or was terminated by Oracle Fail Safe. If the script failed, the error code is also written to the event log.

Regardless of whether the script succeeds or fails, Oracle Fail Safe continues to attempt to bring the single-instance database online as defined in the database restart and failover policies.

7.10.2 Troubleshooting Problems

In most cases, the first step in troubleshooting a problem is to run the `Verify Cluster`, `Verify Group`, or `Verify Standalone Database` command. These tools are described in general in Chapter 6. If the `Verify` commands do not reveal the source of the problem, review the Windows application event log to see if any error messages have been posted. When the operation fails, Oracle support may also ask to have tracing enabled. See Appendix B for more information on enabling Oracle Fail Safe tracing.

When you run a `Verify Group` command on a group containing a single-instance database, Oracle Fail Safe performs the following tasks:

- Queries each database in the group to determine which disks it uses. Then, it validates that the disks are cluster disks and have been added to the group. If the disk validation fails (for example, because a disk has been added to the database since it was configured for high availability), then the `Verify Group` operation prompts you before fixing the problem.

- Detects disk drive changes and updates resource dependencies, if necessary.

- Validates that the network name pings the correct IP address.

- Ensures that the Oracle Net configurations are correct.

- Repairs any misconfigured resources in the group.

You can run the `Verify Group` operation at any time. However, you *must* run it when any of the following occurs:

- A group or resource in a group does not come online.

- Failover or failback does not perform as you expect.

- You add more disks to a single-instance database that is configured in a group.

- A new node is added to the cluster.

For example, assume that you add a new disk to a single-instance database, but you do not use Oracle Fail Safe Manager to update the cluster configuration. If a server node subsequently shuts down, failover does not occur correctly because the cluster software was never notified that there was a change in the configuration. To prevent this from happening, you must verify the group containing a single-instance database whenever you make a structural change to the database. When you verify the group, Oracle Fail Safe automatically detects changes and updates the cluster configuration for you. In the previous example, Oracle Fail Safe would add the new disk to the group for you.

If any problems are found during the group verification, then Oracle Fail Safe prompts you to fix them or returns an error message that further describes the problem.

7.10.3 Problems Adding a Database to a Group

To troubleshoot problems when adding a single-instance database to a group:

- Run the `Verify Standalone Database` operation (described in Section 6.1.3).

    Running the `Verify Standalone Database` operation verifies that the database is a valid working standalone database.

    For example, if you try to add a standalone database to a group and it fails during the Oracle Net configuration, then Oracle Fail Safe rolls back the clusterwide operation and the database remains as a standalone database. To fix this problem, perform the following steps:

    1. Run the `Verify Cluster` operation to ensure that the cluster network configuration is correct.

    2. Run the `Verify Standalone Database` operation to ensure that the network (Oracle Net) is working.

    3. Attempt to add the standalone database to a group.

    4. If the Add Database to Group operation fails, then check the Oracle Net rollback file as described in Section 7.10.8.4.

- Ensure that the following conditions are met:

    – The single-instance database files are located on shared cluster disks

        For each disk in the configuration, Oracle Fail Safe determines if the disk resides on a shared storage interconnect. If database files are on nonclustered disks, then you must move the database files so that they are located on a shared cluster disk.

    – The following information must be correctly specified in the Add Resource to Group Wizard:

        * User name and password used to access the database

        * Database parameter file

        * (Net) service name

        * Database name

        * Instance name

7.10.4 Problems Placing a Group Online

If there is a problem placing a group that contains a single-instance database online, then try the following:

- Verify the group.

    When you use the `Verify Group` command (from the Oracle Fail Safe Manager Troubleshooting menu), Oracle Fail Safe checks the group configuration and attempts to fix any problems that it finds. If the `Verify Group` command cannot fix the problem, then it returns an error message that should help you to resolve the problem manually.

    If the `Verify Group` command finds a problem, then it prompts you on how to proceed.

- Check the Oracle Net listener log.

  Oracle Net logs an entry to the listener log file every time an error is encountered or a database is accessed through the listener. Check for errors in the log file that may help you to identify the problem.

- Check the net service name of the single-instance database.

  The Oracle Fail Safe database resource DLL accesses each database in a group at the Is Alive interval. (The Is Alive interval appears on the Failover property page for the database in Oracle Fail Safe Manager.) It uses the database connection information to access the database. If the database access information changes, then Oracle Fail Safe fails to access the database. Hence, MSCS does not consider the database resource to be *alive*.

- Check the Oracle Net configuration data.

  The `listener.ora` file on the server system and the `tnsnames.ora` file on both the client and server systems must contain valid virtual server addresses for the groups in your cluster.

- Bring each resource in the group online individually.

  If multiple single-instance database is in the group, then this helps you to identify the database causing the problem.

- Ensure that the single-instance database `Pending Timeout` value is sufficient.

  If a group containing a database fails to come online or frequently fails over, then check that the `Pending Timeout` value is set correctly. Failure to come online and frequent failovers occur if the `Pending Timeout` value for the database is set too low.

  Set the `Pending Timeout` value to specify the length of time you want the cluster software to allow for the database to be brought online (or taken offline) before considering the operation to have failed. Set the value high enough to prevent a cluster system from mistaking slow response time for unavailability, yet low enough to minimize the failover response time when a failure does occur.

  You can set the `Pending Timeout` value by modifying the database properties, as follows:

  1. In the Oracle Fail Safe Manager tree view, select the database name.

  2. Click the **Policies** tab.

  3. In the **Pending Timeout** box, modify the `Pending Timeout` value.

- If users use the `SYS` account to access the database, then ensure that the initialization parameter `REMOTE_LOGIN_PASSWORDFILE` in the database initialization parameter file (`init`*database-name*`.ora`) is set to `SHARED` or `EXCLUSIVE`.

- If users access the database using operating system authentication only, then ensure that the initialization parameter `REMOTE_LOGIN_PASSWORDFILE` in the database initialization parameter file is set to `NONE`.

- If the account password that Oracle Fail Safe uses to access a database has changed, then update that change in Oracle Fail Safe Manager.

  If the password for the account through which Oracle Fail Safe accesses a database changes and you do not update the information through Oracle Fail Safe Manager, then the attempts at polling the database will fail. See Section 7.5.2 for information about how to update database password changes for Oracle Fail Safe.

7.10.5 Group Fails Over During Processing-Intensive Operations

Sometimes, processing-intensive operations (such as an Import operation) can cause Is Alive polling to fail and may result in an undesired group failover. In such cases, you can disable Is Alive polling for the database by issuing the `FSCMD DISABLEISALIVE` command. However, be aware that when you disable Is Alive polling, Oracle Fail Safe suspends monitoring the instance until Is Alive polling is reenabled. You reenable Is Alive polling with the `FSCMD ENABLEISALIVE` command.

Oracle recommends that you run these `FSCMD` commands from within a script so that you can ensure that Is Alive polling is reenabled when the processing-intensive operation completes.

For information about the `FSCMD` commands, see Chapter 5.

7.10.6 Database Authentication

If there is a problem when Oracle Fail Safe tries to bring a single-instance database online or offline, then the problem may be caused by the way you have set up database authentication. Try the following to solve the problem:

- If you selected "Use this account" on the Authentication page in the Add Resource to Group Wizard, then ensure that the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter in the database initialization parameter file (`init`*database-name*`.ora`) is set to `SHARED` or `EXCLUSIVE`.

  Section 7.5 describes how to correctly set up this parameter for database authentication.

- Ensure that Oracle Fail Safe has access to the databases in the group.

  For some operations that Oracle Fail Safe performs, such as a group verification and polling the database to ensure that it is online, Oracle Fail Safe must have access to the databases in a group. If the database account password has changed, then you must update it in Oracle Fail Safe Manager. Otherwise, Oracle Fail Safe cannot monitor the database using Is Alive polling. This situation will be logged to the Windows Event Viewer.

  Section 7.5.2 describes how to correctly update the database password.

7.10.7 Problems with Sample Databases

If you receive errors when you create or delete a sample database, then check the following:

- If the `Create Sample Database` command cannot open the `Create Sample Database` script files, then reinstall or repair your Oracle Fail Safe installation and try the `Create Sample Database` operation again.

- If the `Delete Sample Database` command fails, then you may have selected a database that is not a sample database. Ensure that the database you selected is a sample database and try the delete operation again.

  Oracle Fail Safe stores information about sample databases in the Windows registry under the `HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\FailSafe\SampleDB` key.

  If a sample database that you have deleted is listed in the Oracle Fail Safe Manager tree view, then ensure that the Windows registry entries for that database have been cleared from the registry.

7.10.8 Problems with Virtual Server Configurations

If you encounter problems when trying to establish a connection to either a standalone database or a database configured in a group, then you must check the Oracle Net configuration for the database.

Oracle Fail Safe provides the `Verify Group` and `Verify Standalone Database` operations to help you verify and repair the Oracle Net configuration. See Section 6.1.2 and Section 6.1.3 for details.

7.10.8.1 Problems Configuring the Virtual Address

Oracle Fail Safe changes the `listener.ora` and `tnsnames.ora` files, and stops and starts listeners when configuring the virtual address information. The following list describes potential problems and the action you can take to correct each problem:

- `FS-10070 Oracle Net:` *name*

  This message code reports any problems parsing (reading or updating) the Oracle Net `listener.ora` and `tnsnames.ora` files:

  – If these files are no longer valid due to improper update or file damage, then Oracle Fail Safe cannot use these to configure virtual server information. You must retrieve a valid version of these files or re-create the files using Oracle Net Assistant.

  – If these files are valid, then check that the net service name, the database `SID`, and the network name of the group used in the operation are correct. Incorrect information may cause the virtual server configuration to fail. You must ensure that a database `SID` is not included in multiple listeners. On systems with multiple Oracle homes, check all of the `listener.ora` files.

- `FS-10066 Failed to start Windows service` *name* `for the Oracle Net listener`

  Oracle Fail Safe starts a listener after changing the definition of a listener or creating the definition of a new listener.

  The most common reason for this error is that another listener is already listening for a database. There can be only one listener on the system listening for a particular address or database `SID`. For example, if `LISTENER_A` has the following definition, then no other listener on the system can listen for key `ORCL` using the IPC protocol, or port 1521 on host server_A using the TCP protocol, or `ORCL SID` name:

```
LISTENER =
    (ADDRESS_LIST=
        (ADDRESS=
                (PROTOCOL=IPC)
                (KEY=ORCL)
        )
        (ADDRESS=
            (PROTOCOL=TCP)
            (Host=server_A)
            (Port=1521)
        )
    )

SID_LIST_LISTENER =
    (SID_LIST=
        (SID_DESC=
            (SID_NAME=ORCL)
```

```
        )
    )
```

Any other listeners that try to use the same address or database SID as LISTENER_A will fail to start.

When you encounter this problem, Oracle Fail Safe saves the listener.ora and tnsnames.ora files that contain definitions of the updated listeners and net service names as rollback files. The file names of the rollback files are *filename_rlb.ora*.

Read the rollback files to find the definition of the listener and check it against the definition of all other listeners running on the system. There must be no overlapping addresses or database SID names among the listeners. Again, check for all listener.ora files on systems using multiple Oracle homes.

■ Another common cause for failing to start a listener is the virtual address. The virtual address used by the listener must be active on the node where Oracle Fail Safe tries to start the listener.

See the Oracle Net documentation (including information about the log directory) for additional information about troubleshooting problems with the network configuration.

7.10.8.2 Problems Creating Listeners

Oracle Fail Safe Manager uses the Listener Control Utility (LSNRCTL) to create new listeners, and captures the output in a file located in your Oracle home.

For example, if the Oracle home and network directory path is C:\ORANT\NETWORK\ADMIN, and the virtual address on which the listener listens is ntclu-155, then the listener output files will be written to the following directory and file:

```
C:\ORANT\NETWORK\LOG\fslntclu-155.out
```

Each listener has its own output file that is named using the listener name and the .out extension. (In the example, the listener name is fslntclu-155.) If you experience difficulties when creating a new listener, then you can use the output file to help you diagnose the problem.

7.10.8.3 Archived listener.ora or tnsnames.ora Files

Whenever Oracle Fail Safe makes changes in the listener.ora or tnsnames.ora files, the original version of the file is archived. If you need to reference an Oracle Net net service name definition or a listener definition as it was before Oracle Fail Safe changed the definition, then you can look at the archived versions of the configuration files.

Oracle Fail Safe keeps up to two archived versions of configuration files. The file name of the archived version has a format of *filename_000.ora* and *filename_001.ora*. Note that *filename_000.ora* is the most recent file.

7.10.8.4 Rollback Files

Whenever Oracle Fail Safe encounters an error during an operation after Oracle Net configuration files have been changed, the updated version of the file is saved as *filename_rlb.ora*. Then, the original version of the file is restored.

The rollback version of the file may be useful for problem diagnosis.

7.10.9 Security Access and Authentication Problems

Access and authorization problems occur most frequently when you are attempting to perform operations through Oracle Enterprise Manager.

The following list addresses some typical authentication problems:

- From Oracle Enterprise Manager, the following error is returned when you are trying to connect to Oracle Fail Safe:

  FS-10101: Failed to authenticate the user *username* on the cluster.

  In Oracle Enterprise Manager, ensure that the User Credentials for the cluster are those of a Windows Administrator on all cluster nodes and that the user name and domain are specified correctly.

- Jobs that you submit to Oracle Fail Safe from Oracle Enterprise Manager fail with the error `Failed to authenticate user`.

  - Ensure that you have a Windows account that was set up with "Log on as batch user" access rights on each node in the cluster.

  - In Oracle Enterprise Manager, ensure that the User Credentials for each node in the cluster match the user name and password for your local account on each node in the cluster.

- When you try to perform an operation on or access a database that is configured in a group, the `ORA-01031: Insufficient privileges` error is returned.

  - When you create a sample database or add a database to a group, ensure that the authorization information for the database uses the `SYS` account with a password.

  - If you are attempting to access the database from Oracle Enterprise Manager, then ensure that the User Credentials for each database match the database `SYS` account.

7.10.10 Clients Cannot Access a Database

If users and client applications are unable to access a database that is configured in a group, then perform the following steps to fix the problem:

1. Update the `tnsnames.ora` file to use the virtual server for the group.

2. Run the `Verify Group` command to validate the network (Oracle Net) configuration.

7.11 Using Highly Available Databases with Oracle Data Guard

While Oracle Fail Safe provides high availability to single-instance Oracle databases, Oracle Data Guard provides disaster tolerance. For example, Oracle Fail Safe can ensure nearly continuous high availability for a given system, but does not protect against a disaster that incapacitates the site where that system resides. Similarly, while Oracle Data Guard provides excellent disaster recovery features, the time required to switch operations from the primary site to a physically separate site can range from several minutes to hours. By combining Oracle Fail Safe with Oracle Data Guard, your databases can be highly available and disaster tolerant.

If you have an Oracle Support contract, then you can find information about using Oracle Data Guard with Oracle Fail Safe, by logging into My Oracle Support and searching for note 259902.1 at

```
https://metalink.oracle.com
```

8

Configuring Generic Services for High Availability

A **generic service** is a Windows service that is supported by the generic service resource DLL provided with Microsoft Cluster Server (MSCS). Oracle Fail Safe support for configuring generic services for high availability lets you:

- Make your own applications highly available
- Make other applications highly available for which Oracle Fail Safe does not currently provide specific support

    For example, Oracle Fail Safe currently provides specific support for making Oracle single-instance databases and Oracle Management Agent highly available. Using the generic service support that Oracle Fail Safe provides, you can configure other Windows services to be highly available. (Section 8.3 describes how this can be performed.)

The following topics are discussed in this chapter:

- Introduction
- Discovering Standalone Generic Services
- Adding Generic Services to a Group
- Security Requirements for Generic Services
- Troubleshooting Problems with Generic Services

8.1 Introduction

The difference between configuring a resource for high availability that is specifically supported by Oracle Fail Safe and configuring a generic service for high availability, is in the level of assistance that the Add Resource to Group Wizard provides. For resources that are specifically supported, the Add Resource to Group Wizard in Oracle Fail Safe Manager requests configuration information that is targeted at that specific resource. For generic resources, the Add Resource to Group Wizard cannot know what the configuration information will be, and thus the requested data is less well defined. Therefore, you must be more aware of the resources upon which your generic service depends, the Windows registry entries required, and so on.

8.1.1 Advantages of Using Oracle Fail Safe

This section lists some advantages of using Oracle Fail Safe instead of MSCS for configuring generic services:

- Oracle Fail Safe can configure an existing service for high availability or it can create and configure a generic service for you as part of the Add Resource to Group operation; MSCS can configure only existing generic services for high availability.

- The Oracle Fail Safe Add Resource to Group Wizard provides more questions to help you configure a generic resource for high availability. For example, it lets you specify the disks required by the generic resource, Windows registry entries that need to be replicated across the cluster, and so on. The MSCS wizard does not provide as many questions to guide you.

- As part of the Add Resource to Group operation, Oracle Fail Safe tests your configuration (as it does for all types of resources it configures for high availability). Oracle Fail Safe tests the network, failover and failback, and ensures that the resource can be started on all cluster nodes that are possible owners of the resource.

- Oracle Fail Safe sets the startup type for the resource on each cluster node to manual, as is required in a cluster environment. MSCS only sets the startup type on the node that currently owns the resource; you must remember to set the startup type to manual on the other cluster nodes that are possible owners.

Use the generic service resource type in the Add Resource to Group Wizard to configure a service for high availability if Oracle Fail Safe does not provide a customized wizard for that service. You can determine the services for which Oracle Fail Safe provides customized wizards by selecting the cluster in the Oracle Fail Safe Manager tree view and then clicking the Resources tab. Resources for which Oracle provides a customized wizard are listed on this property page.

8.1.2 Generic Resources That Must Not Be Configured for High Availability

When you consider configuring a generic service for high availability, keep in mind that after you do so, the service runs on only one cluster node at a time. Services you run on cluster nodes concurrently must not be configured for high availability. For example, consider the Windows Event Log. The Windows Event Log is a file to which all services on a given system can write informational messages, error messages, and so on. It is a means for the service to communicate conditions to the administrator.

If you make the Windows Event Log service highly available, then the service would run on only one cluster node at a time. Messages returned by services on the other cluster nodes would not have access to the Event Log on the cluster node running the Event Log. Therefore, it would be unwise to configure the Windows Event Log as a cluster resource.

8.2 Discovering Standalone Generic Services

Oracle Services for MSCS discovers generic resources by searching for them in the Windows service manager. During the discovery process, Oracle Fail Safe locates services in the Windows service manager on each node in the cluster and then displays the newly discovered services in the Oracle Fail Safe Manager Add Resource to Group Wizard.

To ensure that the properties of standalone and cluster resources are discovered and displayed correctly by Oracle Fail Safe Manager and Oracle Enterprise Manager, each resource must have a unique name within the cluster. It may be necessary to specify names that are different from default values or to change the default names of resources.

8.3 Adding Generic Services to a Group

To configure a generic service for high availability, you must add it to a group using the Oracle Fail Safe Manager Add Resource to Group Wizard. You can either add an existing generic service to a group, or you can specify Oracle Fail Safe to create the generic service. The following sections describe the configuration steps and the data that is needed to complete the Add Resource to Group Wizard for a generic service.

8.3.1 Configuration Steps

Table 8–1 provides a quick reference to the tasks needed to configure a generic service for high availability. For step-by-step instructions about any particular task, refer to the Oracle Fail Safe online help. From the Oracle Fail Safe Manager menu bar, select **Help,** then **Search for Help on.**

Table 8–1   Steps for Configuring a Generic Service

Step	Procedure	Comments
1	Ensure that the generic service executable file is installed on a private disk on each cluster node that will be a possible owner for the generic service.	This is not required, but is strongly recommended. Typically, several service instances use the same executable file. If the executable file is installed on a shared cluster disk, then all services that use that executable file must run on the cluster node that currently hosts that disk.
2	Copy files required by the generic service to a cluster disk.	If data files are required by the generic service, then they must be located on the cluster disks on the shared storage interconnect.
3	Start Oracle Fail Safe Manager.	From the Windows **Start** menu, select **Oracle - *Oracle_Home,*** then **Oracle Fail Safe Manager.**
4	Verify the cluster.	Select **Troubleshooting,** then **Verify Cluster** to run a procedure that validates the cluster hardware and software configurations.
5	Create a group.	Select **Groups,** then **Create** to run the Create Group Wizard. The wizard helps you to set up failover and failback policies and automatically opens the Add Resource to Group Wizard to let you add a virtual address to the group. Oracle Fail Safe does not require you to add a virtual address to a group before you add a generic service. However, the resources on which the generic service depends may require a virtual address. See Section 8.3.2.5.2 for details.
6	If needed, add one or more virtual addresses to the group.	Select **Resources,** then **Add to Group** to run the Add Resource to Group Wizard. The wizard helps you to create and configure the virtual server address for high availability.
7	Add resources upon which the generic service depends.	Select **Resources,** then **Add to Group** to open the Add Resource to Group Wizard.
8	Add the generic service to the group.	select **Resources,** then **Add to Group** to open the Add Resource to Group Wizard. The wizard helps you configure the generic service into a group. You can create a new generic service or specify an existing generic service.
9	Verify the group.	Select **Troubleshooting,** then **Verify Group** to check for and fix any problems with the group, virtual addresses, resources, or the failover configuration.

8.3.2 Configuration Data for Generic Services

To configure a generic service for high availability, you add it to a group. Oracle Fail Safe can create and add a new generic service to group, or you can add an existing generic service to a group. In either case, when you use the Oracle Fail Safe Manager Add Resource to Group Wizard, you need the following data:

- Possible owner nodes for the generic service, if the cluster consists of more than two nodes, or if one node is not available in a two-node cluster

- Identity (node name, display name, service name, and image path) of the generic service

- Account under which the generic service is run and its startup parameters

- Disks, if any are used by the generic service

- Other resources upon which the generic resource depends

- Windows registry key values that the generic service uses

Unlike most other resources that you configure for high availability, you are not required to add a virtual address to a group before adding a generic service. You must determine, based on the use of the generic service, if a virtual address is needed. The following sections examine the issues you must consider to determine whether you must add a virtual address to the group and the configuration information needed to add a generic resource to a group.

8.3.2.1 Choose Nodes

If you are adding a generic service to a group and the cluster consists of more than two nodes, then you are asked to specify the nodes which must be possible owners for the generic service by specifying a list of selected nodes, as shown in Figure 8–1. To specify that a particular node must not be a possible owner for the generic service, select the node from the Selected Nodes list and click the left arrow.

Section 2.6.7 describes in detail the concept of the possible owner nodes list.

Figure 8–1   Choose Nodes Wizard Page When All Nodes Are Available
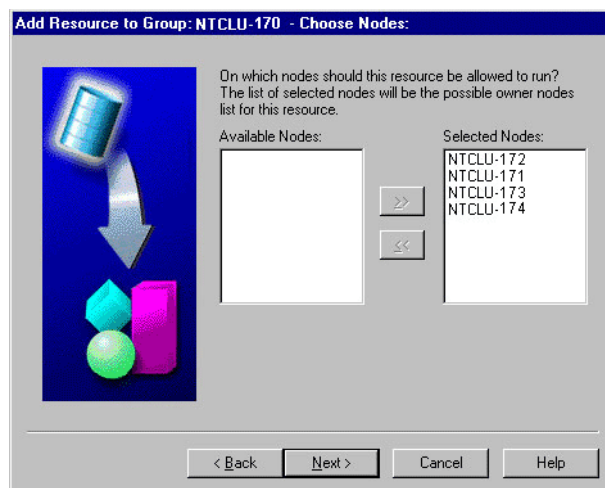
If you are adding a generic service to a group and the cluster consists of two or more nodes, but one or more nodes are unavailable, then you are also asked to specify which nodes must be possible owners for the generic service. In this case, the wizard page displays the nodes that are unavailable and why, as shown in Figure 8–2.

Figure 8–2 Choose Nodes Wizard Page When Any Node Is Unavailable
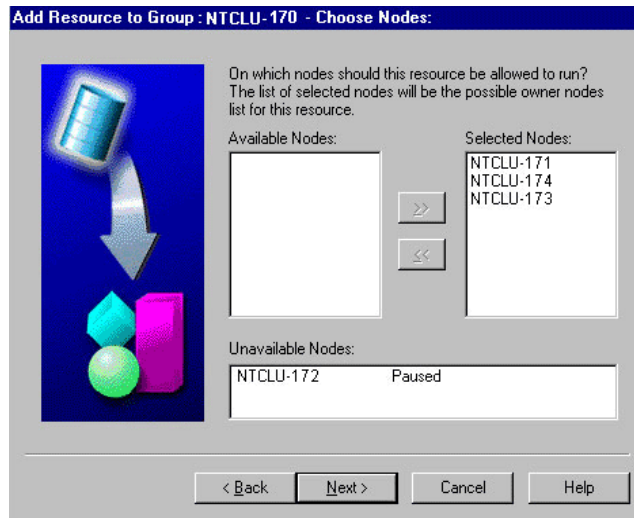
8.3.2.2 Generic Service Identity

When you configure a generic service for high availability, you must provide some basic information that Oracle Fail Safe can use to uniquely identify and locate the executable files for the generic service. In particular, the Add Resource to Group Wizard requests the following about the generic service identity:

- Node name

  For an existing generic service, Oracle Fail Safe must know on which cluster node the generic service currently exists. If it exists on multiple nodes, then specify any one of them in the Add Resource to Group Wizard. If the service does not already exist, then select any node that is a possible owner for the generic service.

- Display name

  The display name is used to describe the service in more detail than the service name. It can contain both spaces and up to 256 characters. The display name is shown in the Windows Services dialog box.

  The display name is also the name used by Oracle Fail Safe to refer to the service in the Oracle Fail Safe Manager tree view.

- Service name

  The service name, sometimes referred to as a short name, labels the Windows registry subkey that contains the configuration information for the service. It must not contain spaces and is typically shorter than the display name.

- Image name

  This is the path and file name for the generic service executable file. The executable file for a generic service must be installed on the same private disk and directory on all cluster nodes that are possible owners of the generic service. This ensures that if the generic service fails over, the executable files upon which it depends are available on the other cluster nodes.

  Oracle recommends that you do not install the generic service executable file on a shared cluster disk. Typically, several service instances use the same executable file. If the executable file is installed on a shared cluster disk, then all services that use that executable file must run on the cluster node that currently hosts that disk.

When you install the executable file at the same location on each cluster node, each cluster node can host different service instances that access that same executable file. For example, you have two services, Service_A and Service_B, which use the same executable file. If the executable file is installed at the same location on each cluster disk, then Service_A can belong to Group_A, whose primary node is Node_1; and Service_B can belong to Group_B, whose primary node is Node_2. If you install the executable file on a shared cluster disk that belongs to Group_C, then the service can run only on the cluster node that is currently hosting Group_C.

Figure 8–3 shows the page in the Add Resource to Group Wizard on which you specify the generic service identity. If you enter an existing service in the Service Name box, then the Status box displays the status of the service. If you enter a new service in the Service Name box, then the Status box is empty. Oracle Fail Safe Manager presents the status for your information. You can add an existing generic resource to a group regardless of whether it is running or stopped.

Figure 8–3   Generic Service Identity Wizard Page
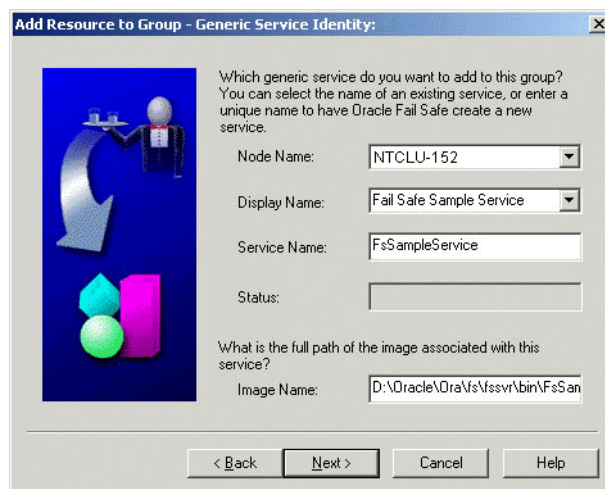
8.3.2.3 Generic Service Startup Parameters

The Add Resource to Group Wizard asks for the following details about how the generic service must be started:

- Startup parameters

  You specify startup parameters that you want Oracle Fail Safe to pass to the Windows Service Control Manager. These parameters are the same as those you would specify if you were using the Windows Services dialog box, for example, -t. Oracle Fail Safe passes parameters unchanged to the Service Control Manager.

- Log on as system account or user account

  You specify the account under which you want the service to run: the system account or a user account. Log on as System Account is selected by default. To log on as a user account, select This Account under "Log on as:" The account under which the service runs defines the security context for the generic service. When the service logs on as a system account (LocalSystem), it has access to all files on the local system, but no access to files across the network. When the service logs on as a user account, it can have access both to files on the local system and across the network, depending on which privileges it has. For example, Oracle Fail Safe

itself runs under a user account (which you specify when you install Oracle Fail Safe) because it should be able to access files on all cluster nodes.

The Add Resource to Group Wizard does not allow you to change the account under which an existing service runs. To change the account under which an existing generic service runs, use the Windows Services dialog box to change the account before attempting to add it to a group. (See Section 4.3.1 for information about changing the account under which Oracle Services for MSCS runs.)

You will notice that Oracle Fail Safe does not request startup type information (Automatic, Manual, or Disabled) for the generic service. The startup type for all resources configured for high availability using Oracle Fail Safe is set to Manual. In a cluster environment, the service must only run on one node at a time. By setting the startup type to Manual, Oracle Fail Safe ensures that the resource runs on one node at a time and is started by MSCS only.

Figure 8–4 shows the page in the Add Resource to Group Wizard on which you specify the generic service startup parameters and account.

Figure 8–4   Generic Service Account Wizard Page
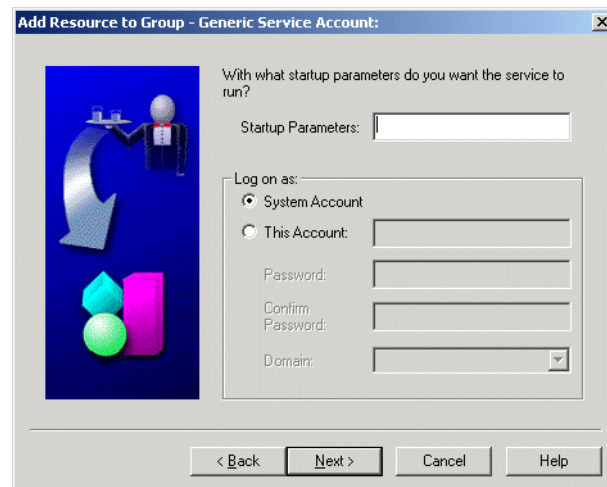
8.3.2.4 Disks Used by a Generic Service

Oracle Fail Safe requires that data files needed by a highly available generic service be available on the cluster node currently running the service. This is accomplished in one of two ways:

- You place the data files required by the service on a shared cluster disk that is included in the same group as the resource.

  In a failover, the disk fails over with the service so that the files are still available to the service.

- You place the same file on the same private disk and directory on all cluster nodes that are possible owners of the generic service.

  In a failover, the service uses the same path to find the file on the private disk. Because the path to the file is the same on each cluster node, the resource can locate it, regardless of which cluster node is hosting the resource.

Typically, the service executable file is installed on a private disk on each cluster node, and the data files are placed on shared cluster disks. See Section 8.3.2.2 for more information about the placement of the executable file.

You may decide to place the data files on the same private disk and directory on each cluster node if you specifically want the contents of the files to be different depending upon the cluster node on which the generic service is running. For example, suppose Node_1 has twice the CPU and memory that Node_2 has. If your generic service uses a file to specify the maximum number of users that can access it concurrently, then you may want to set that number to 100 on Node_1 and 50 on Node_2.

However, Oracle recommends that data files be placed on shared cluster disks whenever possible. If you intend to follow this recommendation, then you must move any data file that the generic resource uses to a shared cluster disk prior to running the Add Resource to Group Wizard.

Figure 8–5 shows the page in the Add Resource to Group Wizard on which you specify the disk dependency.

Figure 8–5   Generic Service Disks Wizard Page
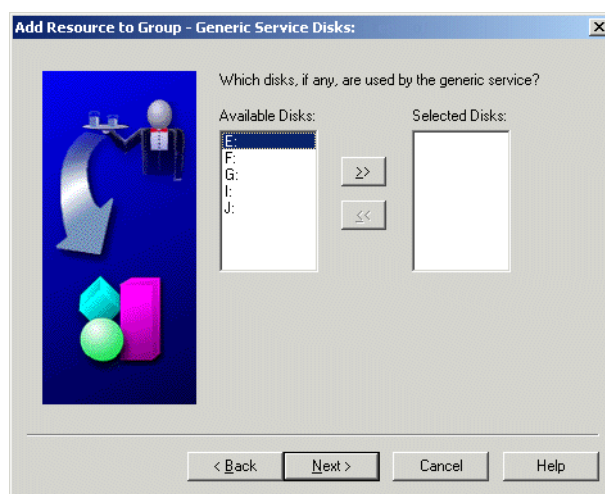
8.3.2.5 Generic Service Dependencies

Because you are configuring a resource about which neither Oracle Fail Safe Manager nor MSCS has detailed information, the process for configuring a generic service for high availability is less automated than for resource types about which Oracle Fail Safe does have detailed information (such as an Oracle database). For example, when you use Oracle Fail Safe to configure an Oracle database for high availability, Oracle Fail Safe Manager includes the resources upon which the database depends and determines the order in which those resources need to be brought online. When you configure a generic resource for high availability, you need to provide this dependency information.

8.3.2.5.1   Specifying Generic Service Dependencies You provide dependency information for a generic resource by the order in which you add resources to a group. For example, suppose you want to make Service_A highly available, but in order for Service_A to come online, Service_B and Service_C must be online already. In other words, Service_A has a dependency on Service_B and Service_C. Furthermore, in order for Service_B to come online successfully, Service_C must be online already. This chain of dependencies can be illustrated by a tree, as shown in Figure 8–6.

Figure 8–6   Dependency Tree
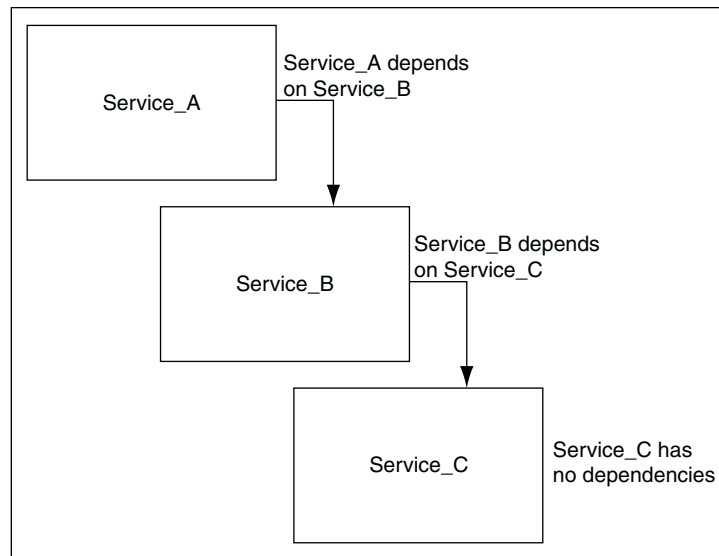
In this scenario, you add Service_C to the group first. Then you add Service_B to the group and specify Service_C as a dependency. Finally, you add Service_A to the group and specify Service_B as a dependency. In effect, you build the dependency tree one resource at a time. Each time you add a resource to a group, you can specify only one dependency level. Therefore, if the dependency tree is two levels deep (or more), then the order in which you add the resources to the group is important.

Figure 8–7 shows the page in the Add Resource to Group Wizard on which you specify resource dependencies.

Figure 8–7   Generic Service Dependencies Wizard Page
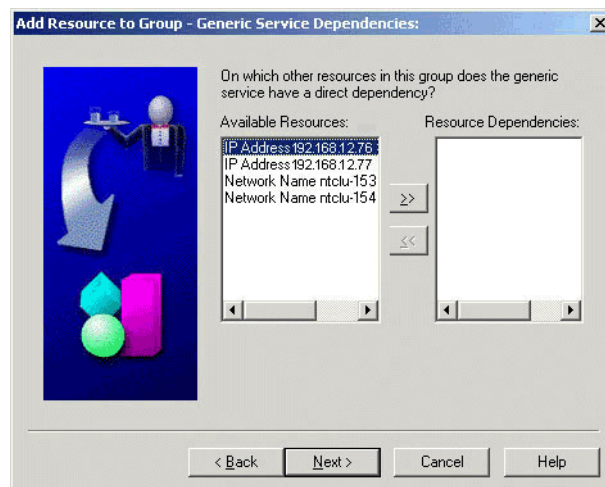
8.3.2.5.2   Generic Services and Virtual Address Dependencies Oracle Fail Safe does not require you to add a virtual address to a group before you add a generic service to the group. A virtual address specifies the network address at which a resource can be found by clients or other services. If neither clients nor other services attach to the generic service, then you do not need to add a virtual address to the group before you add the generic service.

You may need to add a virtual address to the group, however, if the generic service is accessed by clients or other services.

8.3.2.6 Generic Service Registry Keys

If your generic service uses Windows registry entries to store information, then you can specify these entries in the Add Resource to Group Wizard. By specifying them in the wizard, you ensure that the Windows registry entries for your service are consistent across the cluster nodes that are possible owner nodes for the generic resource. This is important so that in a failover, your service can run correctly on any cluster node that is a possible owner of the generic resource.

For example, if you were to manually configure an Oracle Forms Server as a generic service, then you would specify the `FORMS60_PATH` registry variable.

The root for the registry keys you specify is assumed to be `HKEY_LOCAL_MACHINE`. This is discussed in more detail in the online help for the Add Resource to Group Wizard.

Figure 8–8 shows the page in the Add Resource to Group Wizard on which you specify Windows registry keys.

Figure 8–8   Generic Service Registry Wizard Page
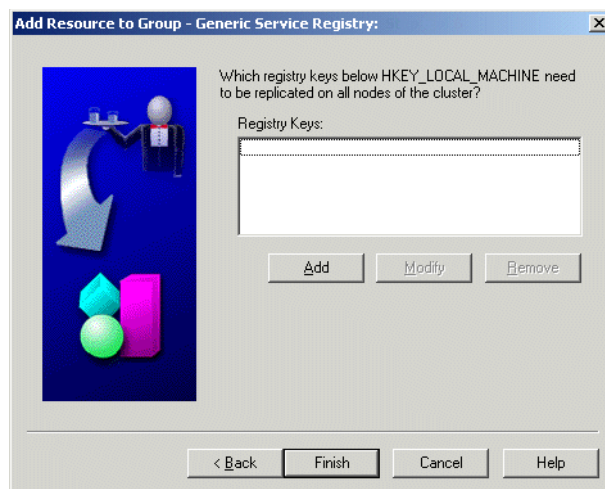
8.4 Security Requirements for Generic Services

By default, a generic service runs under the local system account. If you specify that the generic service must run under a user account, then it must have the "Log on as a service" privilege. When you add a generic service to a group, Oracle Fail Safe checks to see if the account under which the generic service is running has this privilege; if it does not, then Oracle Fail Safe grants it to the specified user account.

In addition, Oracle Fail Safe checks that the user account and password you specified in the Add Resource to Group Wizard are valid. If not, then Oracle Fail Safe returns an error.

8.5 Troubleshooting Problems with Generic Services

You can run the `Verify Group` operation at any time. However, you *must* run it when any of the following occurs:

- A group or resource in a group does not come online.

- Failover or failback does not perform as you expect.

- A new node is added to the cluster.

- You accidently delete a generic service from a cluster node.

  When you run the `Verify Group` operation it automatically re-creates the same generic service on all cluster nodes that are possible owners of the service.

General information about the Oracle Fail Safe troubleshooting tools (`Verify Cluster` and `Verify Group`) is in Chapter 6.

9

Configuring Oracle Management Agent for High Availability

You can configure Oracle Enterprise Manager 10*g* Grid Control to monitor databases configured for high availability by using the following procedure:

1. Install Oracle Management Agent.

2. Create an Oracle Management Agent that listens on a virtual address.

3. Add the Oracle Management Agent created in Step 2 to the same group as the Oracle Database (or databases) configured for high availability.

4. Configure the Oracle Management Agent to monitor the database or databases.

The following topics are discussed in this chapter:

- Prerequisites for High Availability

- Procedure for Configuring Oracle Management Agent for High Availability

- Removing Oracle Management Agent from a Group

9.1 Prerequisites for High Availability

You must install the following software on the cluster system before you can configure an Oracle Management Agent for high availability:

- Oracle Database – any release supported by Oracle Enterprise Manager 10*g* Grid Control

- Oracle Management Agent

  You must install the Management Agent on each cluster node, using the same Oracle home on each node.

- Oracle Fail Safe

In addition, the following components must be configured:

- An Oracle database instance must be configured for high availability.

- An Oracle Enterprise Manager Management Server must be configured and available for setup. The Management Server need not reside on the cluster system.

9.2 Procedure for Configuring Oracle Management Agent for High Availability

You must follow the following steps to configure Oracle Management Agent for high availability:

- Step 1, "Make the Management Agent Highly Available"
- Step 2, "Add the Highly Available Database as a Target in Oracle Enterprise Manager"
- Step 3, "Test the Highly Available Management Agent"
- Step 4, "Remove Extraneous Targets from the Oracle Enterprise Manager Environment"

Step 1 Make the Management Agent Highly Available

Use Oracle Fail Safe Manager to add the new Oracle Management Agent to the Fail Safe group that contains the databases it will monitor. Oracle Fail Safe will create a Generic Service resource for the Oracle Management Agent when you select the **Oracle Management Agent** Resource Type in the Add Resource to Group Wizard. Follow these steps:

1. From the Resources menu, click **Add to Group.** The Add Resource to Group Wizard opens.

Figure 9–1   Add Resource to Group Wizard - Resource Page
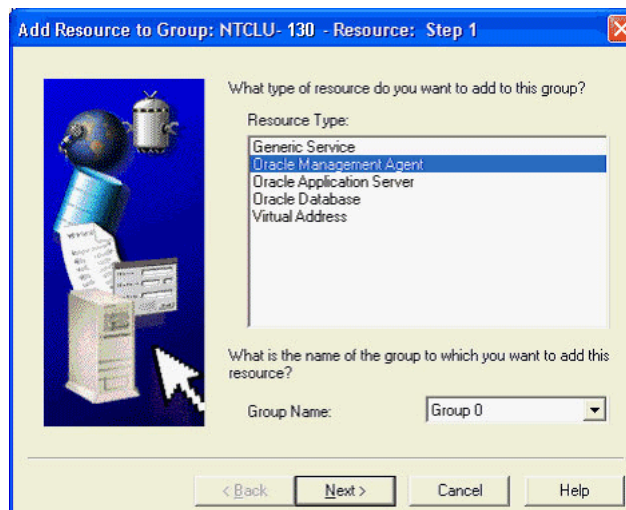
2. In the **Resource Type** box, select **Oracle Management Agent.**

3. In the **Group Name** box, select the group to which you want to add Oracle Management Agent. This must be the group that contains the database you want to monitor with Oracle Management Agent.

4. Click **Next.** The Management Agent Oracle Home page opens.

Figure 9–2   Add Resource to Group Wizard - Management Agent Oracle Home Page
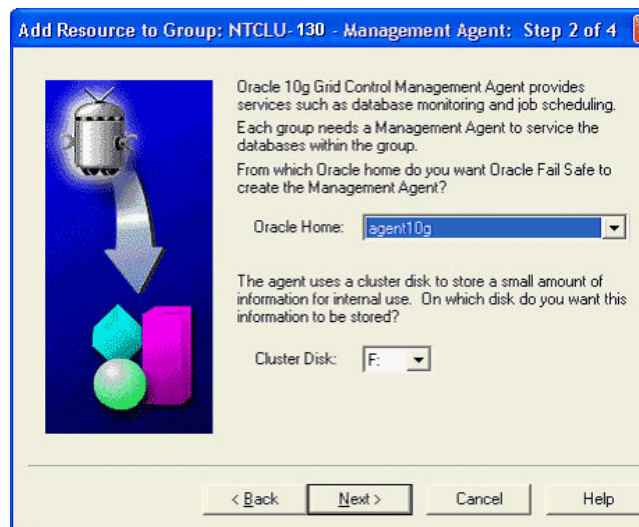
5. In the **Oracle Home** box, select the Oracle home that you want to add to the group.

6. Select the shared disk that will be used to store Management Agent context information.

7. Click **Next.** The Management Agent Virtual Address page opens.

Figure 9–3   Add Resource to Group Wizard - Management Agent Virtual Address Page
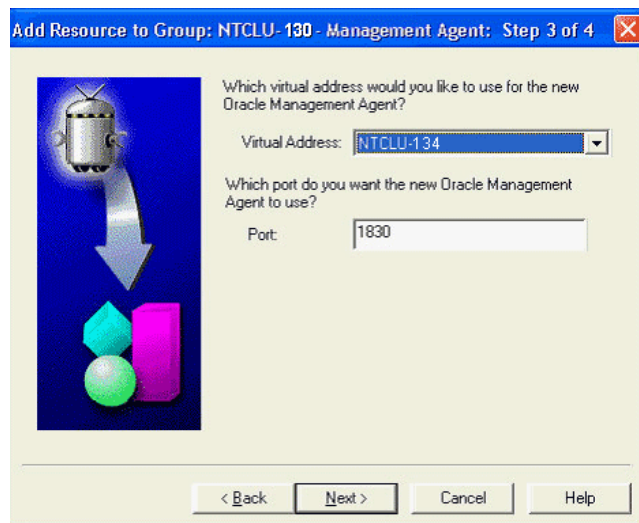
8. In the **Virtual Address** box, select the desired virtual address.

9. In the **Port** box, enter an open port number.

10. Click **Next.** The Management Agent password page opens.

Figure 9–4   Add Resource to Group Wizard - Management Agent Page
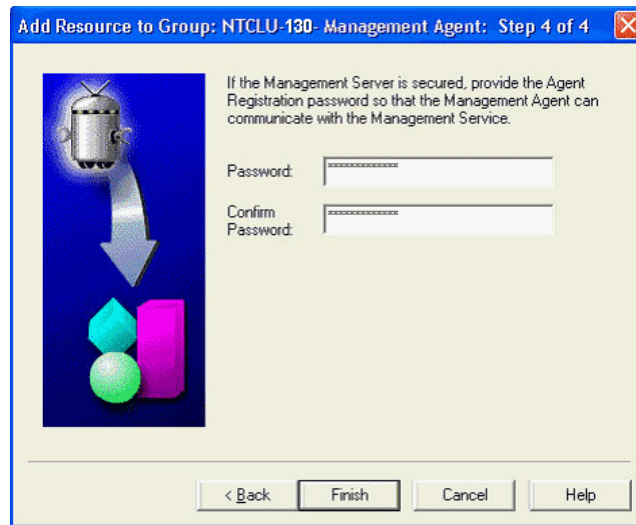
11. Provide the password used for the Oracle Management Server.

12. Click **Finish.** The Finish Adding Management Agent window opens.

13. This window summarizes the properties you have set with the wizard. If the summary information is correct, then click **OK.** Otherwise, click **Cancel** and then click **Back** to find the page or pages in the wizard on which you want to change entries.

    When you click **OK,** a window opens to display the progress of this clusterwide operation. (A clusterwide operation is one that occurs on multiple cluster nodes.)

Step 2 Add the Highly Available Database as a Target in Oracle Enterprise Manager

To configure the highly available database for monitoring through the highly available Management Agent, follow these steps:

1. Log in to the Oracle Enterprise Manager 10*g* Grid Control Console.

2. Click the **Targets** tab. The Hosts page opens.

3. Click the **Databases** secondary tab. The Databases page opens.

4. Click **Add** (in the upper right-hand section of the page). The Add Database Target: Specify Host page opens.

5. Click the **flashlight** icon. The Search and Select Host window opens.

6. Select the virtual host, and then click **Select.** The Search and Select Host window closes and the Host field in the Add Database to Target: Specify Host page contains the selected virtual host.

7. Click **Continue.** The Targets Discovered on Host page opens.

8. Select the highly available database or databases that you want to monitor, then click **OK.** The Database Configuration Results page opens.

9. Click **OK.**

Step 3 Test the Highly Available Management Agent

To test that the Management Agent is highly available, run a job against the highly available database it is monitoring, and follow these steps:

1. Log in to the Oracle Enterprise Manager 10*g* Grid Control Console.

2. Create a SQL job and submit it against the highly available database, as follows:

   a. Click the **Jobs** tab. The Job Activity page opens.

   b. In the **Create Job** box, select **SQL Script,** and then click **Go.** The Create 'SQL Script' Job page opens.

   c. In the **Job Name** box, enter `TEST JOB,` and in the **SQL Script** field, enter `SELECT * FROM ts$.`

   d. In the Databases region of the page, click **Add.** The Add Targets page opens.

   e. Select the highly available database host name (which matches the virtual host name), and then click **Add.** The Create 'SQL Script' Job page opens.

   f. In the Host and Database Credentials portion of the page, specify the database credentials, and then click **Submit.**

3. Ensure that the submitted job completes successfully.

4. Create another job against the same database (by following Step 1 and Step 2 in this list), but schedule it for 10 minutes from the current time.

5. Using Oracle Fail Safe Manager, fail over the group.

6. After 10 minutes pass, check that the second job scheduled ran successfully.

Step 4 Remove Extraneous Targets from the Oracle Enterprise Manager Environment

When you reach this step, the Oracle Enterprise Manager 10*g* Grid Control Console shows three targets for the same database. During the Management Agent installation, the installer automatically discovers Oracle components, including highly available databases, and adds the discovered components as targets.

Because the highly available database instance exists on each cluster node, there are two targets for the database – each monitored by a different Management Agent. In addition, you create a third target when you add the database as a target for the highly available Management Agent listening on the virtual server (in Chapter 10, "Configuring Oracle Application Server Components for High Availability").

You can safely remove the database targets that were discovered when the Management Agent was installed by performing the following steps:

1. Log in to the Oracle Enterprise Manager 10*g* Grid Control Console.

2. Click the **Targets** tab. The Hosts page opens.

3. Click the physical (as opposed to virtual) host name for one of the cluster nodes. The Host page for that physical host opens.

4. Click the **Targets** locator link.

5. Select the database on this host, and then click **Remove.**

6. Repeat Step 2 though Step 5 for each cluster node.

9.3 Removing Oracle Management Agent from a Group

If you decide you no longer want to have an Oracle Management Agent configured for high availability, then remove it from the group that contains it. When you do so, the Oracle Management Agent is deleted from the cluster.

To remove an Oracle Management Agent from a group, perform the following steps:

1. Open Oracle Fail Safe Manager.

2. In the tree view, select the Oracle Management Agent that you want to remove and then, on the **Resources** menu, select **Remove from Group.**

3. The Confirm Remove from Group box opens. Click **Yes** in the **Confirm Remove from Group** box.

10

Configuring Oracle Application Server Components for High Availability

Oracle Fail Safe Manager provides a wizard to help you configure instances of the following Oracle Application Server components for high availability:

- Oracle Process Manager and Notification Server (OPMN)

- Oracle Enterprise Manager 10*g* Application Server Control (Application Server Control) service

- The metadata repository, if it was installed in the same Oracle home as the other Oracle Application Server components

The following topics are discussed in this chapter:

- Prerequisites for High Availability

- Procedure for Configuring Oracle Application Server Components for High Availability

- Removing Oracle Application Server Components from a Group

10.1 Prerequisites for High Availability

If a metadata repository is in the same home where you installed Oracle Application Server, then before adding instances of Oracle Application Server components to a group, you must do the following:

- Add at least one virtual address to the group.

- Ensure that the database files are on cluster disks. If a failover occurs, then any cluster node that is a possible owner of the instance must be able to access the database files.

You need the following information to complete the wizard. Examples are provided in parentheses.

- Oracle home where Oracle Application Server is installed (OAS)

- If the metadata repository was installed in the Oracle home where Oracle Application Server is installed, then you also need the following information:

  - Database instance name (OFS5)

  - Database name (OFS5DB)

  - Database parameter file (`T:\OFSDB\PARA\initofs5.ora`)

  - Database `SYS` account password if the database has a password file

10.2 Procedure for Configuring Oracle Application Server Components for High Availability

To configure Oracle Application Server components for high availability, perform the following steps:

1. Install Oracle Application Server as described in *Oracle Application Server Installation Guide* (10*g* Release 2 (10.1.2) for Microsoft Windows).

2. Open Oracle Fail Safe Manager.

3. On the tree view, select the standalone Oracle Application Server instance that you want to add to a group.

4. From the **Resources** menu, select **Add to Group.** The Add Resource to Group Wizard opens.

5. In the **Resource Type** box, select **Oracle Application Server.**

6. In the **Group Name** box, select the group to which you want to add Oracle Application Server.

7. Click **Next.**

8. Complete the remaining wizard pages. If you need assistance, then click **Help** on the wizard page.

Note the following:

- Unlike a standalone database instance, a standalone Oracle Application Server instance is represented in the tree view on each cluster node. You can select the standalone instance that you want to add to a group from any one of the cluster nodes.

- A standalone Oracle Application Server instance is represented in the tree view as the home into which Oracle Application Server was installed. After it is configured into the group, it is represented by the resources that compose it: OPMN and Application Server Control. Oracle Fail Safe brings these resources online after it has configured them on each cluster node.

- To make an Oracle Application Server metadata repository that is not installed in the same home as Oracle Application Server highly available, you add it to the group to which you intend to add the other Oracle Application Server components. It is important that you add it to the group before you add the other Oracle Application Server components because Oracle Fail Safe sets up a dependency between OPMN and the metadata repository when it configures OPMN into the group.

10.3 Removing Oracle Application Server Components from a Group

If you decide you no longer want to have Oracle Application Server components configured for high availability, then remove them from the group that contains them.

Because the OPMN server has a dependency on the Oracle Application Server Control service, you must remove the Oracle Application Server Control service from the group before you remove the OPMN server. In addition, if the metadata repository was added to the group when you configured the other Oracle Application Server components into the group, then it must be removed after the Oracle Application Server Control service and the OPMN server are removed because of dependencies it has on the other components.

To remove Oracle Application Server components from a group:

1. Open Oracle Fail Safe Manager.

2. In the tree view, select the Oracle Application Server Control service that you want to remove and then, on the **Resources** menu, select **Remove from Group.**

3. The Confirm Remove Resource from Group box opens. Click **Yes** in the **Confirm Remove from Group** box. A window opens to show you the progress of the operation.

4. When the operation to remove the Oracle Application Server Control service completes, select the OPMN server that you want to remove and then, on the **Resources** menu, select **Remove from Group.**

5. The Confirm Remove Resource from Group box opens. Click **Yes** in the **Confirm Remove from Group** box. A window opens to show you the progress of the operation.

6. When the operation to remove the OPMN server completes, select the metadata repository (if it was added when you configured the Oracle Application Server Control service into the group) and then, on the **Resources** menu, select **Remove from Group**.

7. The Confirm Remove Resource from Group box opens. Click **Yes** in the **Confirm Remove from Group** box. A window opens to show you the progress of the operation.

After you remove the Oracle Application Server components from a fail-safe group, they remain on the cluster as standalone resources. Note the following:

- Unlike a standalone database instance, a standalone Oracle Application Server instance is represented in the tree view on each cluster node.

- A standalone Oracle Application Server instance is represented in the tree view as the home into which Oracle Application Server was installed.

A

Network Configuration Considerations

This appendix describes the network considerations for the Microsoft Windows software.

This appendix covers the following topics:

- Registering Host Names and IP Addresses
- Validating Proper Name Resolution in the Cluster
- Changing IP Addresses of Cluster Nodes
- Troubleshooting Problems with Improper Name Resolution

A.1 Registering Host Names and IP Addresses

Each **IP address** and its corresponding **host name** *must* be registered in either the Hosts file (in the `system32\drivers\etc\` directory under the system directory) on each cluster node or in DNS regardless of whether WINS is used. You must register all of the following addresses:

- Cluster node addresses
- Cluster alias address
- Virtual server addresses
- Oracle Fail Safe Manager client systems

If you make changes in any of the name registries (DNS server or Host file), then use the following command to purge the local cache and make the changes effective:

```
nbtstat -R
```

A.2 Validating Proper Name Resolution in the Cluster

This section describes how to ping all systems to validate that you have properly registered the IP addresses and host names. The procedures described in this section ensure that the same IP address is echoed when you ping a host name from the host itself and from any other systems.

Suppose there is a cluster of two systems and a client system. The host name of the systems and the cluster alias are shown in the following table:

System	Host Name
Cluster system 1	ClusterHost1
Cluster system 2	ClusterHost2

System	Host Name
Cluster alias	ClusterAlias
Client system	ClientHost

On each system, Cluster system 1, Cluster system 2, and Client system ping all four host names. For example, on Cluster system 1:

```
ping ClusterHost1
ping ClusterHost2
ping ClusterAlias
ping ClientHost
```

For each host name that you ping, the ping test must echo the same IP address. For example, if ClusterHost1 is registered to map to IP address 192.1.99.202, then pinging ClusterHost1 on Cluster system 1, Cluster system 2, and Client system must echo 192.1.99.202.

If all ping tests are successful, then your network is configured correctly.

If any ping tests do not return the proper IP address, then see Section A.4.

A.3 Changing IP Addresses of Cluster Nodes

A Microsoft Windows cluster system is typically composed of two nodes and two networks: a public network, which is used by clients of the cluster, and a private network called the cluster interconnect, which is used internally by the cluster hardware and software to monitor the state of the two nodes during their operation.

Thus, each node has at least two IP addresses, one on the public network and one on the private network (for example, 192.168.10.1 could be the IP address on the public network and 10.10.0.1 could be the IP address on the private network).

These addresses are specified in the Microsoft Cluster Server software.

The following steps describe how to change the IP addresses of cluster nodes in a Microsoft Windows cluster. The steps assume that there are two cluster nodes: Node A and Node B. The following steps are performed on Node A unless otherwise noted.

1. Open Oracle Fail Safe Manager and take all of the groups offline except the Cluster Group. (In the tree view, right-click a group, then select **Take Offline.** Repeat until all groups except the Cluster Group have been taken offline.)

2. Close Oracle Fail Safe Manager.

3. Open MSCS Cluster Administrator and check that all of the groups are offline except the Cluster Group. Then follow these steps:

   a. In the tree view, select the **Resources** folder.

   b. Right-click **Oracle Services for MSCS** and select **Take Offline**.

   c. Right-click a group and select **Move Group.** Repeat until all groups are moved to Node B.

4. Change the Cluster IP address (for Node A) in MSCS Cluster Administrator, as follows.

   a. Select the **Resources** folder.

   b. Right-click **Cluster IP Address.**

   c. Select **Properties.**

    d. Change the IP Address.

    e. Click **OK.**

5. For each group, change its IP address in MSCS Cluster Administrator, as follows:

    a. Select the **Resources** folder.

    b. Right-click the Network Name *IP address* for a group and select **Properties.**

    c. Change the network Name on the **General** tab and change the IP address Name on the **Parameters** tab.

    d. Click **Apply.**

    e. Repeat steps b through d for all of the groups.

6. Change the IP Address of the network adapter on Node A. (For instructions, see Microsoft Article ID: Q230356 - *Changing the IP Address of Network Adapters in Cluster Server,* available on the Microsoft Corporation Web site.)

7. Restart Node A.

8. After Node A restarts, open MSCS Cluster Administrator.

9. Right-click a group and select **Move Group.** Repeat until all groups are moved to Node A.

10. Change the IP Address of the network adapter on Node B.

11. Restart Node B.

12. After Node B restarts, update the host file on both Node A and Node B (located in `winnt/system32/drivers/etc/host`).

13. Update the DNS Server if necessary.

14. Open MSCS Cluster Administrator and ensure that the Cluster Group and all resources are online.

15. Open Oracle Fail Safe Manager and then perform the following steps:

    a. Ensure that all resources are online.

    b. Select the cluster from the tree view, and on the **Troubleshooting** menu, select **Verify Cluster.** Ensure that no warnings are returned due to a problem with host name to IP address resolution.

    c. Select a group from the tree view, and on the **Troubleshooting** menu, select **Verify Group.** This reports a number of errors in Oracle Net configuration files and Microsoft Windows services and asks if you want the errors to be fixed. Accept all fixes and rerun the Verify Group operation until no more fixes are suggested and no warnings are produced.

    d. Repeat step 15c for each group on the cluster.

    e. Test manual failover of each group on both nodes and then test connectivity from clients. (To test manual failover of a group, right-click the group in the Oracle Fail Safe Manager tree view, then select **Move Group.**)

A.4 Troubleshooting Problems with Improper Name Resolution

If the ping test in Section A.2 does not return the correct addresses, then one of the following is probably causing the discrepancy:

- You have incorrectly registered the IP address, or the host name, or both in the Host file or with the DNS server.

  If you suspect this is the problem, then reexamine Section A.1 to ensure that you have registered the addresses correctly, or contact your network administrator to verify your IP addresses and host names.

- You have multiple network interface cards.

  On a system with multiple network interface cards, multiple IP addresses are assigned to the host name. For example, Cluster system 1 has two network interface cards. Each card is assigned an IP address on a different network or subnet, as follows:

Network	IP Address
Net1	192.1.22.101
Net2	192.1.99.202

  On system Cluster system 1, host name ClusterHost1 maps to both IP addresses, although the ping program echoes only one IP address. If the ping program does not echo the same IP address that is echoed on other systems, then it could mean that there is an ordering problem with the IP addresses for the host name.

  Refer to the Microsoft Windows documentation for information about correcting an ordering problem.

B

Contacting Oracle Support Services

This appendix discusses the following topics:

- Reporting a Problem
- Finding Your Version Information
- Tracing Oracle Fail Safe Problems
- Locating Trace and Alert Files

B.1 Reporting a Problem

Some messages recommend calling Oracle to report a problem. When you call your Oracle Support representative, have the following information available:

- The hardware, operating system, and release number of the operating system on which the Oracle software is running
- The complete release number of Oracle and other product software
- All Oracle programs (with release numbers) in use when the error occurred
- If you encountered one or more error codes or messages, then have the exact code numbers and message texts, in the order that they were displayed
- Provide the exact text of Oracle Fail Safe messages (if any) that were written to the Windows application event log
- The problem severity, according to the following codes:

  1 = Program not usable. Critical impact on operations.
  2 = Program usable. Operations severely restricted.
  3 = Program usable with limited functions. Not critical to overall operations.
  4 = Program circumvented by customer. Minimal effect, if any, on operations.

- Your personal and company information:
  - Name
  - Company name
  - Company Oracle Support ID Number
  - Phone number
- In some cases, Oracle Support Services will request a trace file.

  See Section B.3 for information about using the trace function to log error output to a file.

B.2 Finding Your Version Information

You can find the version of software that you run in the Oracle Fail Safe Manager help menu. Select **Help,** then select **About Oracle Fail Safe Manager.** Version information for Oracle products that are integrated with Oracle Fail Safe is displayed in the output window for the Verify Cluster command.

B.3 Tracing Oracle Fail Safe Problems

Tracing is available to help you track, report, and examine errors that you receive in Oracle Fail Safe by dumping information about the errors to a log file.

You enable tracing for each node.

Follow these steps to enable tracing and set tracing flags on the cluster server nodes:

1. Run the Windows registry editor.

2. Select the following from the Registry tree:

   HKEY_LOCAL_MACHINE, then **SOFTWARE,** then **ORACLE,** then **FailSafe,** and finally, **Tracing**

3. From the Registry Editor menu bar, select **Edit**, then select **Add Value** to open the Add String dialog box.

4. In the **Value Name** field, enter an Oracle Services for MSCS value from Table B–1.

5. In the **Data Type** field, enter **REG_SZ.**

6. Click **OK** to open the String Editor dialog box.

7. In the **String** field, enter one or more of the Oracle Services for MSCS strings shown in Table B–1. Separate multiple entries with commas.

8. Repeat steps 3 through 7 to set additional Oracle Services for MSCS trace flags.

9. Using MSCS Cluster Administrator, shut down Oracle Services for MSCS.

10. In the Windows control panel, select **Services** and stop the Oracle Services for MSCS on each cluster node.

11. Using MSCS Cluster Administrator, restart Oracle Services for MSCS to begin tracing.

Repeat these steps on each cluster node to ensure that tracing is enabled across the cluster.

Table B–1   Trace Flags for Cluster Server Nodes

Value	String	Description
FSR_TRACE_OUTPUT	A path and file name	Specifies the path and file name for the file to which you want tracing information about the Oracle Fail Safe resource DLL to be written. For example: `C:\fsr_tracelog`
FSS_TRACE_OUTPUT	A path and file name	Specifies the path and file name for the file to which you want tracing information about the Oracle Services for MSCS to be written. For example: `C:\fss_tracelog`
FSS_TRACE_FLAGS	AGENT	Logs information related to Oracle Management Agent activity.

Table B–1   (Cont.) (Cont.) *Trace Flags for Cluster Server Nodes*

Value	String	Description
	ALL	Enables logging of all Oracle Fail Safe trace messages. Typically this is the most convenient flag to use. If this flag is enabled, then expect trace files to potentially grow large.
	COM	Logs activity related to the Microsoft DCOM interface.
	COMMAND_RESULT	Logs information related to spawned commands.
	COMMON	Logs information that would be common to all Oracle Fail Safe components, such as error logging or work item processing.
	CREATE_SA	Logs information related to the creation of standalone resources, such as a sample database.
	CR_RES	Logs information related to the creation of cluster resources.
	CLUSTER_MGR	Logs information related to the MSCS cluster interface.
	DB_RES	Logs information related to database access by the server or resource monitor DLL.
	DEL_RES	Logs information related to the deletion of a cluster resource.
	DELETE_SA	Logs information related to the deletion of a standalone resource, such as a sample database.
	HOME	Logs information related to the processing of Oracle homes.
	LOCAL_TRACE	Enables local tracing, which specifies that trace output for a given cluster node be written to the `FSS_TRACE_OUTPUT` file for that node. If this flag is not specified, then trace output for all cluster nodes is written to the `FSS_TRACE_OUTPUT` file on the node where Oracle Services for MSCS is running (the node where the Cluster Group resides).
		You must specify one or more additional `FSS_TRACE_FLAG` strings to specify the type of information that you want to have traced. If you specify only the `LOCAL_TRACE` string, then no trace output is produced.
	SQLNET	Generates detailed internal information related to the Oracle Net configuration performed by Oracle Services for MSCS. Information is logged whenever an operation is performed that requires a change to the Oracle Net configuration. This includes creating and deleting a sample database, or adding and removing a database from a group.
	VERIFY_CLUSTER	Logs information about the VERIFY CLUSTER operation.
	VERIFY_GR	Logs information about the `Verify Group` operation.
	VERIFY_SA	Logs information regarding verification of standalone resources.

Table B–1   (Cont.) (Cont.)   ***Trace Flags for Cluster Server Nodes***

Value	String	Description
	VERIFY_DB	Logs information about the `Verify Standalone Database` operation.
	XML	Logs activity related to the exchange of XML messages between Oracle Fail Safe components.
FSU_TRACE_OUTPUT	A path and file name	Specifies the path and file name for the file to which you want tracing information about the Oracle Services for MSCS surrogate to be written. For example: `C:\fsu_tracelog.log`
		Note: `FSU_TRACE_OUTPUT` file is always appended to and never overwritten. It means the file will continually grow until the file is deleted, or until the `FSU_TRACE_OUPUT` registry entry is deleted or redefined. Oracle recommends that the file be monitored to ensure that it does not grow too large and that tracing be enabled only for short periods of time.

B.4 Locating Trace and Alert Files

Oracle Fail Safe trace files must be directed to a private disk.

Database trace and alert files can be located on either a cluster disk or a private disk:

- If you use a cluster disk, then trace and alert files contain complete information about the operation. However, information about the node hosting the database is not recorded. The cluster disk used for these files must be one of the disks used for the archive log files or the database data files (where `Create Sample Database` places them, for example); otherwise, they will not be added to the group.

- If you use a private disk, then trace and alert files each contain node-specific information about the operation. However, you must view files from each cluster node at the same time to obtain complete chronological information if the database has failed over or been moved. Use a path name that is valid on each node so that data can be written to these files correctly. Files on private disks are never added to a group.

Glossary

24x365

24 hours a day, 365 days a year.

active/active configuration

A cluster configuration in which all cluster nodes perform work. If one node becomes unavailable, then one or more other nodes take over the workload of the node that is no longer available.

active/passive configuration

A cluster configuration in which one node usually stands idle in anticipation of a failover from another node.

availability

The measure of the ability of a system or resource to provide the desired service when required. Availability is measured in terms of the percentage of time the device is accessible out of the total time it is needed. Businesses that require uninterrupted computing services have an availability goal of 24x365.

bequeath protocol

A protocol that enables clients to retrieve information from an Oracle database without using the network listener. The bequeath protocol internally spawns a server thread for each client application. In a sense, it does the same operation that a remote network listener does for a database connection, but locally.

client application

The application that provides all user-oriented activities, such as character or graphical user display, screen control, data presentation, application flow, and other application-specific tasks.

cluster

A group of two or more independent computing systems that operate as a single virtual system.

cluster alias

A node-independent network name that identifies a cluster and is used for cluster-related system management.

cluster node

A Windows system that is a member of a cluster.

cluster resource

A resource that is configured and managed on a cluster node. *See also* **resource** and **standalone resource.**

data file

A file that contains the contents of logical database structures, such as tables and indexes. One or more data files form a logical unit of storage called a tablespace. A data file can be associated with only one tablespace, and only one database.

downtime

The measure of the inability of a system or resource to provide the desired service when required. Downtime is measured in terms of the percentage or amount of time the device is not accessible out of the total time it is needed.

failback

The process of intentionally returning a group of cluster resources to a preferred owner node from the failover node after the preferred owner node returns to operational status.

failback policy

See **group failback policy.**

failover

The process of taking cluster resources offline on one node and bringing them back online on another node. This process can either be planned (for upgrades and maintenance, for example) or unplanned (due to system or resource failure, for example).

failover node

The server node that takes over the workload of an unavailable node.

failover period

A user-specified time period in which the cluster software must continue to try to move cluster resources from one node to another before discontinuing the failover process and taking the resources offline. *See also* **group failover policy.**

failover policy

See **group failover policy** or **resource failover policy.**

failover threshold

The maximum number of times the cluster software must attempt to move resources from one node to another during the time period (failover period) that you specify. After reaching the specified failover threshold, the cluster software will stop the failover process and take the resources offline. *See also* **group failover policy.**

fail-safe resource

A resource that has been configured for high availability.

failure

The inability of a computing component to perform its function correctly.

generic service

A Windows service that is supported by the generic service resource DLL provided with Microsoft Cluster Server (MSCS). The generic service resource DLL is used to configure standard Windows services (such as IP addresses, physical disks, and some applications) as resources in a cluster.

group

A logical collection of cluster resources that forms a minimal unit of failover. In a failover situation, the group of resources is moved together to a failover node. A group resides on only one cluster node at a time. It is also referred to as service or application.

group failback policy

A user-specified plan that determines when and if cluster resources must fail back to the preferred owner node from the failover node.

group failover

The process of taking a group of cluster resources offline on one node and attempting to bring them back online on another node. This process can either be planned (for upgrades and maintenance, for example) or unplanned (due to system or resource failure, for example).

group failover policy

A user-specified plan that determines two parameters: the time period in which the cluster software must continue to move resources from one node to another (failover period), and the maximum number of times failover must occur during the failover period (failover threshold). *See also* **failover period** and **failover threshold.**

heartbeat connection

See **private interconnect.**

host name

A name that represents the specific IP address on a network. In Microsoft Cluster Server (MSCS), the host name is mapped to a network name resource. *See also* **network name.**

instance

A combination of System Global Area (SGA) and one or more Oracle database processes. When a database is started, Oracle allocates SGA and starts one or more Oracle processes. The memory and processes of an instance efficiently manage the associated database's data and serve the database users. Each instance has a unique Oracle System Identifier (SID), instance name, instance number, rollback segments, and thread ID.

internode network connection

See **private interconnect.**

IP address

The Internet Protocol (IP) address. An IP address takes the form n.n.n.n, for example, 138.2.134.113.

listener

A service that receives requests by clients and redirects them to the appropriate server.

Microsoft Cluster Services (MSCS)

Microsoft Corporation software that provides the capability to cluster individual nodes that run supported Windows operating systems. See *Oracle Fail Safe Release Notes* for a list of the supported operating system releases. *See also* **Microsoft Server Clusters** and **Microsoft Failover Clusters (MSFC)**.

Microsoft Server Clusters

In Windows Server 2003, Microsoft Cluster Services (MSCS) is known as Microsoft Server Clusters. See also **Microsoft Cluster Services (MSCS)** and **Microsoft Failover Clusters (MSFC)**.

Microsoft Failover Clusters (MSFC)

In Windows Server 2008, Microsoft Cluster Services (MSCS) is known as Microsoft Failover Clusters. See also **Microsoft Cluster Services (MSCS)** and **Microsoft Server Clusters**.

mission-critical application

A type of business function that is critical to the company and requires high availability.

net service name

Network information that describes the network and connection data of an Oracle database. More than one net service name can be defined for an Oracle database.

network name

The Microsoft Cluster Server (MSCS) term for a NetBIOS name, which translates into a specific IP address on a network. *See also* **host name.**

node

A computing system that is a member of a cluster.

planned group failover

The process of intentionally taking client applications and cluster resources offline on one node and bringing them back online on another node. For example, the *Oracle Fail Safe Installation Guide* describes how to perform a planned failover to perform a rolling upgrade (you fail over all resources to one cluster node as you sequentially upgrade software or hardware on another node). *See also* **unplanned group failover.**

possible owner node

A node capable of running a specified resource based on the following qualities:

- The resource DLL for the specified resource has been installed on the node.

- The resource has been configured to run on the node.

- You have not manually removed the node from the possible owner nodes list for the resource or the group containing the resource.

In a two-node cluster, both nodes must be possible owner nodes for all resources in a group if you want that group to be able to fail over.

possible owner nodes list

The set of all nodes on which the resource DLL for the specified resource has been installed and configured to run, less any nodes that you explicitly remove from the set.

preferred owner node

The node on which you want a group to reside when all cluster nodes that are possible owners are up and running. *See also* **failover node.**

primary nodes

In an active/passive configuration, the nodes that perform work. *See also* **active/passive configuration.**

private interconnect

A network connection that is dedicated to intracluster communication. The private interconnect is also referred to as a heartbeat connection, because it allows one node to detect the availability or unavailability of another node. The private interconnect is distinct from the public interconnect. *See also* **public interconnect.**

public interconnect

A network connection (such as a LAN or WAN) that connects clients to the cluster. *See also* **private interconnect.**

redundant components

Duplicate or extra computing components that safeguard the integrity of a computing system.

reliability

The ability of a computing system to operate without failing.

resource

A physical or logical component that is available to a computing system. For example, a resource can be a disk, a network IP address, an Oracle database, or a listener. *See also* **cluster resource** and **standalone resource.**

resource dependencies

Relationships between resources in a group that define the order in which the cluster software brings those resources online and offline.

resource failover policy

A policy that specifies whether a resource failure must result in a group failover.

resource restart policy

A policy that specifies whether the cluster software must attempt to restart a failed resource on its current node, and if so, how many attempts within a given time period must be made to restart it.

rolling upgrade

A software installation technique that allows a cluster system to continue to provide service while the software is being upgraded to the next release. This process is called a rolling upgrade because each node is upgraded and restarted in turn, until all cluster systems and client nodes have been upgraded. While each node is temporarily offline, another node takes over the workload of the node being upgraded.

sample database

An optional, preconfigured starter database that is provided with Oracle Fail Safe so you can try out the functions of Oracle Fail Safe before using them on your production database.

secondary node

In an active/passive configuration, a node that stands by to accept the work of a primary node in a failover. *See also* **active/passive configuration** and **primary nodes.**

service name entry

See **net service name.**

shared-nothing configuration

A cluster configuration in which all cluster nodes are cabled physically to the same disks, but only one node can access a given disk at a time for either read or write activity.

shared storage interconnect

An I/O connection on which the cluster disks are accessible from all nodes in a cluster.

silent mode

An installation method that lets you install software by supplying input to Oracle Universal Installer with a response file.

standalone resource

A resource that is not contained in a group. A standalone resource is hosted by a specific cluster node. *See also* **cluster resource** and **group.**

standby node

A node in an active/passive architecture that is ready to pick up application processing if a preferred owner node fails. *See also* **active/passive configuration** and **preferred owner node.**

subnet mask

A 32-bit value that indicates how many bits in an address are being used for the network ID.

transparent application failover

The ability of client applications to automatically reconnect to a database and resume work after a failover occurs.

unplanned group failover

A software-initiated failover process that is triggered automatically in response to a software or hardware failure. *See also* **planned group failover.**

virtual address

A network address at which resources in a group can be accessed, regardless of the cluster node hosting those resources. A virtual address on an MSCS cluster consists of a network name and associated IP address.

virtual directory

A name that maps to a physical directory specification. You specify a virtual directory to hide your file structure from users. If the physical directory changes, then the URL specified by users does not change. For example, if your virtual address is `Company`, and you have mapped the virtual directory `Sales` to `U:\SalesInfo\Webfiles`, then users will access sales information by entering the URL `http://Company/Sales`.

virtual server

A group with one or more virtual addresses.

Index