

Oracle® Fail Safe

Concepts and Administration Guide



Release 4.2.1 for Microsoft Windows

E66236-03

June 2018

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Fail Safe Concepts and Administration Guide, Release 4.2.1 for Microsoft Windows

E66236-03

Copyright © 2015, 2018, Oracle and/or its affiliates. All rights reserved.

Primary Author: Tanaya Bhattacharjee

Contributing Authors: Paul Mead

Contributors: Janelle Simmons

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	viii
Documentation Accessibility	viii
Related Documents	viii
Conventions	ix

1 Introduction to Oracle Fail Safe

1.1 Overview of Oracle Fail Safe	1-1
1.2 Benefits of Oracle Fail Safe	1-2
1.2.1 Highly Available Resources and Applications	1-2
1.2.2 Ease of Use	1-3
1.2.3 Product Accessibility	1-5
1.2.4 Ease of Integration with Applications	1-6
1.3 Overview of a Typical Oracle Fail Safe Configuration	1-6
1.4 Deploying Oracle Fail Safe Solutions	1-8

2 Cluster Concepts

2.1 Cluster Technology	2-1
2.1.1 About Clusters Providing High Availability	2-2
2.1.2 About System-Level Configuration	2-2
2.1.3 About Disk-Level Configuration	2-3
2.2 Resources, Groups, and High Availability	2-3
2.2.1 About Resources	2-3
2.2.2 About Groups	2-4
2.2.3 About Resource Dependencies	2-4
2.2.4 About Resource Types	2-5
2.3 Groups, Network Names, and Virtual Servers	2-6
2.4 Allocating IP Addresses for Network Names	2-7
2.5 Cluster Group and Cluster Alias	2-8
2.6 About Failover	2-9
2.6.1 Unplanned Failover	2-9

2.6.1.1	Unplanned Failover Due to a Resource Failure	2-10
2.6.1.2	Unplanned Failover Due to Node Failure or Unavailability	2-11
2.6.2	Planned Group Failover	2-12
2.6.3	Group and Resource Policies That Affect Failover	2-13
2.6.4	Detecting a Resource Failure	2-14
2.6.5	About Resource Restart Policy	2-14
2.6.6	About Resource Failover Policy	2-15
2.6.7	Resource Possible Owner Nodes List	2-15
2.6.8	About Group Failover Policy	2-18
2.6.9	Effect of Resource Restart Policy and Group Failover Policy on Failover	2-19
2.6.10	About Group Failover and the Preferred Owner Nodes List	2-20
2.6.11	Determining the Failover Node for a Group	2-21
2.7	About Failback	2-22
2.7.1	Group Failback and the Preferred Owner Nodes List	2-23
2.7.2	Client Reconnection After Failover	2-24

3 Designing an Oracle Fail Safe Solution

3.1	Customizing Your Configuration	3-1
3.1.1	Active/Passive Configuration	3-1
3.1.2	Active/Active Configuration	3-3
3.2	Integrating Clients and Applications	3-5

4 Management for High Availability

4.1	Configuring Resources for Failover	4-1
4.2	How Does Oracle Fail Safe Use the Wizard Input?	4-2
4.3	Managing Cluster Security	4-4
4.3.1	Oracle Fail Safe Services	4-4
4.3.1.1	Changing the Oracle Fail Safe Server Account	4-5
4.3.2	Oracle Fail Safe Manager	4-6
4.4	Discovering Standalone Resources	4-6
4.5	Renaming Resources	4-7
4.6	Using Oracle Fail Safe in a Multiple Oracle Homes Environment	4-7
4.7	Configurations Using Multiple Network Names	4-7
4.8	Adding a Node to an Existing Cluster	4-8

5 PowerShell Commands

5.1	Getting Started	5-1
5.2	About Common Parameters	5-3
5.3	Using the Oracle Fail Safe cmdlets in Scripts	5-3

5.4	FSCMD Equivalent cmdlets	5-4
5.5	Examples	5-5

6 Validating Actions

6.1	Validating Operations	6-1
6.1.1	Validating Cluster	6-1
6.1.2	Validating the Configuration of Oracle Resources	6-3
6.1.3	Validating Standalone Database	6-5
6.2	Dumping Cluster	6-7
6.3	Verifying Security Parameters	6-8
6.4	Additional Troubleshooting Information	6-9

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

7.1	Discovering Standalone Single-Instance Databases	7-1
7.2	Oracle Net Configuration for Standalone Single-Instance Databases	7-1
7.2.1	Listener Must Use IP Address for Local Host, Not Host Name	7-2
7.2.2	Shared Server Configuration and a Standalone Database	7-2
7.2.3	SID List Entries	7-3
7.2.4	Configuring Oracle Net on Nodes with Multiple Listeners	7-3
7.3	Adding Single-Instance Oracle Databases to a Group	7-4
7.3.1	Prerequisites	7-4
7.3.2	Configuration Steps	7-4
7.3.3	Configuration Data for Oracle Databases	7-5
7.3.3.1	Naming a Cluster Resource	7-6
7.3.3.2	Choosing Nodes	7-6
7.3.3.3	Selecting Network Names	7-8
7.3.3.4	Identifying Database Parameters	7-9
7.3.3.5	Database Authentication	7-11
7.3.3.6	Database Password	7-13
7.3.3.7	Database Resource Addition Confirmation	7-14
7.4	About Oracle Net Listener Resource Creation and Configuration	7-15
7.4.1	Client Connections to Highly Available Single-Instance Databases	7-16
7.4.2	Updating Oracle Net Configuration After Adding a Database to a Group	7-16
7.4.2.1	Updates That Oracle Fail Safe Makes to the tnsnames.ora File	7-16
7.4.2.2	Updates That Oracle Fail Safe Makes to the listener.ora File	7-17
7.4.2.3	Updates That Oracle Fail Safe Makes to the sqlnet.ora File	7-18
7.4.3	Using External Procedures with Databases Configured for High Availability	7-18
7.4.4	Support for Databases Using Shared Servers	7-18

7.4.4.1	Shared Servers for Oracle8 Databases	7-18
7.4.4.2	Shared Servers for Databases	7-20
7.5	Security Requirements for Single-Instance Databases	7-20
7.5.1	Synchronizing Password Files on Cluster Nodes	7-21
7.5.2	Changing the SYS Account Password	7-22
7.5.3	Upgrading a Fail-Safe Database with the Oracle Database Upgrade Assistant	7-23
7.6	Configuring a ctxsrv Server for High Availability	7-24
7.7	Integrating with Oracle Enterprise Manager	7-25
7.7.1	Adding an Oracle Management Agent to a Group	7-25
7.8	Optimizations for Single-Instance Database Recovery	7-27
7.9	Performing Administrative Tasks on a Single-Instance Fail-Safe Database	7-27
7.10	Database Homes	7-28
7.11	Configuring Transparent Application Failover (TAF)	7-29
7.12	Handling Errors and Troubleshooting Problems with Databases	7-30
7.12.1	Handling Errors That Occur When Bringing a Database Online	7-30
7.12.2	Troubleshooting Problems	7-30
7.12.3	Problems Adding a Database to a Group	7-31
7.12.4	Problems Placing a Group Online	7-31
7.12.5	Group Fails Over During Processing-Intensive Operations	7-33
7.12.6	Problems with Database Authentication	7-33
7.12.7	Problems with Virtual Server Configurations	7-33
7.12.7.1	Problems Configuring the Network Name	7-34
7.12.7.2	Archived listener.ora or tnsnames.ora Files	7-35
7.12.8	Security Access and Authentication Problems	7-35
7.12.9	Clients Cannot Access a Database	7-35
7.13	Starting and Stopping Pluggable Databases	7-35

A Contacting Oracle Support Services

A.1	Reporting a Problem	A-1
A.2	Finding the Version of Oracle Software	A-2
A.3	Viewing Error Information	A-2
A.4	Tracing Oracle Fail Safe Problems	A-2
A.5	Locating Trace and Alert Files	A-5

Glossary

Index

List of Tables

1-1	Comparing Oracle Fail Safe and Oracle Parallel Server	1-6
2-1	Example of Possible Owners for Resources in Group Test_Group	2-21
4-1	Permissions and Privileges	4-4
7-1	Steps for Configuring Databases	7-5
A-1	Trace Flags for Cluster Server Nodes	A-3

Preface

This guide describes how to use Oracle Fail Safe running on a Microsoft cluster system to configure Oracle databases.

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

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

Access to Oracle Support

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

Related 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 Microsoft Windows*.
- For installation, deinstallation, and upgrade instructions, see *Oracle Fail Safe Installation Guide for Microsoft Windows*.
- For online assistance, Oracle Fail Safe Manager provides help topics online. To access the online help topics, select **Help Topics** from the **Help** 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.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Introduction to 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 Windows Failover Clusters](#).

This chapter discusses the following topics:

- [Overview of Oracle Fail Safe](#) (page 1-1)
- [Benefits of Oracle Fail Safe](#) (page 1-2)
- [Overview of a Typical Oracle Fail Safe Configuration](#) (page 1-6)
- [Deploying Oracle Fail Safe Solutions](#) (page 1-8)

1.1 Overview of Oracle Fail Safe

Oracle Fail Safe is a user-friendly software that works with Microsoft Windows Failover Clusters 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 Microsoft Windows Failover Clusters 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 have been configured using the Microsoft Windows Failover Cluster Manager. 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 Fail Safe Server and Oracle Fail Safe Manager:

- Oracle Fail Safe Server works with the Microsoft Windows Failover Clusters 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 Fail Safe can attempt to restart a failed software resource so that a failover from one cluster node to another may not be required.

Note:

Oracle Fail Safe Server was referred to as Oracle Services for MSCS in previous releases.

- Oracle Fail Safe Manager provides a user-friendly interface and wizards that help configure and manage cluster resources, and troubleshooting tools that help 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:

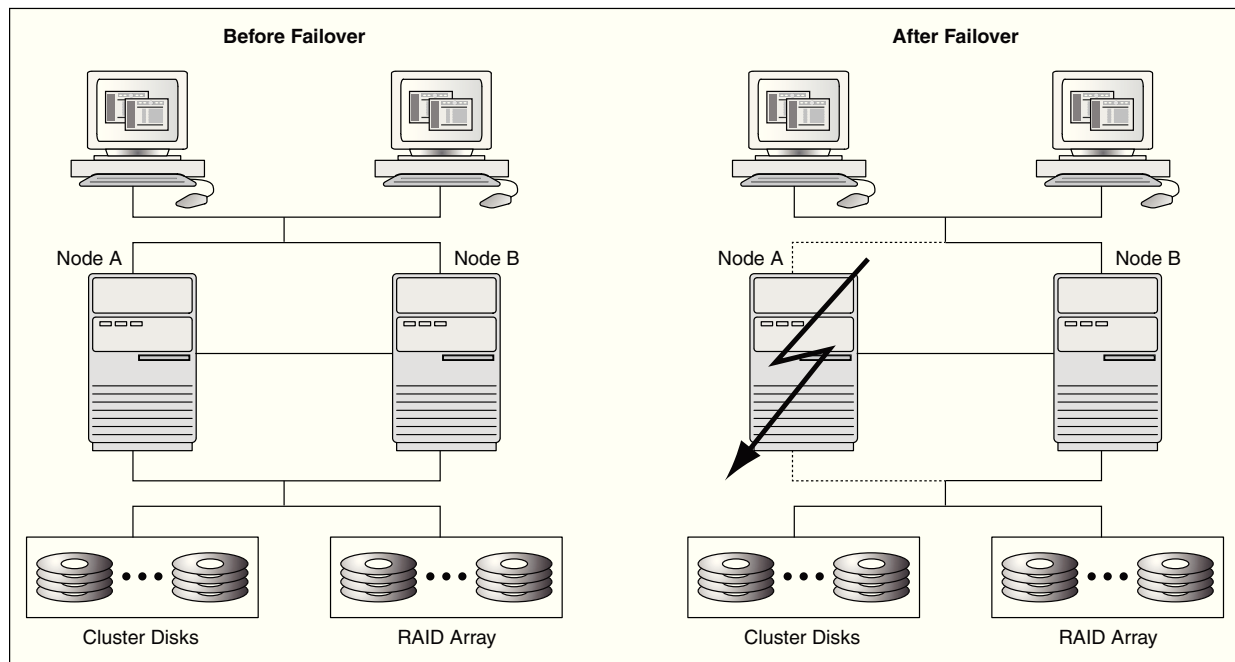
- [Highly Available Resources and Applications](#) (page 1-2)
- [Ease of Use](#) (page 1-3)
- [Product Accessibility](#) (page 1-5)
- [Ease of Integration with Applications](#) (page 1-6)

1.2.1 Highly Available Resources and Applications

Oracle Fail Safe works with Microsoft Windows Failover Clusters 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 [network name](#), without requiring any knowledge of the underlying cluster. If one node in the cluster becomes unavailable, then Microsoft Windows Failover Clusters moves the workload of the failed node (and client requests) to another node.

For example, the left side of [Figure 1-1](#) (page 1-3) 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, Microsoft Windows Failover Cluster ensures that each disk can be owned and accessed by only one node at a time. However, Microsoft Windows Failover Cluster also provides the ability to concurrently access a disk from multiple cluster nodes. This feature is known as Cluster Shared Volumes (CSV).

The right side of [Figure 1-1](#) (page 1-3) 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 can then access the database through Node B using the same network name 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, install Oracle Fail Safe either interactively or in silent mode. With the [silent mode](#) installation method, install software by supplying input to Oracle Universal Installer with a response file. Also, perform [rolling upgrades](#) 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 Microsoft Windows* 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, 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 help
- A command-line interface (PowerShell) for managing the cluster through batch programs or scripts

Figure 1-2 (page 1-4) shows an Oracle Fail Safe Manager window. The left pane displays a tree view showing the Microsoft Windows Failover Cluster Manager and the Oracle Fail Safe Manager. The Oracle Fail Safe Manager has a cluster, that includes the Oracle resources and a group (A group is sometimes referred to as a "service or application" or "clustered role"). The right pane displays the actions associated with Oracle Resources and Available Oracle Resources. The actions listed at the top of the **Actions** menu are relevant to the currently selected item in the tree view pane on the left. While the actions listed at the bottom of the **Actions** menu are related to the selected list item (if any) in the middle pane.

Figure 1-2 Oracle Fail Safe Manager

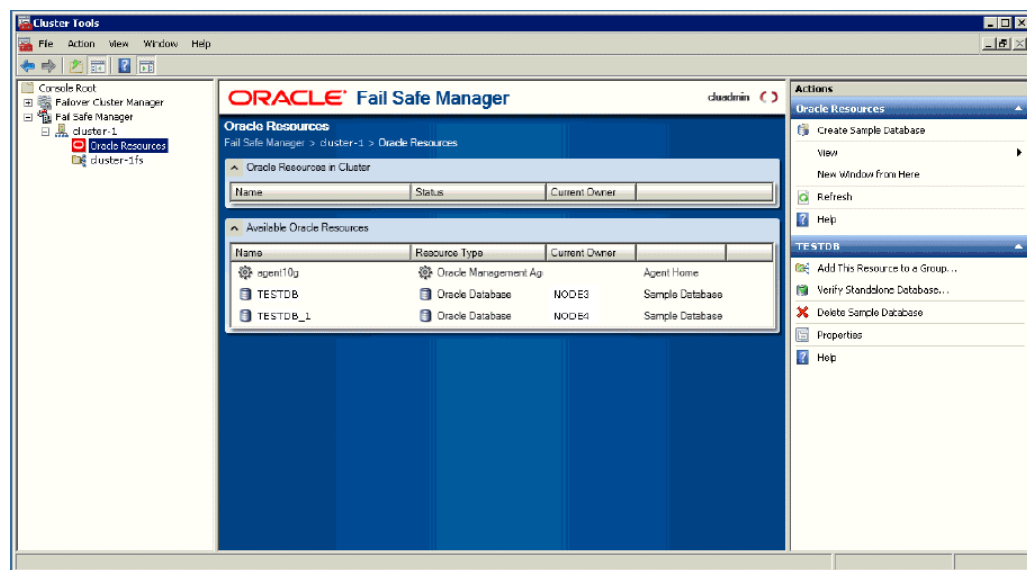
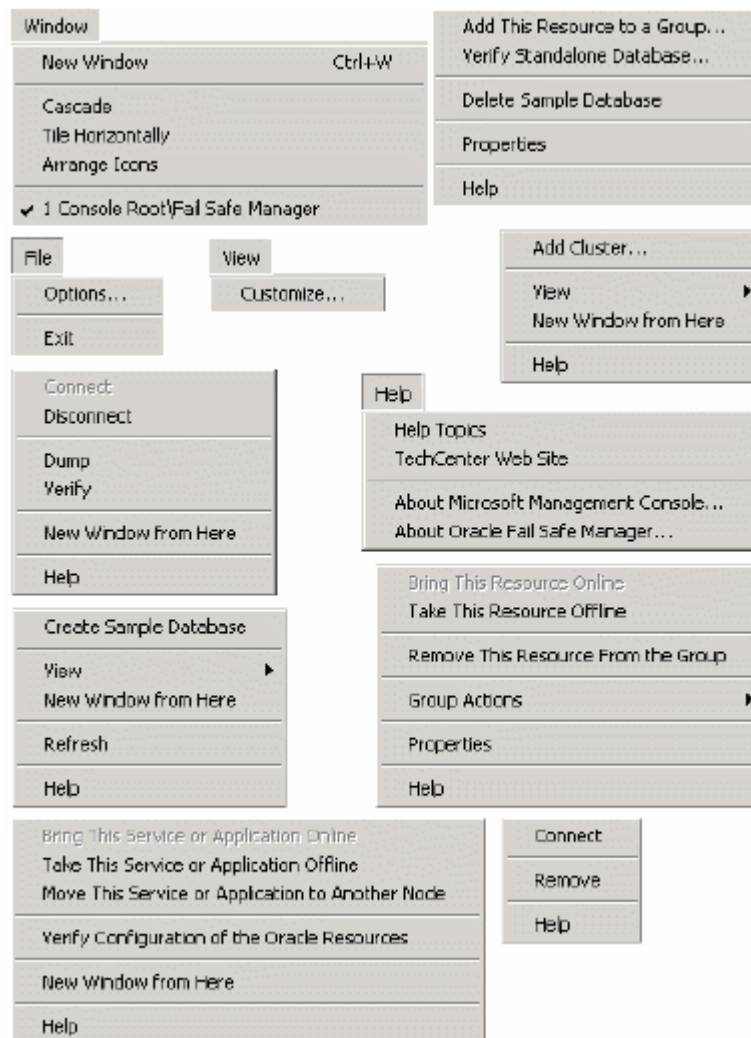


Figure 1-3 (page 1-5) shows the Oracle Fail Safe menus and the items within each menu.

Figure 1-3 Oracle Fail Safe Manager Menus and Contents



1.2.3 Product Accessibility

Oracle Fail Safe has two user interfaces: the PowerShell cmdlets 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 the following three panels:

- A navigation tree in the left panel
- The middle panel representing the selected tree view item
- The right panel showing actions for the selected tree view at the top of the **Actions** menu list, and actions for the selected list item (if any) chosen from the middle panel, at the bottom of the **Actions** menu list.

Wizard pages are displayed when the user selects an action that requires multiple steps, such as adding a resource to a group.

Refer to the Microsoft Management Console (MMC) help topic titled "Accessibility for MMC 3.0" for more information regarding the accessibility features of MMC.

1.2.4 Ease of Integration with Applications

To configure an existing application to access databases or other applications configured with Oracle Fail Safe, few or no changes are required. Because applications always access cluster resources at the same network name, 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 [Configuring Transparent Application Failover \(TAF\)](#) (page 7-29) for more information.

Table 1-1 Comparing Oracle Fail Safe and Oracle Parallel Server

Oracle Fail Safe	Oracle Parallel Server (OPS)
Designed for departmental and workgroup businesses running on Windows NT clusters.	Designed for enterprise-level and corporate database businesses that can benefit from multi-instance databases on Windows NT as well as other operating systems.
Provides high availability for single-instance Oracle databases and for nearly all currently shipping Oracle products that run on Windows NT. ¹	Provides high availability and scalability for multi-instance Oracle databases.
Available on Windows NT only.	Available on all major operating systems.
Currently supports two nodes in an Windows Failover Cluster; node limitation is required by the underlying Windows Failover Clusters.	Currently supports multiple nodes in a cluster; does not layer on Windows Failover Clusters or any other cluster software.
Offers the manageability of a single-instance Oracle database.	Requires administrative skills to tune and manage multi-instance Oracle databases and distributed configurations.
Uses the shared nothing method of disk access.	Uses the shared data method of disk access.

¹ Oracle Developer Forms and Reports, Oracle Application Server, Oracle Management Agent, and Microsoft Internet Information Server (IIS), and all applications that can be configured as a Windows NT service (for example, Oracle Applications Release 11 and Oracle WebDB).

1.3 Overview of a Typical Oracle Fail Safe Configuration

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

Most clusters are configured similarly, differing only in choice of storage interconnect (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](#) (Fibre Channel, or SAN). All [data files](#), log files, and other files that must fail over from one node to another are located on these cluster disks.

 **Note:**

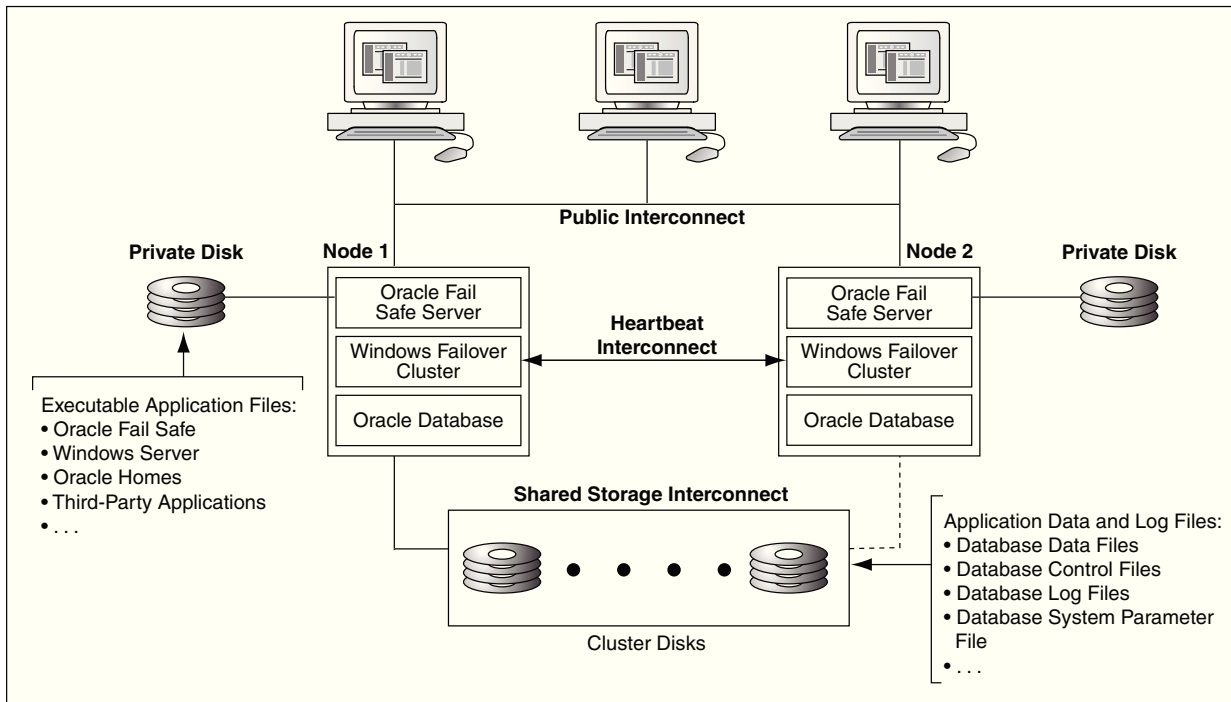
See the documentation for your cluster hardware for information about using redundant hardware, such as RAID, to further ensure high availability.

- Additional [redundant components](#) (UPS, network cards, disk controllers, and so on).
- Software (installed on each node)
 - Microsoft Windows
 - Oracle Fail Safe Server
 - Oracle Fail Safe Manager (installed on one or more cluster nodes, one or more client workstations, or both)
 - Oracle database

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

[Figure 1-4](#) (page 1-8) 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 Microsoft Windows Failover Clusters 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 failovers** (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 failovers** (those which you intentionally start, such as when you upgrade software on the cluster). To 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. [Designing an Oracle Fail Safe Solution](#) (page 3-1) 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 Windows Failover Clusters 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 Windows Failover Clusters](#) 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 Microsoft Windows Failover Clusters concepts.

This chapter discusses the following topics:

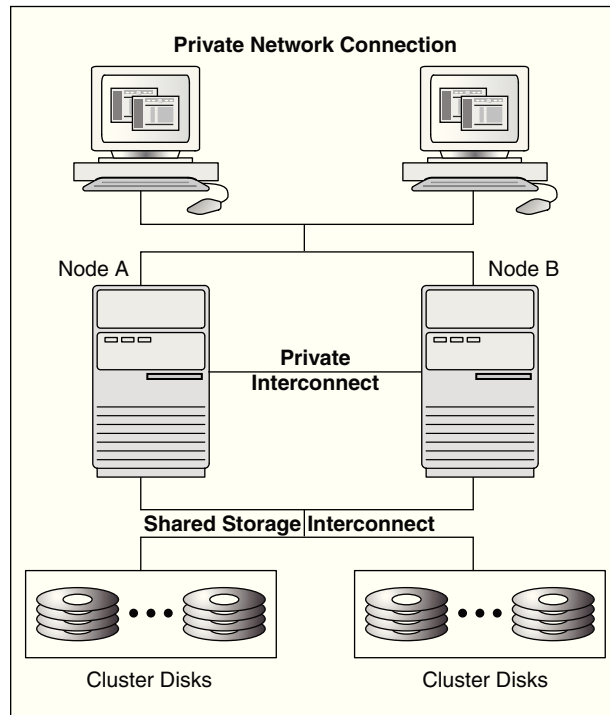
- [Cluster Technology](#) (page 2-1)
- [Resources, Groups, and High Availability](#) (page 2-3)
- [Groups, Network Names, and Virtual Servers](#) (page 2-6)
- [Allocating IP Addresses for Network Names](#) (page 2-7)
- [Cluster Group and Cluster Alias](#) (page 2-8)
- [About Failover](#) (page 2-9)
- [About Failback](#) (page 2-22)

2.1 Cluster Technology

The Windows systems that are members of a cluster are called [cluster nodes](#). 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](#) (page 2-2) shows the network connections in a two-node Microsoft cluster configuration.

Figure 2-1 Microsoft Cluster System

2.1.1 About Clusters Providing 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 Microsoft Windows Failover Clusters software, 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.

2.1.2 About 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 [Designing an Oracle Fail Safe Solution](#) (page 3-1) for information about these configurations.

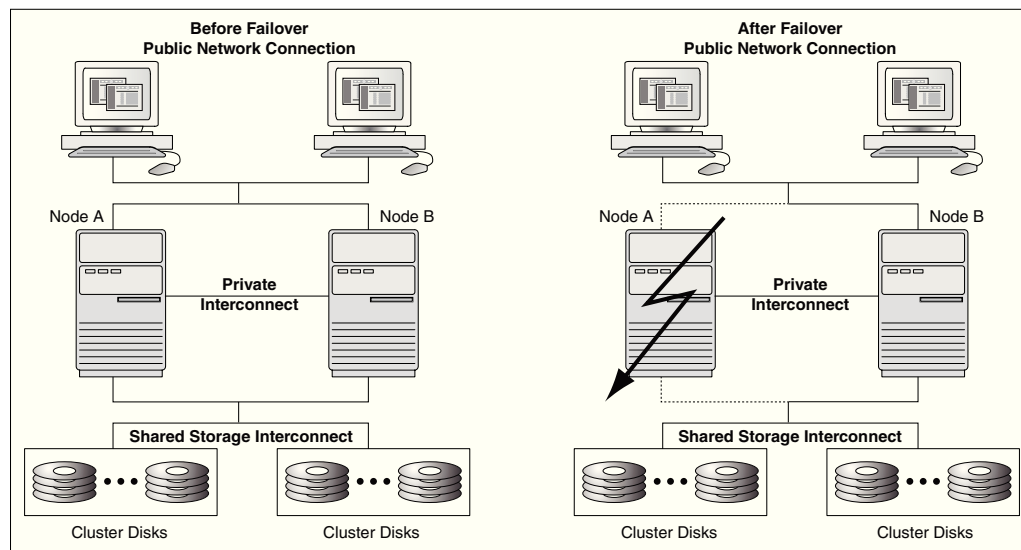
2.1.3 About Disk-Level Configuration

When a Windows Failover 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](#) (page 2-3) 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



Cluster Shared Volumes allow multiple nodes to concurrently access the same disk.

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. A group is sometimes referred to as a "service or application" or "clustered role". The following sections describe resources and groups, and how they are configured for high availability.

2.2.1 About 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 About Groups

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" or "clustered role". 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:

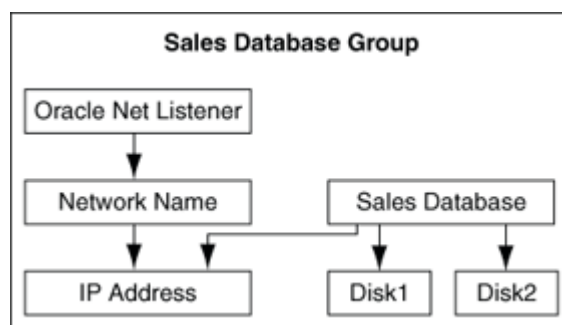
- All disks used by the Oracle Database
- An Oracle Database [instance](#)
- One or more network names, each one consisting of:
 - An IP address
 - A [network name](#)
- An Oracle Net network [listener](#) that listens for connection requests to databases in the group

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.

Microsoft Windows Failover Cluster Manager helps to create groups and add the resources needed to run applications.

2.2.3 About Resource Dependencies

[Figure 2-3](#) (page 2-5) 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; Windows Failover Cluster 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 About 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 Microsoft Windows Failover Clusters resource DLLs as well as custom Oracle resource DLLs. The same resource DLL may support several different resource types.

Microsoft Windows Failover Clusters 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 Microsoft Windows Failover Clusters.)

Oracle Fail Safe uses many of the Windows Failover Cluster 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. Windows Failover Cluster 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 a DLL file to enable Microsoft Windows Failover Clusters to communicate with and monitor Oracle Database resources. `FsResOdb.dll` provides functions that enable Microsoft Windows Failover Clusters to bring an Oracle Database and its listener online or offline and check its status through Is Alive polling.

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 Microsoft Windows Failover Clusters about Oracle cluster resources, Oracle recommends that you use Oracle Fail Safe Manager (or the Oracle Fail Safe PowerShell cmdlets) to configure and administer Oracle Databases and applications.

2.3 Groups, Network Names, and Virtual Servers

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

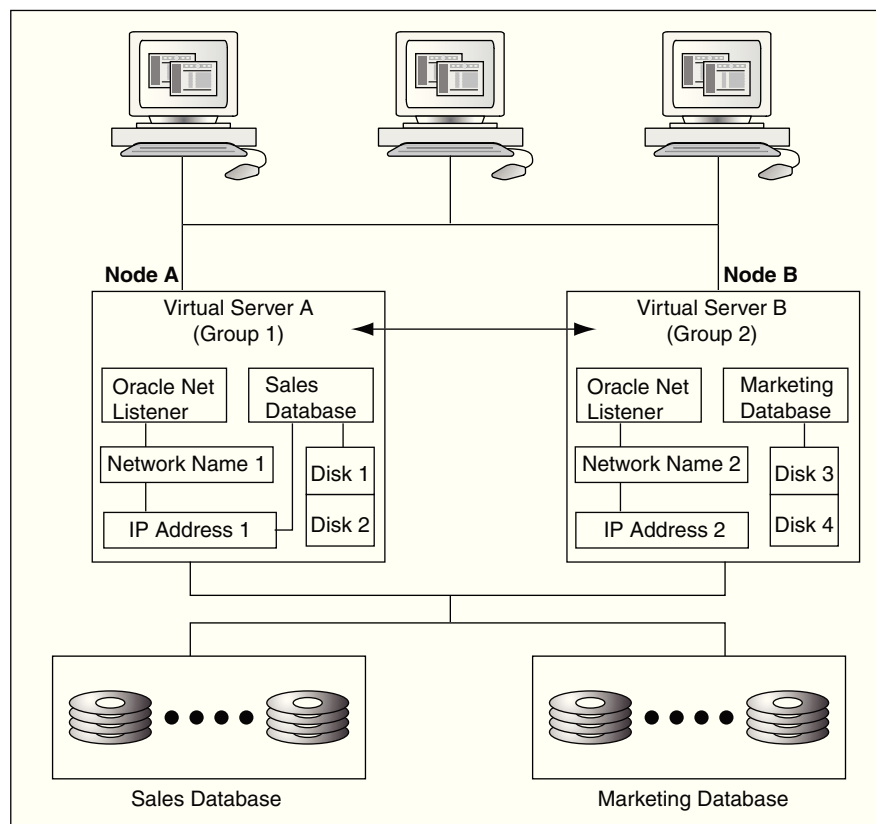
Because groups move from an unavailable node to an available node during a failure, a client cannot connect to an application that uses an address that is identified with only one node. To identify a network name for a group in Oracle Fail Safe Manager, add a unique network name and IP address to a group.

To add a network name resource to a group, use the Microsoft Windows Failover Cluster Manager.

Once you add a network name to a group, the group becomes a **virtual server**. Although at least one network name is required for each group for client access, you can assign multiple network names to a group. You may assign multiple network names 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 network name of the group. To the client, the virtual server is the interface to the cluster resources and looks like a physical node.

Figure 2-4 (page 2-7) 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 network name 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 2-4 Accessing Cluster Resources Through a Virtual Server

2.4 Allocating IP Addresses for Network Names

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 [Cluster Group and Cluster Alias](#) (page 2-8))
- One IP address for each group

For example, the configuration in [Figure 2-4](#) (page 2-7) 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 network names for a group. See [Configurations Using Multiple Network Names](#) (page 4-7) for details.

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. Microsoft Windows Failover Clusters creates a group called the Cluster Group, and the cluster alias is the network name of this group. Oracle Fail Safe is a resource in the Cluster Group, making it highly available and ensuring that Oracle Fail Safe 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 Microsoft Windows Failover Clusters using the cluster alias.

To add a cluster to Oracle Fail Safe Manager cluster list, perform the following steps:

1. Select **Add Cluster** from the **Actions** menu on the right pane of Oracle Fail Safe Manager page.
2. Enter the network name of the cluster in the **Cluster Alias** field as shown in [Figure 2-5](#) (page 2-9).
3. Optionally, you can select **Connect using different credentials**.

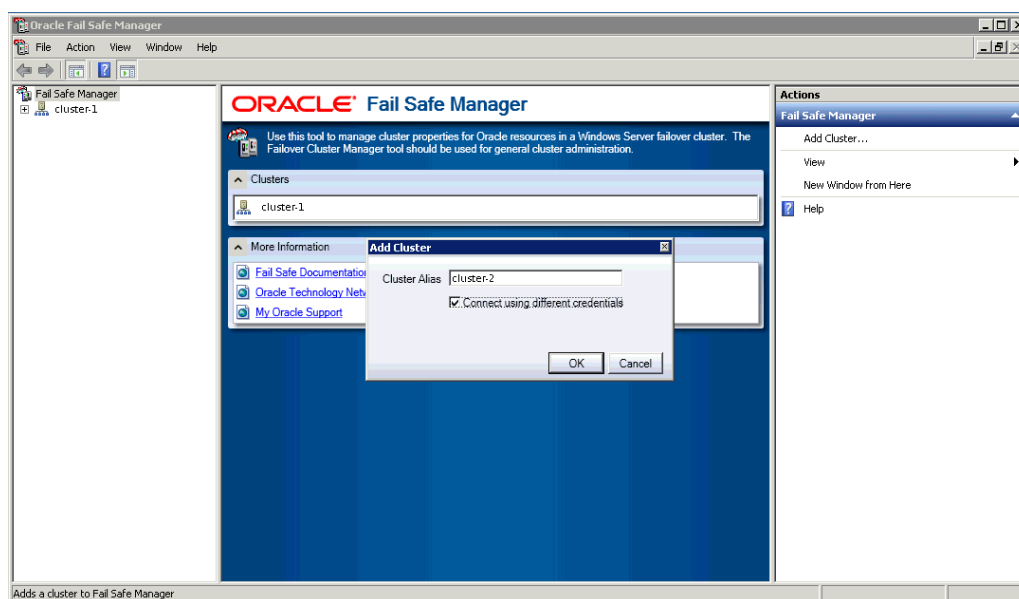
By default, Oracle Fail Safe Manager connects to the cluster using the credentials for your current Windows login session. To use another user's credentials, perform the following steps:

1. Select **Connect using different credentials** option.
2. This opens a Windows Security Cluster Credentials dialog box that allows you to enter new credentials for administering the cluster. Enter the user name and password in the fields provided to continue.
3. To save the credentials, select the **Remember my credentials** option and click **OK**. The credentials will be saved in the Windows credentials cache so that when you connect to the cluster, Oracle Fail Safe Manager will check to see if there are any saved credentials for that cluster and use the same to connect to the cluster.

See Also:

Oracle Fail Safe Tutorial for step-by-step instructions on adding a cluster to the Oracle Fail Safe Manager cluster list and connecting to a cluster

Figure 2-5 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 network names of the group that contains that resource.

2.6 About 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. Windows Failover Cluster 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 failovers**, 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 [Detecting a Resource Failure](#) (page 2-14) 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 2-6](#) (page 2-11) 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 [Resource Possible Owner Nodes List](#) (page 2-15) for more information about the resource possible owner nodes list, and see [About Group Failover and the Preferred Owner Nodes List](#) (page 2-20) 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 [About Group Failover Policy](#) (page 2-18) for more information about the group failover policy.
5. The [fallback 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 [About Fallback](#) (page 2-22) for information about fallback.

In [Figure 2-6](#) (page 2-11), Virtual Server A is failing over to Node B due to a failure of one of the resources in Group 1.

Figure 2-6 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 [Resource Possible Owner Nodes List](#) (page 2-15) for more information about the resource possible owner nodes list, and see [About Group Failover and the Preferred Owner Nodes List](#) (page 2-20) 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 [About Group Failover Policy](#) (page 2-18) 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 [About Failback](#) (page 2-22) for information about failback.

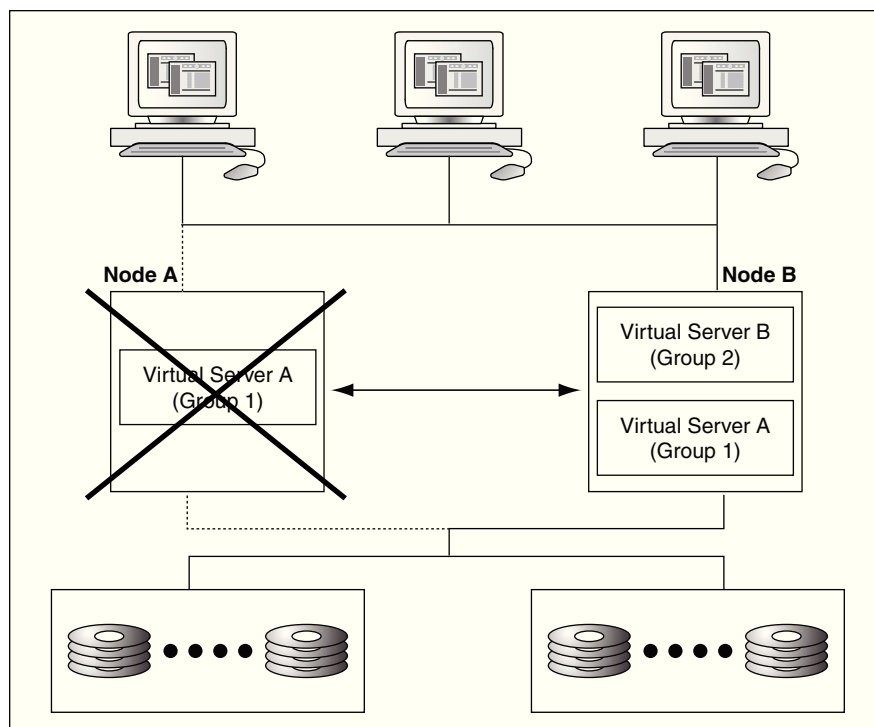
[Figure 2-7](#) (page 2-12) shows Group 1 failing over when Node A fails. Client applications (connected to the failed server) must reconnect to the server after failover occurs. If the application is performing updates to an Oracle Database and

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 [Unplanned Failover Due to a Resource Failure](#) (page 2-10). 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 2-7 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, planned failover helps to balance the load across the nodes in the cluster. In other words, use planned failover to move a group from one node to another. In fact, to implement a planned failover, 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 Fail Safe works with Microsoft Windows Failover Clusters 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 [Configuring Transparent Application Failover \(TAF\)](#) (page 7-29) for information about transparent application failover). Then, take your time to perform 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 an Oracle resource to a group using Oracle Fail Safe Manager and Microsoft Windows Failover Cluster Manager. However, you can reset the values in these policies with the Oracle Resource **Properties** page of Oracle Fail Safe Manager and the Group Failover **Properties** page of Microsoft Windows Failover Cluster Manager. You can set values for the group failback policy at group creation time or later, using the Group Failover **Properties** page and selecting the **Allow Failback** option.

Figure 2–8 shows the page for setting group failover policies. To access this page, select the group of interest in Microsoft Windows Failover Cluster Manager tree view, then select the **Properties** action from the **Actions** menu on the right pane of the screen. Then select the **Failover** tab in the Properties page.

Figure 2–9 shows the page for setting Oracle resource policies. To access this page, either select the **Properties** action menu from the right pane of the screen or select **Properties** from the **Action** item in the menu bar. Then select the **Policies** tab in the Oracle resource Properties page.

Figure 2-8 Group Failover Property Page

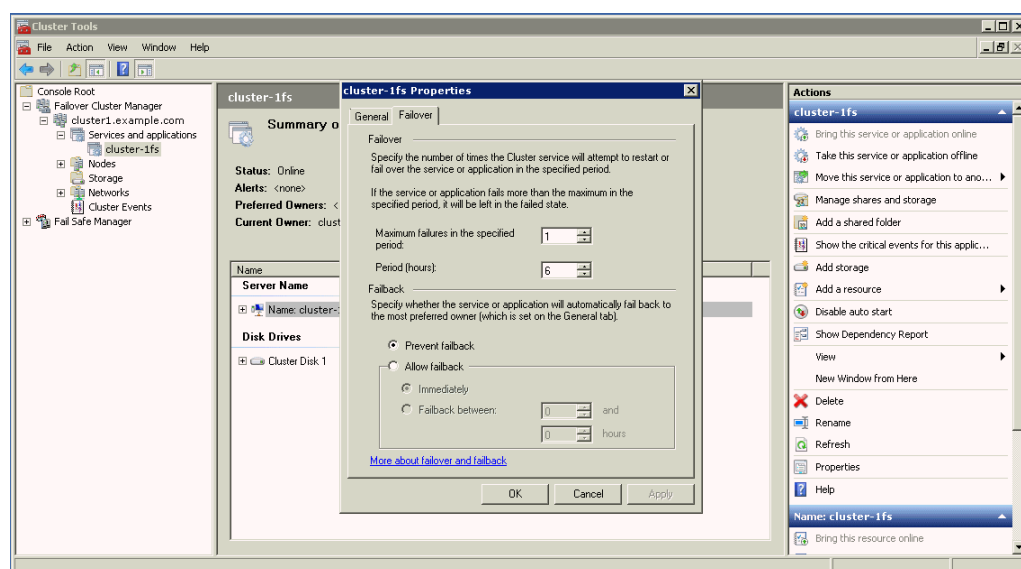
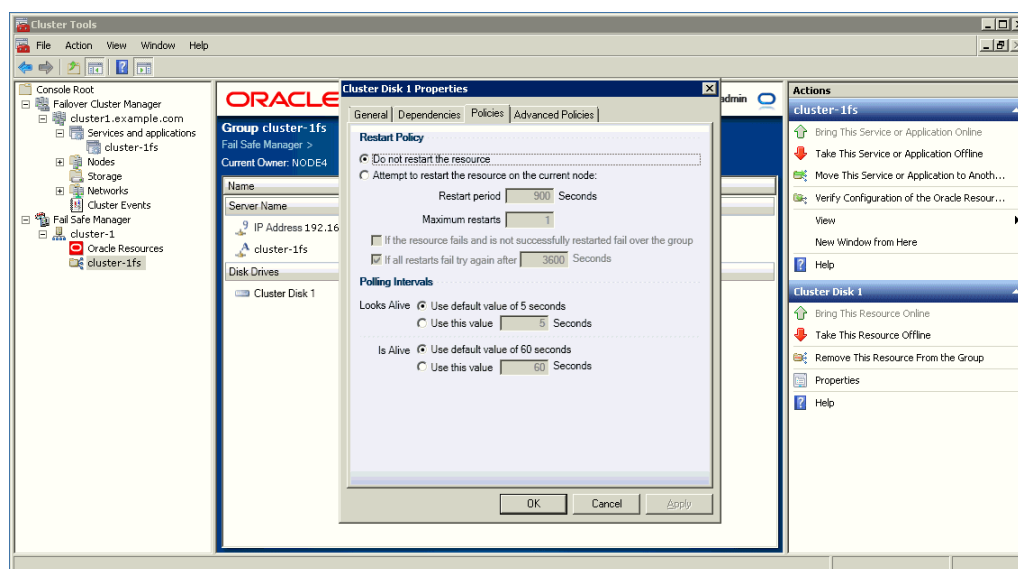


Figure 2-9 Resource Policies Property Page



2.6.4 Detecting a Resource Failure

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. Either use the default value for the resource type or specify a number (in seconds). 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. Either use the default value for the resource type or specify a number (in seconds). This check is less thorough, but also uses fewer system resources, than the check performed during an Is Alive interval.

2.6.5 About 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 2-9](#) (page 2-14):

- 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 About 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 [About Group Failover and the Preferred Owner Nodes List](#) (page 2-20) and [Group Failback and the Preferred Owner Nodes List](#) (page 2-23) 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. Some of the options are:

- The **Advanced Policies** property tab of the resources displays the possible owner nodes.
- To change the node, use the Microsoft Windows Failover Cluster Manager.
- Run the **Validate** group action after changing the possible owner nodes.
- [Figure 2-10](#) (page 2-16), [Figure 2-11](#) (page 2-17), and [Figure 2-10](#) (page 2-16) are examples of the Group property pages in Microsoft Windows Failover Cluster Manager.

Note:

If you select the **Validate** group action, 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 select the **Validate** group command for each group for which the new node is listed as a possible owner. [Validating the Configuration of Oracle Resources](#) (page 6-3) describes the **Validate** group action.

Figure 2-10 Preferred Owner Nodes Property Page

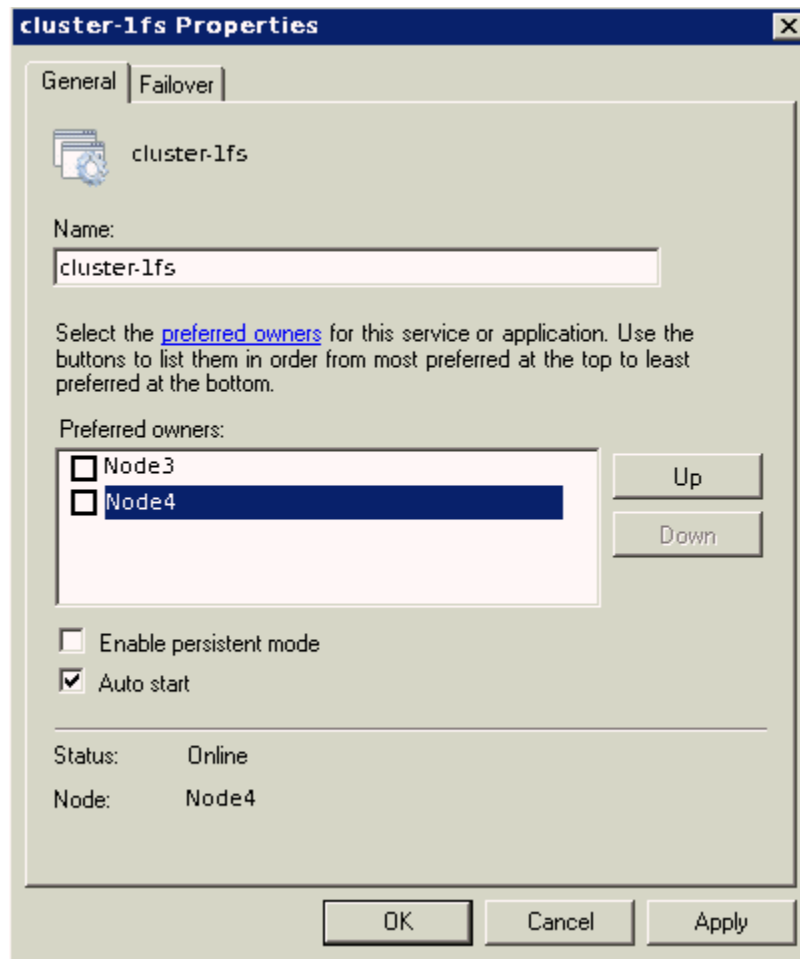


Figure 2-11 Preferred Owner Nodes Property Page

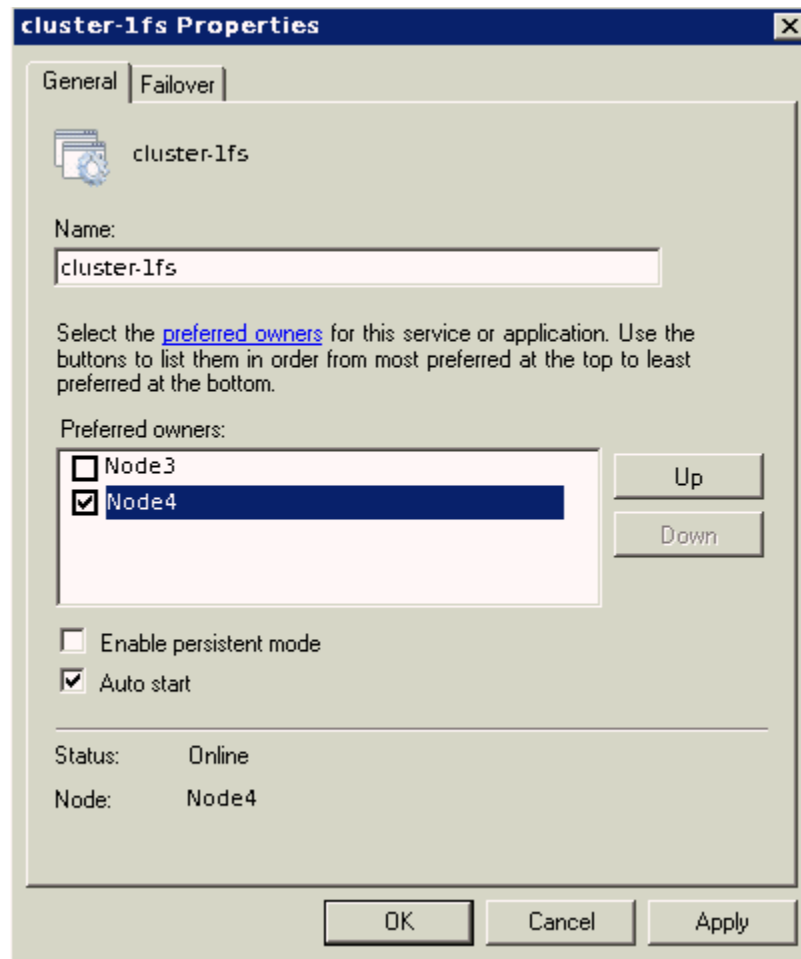
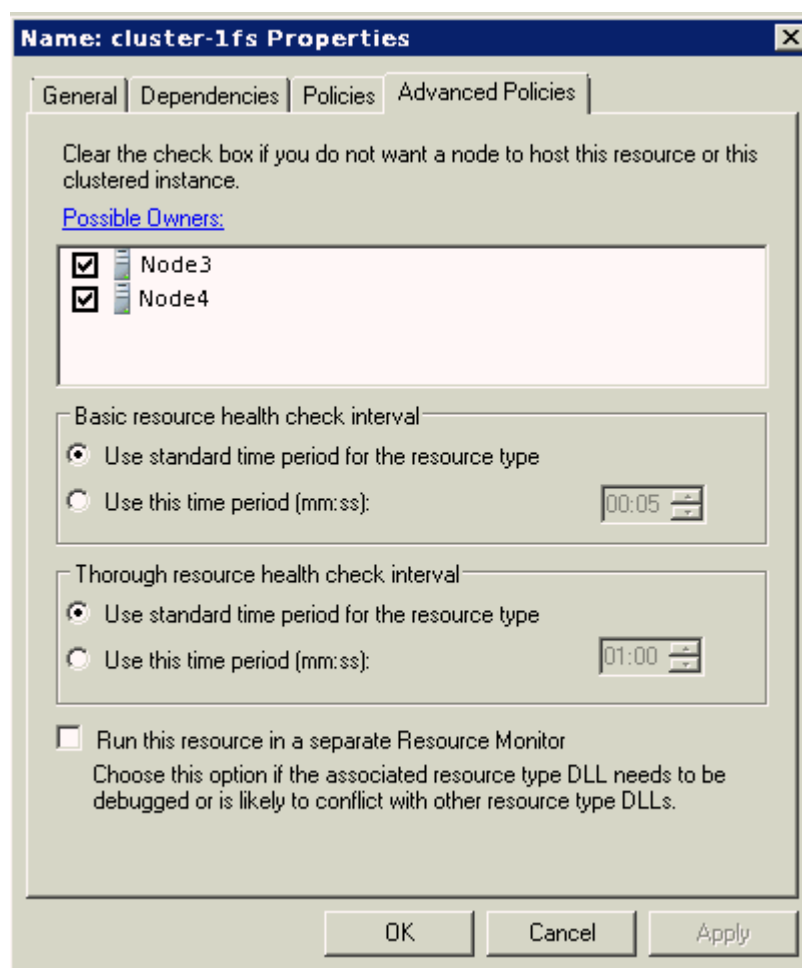


Figure 2-12 Possible Owner Nodes Property Page



2.6.8 About 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.

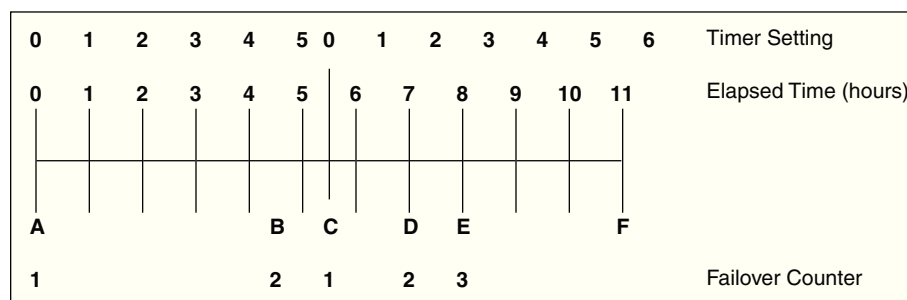
When the first failover occurs, a timer to measure the failover period is set to 0 and a counter to measure the number of failovers is set to 1. The timer is not reset to 0 when the failover period expires. Instead, the timer is reset to 0 when the first failover occurs after the failover period has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

For example, assume again that the failover period is 5 hours and the failover threshold is 3. As shown in [Figure 2-13](#) (page 2-19), 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 2-13 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 About 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.) 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. Node 2 has been omitted 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. [Determining the Failover Node for a Group](#) (page 2-21)

describes in more detail how the cluster software determines the node to which a group fails over.

See [Group Failback and the Preferred Owner Nodes List](#) (page 2-23) 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, 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 [Resource Possible Owner Nodes List](#) (page 2-15).)
- List of preferred owner nodes for the group containing the resources (See [About Group Failover and the Preferred Owner Nodes List](#) (page 2-20).)

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](#) (page 2-21).

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

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

By reviewing [Table 2-1](#) (page 2-21), 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

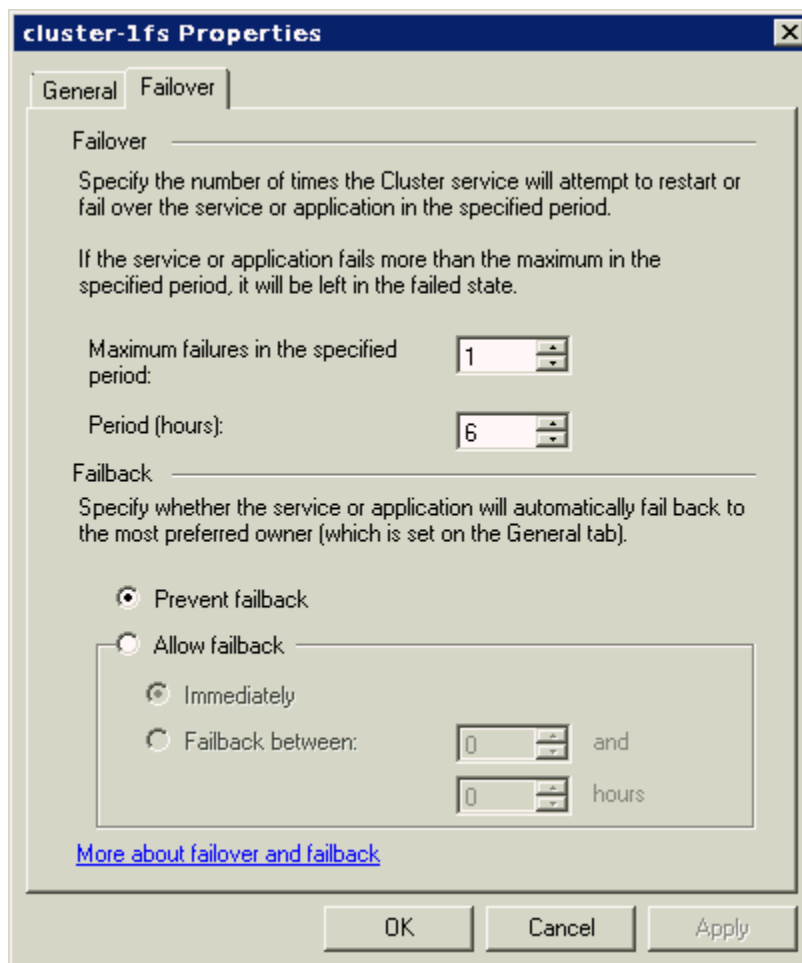
Because Node 3 has failed, the cluster software fails over Test_Group to Node 4. If both Node 3 and Node 4 are not available, then the cluster software fails over Test_Group to Node 1. If Nodes 1, 3, and 4 are not available, then the group fails over to Node 2. However, because Node 2 is not a possible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

2.7 About 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. The following figure shows the property page for setting the failback policy for a group.

Figure 2-14 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 [About Group Failover and the Preferred Owner Nodes List](#) (page 2-20) 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 network name) 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 [Configuring Transparent Application Failover \(TAF\)](#) (page 7-29) 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](#) (page 3-1)
- [Integrating Clients and Applications](#) (page 3-5)

3.1 Customizing Your Configuration

You can deploy highly available solutions using the following configurations:

- [Active/Passive Configuration](#) (page 3-1)
- [Active/Active Configuration](#) (page 3-3)

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 Fail Safe software automatically runs as needed on one or more cluster nodes to ensure proper configuration and failover.

[Figure 1-4](#) (page 1-8) 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](#) (page 3-2) shows a two-node configuration with Oracle Database running on Node 1, and with Node 2 as a standby server. 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

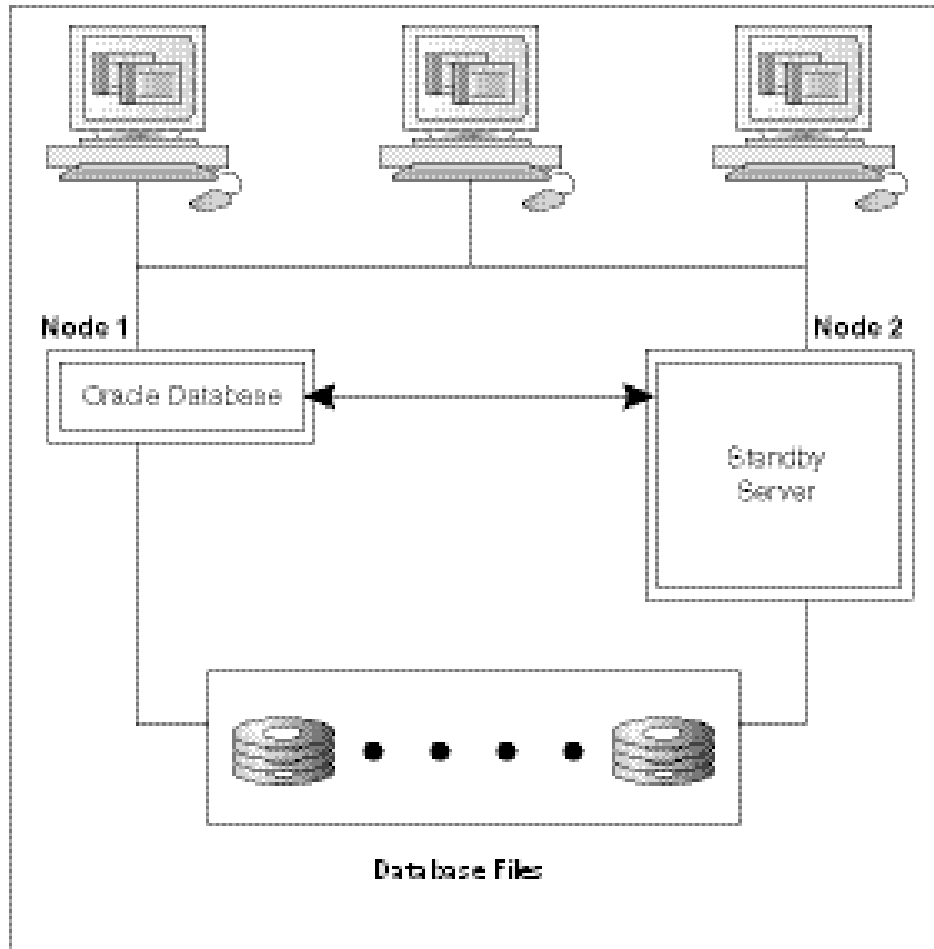
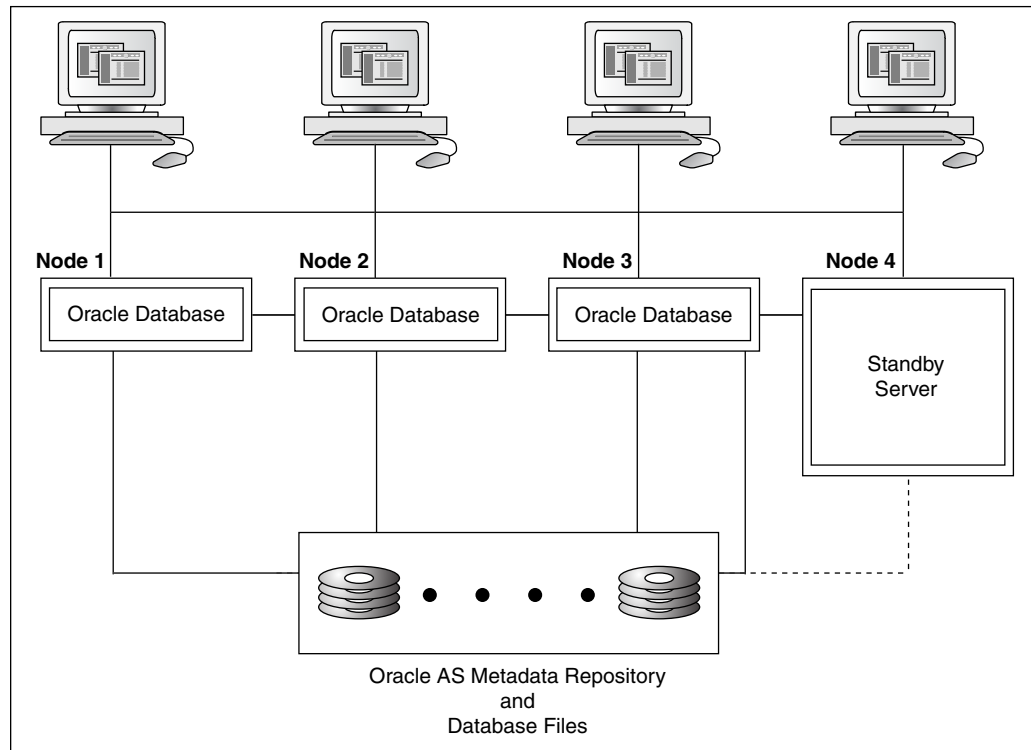


Figure 3-2 (page 3-3) shows a four-node configuration with Oracle Database running on Node 1, Node 2, and Node 3. Node 4 is the [standby node](#). Currently, nothing is running on Node 4. In the event of a failover, Node 4 takes over the failover workload.

Figure 3-2 Active/Passive (Standby) Four-Node Configuration

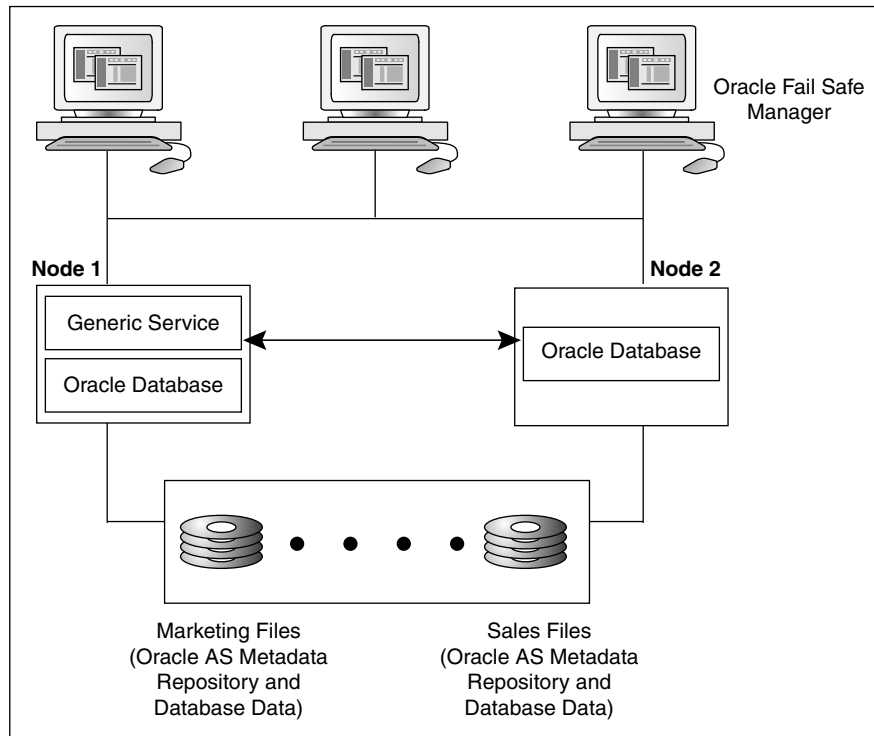


An active/passive configuration is the fastest failover configuration, because the passive standby node has no workload of its own.

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 enables division of the workload to best meet your business needs.

[Figure 3-3](#) (page 3-4) shows a two-node active/active configuration with an Oracle Database running on both cluster nodes. In addition, a generic service is running on Node 1. In [Figure 3-3](#) (page 3-4), 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 throughput) 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 applications](#)
- Use different database parameter files on each node so that fewer system resources are used after a failover
- Configure each component (Oracle Database 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 PowerShell cmdlets described in [PowerShell Commands](#) (page 5-1)) 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.

[Configuring Single-Instance Databases for High Availability and Disaster Tolerance](#) (page 7-1) contains a section specific to how you can integrate clients and applications. [Configuring Single-Instance Databases for High Availability and Disaster Tolerance](#) (page 7-1) 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:

- [Configuring Resources for Failover](#) (page 4-1)
- [How Does Oracle Fail Safe Use the Wizard Input?](#) (page 4-2)
- [Managing Cluster Security](#) (page 4-4)
- [Discovering Standalone Resources](#) (page 4-6)
- [Renaming Resources](#) (page 4-7)
- [Using Oracle Fail Safe in a Multiple Oracle Homes Environment](#) (page 4-7)
- [Configurations Using Multiple Network Names](#) (page 4-7)
- [Adding a Node to an Existing Cluster](#) (page 4-8)

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 [Configuring Single-Instance Databases for High Availability and Disaster Tolerance](#) (page 7-1) in this manual, *Oracle Fail Safe Tutorial for Microsoft Windows*, and online help.

4.1 Configuring Resources for Failover

Using Oracle Fail Safe Manager wizards, you can easily configure failover automatically and with minimal work by a system 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, all of these tasks must be performed 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 for Microsoft Windows*.)
2. Start Oracle Fail Safe Manager.
3. Verify the cluster.
4. Use the **Validate** action to validate the standalone Oracle Database you are adding.
5. Add resources to the group.

6. Verify the group.
7. Update any Oracle Net file (such as the `tnsnames.ora` file) on client systems.



Note:

Depending on the type of resource being configured, 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 Fail Safe (which in turn interacts with Microsoft Windows Failover Clusters) 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:

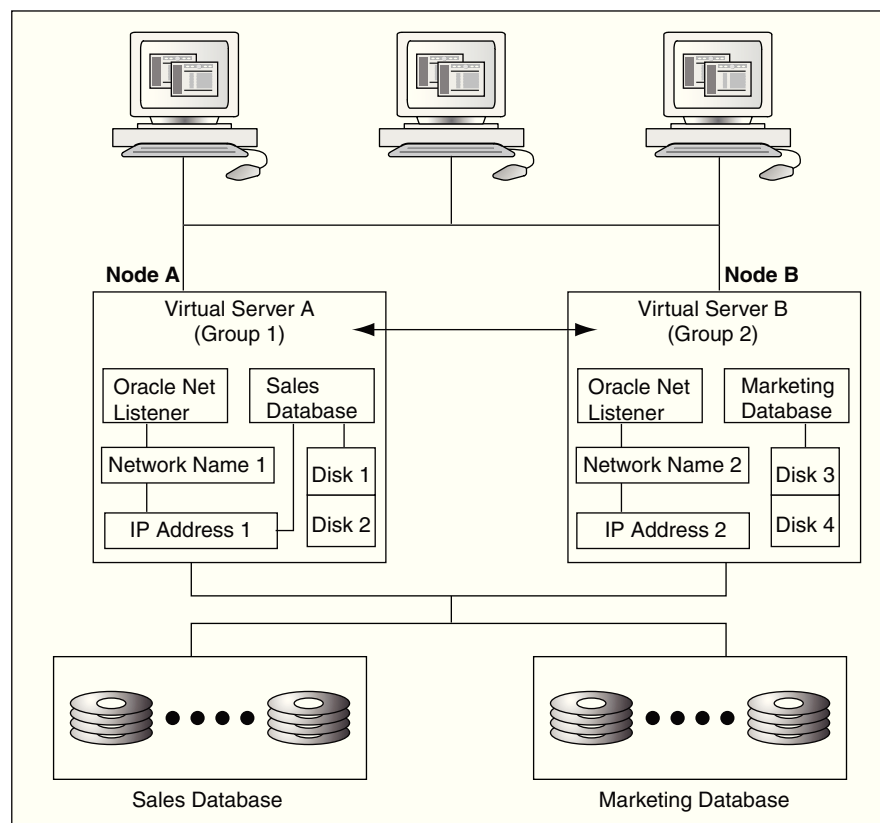
1. Configures access to the database using a network name:
 - a. Configures Oracle Net to use the network name or names associated with the database on all nodes listed in the possible owner nodes list for the database. (On a two-node cluster, this is both cluster nodes. On clusters that consist of more than two nodes, 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. If the group has a preferred owner node, then the group is moved to that node. 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 (page 4-3) 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** (network name), database instance, 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 network name 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](#) (page 4-4) provides a quick reference to the privileges required to use the services 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 Fail Safe	Domain user account that has Administrator privileges on all cluster nodes	Oracle Fail Safe Services (page 4-4)
Oracle Fail Safe Manager	Domain user account that has Administrator privileges on all cluster nodes	Oracle Fail Safe Manager (page 4-6)
Oracle Database	Database administrator account with SYSDBA privileges	Security Requirements for Single-Instance Databases (page 7-20)

4.3.1 Oracle Fail Safe Services

Oracle Fail Safe accesses database resources from two different Windows services: the Cluster Service service and the Oracle Fail Safe 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 Oracle Fail Safe service processes requests from the Oracle Fail Safe clients, such as Oracle Fail Safe Manager or PowerShell cmdlets, that are related to Oracle cluster resources.

Each of these services executes in the context of the Log On As user specified for the particular service. The Oracle Fail Safe service executes under the account provided to the Oracle Fail Safe Security Setup tool during the installation of Oracle Fail Safe.

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 Oracle Fail Safe service, that means that authentication is done using the Log On As account specified for the Oracle Fail Safe service. The Oracle Fail Safe database resource DLL impersonates the Oracle Fail Safe account when connecting to the database. Even though the Cluster service is executing as Local System, database access authentication is done using the Oracle Fail Safe account.

When using operating system authentication, Oracle Fail Safe only attempts to authenticate database access using the account specified for the Oracle Fail Safe service. See [Database Authentication](#) (page 7-11) for more information regarding database authentication.

4.3.1.1 Changing the Oracle Fail Safe Server Account

The Oracle Fail Safe Server service must run under a domain user account that is a member of the Administrators group and has access to all nodes in the cluster. The account must be a member of each node's local Administrators group; domain Administrator privileges are not required. This account is used by Oracle Fail Safe to change the configuration of Oracle resources in the cluster and also used to access Oracle Databases managed by Oracle Fail Safe. During the installation of Oracle Fail Safe, the installation process prompts for an account and password to be used by the Oracle Fail Safe server. To change the account used by Oracle Fail Safe, run the **Set Credentials** tool, and specify a new account to be used by Oracle Fail Safe.

To change the Oracle Fail Safe account, perform the following steps:

1. Select **All Programs** from the Windows **Start** menu.
2. Select **Oracle Fail Safe home** and then select **Set Credentials** to see the introduction screen.
3. Click **Continue** to enter the new credentials.
4. After changing the user account for the Oracle Fail Safe server, the Cluster Service service must be restarted.

 **Note:**

All cluster resources currently owned by the node will fail over to another node in the cluster when the Cluster Service service is stopped

[Figure 4-2](#) (page 4-6) shows the dialog box for Oracle Fail Safe Server Credentials explaining its utility.

Figure 4-2 Oracle Fail Safe Server Credentials

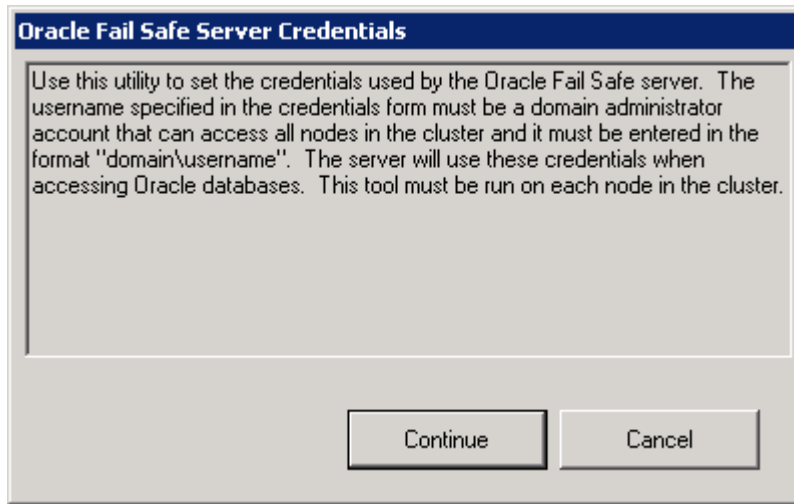
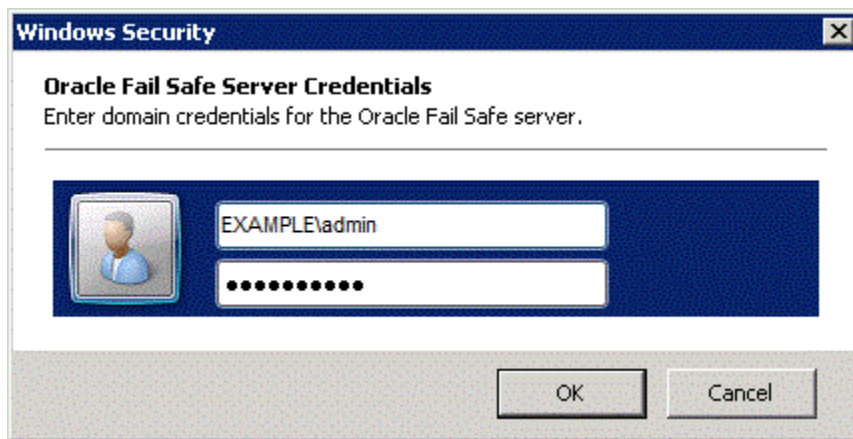


Figure 4-3 Windows Security Settings for the Oracle Fail Safe Server



4.3.2 Oracle Fail Safe Manager

When you run Oracle Fail Safe Manager, you must login to an account that is a domain user account (not a local account) with Administrator privileges on all cluster nodes. The account must be a member of each node's local Administrators group; domain Administrator privileges are not required.

4.4 Discovering Standalone Resources

Oracle Fail Safe automatically discovers (locates) and displays [standalone resources](#) in the Oracle Fail Safe Manager tree view when you select the Standalone Resources folder from the tree view. [Configuring Single-Instance Databases for High Availability and Disaster Tolerance](#) (page 7-1) contains information about how Oracle Fail Safe discovers each type of component that can be configured for high availability with Oracle Fail Safe.

4.5 Renaming Resources

Once a resource is added to a group, the resource name must not be changed. 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 Fail Safe in any one Oracle home on all cluster nodes. Only one version of Oracle Fail Safe 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 Microsoft Windows* for information about the compatibility of various versions on Oracle Fail Safe Manager and the Oracle Fail Safe server component.

 **Note:**

Multiple releases of Oracle Fail Safe Manager can be installed on a system, but each release must be installed in a different Oracle home.

- 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 cluster **Validate** action validates this symmetry. See [Validating Cluster](#) (page 6-1) for information about the cluster **Validate** action.
- All databases and listeners in a group must come from the same Oracle home.

On adding a database to a group, an Oracle Net listener resource is also added to the group. Optionally, you can add an Oracle Management Agent resource to the group.

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

4.7 Configurations Using Multiple Network Names

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

You can add up to 32 network names to a group, prior to adding resources, by starting the Add Resource to Group Wizard. In Microsoft Windows Failover Cluster Manager, select a group, then select **Add a Resource** action from the **Actions** menu in the right pane of the screen to add a network name (also known as client access point) to the group.

Note the following restrictions:

- At least one network name must be added to a group before you can add another resource to the group.
- If the group contains one or more Oracle Databases, then:
 - All network names 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 network names that you specify for the first database that you add to the group. (The set of network names can contain as few as one address.)

See [Selecting Network Names](#) (page 7-8) for more information about configuring multiple network names with Oracle Databases.

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

Multiple network names 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 different network names can be allocated 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 network name 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 a Custom Application), then you can dedicate one network name for users to access the database directly and another for users to access the Custom Application. Alternatively, if there are many database users, then you can have some users access the database using one network name and the others use the other network name, to balance the network traffic.

See the online help in Oracle Fail Safe Manager for information about adding a network name 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 for Microsoft Windows*. Once that task is completed, there is one final step. Select the **Validate** action for 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 **Validate** action, 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 **Validate** operation for each group for which the new node is listed as a possible owner. [Validating the Configuration of Oracle Resources](#) (page 6-3) describes the **Validate** operation. Groups can also be verified using the Oracle Fail Safe PowerShell cmdlet Test-OracleClusterGroup, as described in [PowerShell Commands](#) (page 5-1).

5

PowerShell Commands

In Windows Server 2008 R2, Microsoft introduced a set of failover cluster PowerShell commands (cmdlets) as the preferred scripting tool for managing failover clusters. PowerShell is the new command and scripting language offered by Microsoft and intends to replace the old command (CMD) environment used in the past. The new failover cluster cmdlets have replaced the old command line utility, `CLUSTER.EXE`, which might not be available in the future releases of Windows Server. Oracle now provides a new set of PowerShell cmdlets that has also replaced the old `FSCMD.EXE` utility.

The PowerShell cmdlets can be used on server systems, such as Windows Server 2012 R2, or on client systems, like Windows 8.1. The Microsoft failover cluster cmdlets are added to a server when the failover cluster feature is added to the system. The failover cluster cmdlets can be added to a client system by installing the Remote Server Administration Tools package available in Microsoft. Oracle Fail Safe cmdlets are installed as part of the Oracle Fail Safe Manager installation component.

5.1 Getting Started

To start a basic PowerShell session on a system that already has PowerShell installed, perform the following steps:

1. Click **Start**.
2. Click **Accessories** and then click **Windows PowerShell**.
3. Double-click **Windows PowerShell**. This starts a basic PowerShell session with the standard cmdlets loaded.

A group of related cmdlets are packaged in a container called a module. Before using the cmdlets that are specific to failover clusters, load them into the PowerShell session using the **Import-Module** cmdlet. Use the following command to load the failover clusters module.

```
PS C:\Users\admin> Import-Module FailoverClusters
```

To load the Oracle Fail Safe cmdlets, run the following command:

```
PS C:\Users\admin> Import-Module FailSafe
```

A PowerShell module also contains a link to the descriptions of the commands that it contains. After loading the module, find out what cmdlets are contained in the module and display the help text describing each of the cmdlets.

For example, to see what cmdlets are provided by the FailSafe module, type the following command:

```
PS C:\Users\Admin> Get-Command -Module FailSafe
```

CommandType	Name	Definition
-----	----	-----
Cmdlet	Add-OracleClusterResource	Add-OracleClusterResource [...
Cmdlet	Get-OracleClusterResource	Get-OracleClusterResource [...

```

Cmdlet          Remove-OracleClusterResource  Remove-OracleClusterResource...
Cmdlet          Stop-OracleClusterDatabase    Stop-OracleClusterDatabase [...
Cmdlet          Test-OracleCluster           Test-OracleCluster [[-Name] ...
Cmdlet          Test-OracleClusterAvailableD... Test-OracleClusterAvailableD...
Cmdlet          Test-OracleClusterGroup    Test-OracleClusterGroup [[-N...

```

To get basic information about all cmdlets that contain the string "cluster", use the following command:

```
PS C:\Users\Admin> Get-Help cluster
```

```

Name                Category  Synopsis
----                -
Add-OracleClusterResource  Cmdlet    Adds an Oracle resource to a fai...
Get-OracleClusterResource  Cmdlet    Gets an Oracle resource object
Remove-OracleClusterResource  Cmdlet    Removes an Oracle cluster resour...
Stop-OracleClusterDatabase  Cmdlet    Stops an Oracle database
Test-OracleCluster         Cmdlet    Verifies the installation and co...
Test-OracleClusterAvailableDat... Cmdlet    Verifies the configuration of an...
Test-OracleClusterGroup    Cmdlet    Verifies the configuration of Or...
Add-ClusterDisk           Cmdlet    Make a new disk available for us...
Add-ClusterFileServerRole  Cmdlet    Create a clustered file server (...
.
.
.

```

To get information specific to a particular cmdlet specify the cmdlet name to `Get-Help`.

Example 5-1 Get-OracleClusterResource

```
PS C:\Users\Admin> Get-Help Get-OracleClusterResource
```

```

NAME
    Get-OracleClusterResource

SYNOPSIS
    Gets an Oracle resource object.

SYNTAX
    Get-OracleClusterResource [[-Name] <String>] [-Available] [-InputObject <PS
    Object>] [-Cluster <String>] [-Credential <PSCredential>] [<CommonParameter
    s>]

DESCRIPTION
    This cmdlet will return an Oracle cluster resource object or objects based
    on the input provided.

RELATED LINKS
    Add-OracleClusterResource
    Remove-OracleClusterResource
    Test-OracleClusterAvailableDatabase
    Stop-OracleClusterDatabase

REMARKS
    To see the examples, type: "get-help Get-OracleClusterResource -examples".
    For more information, type: "get-help Get-OracleClusterResource -detailed".
    For technical information, type: "get-help Get-OracleClusterResource -full".

```

5.2 About Common Parameters

Oracle Fail Safe cmdlets displays the same progress information that is shown in the Oracle Fail Safe Manager progress dialog output box for the same operation. However, for the output to be displayed the `-Verbose` switch must be specified.

```
PS C:\Users\Admin > Test-OracleClusterGroup "Test Group" -Verbose
VERBOSE: FS-10371: FSWIN3 : Performing initialization processing
VERBOSE: FS-10371: FSWIN4 : Performing initialization processing
VERBOSE: FS-10373: FSWIN3 : Determining owner node of resource
.
.
.
```

To run a cmdlet on a different cluster, the `-Cluster` parameter must be added to a command. This is mandatory for running a cmdlet on a Windows client system that only has the Oracle Fail Safe Manager component installed. For example, to show the Oracle cluster resources on a cluster named "FinanceCluster" use the following command:

```
PS C:\Users\Admin > Get-OracleClusterResource -Cluster FinanceCluster
```

Name	State	Group
----	----	----
OracleOraDb11g_home1TNSLis...	Online	Receivables
OracleOraDb12c_home1TNSLis...	Online	Payables
ReceivablesDb	Online	Receivables
PayablesDb	Online	Payables

Some commands are not expected to be used as pipeline input. Thus running these commands does not send an object to the pipeline. To force the command to write the target object to the output pipeline stream, use the `-PassThru` switch. For example, the following command does not send any objects to the pipeline.

```
PS C:\Users\Admin> Get-ClusterGroup | Test-OracleClusterGroup
```

Adding the `-PassThru` switch forces the cluster group objects to be sent to the pipeline.

```
PS C:\Users\Admin> Get-ClusterGroup | Test-OracleClusterGroup -PassThru
```

Name	OwnerNode
----	-----
Test Group	Node1
Cluster Group	Node2
Available Storage	Node2

5.3 Using the Oracle Fail Safe cmdlets in Scripts

Some Oracle Fail Safe cmdlets will prompt the user for confirmation before proceeding. For example, the `Test-OracleClusterAvailableDatabase` cmdlet discovers that the specified database instance is not running and asks if the instance can be started. This is not a problem in an interactive session, but when running a script, the confirmation prompt receives the next command line in the script, and that may cause the script to fail. PowerShell provides different methods to address this situation. One way to avoid prompts is to specify the `"-Confirm:$false"` parameter in a command. When that qualifier is present, no confirmation prompt is displayed and the cmdlet

proceeds without interruption. In the following example, the first command causes a confirmation prompt to be displayed while the second command that includes "-Confirm:\$false", proceeds without asking the user if the database may be started.

```
PS C:\Users\Admin> Test-OracleClusterAvailableDatabase TestDb
WARNING: FS-10349: Database instance OFS3 is not alive. Do you want to stop and
restart the database instance?

Confirm
Are you sure you want to perform this action?
Performing operation "Test-OracleClusterAvailableDatabase" on Target "TestDb".
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is
"Y"): N
WARNING: FS-10340: Database instance OFS3 is not started and therefore cannot be
fully verified.
PS C:\Users\Admin> Test-OracleClusterAvailableDatabase TestDb -Confirm:$false
```

Another way to prevent confirmation prompts is to set the default response variable, `$ConfirmPreference`. You can set this variable to the lowest level that requires confirmation ("high", "medium", "low" or "none"). For example, most Oracle Fail Safe cmdlets are known to have a high impact on system operations. Therefore for Oracle Fail Safe cmdlets to be confirmed automatically, you must set the `$ConfirmPreference` variable to "none", as shown below.

```
PS C:\Users\Admin> $ConfirmPreference="none"
PS C:\Users\Admin> Test-OracleClusterAvailableDatabase TestDb
```

Some commands, such as the `Stop-OracleClusterDatabase` command are always expected to prompt for confirmation before proceeding. This cmdlet provides a `-Force` switch that can be used to prevent a confirmation prompt.

```
PS C:\Users\Admin> Stop-OracleClusterDatabase TestDb -Force
```

5.4 FSCMD Equivalent cmdlets

The following table lists `FSCMD.EXE` commands and the Oracle Fail Safe PowerShell cmdlet(s) that can be used to accomplish the same task.

FSCMD.EXE Command	PowerShell Command
DISABLEISALIVE	(Get-OracleClusterResource <db name>).IsAliveEnabled=\$false
ENABLEISALIVE	(Get-OracleClusterResource <db name>).IsAliveEnabled=\$true
MOVEGROUP	Move-ClusterGroup
OFFLINEGROUP	Stop-ClusterGroup
OFFLINERESOURCE	Stop-ClusterResource or Stop-OracleClusterDatabase
ONLINEGROUP	Start-ClusterGroup
ONLINERESOURCE	Start-ClusterResource
VERIFYALLGROUPS	Get-ClusterGroup Test-OracleClusterGroup
VERIFYCLUSTER	Test-OracleCluster
VERIFYGROUP	Test-OracleGroup

5.5 Examples

When operating system (OS) authentication is not enabled for a cluster or a database, you must provide a user name and password for the database. The `Test-OracleClusterAvailableDatabase` cmdlet provides the `-SysPwd` parameter for specifying the password for the database `sys` account. Note that the password is obtained from the user by running the `Read-Host` cmdlet.

```
PS C:\Users\Admin> Test-OracleClusterAvailableDatabase TestDb -SysPwd (Read-Host
-AsSecureString -Prompt "SYS Password")
SYS password: ****
```

When adding an Oracle resource, depending on the resource type, the resource may or may not have a user name and password property. For a database resource, if OS authentication is not being used, the user name and password must be set in the resource before it is added to a cluster group.

```
PS C:\Users\Admin> $testdb = Get-OracleClusterResource TestDb -Available
PS C:\Users\Admin> $testdb.UserName="SYS"
PS C:\Users\Admin> $testdb.Password=Read-Host -AsSecureString -Prompt "SYS
Password"
SYS Password: ****
PS C:\Users\Admin> $testdb | Add-OracleClusterResource -Group FsTutorial
```

By using the pipeline capabilities of PowerShell you can link together various commands.

In the following example, all Oracle resources are first fetched, then the databases selected from that list, and finally the databases are stopped using the immediate mode.

```
PS C:\Users\Admin> Get-OracleClusterResource |
>> where {($_.Type -ieq "Oracle Database")} |
>> Stop-OracleClusterDatabase -Mode Immediate -Force
```

6

Validating Actions

This chapter provides general information about the validation process in Oracle Fail Safe Manager. The following topics are discussed in this chapter:

- [Validating Operations](#) (page 6-1)
- [Dumping Cluster](#) (page 6-7)
- [Verifying Security Parameters](#) (page 6-8)
- [Additional Troubleshooting Information](#) (page 6-9)

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. Find more information about these messages in the *Oracle Fail Safe Error Messages for Microsoft Windows* manual.

6.1 Validating Operations

Oracle Fail Safe provides a family of tools to help you validate 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 validate operation takes the appropriate action to fix any potential or actual problems.

Use the validate commands at any time to validate your cluster, group, or standalone database. If problems are found during validation, 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 validate commands, then fix the errors and then rerun the validate command. Repeat this process until the validate operation runs without errors.

6.1.1 Validating Cluster

The **Validate** cluster action validates the installation of the cluster. You can perform a cluster verification at any time. Select the cluster you want to validate from the list, then select **Validate** from the **Actions** menu in the Cluster view.

The first time you connect to a cluster after installing or upgrading the Oracle Fail Safe software, you are prompted to run **Validate**. You can run the **Validate** action at any time, however, you must run it whenever the cluster configuration changes. The **Validate** action 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

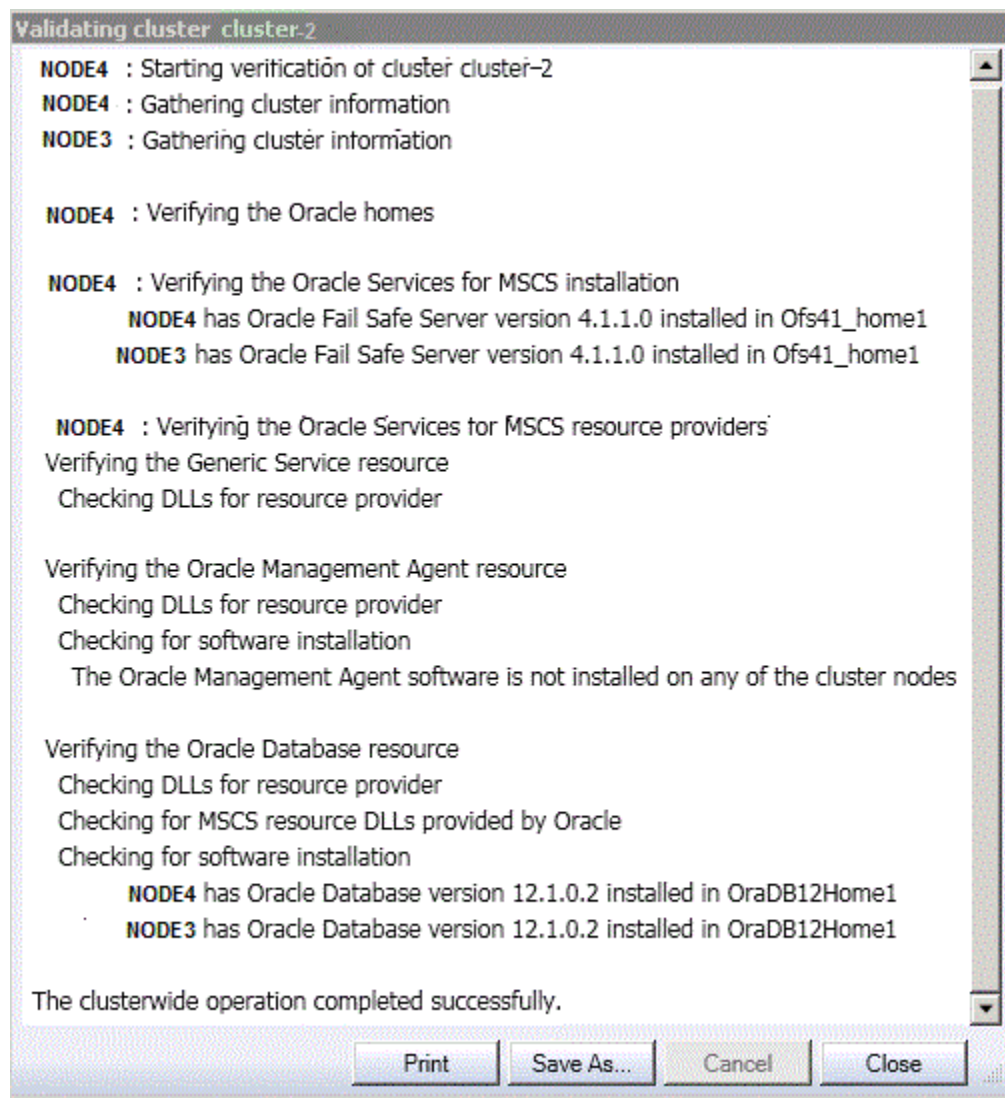
Oracle home name on all nodes in the cluster where the Oracle Database software is installed.

- The Oracle Fail Safe 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

Validate also registers Oracle resource DLLs with Microsoft Windows Failover Clusters. Moreover, if any of the cluster configuration changes, then Oracle recommends that you run the Microsoft Windows Failover Cluster Manager **Validate Cluster** wizard to verify that the cluster configuration is still valid.

Figure 6-1 (page 6-2) shows the output from a typical **Validate** action.

Figure 6-1 Verifying Cluster Progress Window



If you run the **Validate** operation and it does not complete successfully, it may indicate one or more of the following problems:

- A problem exists in the configuration of the hardware, network, or the Microsoft Windows Failover Clusters.
- 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 face problems with Oracle Fail Safe, then the problem is based in the Oracle Fail Safe configuration.

6.1.2 Validating the Configuration of Oracle Resources

The **Validate** action 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.
- Repairs a group that is misconfigured after prompting.

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

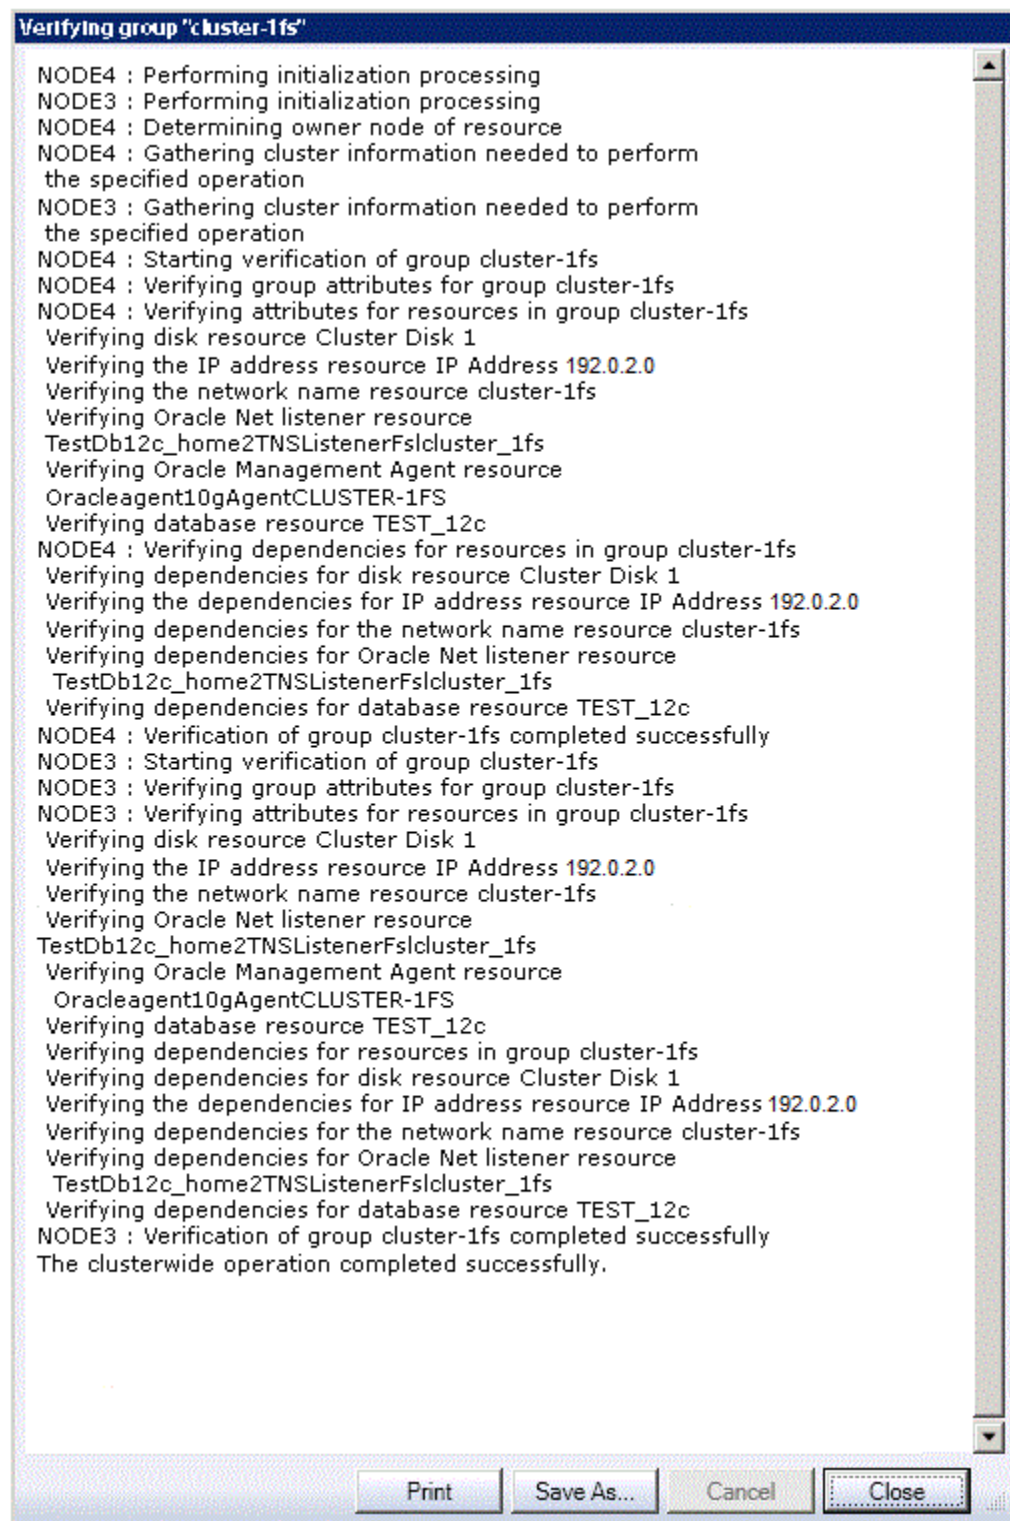
- A group or resource in a group does not come online.
- Failover or failback do not perform as you expected.
- You add a node to the cluster.

Perform the following steps for validating the configuration of Oracle resources:

1. Select a group, then select **Validate** from the **Actions** menu in the Cluster view. Or, you can run the **Validate** action using the PowerShell cmdlet `Test-OracleClusterGroup` command (see [PowerShell Commands](#) (page 5-1)).
2. Run the `Test-OracleClusterGroup` command in scripts as batch jobs.
3. Watch the progress of the **Validate** action and view the status of the individual resources in the group as Oracle Fail Safe verifies the group.

[Figure 6-2](#) (page 6-4) shows the output from a **Validate** action.

Figure 6-2 Verifying Group Progress Window



6.1.3 Validating Standalone Database

A standalone database can be validated at any time by selecting the **Validate** action. Select the database from the Available Oracle Resources list and then run the **Validate** action.

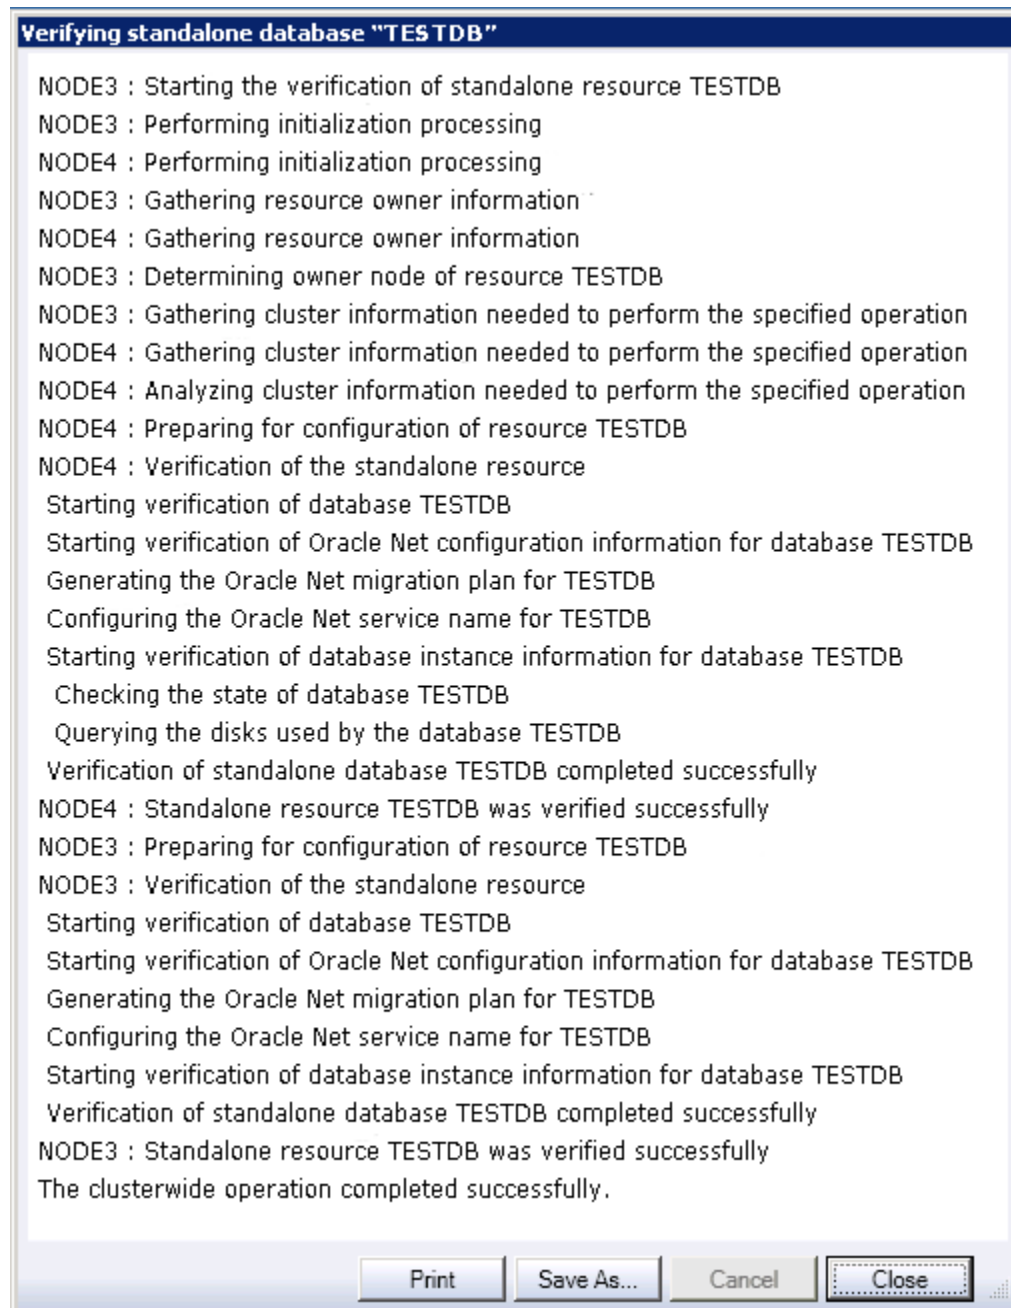
The **Validate** operation performs validation checks to ensure that the standalone database is configured correctly on the node where it resides and to remove any references to the database that may exist on other cluster nodes. (References to the database may exist on other cluster nodes if the database 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 **Validate** 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 Microsoft Windows Failover Clusters 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 for Microsoft Windows*.

[Figure 6-3](#) (page 6-6) shows the output from a typical **Validate** operation in a Clusterwide Operation window.

Figure 6-3 Verifying Standalone Database Progress Window

To verify a standalone database, perform the following steps:

- Select **Oracle Resources** from the tree-view on the left panel of the window.
- Select a resource from the Available Oracle Resources list.
- Select **Validate** action from the **Actions** menu list in the right panel of the window.
- The Verifying standalone database progress window opens. This window shows the different tests run for the standalone database and in case of any errors, a message is displayed. These errors must be resolved before attempting to add the

database to a cluster group. The Oracle Fail Safe Server may be able to resolve some issues, but it will ask for your confirmation before making any changes.

- If a standalone database is open and you select the **Validate** action, then the action 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.
- If any problems are found during verification, then the **Validate** action 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. Run the **Validate** action to fix the network problem and subsequently add the database to a group.

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

6.2 Dumping Cluster

The Dump cluster action allows you to direct Oracle Fail Safe to display cluster data (such as number of cluster nodes, resource types, network information, Oracle Homes, restart action, and so on) in a window. You can then save this data to a file. You can enter this command periodically (and save the output) to maintain a record of changes made to the cluster over time, or you might enter it at the request of customer support so as to provide a snapshot of the cluster environment.

Data displayed when you select the **Dump** cluster action includes:

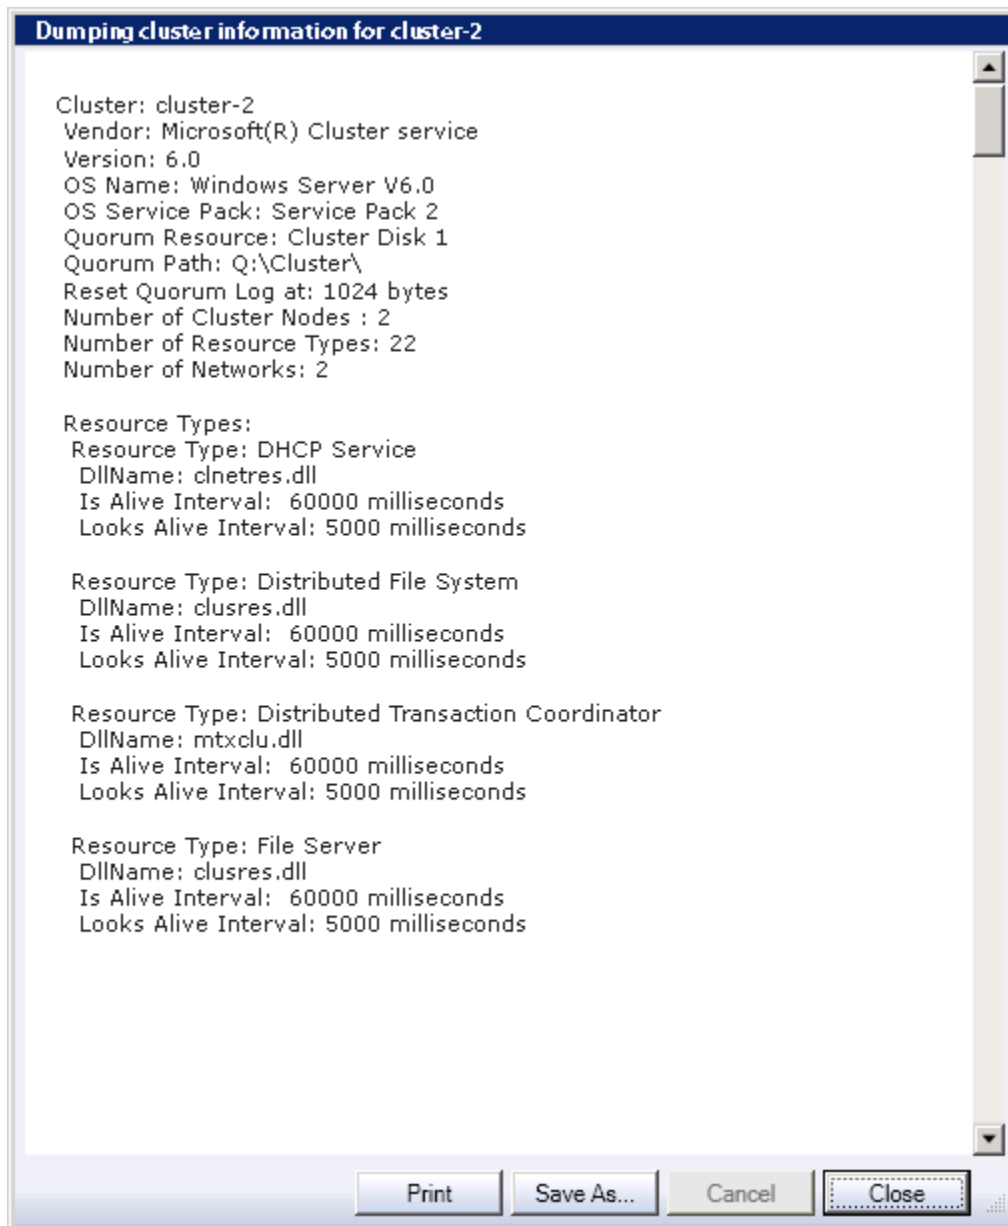
- Information related to the operating system (including the location of the quorum disk)
- Public and private network information
- Resources registered with the cluster
- Group failover and failback policies

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

To run the **Dump** cluster action, select the cluster you want to dump from the list, then select **Dump** from the **Actions** menu in the Cluster view.

[Figure 6-4](#) (page 6-8) shows the portion of the **Dump** cluster command output that provides information about cluster-2 cluster and some of its resources.

Figure 6-4 Dumping Cluster Information Progress Window



6.3 Verifying 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 account with local Administrator
```

privileges. The user account must also have the 'Log on as batch job' privilege.

```
User account specified for OracleMSCSServices is EXAMPLE\ADMIN
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 EXAMPLE\ADMIN
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 Additional Troubleshooting Information

Additional troubleshooting information is available from the following sources:

- Information about troubleshooting a specific component can be found in Chapters 7 and 8, which describe how to configure a particular component for high availability.
- Because Oracle Fail Safe is layered upon Microsoft Windows Failover Clusters software, you may need to refer to the Microsoft Windows Failover Clusters 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 Fail Safe 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 (except Oracle Database Personal Edition) running on Windows clusters configured with Microsoft Windows Failover Clusters.

By making a single-instance Oracle Database highly available, you ensure that even when a cluster node is shut down or fails, applications accessing that database suffer 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 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](#) (page 7-1)
- [Oracle Net Configuration for Standalone Single-Instance Databases](#) (page 7-1)
- [Adding Single-Instance Oracle Databases to a Group](#) (page 7-4)
- [About Oracle Net Listener Resource Creation and Configuration](#) (page 7-15)
- [Security Requirements for Single-Instance Databases](#) (page 7-20)
- [Optimizations for Single-Instance Database Recovery](#) (page 7-27)
- [Performing Administrative Tasks on a Single-Instance Fail-Safe Database](#) (page 7-27)
- [Database Homes](#) (page 7-28)
- [Configuring Transparent Application Failover \(TAF\)](#) (page 7-29)
- [Handling Errors and Troubleshooting Problems with Databases](#) (page 7-30)
- [Starting and Stopping Pluggable Databases](#) (page 7-35)

7.1 Discovering Standalone Single-Instance Databases

Oracle Fail Safe Server discovers standalone single-instance databases (those that are not in a cluster group) by looking for Oracle Database instance Windows services. Any service found on any cluster node that is not currently in a cluster group is displayed in the Oracle Fail Safe Manager's Available Oracle Resources list.

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 Listener Must Use IP Address for Local Host, 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. The local listener also opens any cluster IP addresses causing a cluster group listener failure if it attempts to listen on IP addresses assigned to the group.

To avoid this problem, the listener must use the node IP address for its host entry instead of the host name. Whenever Oracle Fail Safe validates a cluster group or adds a database to a group, it searches ADDRESS entries that have a HOST set to the local node's host name. All HOST entries that use the local node name are changed to use the IP address for the node.

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)
  )
```

7.2.2 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. As a result, 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. Locate the database initialization parameter file of the database and add the LOCAL_LISTENER parameter to the file.

```
LOCAL_LISTENER = network-name
```
2. 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. 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.

7.2.3 SID List Entries

Oracle Fail Safe does not attempt to maintain listener SID lists. If you have an application that requires a cluster database to be in the listener's SID list, then manually edit the appropriate `listener.ora` file on each node of the cluster.

7.2.4 Configuring Oracle Net on Nodes with Multiple Listeners

When Oracle Fail Safe searches for a standalone database listener, it scans the listener Windows services to find one listening on the network address used by the database. If there are multiple listeners on a network address, then Oracle Fail Safe selects the running listener service. If none of the listeners are started, then Oracle Fail Safe chooses the first listener found that is listening on the network address.

Note:

To prevent network configuration errors, 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.

7.3 Adding Single-Instance Oracle Databases to a Group

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

- One or more network names, 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 network name (or names) of the group for connection requests to the databases in the group

7.3.1 Prerequisites

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

- 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 [Identifying Database Parameters](#) (page 7-9) for more information about the placement of the initialization parameter file.
- Resources must belong to one group only. 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 network name.
- Database service names must be unique across the cluster.
- The listener and all the databases in the group must use the same Oracle home.
- Oracle Fail Safe does not support the use of mounted folders (mount points) or symbolic links for files used by an Oracle database. For example, a control file or server parameter file cannot be referenced by any filename other than the actual filename that represents the file. Files that are stored on a cluster shared volume, and thus use the `C:\ClusterStorage` root folder are supported.
- A text initialization parameter file (`PFILE`) must be used to start the database. Use the `oradim` utility to set the location of the `PFILE`.

7.3.2 Configuration Steps

[Table 7-1](#) (page 7-5) 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. To access online help, select **Help**

from the **Actions** menu on the right pane of Oracle Fail Safe Manager window. Or select **Fail Safe Documentation** in the middle pane of Oracle Fail Safe Manager, then select the **HTML** or **PDF** version of the **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 network names.	<ol style="list-style-type: none"> 1. In the Actions menu on the right pane of Microsoft Windows Failover Cluster Manager window, select Configure a Service or Application, and the High Availability wizard opens. 2. Select Other Server from the Select Service or Application page and click Next. 3. In the Client Access Point page, set the network address and click Next. 4. Choose a cluster disk from the list in the Select Storage page and click Next. Do not select any resource types and click Next to confirm your choices.
3	Create a sample database, if desired.	From Oracle Resources view, choose Create Sample Database action from the Actions menu in the right pane of the screen. You can use this sample database to 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.	Select the resource that you want to verify from the Available Oracle Resources list, then select Validate from the Actions menu of the Oracle Resources view. This 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.
5	Add the Oracle Database to the group.	Select the resource that you want to add from the Available Oracle Resources list, then select Add Resource from the Actions menu of the Oracle Resources view. This 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 <code>tnsnames.ora</code> file on each client system using a network configuration tool) to recognize the virtual server. See About Oracle Net Listener Resource Creation and Configuration (page 7-15) 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 network names currently in the group, and the number of nodes in the cluster.

Typically, each group has one network name, but more complex configurations may have multiple network names. 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 instance name and specification for the database initialization parameter file
- The database `sys` password, if OS authentication is not used

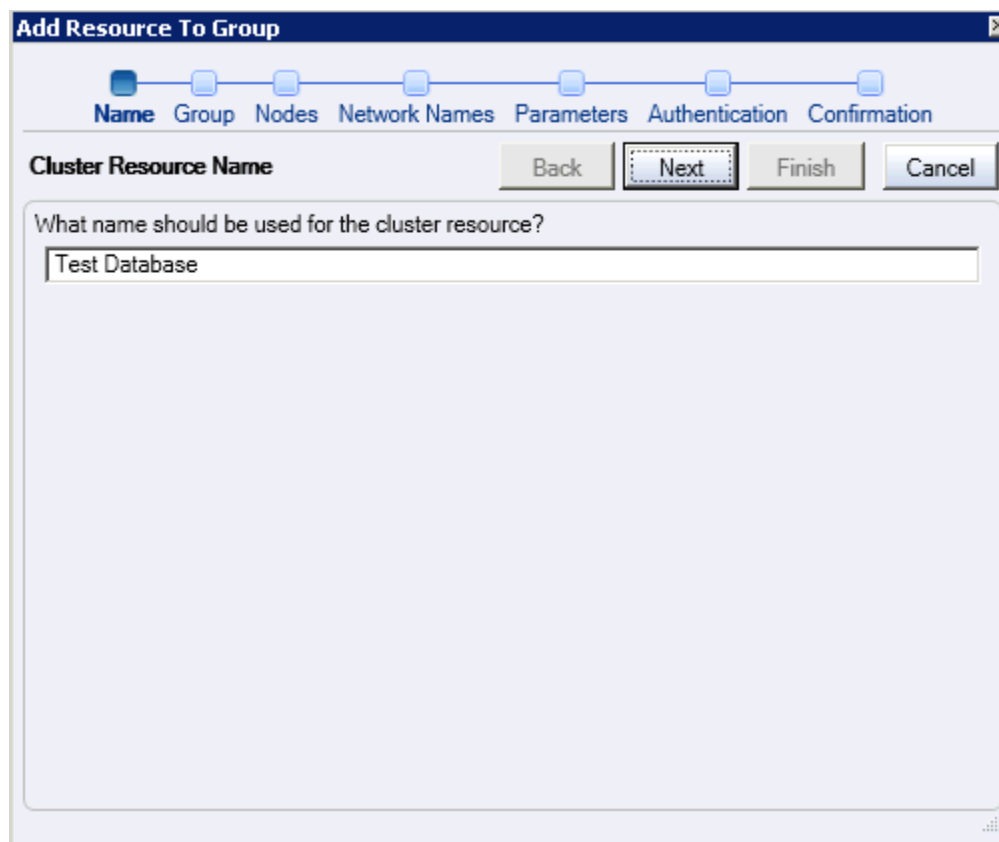
If you add a database to a group that currently contains multiple network names, then you are also asked to specify the network name or names for the listener.

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

7.3.3.1 Naming a Cluster Resource

Microsoft failover clusters allow you to use any text string for the name of a resource. By default, Oracle Fail Safe uses the instance ID for the database. You can change the name to something more meaningful, if desired. For example, the cluster resource name is changed to Test Database here.

Figure 7-1 Add Resource to Group Cluster Resource Name Wizard Page

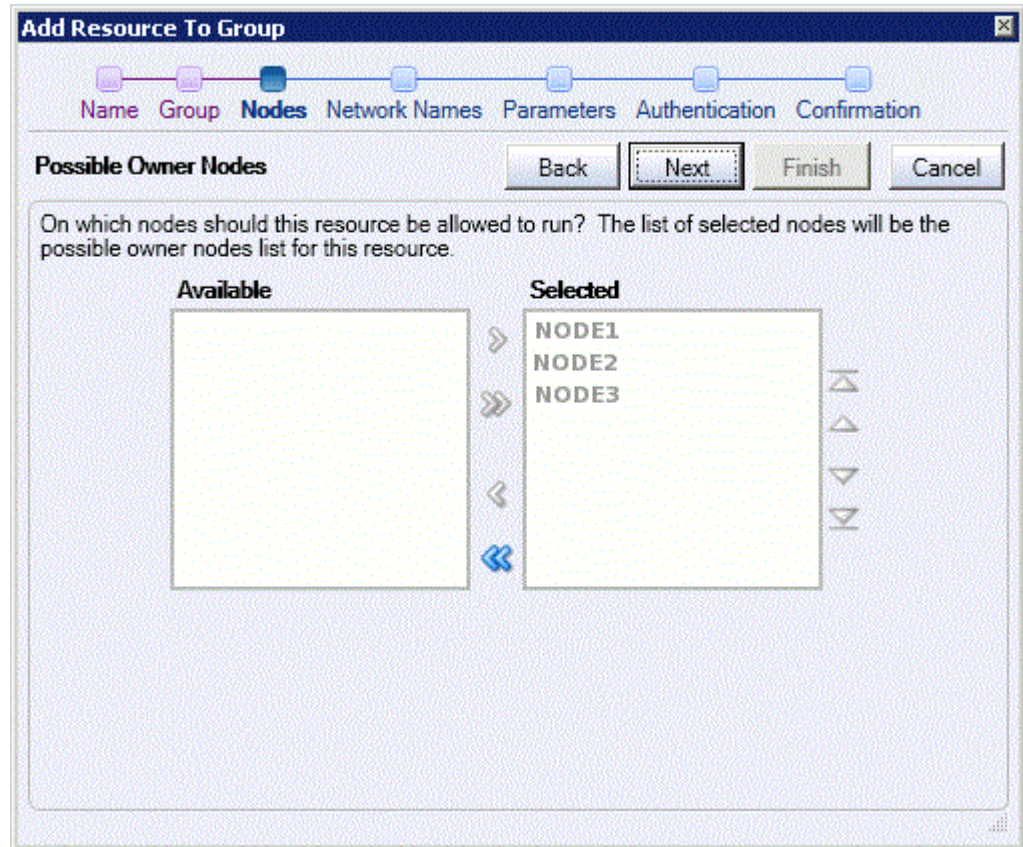


7.3.3.2 Choosing Nodes

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 7-2](#) (page 7-7). 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.

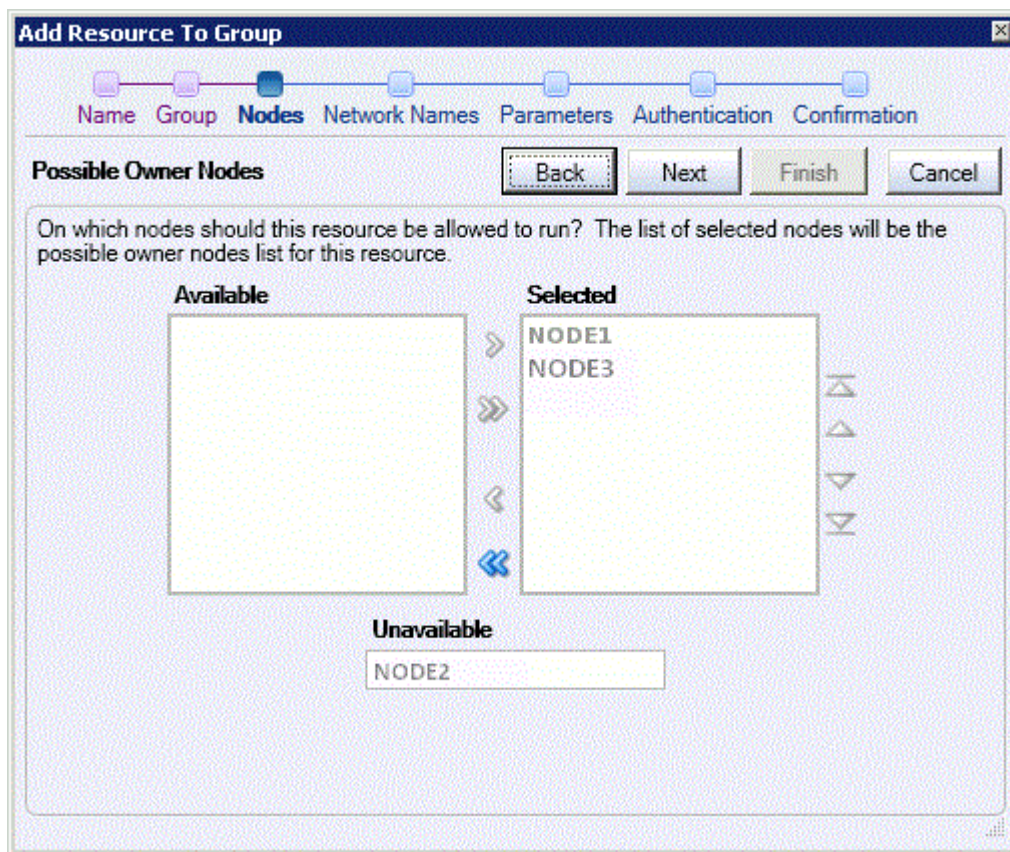
[Resource Possible Owner Nodes List](#) (page 2-15) describes in detail the concept of the possible owner nodes list.

Figure 7-2 Add Resource to Group 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, as shown in [Figure 7-3](#) (page 7-8).

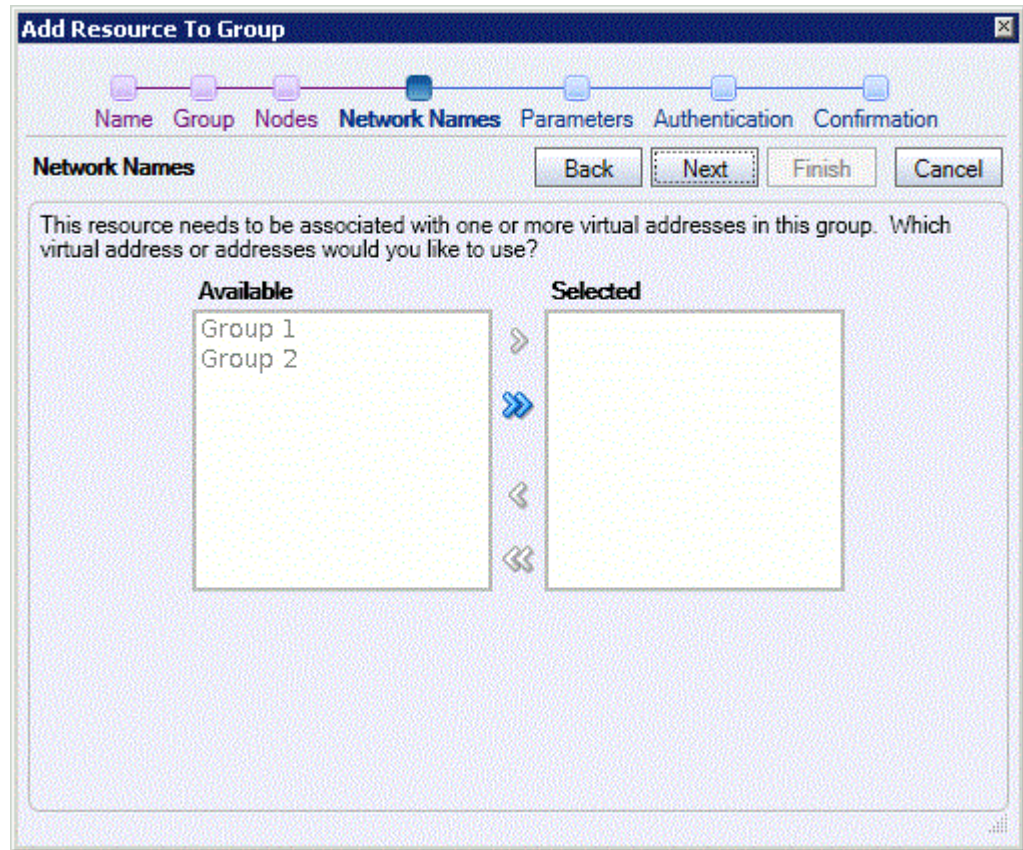
Figure 7-3 Add Resource to Group Wizard Page When Any Node Is Unavailable



7.3.3.3 Selecting Network Names

If the group to which you are adding a single-instance database contains multiple network names, then the Add Resource to Group Wizard asks you to specify the network name or names for the listener, as shown in [Figure 7-4](#) (page 7-9). This page is not displayed if the group to which you are adding a database contains only one network name.

Figure 7-4 Add Resource to Group Network Name Wizard Page



Oracle Fail Safe includes support for multiple network names in a group. All databases in a group must use the same network names, and the network names 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 network names to the group.
3. Add one or more single-instance databases to the group.

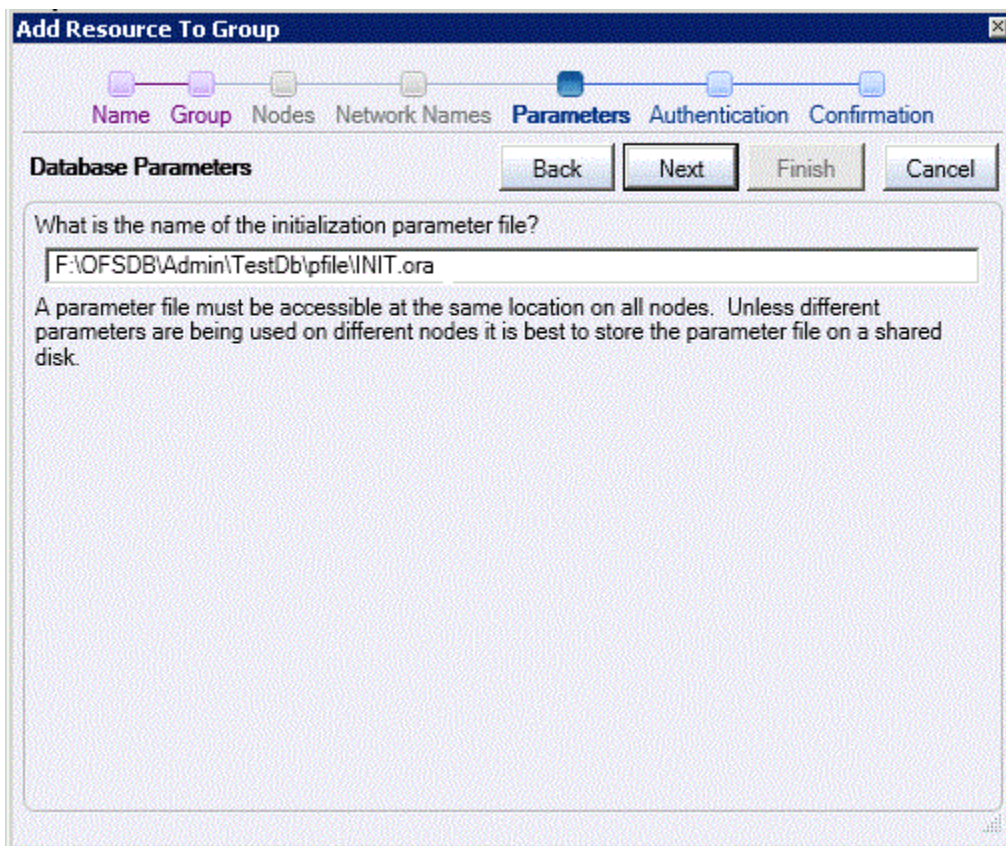
For example, if a group contains a database that is using two network names and you add a second database to the group, then the second database must use the same network names as the first database that was configured into the group. Oracle Fail Safe Manager checks to ensure that the same network names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

See [Configurations Using Multiple Network Names](#) (page 4-7) for information about configuring a resource in a group with multiple network names.

7.3.3.4 Identifying Database Parameters

The Add Resource to Group Wizard requests database parameters information to uniquely identify the single-instance database that is being configured for high availability, as shown in [Figure 7-5](#) (page 7-10).

Figure 7-5 Database Parameters 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 Microsoft Windows Failover Clusters, 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 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, 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, a copy of the initialization parameter file can be placed 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.

 **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.

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. The `SPFILE` must reside on a shared disk that is a member of the cluster group where the database resides. For example, the contents of the `PFILE` may include the following parameters:

```
SPFILE=F:\OFSDB\oradata\OFS1\spfileTestDboradb.ora
```

(If you specify an `SPFILE` in the `PFILE` that Oracle Fail Safe uses, then use caution if 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. Ensure that you 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.)

All Oracle database instances on each node of the cluster must use the same `SPFILE` and the file must be on shared storage. If the `SPFILE` is not currently stored on a shared disk, then create a copy using `SQL*PLUS` as follows:

```
CREATE SPFILE=shared disk path\spfiledb_unique_name.ora
```

Create a `PFILE`, `ORACLE_HOME\dfs\initsid.ora` that contains the name `SPFILE=shared disk path\spfiledb_unique_name.ora`.

7.3.3.5 Database Authentication

The Authentication page is presented if the account where Oracle Fail Safe was installed is not in one of the following Windows operating system groups:

- `ORA_DBA` group
- `ORA_SID_DBA` group
- or the `ORA_home_DBA` group associated with the database.

When the account where Oracle Fail Safe was installed in any of these three groups, it can use the operating system authentication to access the database. If the account is not a member of any of these three groups, then it must use the `SYS` account to access the database.

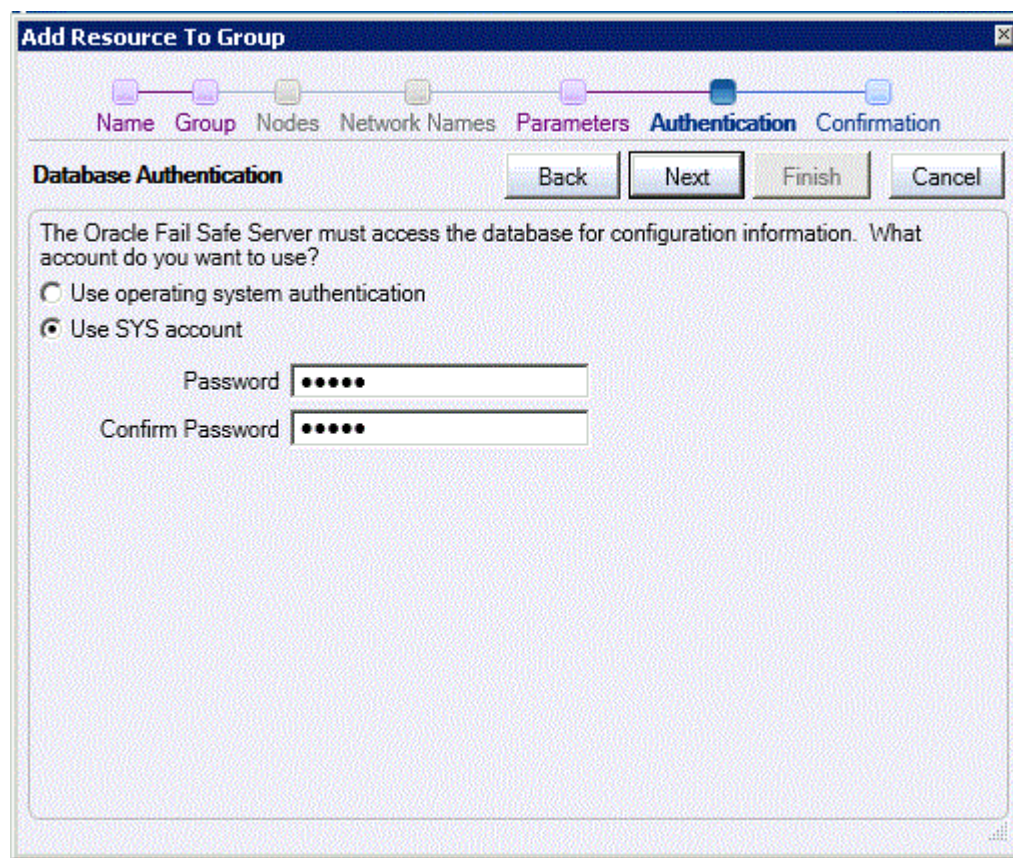
When an Oracle database is added to a cluster group, Oracle Fail Safe checks to see if the `ORA_sid_DBA` and `ORA_sid_OPER` local user groups exist on the cluster node that owns the database. If those local groups exist, then they are replicated to the other nodes in the cluster. Fail Safe will not copy any group members that are specific to that node, such as a local user name. It will copy built-in members. For example, the Administrators member will be copied to other nodes. If a local group has no members then it is ignored.

This page lets you specify whether Oracle Fail Safe should use operating system authentication or the `sys` account to access the database and its instances, as shown in [Figure 7-6](#) (page 7-12).

If operating system authentication is enabled for all databases in the cluster (the Fail Safe server username is a member of the `ORA_DBA` user group), then the **Use operating system authentication** option will be selected by default and cannot be changed.

If cluster-wide operating system authentication is not enabled, then you can choose to **Use operating system authentication** for a specific database (the Fail Safe server user will be added to the `ORA_SID_DBA` user group) or select **Use SYS account** to specify that Fail Safe should authenticate using the database `sys` account and its associated password.

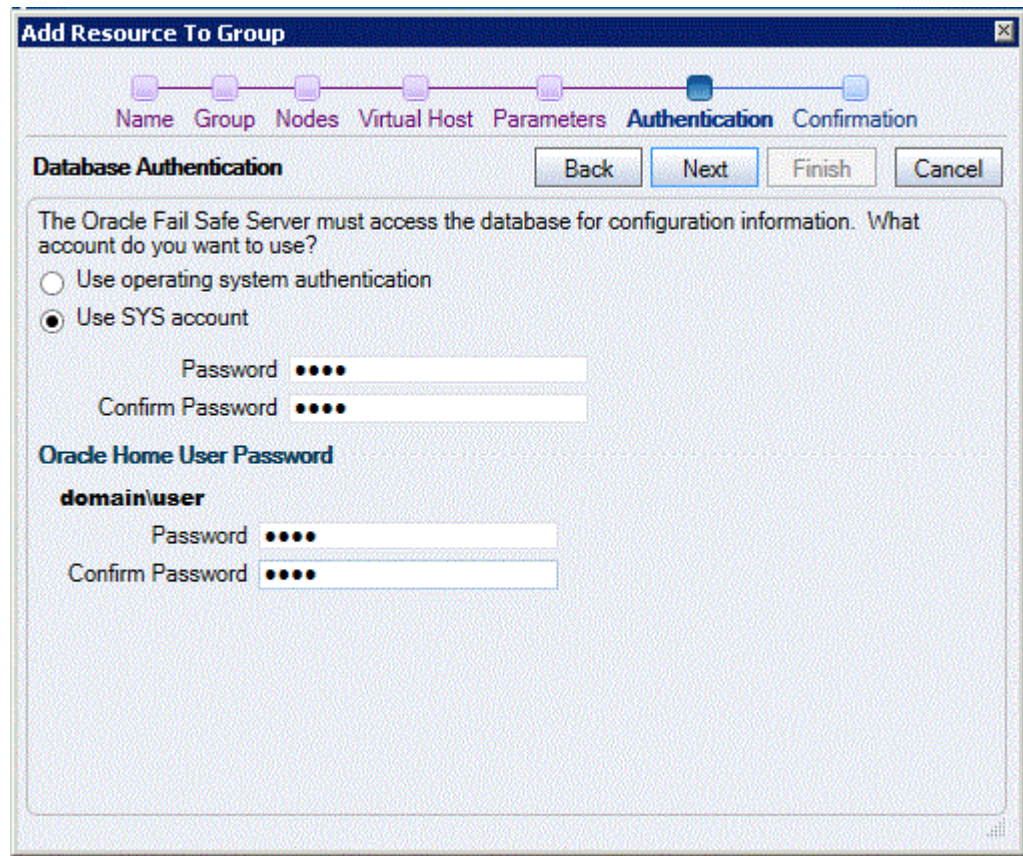
Figure 7-6 Database Authentication Page



The screenshot shows the 'Add Resource To Group' wizard window. The title bar reads 'Add Resource To Group'. At the top, there is a progress bar with seven steps: Name, Group, Nodes, Network Names, Parameters, Authentication, and Confirmation. The 'Authentication' step is currently selected and highlighted in blue. Below the progress bar, the 'Database Authentication' section is visible. It contains the text: 'The Oracle Fail Safe Server must access the database for configuration information. What account do you want to use?'. There are two radio button options: 'Use operating system authentication' (which is unselected) and 'Use SYS account' (which is selected). Below these options are two password input fields labeled 'Password' and 'Confirm Password', both containing five dots to indicate masked text. At the top right of the main content area, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

If an Oracle Home User is configured, then Oracle Fail Safe displays an additional set of password fields for the Oracle Home User. Ensure that you provide the Oracle Home User password too.

Figure 7-7 Database Authentication Page



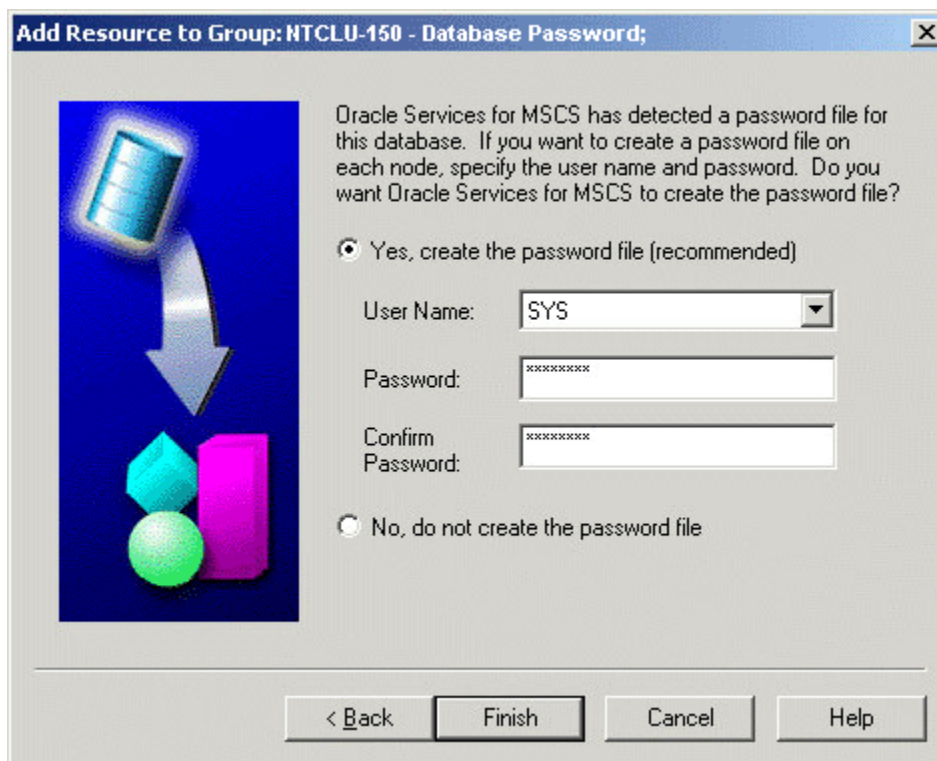
7.3.3.6 Database Password

If Oracle Fail Safe detects that the standalone database has an associated password file, then the wizard asks if you want Oracle Fail Safe to create the password file on all possible owner nodes for the fail-safe database, as shown in [Figure 7-8](#) (page 7-14).

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.

Figure 7-8 Database Password Wizard Page



Oracle Fail Safe makes the following adjustments to the database initialization parameter file, depending on your choice to have Oracle Fail Safe create the password file on all cluster nodes that are possible owner nodes for the database:

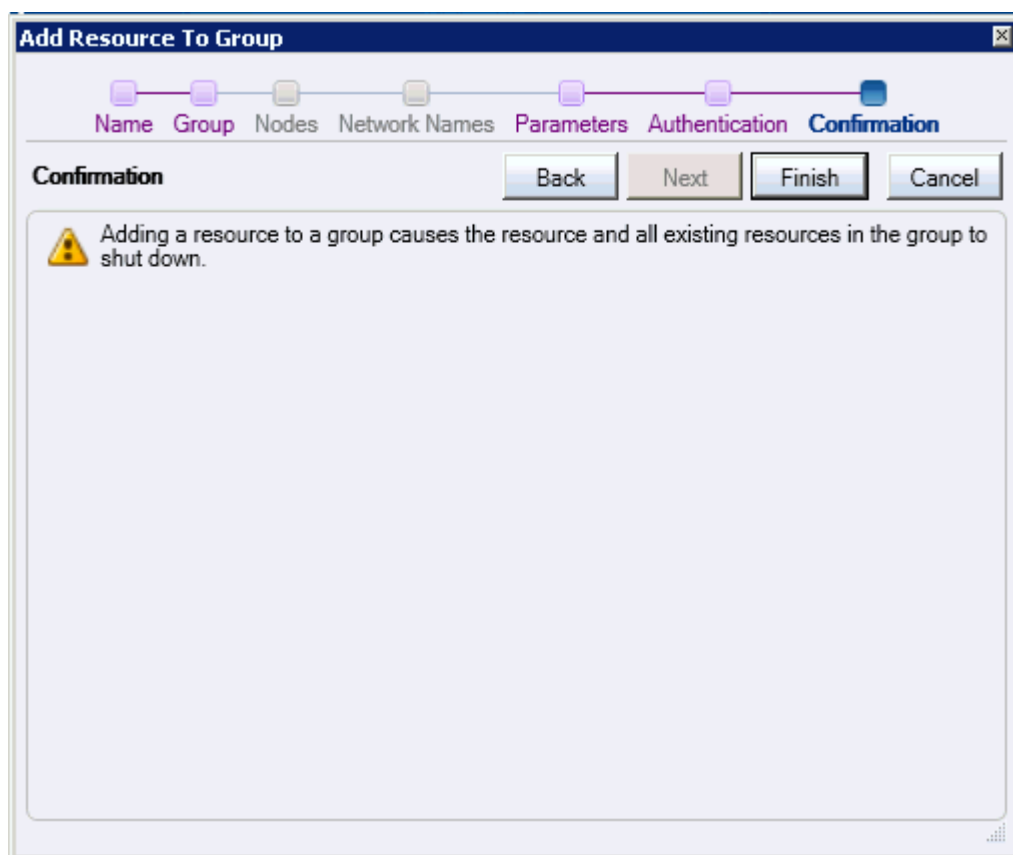
- Yes, create the password file
Sets the `REMOTE_LOGIN_PASSWORDFILE` parameter to `EXCLUSIVE`.
- No, do not create the password file
Sets the `REMOTE_LOGIN_PASSWORDFILE` parameter to `NONE`.

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 [Changing the SYS Account Password](#) (page 7-22) for information about how to update the password for this account after the database has been added to a group.

7.3.3.7 Database Resource Addition Confirmation

Finally, the Add Resource wizard asks you to confirm the operation. Note that the cluster group will be taken offline during the Add operation. The database and any resources in the group will be unavailable while Oracle Fail Safe adds the database to the group. Click **Finish** to complete the task of adding the Oracle Database to group.

Figure 7-9 Database Resource Addition Confirmation Page



7.4 About 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. The new group listener configuration is based on the listener that the standalone database is using when it is being added to the group. The new listener will be given the same parameters as the original listener and it will use the same port numbers in its address list.

During normal operations, the cluster will periodically poll the listener to verify that the Windows service is started and that the listener responds to status commands. If those checks fail, then the listener is terminated and the cluster starts its failover policies to determine if the listener resource should be restarted or if the group should be failed over to a different node. Any resource failure discovered by the Oracle cluster resource control manager will be logged in the Windows application event log.

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 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 network name information, it updates the `tnsnames.ora` files in all Oracle homes on cluster nodes that are possible owners for the database. This enables Oracle Fail Safe to access the database instance using the updated configuration.

7.4.2 Updating 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.

If the `TNS_ADMIN` environment variable exists, then Oracle Fail Safe will update the network configuration files in the directory pointed by the `TNS_ADMIN` environment variable instead of Oracle home.

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

When you add a single-instance database to a group, Oracle Fail Safe ensures that all net service descriptors in the `tnsnames.ora` file are updated to use the network name or names used by the cluster group.

- First, Oracle Fail Safe scans the `tnsnames.ora` file for any existing net service descriptors that reference the database. Any existing descriptors are changed to use the address list of the group's TNS listener.
- Then, for each service name found in the database's `service_names` parameter, Oracle Fail Safe ensures that there is a net service descriptor in the `tnsnames.ora` file.
- If no net service descriptor is found, Oracle Fail Safe creates a new net service descriptor that contains an address list that matches the group's listener address list. If there are multiple Oracle homes on the node and the `TNS_ADMIN` environment variable is not set, then all net service descriptors for the database are duplicated to the `tnsnames.ora` files in the other Oracle homes.
- Similarly, the new net service descriptors are duplicated to all `tnsnames.ora` files on all other nodes in the cluster.

When adding a single-instance database to a group, if you do not specify a domain name in the Oracle Net service name 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 12c 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 12c 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`.

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. On each cluster node, edit or add the service name entry to the `tnsnames.ora` file manually.

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

All client systems that connect to the database must have their `tnsnames.ora` files updated to use the cluster group's network name for the `HOST` parameter in each network service descriptor's address list that references the database. Edit each client's local `tnsnames.ora` file manually or use a network configuration tool.

7.4.2.2 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. Creates a new Oracle Fail Safe listener that is configured to listen on the network name associated with the single-instance database
2. Stops and restarts the standalone database listener to accept the changes that have been made
3. Starts the new Oracle Fail Safe listener

When a new cluster group listener is created, Oracle Fail Safe duplicates the port numbers from the original listener to the new listener. For example, if the original listener had `ADDRESS` entries with ports 1521 and 1522 in the `ADDRESS_LIST`, then the new listener will create an `ADDRESS_LIST` that contains the same port numbers.

When a new group listener is created, Oracle Fail Safe forces all databases in the cluster group to use secure registration through the IPC protocol. So, Oracle Fail Safe creates a parameter, `SECURE_REGISTER_group_listener_name` with the value, `IPC`.

When a database is added to a cluster group and there is no listener configured for the group, then Oracle Fail Safe will copy the parameters from the database's current listener to the new group listener. For example, if the database is currently using the default listener named "listener", and that listener has the parameter `INBOUND_CONNECT_TIMEOUT_LISTENER` in the `listener.ora` file, then Oracle Fail Safe will create the parameter `INBOUND_CONNECT_TIMEOUT_group_listener_name` for the new listener and assign it the value used for the `INBOUND_CONNECT_TIMEOUT_LISTENER` parameter.

7.4.2.3 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).

7.4.3 Using External Procedures with Databases Configured for High Availability

Oracle Fail Safe does not create external procedure parameters for new group listeners. If your application uses external procedures, then you must manually edit the `listener.ora` and `tnsname.ora` files on each node in the cluster and add the parameters needed for the external procedures used by your application.

7.4.4 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, ensure that Oracle Fail Safe can continue to use a dedicated server connection for its internal operations. 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.)

7.4.4.1 Shared Servers for Oracle8 Databases

Oracle Fail Safe supports single-instance Oracle8 databases that use a shared servers configuration. However, Oracle Fail Safe does not automatically update the database initialization file where the shared servers configuration is defined.

You can configure a standalone single-instance database or a single-instance database that is currently a resource in a group to use shared servers. In both cases, you must update the database initialization file by performing the following steps:

1. Determine the listener parameters for the group containing the single-instance database using the shared servers configuration, as follows:
 - a. Find the `listener.ora` file in the Oracle Net configuration directory in the Oracle home where the database resides.
 - b. Search the `listener.ora` file for the SID of the database and find the first listener address for the group using the TCP protocol.

For example, the boldface text in the following `listener.ora` file shows the first listener address of the group:


```

LISTENER =                               (Entries for default Listener)
  (ADDRESS_LIST =
    .
    .
    .
Fslvirtualnode =                          (Entries for Fail Safe Listener)
  (ADDRESS_LIST=
    (ADDRESS=
      (PROTOCOL=IPC)
      (KEY=OFS1)
    )
    (ADDRESS=
      (PROTOCOL=IPC)
      (KEY=805mts.world)
    )
    (ADDRESS=
      (PROTOCOL=TCP)
      (Host=virtualnode)
      (Port=1521)
    )
    (ADDRESS=
      (PROTOCOL=TCP)
      (Host=virtualnode)
      (Port=1526)
    )
  )
)

SID_LIST_Fslvirtualnode =
  (SID_LIST=
    (SID_DESC=
      (SID_NAME=OFS1)
    )
  )

```

- c. Find the first address in the file that includes the line (PROTOCOL=TCP), and format the address parameters into a single line. For example:

```
(ADDRESS=(PROTOCOL=TCP)(Host=virtualnode)(Port=1521))
```

2. Update the database initialization file (for example, initofs1.ora) to use the listener parameters for the group. To do this, perform the following steps:
 - a. Open the database initialization parameter file.

Note that the initialization parameter file for the database may reside on a disk on the shared interconnect, for example:

```
H:\OFSDB\OFS1\PARAM\initofs1.ora
```

If each node of the cluster has its own copy of the file on its private disk, then you need to update all copies.

- b. In the database initialization parameter file, search for the line containing the following parameter:

```
mts_listener_address
```

- c. Replace the value of the mts_listener_address parameter with the listener address that you formatted in step 1c.

For example, assume the original mts_listener_address parameter contains the following value:

```
mts_listener_address = "(ADDRESS=(PROTOCOL=TCP)(HOST=node1)(PORT=1521))"
```

Replace the line, as follows:

```
mts_listener_address = "(ADDRESS=(PROTOCOL=TCP)(HOST=virtualnode)
(PORT=1521))"
```

- d. Save the database initialization file.
3. Check that value of the `mts_service` parameter is the database SID.

The Database Configuration Assistant might use the database name for the `mts_service` parameter value. If so, change the value to the database SID.

4. Stop and restart the resources in the group.

To have your changes take effect, use Oracle Fail Safe Manager or the `FSCMD` command to take the group that contains the database offline and then place it back online. This stops and restarts all resources in the group.

7.4.4.2 Shared Servers for Databases

To use a shared server configuration, it may be necessary 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. Do 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 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 [Shared Server Configuration and a Standalone Database](#) (page 7-2).

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 (`initdatabase-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.

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 Fail Safe 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 (page 7-21).
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 steps 2 (page 7-21) and 3 (page 7-21) 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 SYS Account Password

The password for the SYS account is normally stored in a password file that is located in the Oracle home associated with the database. Since each cluster node has an Oracle home that is used for a database, that means that there are multiple password files to be maintained when a database is a cluster resource. To change the password for the SYS account use the Oracle Fail Safe Manager utility so that the change can be propagated to each Oracle home in the cluster that is associated with the database. Do not change the SYS account password using SQL*Plus or any other utility since that will interfere with the password maintenance strategy used by Oracle Fail Safe.

Use the Properties page to change a password for a database resource. In Oracle Fail Safe Manager, perform the following steps:

1. Select the database from the cluster resource list and click **Properties** in the **Actions** menu.
2. The resource properties page displays.
3. If the operating system authentication is not used, then the password field displays.

TESTDB Properties

General | Dependencies | Policies | Advanced Policies

Resource Name TESTDB
Resource Type Oracle Database
Status Online

Instance Name INS4
Database Name INS4
Parameters File G:\TESTDB\INS4\Admin\INS4\pfile\INITINS4.ora
Home Name OraDb12c_home1
Database Role Primary

Enable Is Alive polling

Authentication Method SYS password

Old Password
New Password
Confirm Password

OK Cancel Apply

For most databases, the password change takes effect immediately. However, if the database password file is being shared, then the database must be re-opened before the password change takes effect. You can do this by taking the database cluster resource offline and then bringing them back online. Alternatively you can move the cluster group that owns the database to another node in the cluster.

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.
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. Specify the location of the converted database files as asked by the Oracle Database Upgrade Assistant. 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. If you are upgrading an Oracle7 database to an Oracle8 database, the Oracle Database Upgrade Assistant creates a new data file called `<Oracle_Home>\database\mig<SID>.ora` on a private disk, where SID is the database instance name. Move this new data file to a cluster disk with your other data files, as follows:

- a. Use SQL*Plus to connect to the database, then shut it down.

```
SQL> SHUTDOWN
```

- b. Copy the `<Oracle_Home>\database\mig<SID>.ora` file to a cluster disk where `<SID>` is the database instance name.

- c. Use SQL*Plus to connect to the database and then execute the following commands:

```
SQL> STARTUP PFILE=init<SID>.ora MOUNT
SQL> ALTER DATABASE RENAME FILE
'<Oracle_Home>\database\mig<SID>.ora' TO
'cluster_disk\mig<SID>.ora';
SQL> SHUTDOWN
SQL> EXIT
```

6. 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. To add an available resource to a group, select the resource you want to add to a group, then select **Add Resource** from the **Actions** menu of the Oracle Resources view.
- b. The Add Resource to Group guided process wizard opens to assist in the configuration of the cluster resource.

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 use the same Oracle home.

7.6 Configuring a ctxsrv Server for High Availability

For releases of the Oracle database server prior to Oracle9i, the ctxsrv server processes background data manipulation language (DML) for indexing, searching, retrieving, and viewing documents. (Beginning with Oracle9i, you can index, search, retrieve, and view documents with standard SQL or PL/SQL procedures.)

If you are using a ctxsrv server with a single-instance Oracle database server, you can configure the ctxsrv server for high availability, as follows:

1. Create a batch file to start the ctxsrv server and specify the personality mask. For example, create a file named context.bat that contains the following command line:

```
ctxsrv -user CTXSYS/CTXSYS -personality QDM
```

2. Open Windows Failover Cluster Manager and configure the ctxsrv server as a highly available generic application, as follows:
 - a. On the **File** menu, click **New**, then **Resource**.
 - b. On the **New Resource** page:
 - In the **Name** field, enter the name of the ctxsrv server.
 - In the **Description** field, enter a description of the ctxsrv server, if desired.
 - In the **Resource type** field, select **Generic Application**.
 - In the **Group** field, select the group that contains the database with which the ctxsrv server is associated.
 - c. On the **Possible Owners** page, specify the nodes in the cluster on which the ctxsrv server can be brought online. These should be the same as the possible owners for the single-instance database with which the ctxsrv server is associated.
 - d. On the **Dependencies** page, specify the single-instance database with which the ctxsrv server is associated as a resource dependency. If the ctxsrv server has a disk dependency (other than those that the database requires), specify the disk or disks as resource dependencies also.
 - e. On the **Generic Application Parameters** page:
 - In the **Command line** field, enter the file specification for the batch file you created earlier. For example, context.bat.
 - In the **Current directory** field, enter the directory where the ctxsrv server was installed (for example, D:\Orant\bin).

- f. On the **Registry Replication** page, you need not enter any registry keys; click **Finish**.

After you complete the configuration, the cluster service starts up the .bat file and opens a command window to display ctxsrv server logging information. If someone closes the command window, another command window opens immediately and the ctxsrv server continues to search, as usual. If the group containing the ctxsrv server fails over, operations and the searching function continue as usual.

7.7 Integrating with Oracle Enterprise Manager

You can use Oracle Enterprise Manager to manage and monitor single-instance databases in an Oracle Fail Safe environment. For example, you can use Oracle Enterprise Manager to:

- Discover clusters

For Oracle Enterprise Manager to discover Oracle Fail Safe clusters, you must edit the `nmiconf.lst` file and add `fs_discover.tcl` as the first entry in the list. This must be performed on all nodes of the cluster. If you are using Oracle Intelligent Agents from multiple Oracle homes, then add the `fs_discover.tcl` entry in one of the Oracle homes. The `nmiconf.lst` file is located in `Oracle_Home\NETWORK\AGENT\CONFIG`.

- Discover groups (virtual servers)

Note:

You must perform discovery on each group for Oracle Enterprise Manager to see the resources configured in that group. Once discovered, each group appears as a node in the Oracle Enterprise Manager nodes list, and you can manage the resources in the group as you would manage any standalone resource.

- Create and register jobs and events for a database in a group as you would for a standalone database
- Create and register jobs and events for a group as if it were a physical node (host)
Each group appears to Oracle Enterprise Manager as a host.

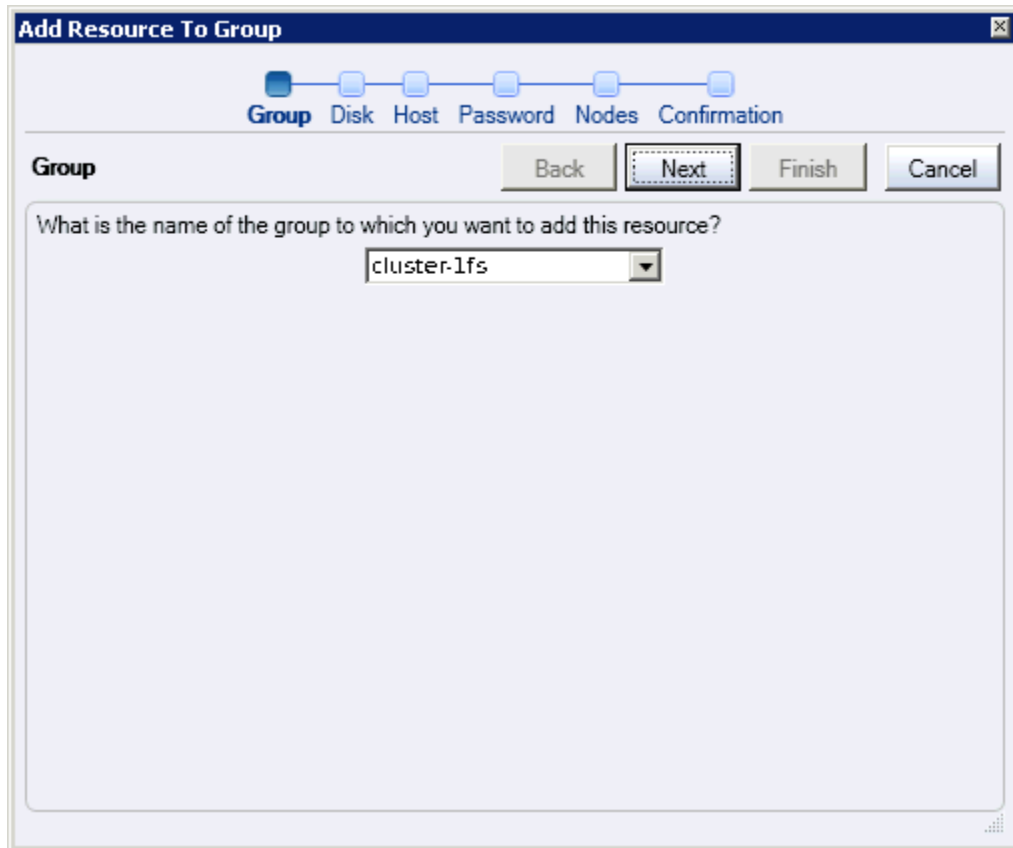
See Also:

[Security Access and Authentication Problems](#) (page 7-35) for information about troubleshooting problems related to integrating with Oracle Enterprise Manager.

7.7.1 Adding an Oracle Management Agent to a Group

After you add a single-instance database to a group, run the Add Resource to Group Wizard again to add an Oracle Management Agent to that group. The wizard displays a dialog box similar to the one shown in [Figure 7-10](#) (page 7-26).

Figure 7-10 Add Resource to Group Wizard - Resource Page



You add only one Oracle Management Agent to a group, regardless of the number of databases in the group. However, the group must contain at least one database resource before you can add the Oracle Management Agent. Similarly, you cannot remove the last database resource from a group without first removing the Oracle Management Agent.

When you specify that you want to add an Oracle Management Agent to a group:

- Oracle Fail Safe creates a new Management Agent
- The new Management Agent uses one of the cluster disks (specified when you ran the Add Resource to Group Wizard to add the Oracle Management Agent) to store its jobs and events information
- Oracle Fail Safe configures the new Management Agent to be a part of the group:
 - The Oracle Management Agent is configured to listen on the virtual address associated with the database.
 - The Oracle Management Agent fails over with the group.

See the Oracle Fail Safe Help for information about scheduling jobs for Oracle databases configured in a cluster and for monitoring events (such as failovers) using Oracle Enterprise Manager.

7.8 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 PowerShell cmdlets) to carry out a planned failover, Oracle Fail Safe 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 Microsoft Windows Failover Clusters to carry out a planned failover, then it does not checkpoint the database before shutting it down.)

 **Note:**

If you use a tool other than Oracle Fail Safe Manager, Oracle Fail Safe PowerShell cmdlet, or Windows failover clusters 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.

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

Perform administrative tasks on a database configured for high availability as you would for any database, with one exception. Use Oracle Fail Safe Manager or the PowerShell cmdlets command-line interface (see [PowerShell Commands](#) (page 5-1)) 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 cold backup operations, administrative operations that must be performed while users continue to access the database, and operations that could affect response times during the periodic Is Alive polling by Microsoft Windows Failover Clusters.

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 Fail Safe PowerShell cmdlets 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. Use a tool such as SQL*Plus to shut down the database once you complete the administrative tasks.
4. Use Oracle Fail Safe Manager or the Fail Safe PowerShell cmdlets to place the database online again. The cluster will resume monitoring the database.

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 **Validate** group operation. Adding a new data file can introduce a new disk dependency in the group. When you run the **Validate** 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.

7.10 Database Homes

Starting with Oracle Database 12c Release 1 (12.1), Oracle Database supports the use of Oracle Home User specified at the time of installation. Oracle Home User must be a domain user account. Oracle Home User is associated with an Oracle home. Ensure that all nodes in a cluster that use the same Oracle home use the same Oracle Home User.

When Oracle Fail Safe accesses a database, it usually uses the same Oracle Database home to access any database on the system. Oracle Fail Safe selects the database home 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 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 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 its `\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 (FsResOdb.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 when a database or a 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 [Oracle Fail Safe Services](#) (page 4-4) for information regarding the user accounts used by Oracle Fail Safe components when accessing databases.

7.11 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 applications](#) 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.

7.12 Handling Errors and Troubleshooting Problems with Databases

The following sections describe how to handle errors that occur when Oracle Fail Safe tries to bring a highly available single-instance database online. They also describe troubleshooting specific problems with single-instance Oracle Databases configured for high availability.

7.12.1 Handling Errors That Occur When Bringing a Database Online

You can create 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\FailSafe\Server\scripts
```

If Oracle Fail Safe cannot bring a single-instance database online, first 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 running the script waits for the script to finish within the 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.

7.12.2 Troubleshooting Problems

In most cases, the first step in troubleshooting a problem is to select the **Validate** cluster, **Validate** group, or **Validate** standalone database action. These tools are

described in general in [Validating Actions](#) (page 6-1). If the **Validate** actions 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 [Contacting Oracle Support Services](#) (page A-1) for more information on enabling Oracle Fail Safe tracing.

When you select **Validate** action for 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 **Validate** group action prompts you before fixing the problem.
- Detects disk drive changes and updates resource dependencies, if necessary.
- Ensures that the Oracle Net configurations are correct.
- Repairs any misconfigured resources in the group.

You can select the **Validate** group action 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.

7.12.3 Problems Adding a Database to a Group

While adding a database to a group, detailed error information might not be displayed by the Oracle Fail Safe Manager when a listener or database resource fails to come online. The Oracle Fail Safe resource control manager will log error information in the Windows application event log. For additional error information, refer to [Contacting Oracle Support Services](#) (page A-1).

7.12.4 Problems Placing a Group Online

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

- Validate the group.

When you select **Validate** (from the Oracle Fail Safe Manager group **Actions** menu), Oracle Fail Safe checks the group configuration and attempts to fix any problems that it finds.

If the **Validate** group action finds a problem, then it returns an error message that should help you resolve the problem manually.

- 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 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 databases are in the group, then this helps you to identify the database causing the problem.

- Ensure that the single-instance database **Pending Timeout** value specified in the Advanced Policies property page 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.

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 **Advanced 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 (`initdatabase-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 [Changing the SYS Account Password](#) (page 7-22) for information about how to update database password changes for Oracle Fail Safe.

7.12.5 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 `(Get-OracleClusterResource dbname).IsAliveEnabled=$false` command. However, be aware that when you disable Is Alive polling, Oracle Fail Safe suspends monitoring the instance until Is Alive polling is reenabled. You reenable Is Alive polling with the `(Get-OracleClusterResource dbname).IsAliveEnabled=$true` command.

Oracle recommends that you run these PowerShell cmdlets 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 PowerShell cmdlets commands, see [PowerShell Commands](#) (page 5-1).

7.12.6 Problems with 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 database authentication has been set up. Try the following to solve the problem:

- If you select **Use SYS account** for authenticating database in the General property page of Oracle Database, then ensure that the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter in the database initialization parameter file (`initdatabase-name.ora`) is set to `SHARED` or `EXCLUSIVE`.

[Security Requirements for Single-Instance Databases](#) (page 7-20) 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 it must be updated 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 Application Event log.

[Changing the SYS Account Password](#) (page 7-22) describes how to correctly update the database password.

7.12.7 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 **Validate** group and **Validate** standalone database operations to help you verify and repair the Oracle Net configuration. See [Validating the Configuration of Oracle Resources](#) (page 6-3) and [Validating Standalone Database](#) (page 6-5) for details.

7.12.7.1 Problems Configuring the Network Name

Oracle Fail Safe changes the `listener.ora` and `tnsnames.ora` files, and stops and starts listeners when configuring the network name 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. 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-10436 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 an address. 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)
    )
  )
```

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

- Another common cause for failing to start a listener is the network name. The network name 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.

7.12.7.2 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 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 retains previous versions of the configuration files. When a file is updated, the previous version of the file is renamed to `filename.timestamp.ora`. Note that `filename.ora` is the most recent file.

7.12.8 Security Access and Authentication Problems

Most authentication problems are due to inconsistencies in copies of the database password file between the cluster nodes. If Fail Safe is able to access the database on one node but not on another, then the password files have lost synchronization. You can resolve this issue in a number of ways:

- Copy the password file from the node which works to the node that does not work.
- Remove the database from the cluster using the Oracle Fail Safe Manager. Remove the command and add it again using the Add Resource command.
- Manually update the SYS password on the failing nodes by using the ORAPWD utility.

If operating system authentication is being used, then the Oracle authentication local user groups may be inconsistent. To verify the user groups, run Oracle Fail Safe Manager and select the Validate command for the cluster. Oracle Fail Safe will compare the contents of all the Oracle authentication local user groups on all nodes of the cluster and will report any inconsistencies found.

7.12.9 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. Select the **Validate** group action to validate the network (Oracle Net) configuration.

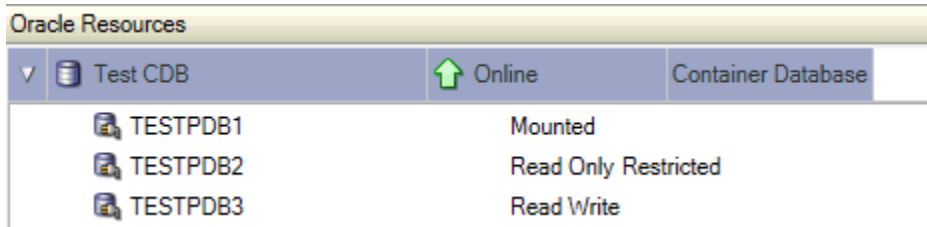
7.13 Starting and Stopping Pluggable Databases

Oracle Fail Safe will recognize that a database is a root container and will start and stop individual pluggable databases owned by the container database. When a database is failed over or moved to another node in the cluster, Oracle Fail Safe will start each pluggable database using the state that was saved by the last `SQL ALTER PLUGGABLE DATABASE ALL SAVE STATE` command. Oracle database 12c patch set 1 (12.1.0.2) is required to have the ability to save the state of the pluggable databases.

You can use the Oracle Fail Safe Manager to open and close individual pluggable databases in a container database. Click on the expansion icon next to the name of the container database to list the pluggable databases owned by the container

database. To open or close a pluggable database, click on the desired database in the list, then choose the open or close action in the Actions list.

Figure 7-11 Example of a Pluggable Database



Oracle Resources	
Test CDB	Online Container Database
TESTPDB1	Mounted
TESTPDB2	Read Only Restricted
TESTPDB3	Read Write

 **Note:**

The container database cluster resource must be online before the Oracle Fail Safe Manager displays any information about pluggable databases owned by the container database.

Oracle Fail Safe does not support the use of triggers to open or close pluggable databases. For example, a trigger similar to the following should not be used for a database in a failover cluster:

```
CREATE OR REPLACE TRIGGER sys.after_startup
AFTER STARTUP ON DATABASE
BEGIN
EXECUTE IMMEDIATE 'ALTER PLUGGABLE DATABASE ALL OPEN';
END after_startup;
```

A

Contacting Oracle Support Services

This appendix discusses the following topics:

- [Reporting a Problem](#) (page A-1)
- [Finding the Version of Oracle Software](#) (page A-2)
- [Viewing Error Information](#) (page A-2)
- [Tracing Oracle Fail Safe Problems](#) (page A-2)
- [Locating Trace and Alert Files](#) (page A-5)

A.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.
- 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 [Tracing Oracle Fail Safe Problems](#) (page A-2) for information about using the trace function to log error output to a file.

A.2 Finding the Version of Oracle Software

To find the version of software that you run in the Oracle Fail Safe Manager help menu, select **Help** in the menu bar, 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.

A.3 Viewing Error Information

Oracle Fail Safe Manager error messages are saved in three ways. They are as follows:

- **Progress Window:** This window displays the error messages to the user. Select **Save As** button to save the contents of the output window to a file that has more details, such as error numbers, timestamps, and version information.
- **Windows Application Event Log:** Oracle Fail Safe resource monitor -- the cluster component that starts, stops, and monitors Oracle cluster resources -- posts error information in the Windows Application Event Log. Check that log if errors are encountered related to starting, stopping, or Is Alive polling of Oracle cluster resources.
- **Oracle Fail Safe trace files:** Oracle Fail Safe logs more detailed information in these files that may provide clues to help determine the cause of errors.

A.4 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.

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 Fail Safe value from [Table A-1](#) (page A-3).
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 Fail Safe strings shown in [Table A-1](#) (page A-3). Separate multiple entries with commas.
8. Repeat steps 3 through 7 to set additional Oracle Fail Safe trace flags.

9. Restart the Cluster Service on each node if you set the `FSR_TRACE_OUTPUT` value. Note that, stopping the Cluster Service will cause all cluster resources to fail over to another node in the cluster. Restart Oracle Fail Safe cluster resource if you set the `FSS_TRACE_OUTPUT` value.

Table A-1 Trace Flags for Cluster Server Nodes

Value	String	Description
<code>FSR_TRACE_OUTPUT</code>	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. Oracle Fail Safe resource DLL starts, stops, and monitors Oracle resources in a cluster. For example: <code>C:\ORACLE_BASE\diag\FailSafe\fsr_trace.log</code>
<code>FSS_TRACE_OUTPUT</code>	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 server. Server errors that occur when executing commands sent by the Oracle Fail Safe Manager client are written to this file. For example: <code>C:\ORACLE_BASE\diag\FailSafe\fss_trace.log</code>
<code>FSS_TRACE_FLAGS</code>	<code>AGENT</code>	Logs information related to Oracle Management Agent activity.
<code>FSS_TRACE_FLAGS</code>	<code>ALL</code>	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.
<code>FSS_TRACE_FLAGS</code>	<code>COM</code>	Logs activity related to the Microsoft DCOM interface.
<code>FSS_TRACE_FLAGS</code>	<code>COMMAND_RESULT</code>	Logs information related to spawned commands.
<code>FSS_TRACE_FLAGS</code>	<code>COMMON</code>	Logs information that would be common to all Oracle Fail Safe components, such as error logging or work item processing.
<code>FSS_TRACE_FLAGS</code>	<code>CREATE_SA</code>	Logs information related to the creation of standalone resources, such as a sample database.
<code>FSS_TRACE_FLAGS</code>	<code>CR_RES</code>	Logs information related to the creation of cluster resources.
<code>FSS_TRACE_FLAGS</code>	<code>CLUSTER_MGR</code>	Logs information related to the Microsoft Windows Failover Clusters cluster interface.
<code>FSS_TRACE_FLAGS</code>	<code>DB_RES</code>	Logs information related to database access by the server or resource monitor DLL.
<code>FSS_TRACE_FLAGS</code>	<code>DEL_RES</code>	Logs information related to the deletion of a cluster resource.
<code>FSS_TRACE_FLAGS</code>	<code>DELETE_SA</code>	Logs information related to the deletion of a standalone resource, such as a sample database.

Table A-1 (Cont.) Trace Flags for Cluster Server Nodes

Value	String	Description
FSS_TRACE_FLAGS	HOME	Logs information related to the processing of Oracle homes.
FSS_TRACE_FLAGS	LOCAL_TRACE	<p>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 Fail Safe is running (the node where the Cluster Group resides).</p> <p>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.</p>
FSS_TRACE_FLAGS	SQLNET	Generates detailed internal information related to the Oracle Net configuration performed by Oracle Fail Safe. 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.
FSS_TRACE_FLAGS	VERIFY_CLUSTER	Logs information about the VERIFY_CLUSTER operation.
FSS_TRACE_FLAGS	VERIFY_GR	Logs information about the Verify Group operation.
FSS_TRACE_FLAGS	VERIFY_SA	Logs information regarding verification of standalone resources.
FSS_TRACE_FLAGS	VERIFY_DB	Logs information about the Validate operation.
FSS_TRACE_FLAGS	XML	Logs activity related to the exchange of XML messages between Oracle Fail Safe components.
FSU_TRACE_OUTPUT	A path and file name	<p>Specifies the path and file name for the file to which you want tracing information about the Oracle Fail Safe surrogate to be written. Server errors that occur when executing commands sent by the Oracle Fail Safe Manager client are written to this file. This file is used on the nodes that do not own the Cluster Group. For example:</p> <pre>C:\ORACLE_BASE\diag\FailSafe \fsu_trace.log</pre> <p>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.</p>

 **Note:**

Oracle recommends using `ALL` for `FSS_TRACE_FLAGS`.

A.5 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 is specified. 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 Windows Failover Clusters. 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" or "clustered role".

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 Windows Failover Clusters, 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 Windows Failover Clusters

Microsoft Windows Failover Cluster provides the capability to cluster individual nodes that run supported Windows operating systems. See *Oracle Fail Safe Release Notes for Microsoft Windows* for a list of the supported operating system releases. See also [Microsoft Windows Failover Clusters](#) and [Microsoft Windows Failover Clusters](#).

Microsoft Windows Failover Clusters

In Windows Server 2003, Microsoft Cluster Services (MSCS) is now known as Microsoft Windows Failover Clusters. See also [Microsoft Windows Failover Clusters](#) and [Microsoft Windows Failover Clusters](#).

Microsoft Windows Failover Clusters

In Windows Server 2008, Microsoft Cluster Services (MSCS) is now known as Microsoft Windows Failover Clusters. See also [Microsoft Windows Failover Clusters](#) and [Microsoft Windows Failover 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.

partitioned workload

A configuration in which each node of the cluster performs different work. For example, in a two-node cluster, one node could serve as a database server and another node could serve as an application server.

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 for Microsoft Windows* 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](#).

network name

A network address at which resources in a group can be accessed, regardless of the cluster node hosting those resources. A network name on Windows Failover Cluster consists of a network name and associated IP address. A network name is sometimes referred to as a client access point.

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 network name 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 network names.

Index

A

access to resources, [2-4](#)

accounts

- adding to database password file, [7-21](#)
- privileges and permissions, [4-4](#)
- to manage a single-instance Oracle Database, [7-20](#)

active/active configurations, [3-3](#)

active/passive configurations, [3-1](#)

Add Resource to Group Wizard

- Oracle Intelligent Agents, [7-25](#)
- Oracle single-instance databases, [7-5](#)

adding a ctxsrv server to a group, [7-24](#)

adding a single-instance Oracle database to a group

- password, [7-13](#)

adding a single-instance Oracle Database to a group, [7-4](#)

- authentication information, [7-11](#)
- network names, [7-8](#)
- Oracle Net listener, [7-16](#)
- possible owner nodes, [7-6](#)
- prerequisites, [7-4](#)
- steps, [7-4](#)
- tnsnames.ora file, [7-5](#)

adding an Oracle Intelligent Agent to a group, [7-25](#)

Administrator privileges

- logging in to a cluster, [4-6](#)

alert files, [A-5](#)

allocating IP addresses for a cluster, [2-7](#)

application log

- troubleshooting startup problems, [6-9](#)

application software

- rolling upgrade of, [1-3](#)

applications

- automatic reconnection after failover, [1-6](#)
- failover, [1-6](#), [2-24](#)
- mission-critical, [3-4](#)

archived files, [7-35](#)

authentication

- permissions and privileges required, [4-4](#)
- REMOTE_LOGIN_PASSWORDFILE initialization parameter, [7-21](#)

authentication (*continued*)

- SYSDBA role, [7-21](#)
- SYSOPER privileges, [7-21](#)
- troubleshooting DBA authentication, [7-33](#)

B

bandwidth

- increasing with multiple network names, [2-6](#)

benefits

- of Oracle Fail Safe, [1-2](#)

C

checkpoints

- prior to database failover, [2-13](#)

client connections

- to single-instance Oracle Databases, [7-16](#)

clients

- access to resources, [2-6](#)
- accessing groups with multiple network names, [4-8](#)
- cluster connections, [2-24](#)
- connecting to resources using multiple virtual addresses, [4-7](#)
- reconnecting to a single-instance database after failover, [7-29](#)

cluster alias

- definition, [2-8](#)
- Oracle Fail Safe Manager and, [2-8](#)
- use of, [2-8](#)

cluster disks, [7-4](#)

- Oracle single-instance databases and, [7-4](#)
- RAID hardware and, [1-7](#)
- redundancy and, [2-2](#)
- storing jobs and events information, [7-26](#)
- See also* disks

Cluster Group, [2-8](#)

- cluster alias network name, [2-8](#)
- resources in, [2-8](#)

cluster nodes, [2-1](#)

- adding to an existing cluster, [4-8](#)
- defined, [1-1](#)
- shared servers configuration and, [7-18](#)

cluster registry

- cluster registry (*continued*)
 - run Validate group to correct, [7-28](#)
 - cluster resources
 - defined, [2-3](#)
 - clusters, [1-1](#)
 - adding a new node to, [4-8](#)
 - cluster alias, [2-8](#)
 - configuration, [2-1](#)
 - connecting to, [2-8](#)
 - definition, [2-1](#)
 - different configurations, [2-2](#)
 - discovering in Oracle Enterprise Manager, [7-25](#)
 - disk configurations, [2-3](#)
 - disks
 - See cluster disks, [7-4](#)
 - introduction, [2-1](#)
 - members of, [2-1](#)
 - nodes, [1-2](#), [2-1](#)
 - synchronizing password files on, [7-21](#)
 - typical configuration, [1-6](#)
 - verifying the configuration, [6-1](#)
 - clusterwide operations
 - Verify Group operation, [6-3](#)
 - communications
 - lost when a system fails, [7-35](#)
 - configurations, [3-1](#)
 - active/active, [3-3](#)
 - active/passive, [3-1](#)
 - customizing, [3-1](#)
 - disk-level, [2-3](#)
 - multiple virtual addresses in, [4-7](#)
 - shared-nothing, [2-3](#)
 - system-level, [2-2](#)
 - typical, [1-6](#)
 - using wizard input, [4-2](#)
 - verifying, [6-3](#), [6-5](#)
 - connect to cluster
 - domain user account, [4-6](#)
 - corruption
 - tnsnames.ora and listener.ora files, [7-34](#)
 - ctxsrv servers
 - configuring as a generic application, [7-24](#)
 - configuring for high availability, [7-24](#)
 - custom resource types, [2-5](#)
 - customizing configurations, [3-1](#)
- D**
-
- data files
 - adding new, [7-28](#)
 - database administrator
 - changing the password, [7-33](#)
 - connecting as internal, [7-33](#)
 - DBA_AUTHORIZATION parameter, [7-21](#)
 - database password file
 - adding accounts to for a database, [7-21](#)
 - database recovery
 - optimization for single-instance databases, [7-27](#)
 - databases, [7-1](#)
 - adding a new data file, [7-28](#)
 - backup operations on, [7-27](#)
 - changing DBA passwords, [7-33](#)
 - checkpointing and, [2-13](#)
 - clients lose single-instance database connection, [7-35](#)
 - configuration data, [7-5](#)
 - configuring database resource to use shared servers, [7-18](#)
 - configuring using wizard input, [4-2](#)
 - EXCLUSIVE access, [7-21](#)
 - group resource, [2-4](#)
 - identity, [7-9](#)
 - Is Alive polling, [7-33](#)
 - LOCAL_LISTENER parameter in the initialization file, [7-20](#)
 - network names, [7-8](#)
 - Oracle7
 - upgrading from, [7-23](#)
 - Oracle8 using shared servers, [7-18](#)
 - Oracle8i or later using shared servers, [7-20](#)
 - security, [7-20](#)
 - SHARED access, [7-21](#)
 - shared servers and Windows Failover Cluster nodes, [7-18](#)
 - steps for configuring, [7-4](#)
 - synchronizing password files, [7-21](#)
 - taking offline, [7-20](#), [7-28](#)
 - due to change in password, [7-33](#)
 - for cold backup operation, [7-27](#)
 - Pending Timeout, [7-32](#)
 - problems, [7-33](#)
 - to stop and restart the group, [7-20](#)
 - testing shutdown on the secondary node, [4-3](#)
 - upgrading with Oracle Database Upgrade Assistant, [7-23](#)
 - using shared servers, [7-18](#)
 - using SQL*Plus for cold backup operations, [7-28](#)
 - verifying, [6-3](#), [6-5](#)
 - See also resources
 - DBA_AUTHORIZATION parameter, [7-21](#)
 - dependencies
 - among resources, [2-4](#), [2-5](#)
 - DISABLEISALIVE parameter, [7-33](#)
 - discovering
 - resources, [4-6](#)
 - standalone single-instance databases, [7-1](#)
 - discovery

discovery (*continued*)
 with Oracle Enterprise Manager, 7-25

disk devices
 detecting changes with Verify Group, 7-31
 resources in a group, 2-4
 verification after adding more, 7-31

disks, 7-4
 configuration, 2-3
 resource type, 2-5
 See also cluster disks

DISPATCHERS parameter
 specifying full listener information in, 7-20
 specifying information in, 7-20

domain account
 for Oracle Fail Safe, 4-4
 for Oracle Fail Safe Manager, 4-4

dynamic-link libraries (DLLs)
 custom, 2-5
 FsResOdbbs.dll, 2-5
 generic services, 2-5
 IP addresses, 2-5
 managing resources, 2-5
 physical disks, 2-5

E

ENABLEISALIVE parameter, 7-33

errors
 FS-10070, 7-34
 FS-10436, 7-34
 tracking information in a trace file, A-2
 written to the Oracle Net listener log, 7-32

events
 storing information about, 7-26

F

failback
 effect of preferred owner nodes list on, 2-23
 move group operation effect on, 2-24

failback policy
 dumping information about, 6-7
 specifying, 2-22
 verifying with Validate group action, 6-3
 verifying with Verify Group operation, 7-31

failover period, 2-18

failover policy
 for groups, 2-18

failover threshold, 2-18

failovers, 2-9
 applications, 1-6
 automatic application reconnection, 1-6
 checking with the Verify command, 4-3
 client applications and, 2-24
 database checkpointing and, 2-13

failovers (*continued*)
 definition, 2-1
 disabling during backup operations, 7-27
 due to resource failure, 2-10
 dumping information about policies, 6-7
 effect of possible owner nodes list on, 2-21
 effect of preferred owner nodes list on, 2-21
 effect of resource restart policy on, 2-19
 fastest, 3-3
 for load-balancing, 2-9
 groups and, 2-4
 in an active/active configuration, 3-3
 limiting repeated, 2-18
 node failures and, 2-11
 node to which group moves during, 2-21
 Oracle Intelligent Agent resource, 7-26
 planned, 2-9, 2-12
 replaying transactions and, 1-6
 troubleshooting, 7-33
 unit of, 2-4
 unplanned, 2-9
 verifying with Validate group action, 6-3
 verifying with Verify Group operation, 7-31

failures
 unplanned, 2-9
 verifying resources to prevent, 6-3, 6-5

Fibre Channel
 shared storage interconnect, 1-6

file corruption
 in tnsnames.ora and listener.ora files, 7-34

files
 alert, A-5
 archived, 7-35
 trace, A-5

FSCMD command, 1-4
 DISABLEISALIVE parameter, 7-33
 ENABLEISALIVE parameter, 7-33

FSR_TRACE_OUTPUT value, A-3

FsResOdbbs.dll file
 functions, 2-5

FSS_TRACE_OUTPUT value, A-3

G

generic services
 defined, 2-5
 for Windows services, 2-5
 resource type, 2-5

GR_VERIFY string, A-4

groups
 added bandwidth for, 2-6
 adding a single-instance database to, 7-4
 adding an Oracle Intelligent Agent to, 7-25
 adding network names to, 2-6
 correcting the network name, 7-34

groups (*continued*)

- DBA password mismatches, [7-33](#)
- definition of, [2-4](#)
- discovering in Oracle Enterprise Manager, [7-25](#)
- dumping information about, [6-7](#)
- examples of resources, [2-4](#)
- failover, [2-4](#)
- failover period, [2-18](#)
- failover policy, [2-18](#)
- failover threshold, [2-18](#)
- multiple virtual addresses in, [4-7](#)
- network names and, [7-5](#)
- Oracle homes, [4-7](#)
- ownership of, [2-4](#)
- populating, [4-2](#)
- preferred owner nodes list, [2-20](#)
- resource dependencies, [2-5](#)
- verifying resources in, [6-3](#)
- with multiple virtual addresses, [4-8](#)

H

hardware

- configuration, [1-6](#)
- RAID, [1-7](#)

heartbeat connection, [2-1](#)

- definition of, [1-7](#)

I

initialization parameter file,

- deleting the LOCAL_LISTENER parameter, [7-20](#)
- exporting SPFILE and, [7-11](#)
- location of, [7-10](#)
- PFILE and, [7-11](#)
- REMOTE_LOGIN_PASSWORDFILE
 - initialization parameter, [7-21](#), [7-33](#)
- restrictions on specifying SPFILE and, [7-11](#)

installation

- interactive, [1-3](#)
- of Oracle Fail Safe, [1-3](#)
- silent mode, [1-3](#)
- using a response file with, [1-3](#)

interconnects

- private internode connection, [2-1](#)
- shared storage, [2-1](#)

intracluster communications, [1-7](#)

IP addresses

- allocating for a cluster, [2-7](#)
- determining number needed, [2-7](#)
- resource type, [2-5](#)

Is Alive polling

- disabling, [7-33](#)

Is Alive polling (*continued*)

- DLL file function, [2-5](#)
- effect during backup operations, [7-27](#)
- enabling, [7-33](#)

J

jobs

- storing information about, [7-26](#)

L

listener.ora files

- archived, [7-35](#)
- example, [7-34](#)
- modifying, [7-34](#)
- sample definition, [7-34](#)

listeners

- changes for shared servers, [7-2](#)
- configuring, [4-2](#)
- configuring Oracle Net on nodes with multiple listeners, [7-3](#)
- created for single-instance databases added to a group, [7-16](#)
- log file, [7-32](#)
- network resources, [2-4](#)
- problems after changing or creating definition of, [7-34](#)
- specifying in the LOCAL_LISTENER parameter, [7-20](#)
- specifying information in the DISPATCHERS parameter, [7-20](#)
- troubleshooting, [7-32](#)

load balancing, [2-9](#)

- in an active/active configuration, [3-5](#)
- static, [2-12](#)

local area network, [1-7](#)

LOCAL_LISTENER parameter

- adding to single-instance database parameter file, [7-2](#)
- deletion after removing a single-instance database from a group, [7-20](#)
- specifying listeners in, [7-20](#)
- updating information for the group, [7-20](#)

M

management operations

- single-instance database, [7-27](#)

managing

- cluster security, [4-4](#)
- integrating with Oracle Enterprise Manager, [7-25](#)
- tnsnames.ora file on client systems, [7-5](#)

managing (*continued*)
 using a cluster alias, [2-8](#)

Microsoft Windows Failover Clusters
 resource types, [2-5](#)

mission-critical applications
 in active/active configurations, [3-4](#)

move group operation
 effect on failback, [2-24](#)

multiple Oracle homes, [4-7](#)
 configuring multiple listeners, [7-3](#)
 single-instance database SID in listener.ora
 files, [7-34](#)

multithreaded servers
 See shared servers, [7-18](#)

N

net service name
 definition
 referencing in archived configuration
 files, [7-35](#)
 entry after adding a single-instance database
 to a group, [7-16](#)
 when configuring the virtual server, [7-34](#)

network name, [1-2](#)
 cluster alias, [2-8](#)
 correcting for group, [7-34](#)

network names
 groups and, [7-5](#)
 identifying, [2-6](#)

networks
 configuration
 for highly available single-instance
 databases, [7-16](#)
 Oracle Net on nodes with multiple
 listeners, [7-3](#)
 verify with Verify Group, [7-31](#)
 dumping public and private information, [6-7](#)
 problems while configuring virtual server,
[7-34](#)
 protocol information in tnsnames.ora file, [4-3](#)
 tracing configuration information, [A-4](#)

node
 to which a group fails over, [2-21](#)

node failures
 effects of, [2-24](#)
 failover and, [2-11](#)

nodes
 adding to an existing cluster, [4-8](#)
 cluster, [1-2](#)
 definition of, [2-1](#)
 possible owner, [2-15](#)
 preferred, [2-22](#)
 preferred ownership of groups and, [2-20](#)

O

operating system authentication, [7-11](#)

ORA_DBA
 Windows operating system group, [7-11](#)

ORA_sid_DBA
 Windows operating system group, [7-11](#)

Oracle Database Upgrade Assistant
 upgrading single-instance databases, [7-23](#)

Oracle Enterprise Manager
 discovering clusters in, [7-25](#)
 discovering groups in, [7-25](#)
 integrating with, [7-25](#)
 use in an active/active configuration, [3-5](#)

Oracle Fail Safe,
 installation of, [1-3](#)
 SID list entries, [7-3](#)

Oracle Fail Safe Manager, [1-1](#)
 cluster alias and, [2-8](#)
 compatibility with prior releases of Oracle Fail
 Safe Server, [4-7](#)
 introduction to, [1-3](#)
 tree view, [1-3](#)
 verify tools, [1-4](#)
 wizards, [1-4](#)

Oracle Fail Safe Server, [1-1](#)
 described, [1-1](#)
 Oracle Fail Safe Manager releases and, [4-7](#)
 See also Oracle Fail Safe

Oracle homes
 groups and, [4-7](#)
 multiple, [4-7](#)
 name of, [6-1](#)

Oracle Intelligent Agent
 adding to a group, [7-25](#)
 virtual addresses, [7-26](#)

Oracle Net Assistant
 using to re-create network files, [7-34](#)

Oracle Net listener
 behavior, [7-2](#)
 configuration of, [7-15](#)
 creation of, [7-15](#)

Oracle Net network
 archived configuration files, [7-35](#)
 resource in a group, [2-4](#)
 single-instance database configuration, [7-16](#)
 tracing configuration information, [A-4](#)
 verifying the configuration, [7-31](#)

Oracle Support
 calling, [A-1](#)

Oracle7 databases
 See databases

Oracle8i databases
 See databases

owner nodes

See possible owner nodes list
ownership
of group, [2-4](#)

P

parameter file
See initialization parameter file
password files
synchronizing, [7-21](#)
passwords
changes cause group problems, [7-32](#)
changing for the DBA account, [7-33](#)
problems with
 REMOTE_LOGIN_PASSWORDFILE
 parameter, [7-33](#)
Pending Timeout value
setting, [7-32](#)
troubleshooting, [7-32](#)
performance
in an active/active configuration, [3-4](#)
PFILE
See initialization parameter file
planned failover
static load balancing, [2-12](#)
planned group failover, [2-12](#)
planned maintenance, [2-12](#)
policies
dumping failover and failback information,
 [6-7](#)
polling
failures, [7-32](#), [7-33](#)
possible owner nodes list, [2-15](#)
effect on failover, [2-21](#)
effect on resources, [2-15](#)
for a single-instance Oracle Database, [7-6](#)
in determining failover node, [2-21](#)
resources and, [2-15](#)
Validate group and, [2-15](#)
Verify Group command and, [4-8](#)
PowerShell cmdlets
use in an active/active configuration, [3-5](#)
preferred owner nodes list
effect on failback, [2-23](#)
effect on failover, [2-20](#), [2-21](#)
move group operation and, [2-24](#)
private interconnect, [1-7](#), [2-1](#)
privileges
granting on each cluster node, [7-21](#)
required, [4-4](#)
protocol address, [7-16](#)
public interconnect, [1-7](#)

Q

quorum disk
dumping information about, [6-7](#)

R

RAID
hardware, [1-7](#)
redundancy
servers and, [2-2](#)
registering
jobs and events, [7-25](#)
registration
dumping information about resource, [6-7](#)
REMOTE_LOGIN_PASSWORDFILE initialization
parameter, [7-33](#)
setting, [7-32](#)
resource types
custom, [2-5](#)
generic service, [2-5](#)
IP addresses, [2-5](#)
Microsoft Windows Failover Clusters, [2-5](#)
Oracle MTS Service, [2-5](#)
physical disks, [2-5](#)
resources, [1-1](#)
access to, [2-4](#)
accessing through a virtual server, [2-6](#)
client access to, [2-1](#)
defined, [2-3](#)
dependencies, [2-4](#), [2-5](#)
discovering, [4-6](#)
disks used by the database configuration, [4-2](#)
DLLs, [2-5](#)
effect of group Nodes property page changes
on, [2-15](#)
examples of, [2-4](#)
failover policy effect on groups, [2-18](#)
failure of, [2-10](#)
Oracle Intelligent Agent, [7-25](#)
possible owner nodes list and, [2-15](#)
renaming, [4-7](#)
repairing misconfigured, [7-31](#)
restart policy, [2-19](#)
returning to a preferred owner node, [2-22](#)
verifying configuration of, [6-3](#), [6-5](#)
rolling upgrade
of application software, [1-3](#)

S

scripts
PowerShell cmdlets in an active/active
configuration, [3-5](#)

scripts (*continued*)
 to handle errors when single-instance database cannot be placed online, [7-30](#)

secondary node
 creating the database instance on, [4-2](#)
 testing database shutdown, [4-3](#)

security
 increasing with multiple network names, [2-6](#)
 requirements
 for single-instance databases, [7-20](#)
 synchronizing password files on nodes, [7-21](#)

server nodes
 attempting to restart, [2-24](#)
 cluster polling failure, [7-33](#)
 losing client connections, [7-35](#)

server redundancy, [2-2](#)

shared servers
 configuring for Oracle8 databases, [7-18](#)
 configuring for Oracle8i or later databases, [7-20](#)

shared storage interconnect, [1-2](#), [2-1](#)
 Fibre Channel, [1-6](#)

shared-nothing configuration, [2-3](#)

SID
 entering for single-instance database, [7-34](#)

silent mode installation, [1-3](#)

single-instance databases
 clients lose connection to, [7-35](#)
 configuring database resource to use shared servers, [7-18](#)
 handling errors for failures to place online, [7-30](#)
 management operations and, [7-27](#)
 possible owner nodes list for, [7-6](#)
 prerequisites to adding to a group, [7-4](#)
 upgrading, [7-23](#)

software configuration, [1-7](#)

SPFILE
 See initialization parameter file

SQL*Plus
 performing administrative tasks, [7-28](#)

SQLNET string, [A-4](#)

standalone databases
 discovering, [7-1](#)
 shared servers and, [7-2](#)
 Validate, [4-1](#), [6-5](#)
 verifying, [4-1](#), [6-5](#)

standalone resources
 discovering, [4-6](#)
 See also index entries for each resource, [4-6](#)

standby configurations, [3-1](#)

static load balancing, [2-12](#)

strings

strings (*continued*)
 Oracle Fail Safe tracing, [A-3](#), [A-4](#)
 SYSDBA privileges, [7-21](#)
 SYSOPER privileges, [7-21](#)

T

temporary tables
 failover and, [7-4](#)

tnsnames.ora files, [4-3](#)
 archived, [7-35](#)
 checking Oracle Net configuration data, [7-32](#)
 modifying for highly available single-instance databases, [7-16](#)
 problems configuring the virtual server, [7-34](#)
 single-instance Oracle Databases and, [7-5](#)
 updating, [7-16](#)

trace files, [A-5](#)

tracing
 enabling, [A-2](#)
 flags, [A-2](#)

transactions
 failover and, [1-6](#)

transparent application failover, [7-29](#)

tree view, [1-3](#)
 populating with clusters, [2-8](#)

troubleshooting
 databases, [7-30](#)
 startup problems, [6-9](#)
 unable to place a group online, [7-31](#)
 verification tools, [6-1](#)

U

unplanned failover, [2-9](#)

upgrading a single-instance database, [7-23](#)

user
 domain account for Oracle Fail Safe Manager, [4-6](#)

V

Validate group action
 purpose, [6-3](#)

Validate group operation
 overview, [6-3](#)
 possible owner nodes list and, [2-15](#)
 updating information in the cluster registry, [7-28](#)

Validate operation, [6-5](#)
 tracing, [A-4](#)

Verify Cluster operation
 purpose, [6-1](#)

Verify Group operation

Verify Group operation (*continued*)
 logging information about, [A-4](#)
 possible owner nodes list and, [4-8](#)
 updating resource dependencies, [7-31](#)

verify tools, [1-4](#)

VERIFY_DB string, [A-4](#)

verifying
 with groups, [7-32](#)

virtual addresses
 cluster alias, [2-8](#)
 definition, [2-6](#)
 for Oracle Intelligent Agents, [7-26](#)
 for single-instance Oracle Databases, [7-8](#)
 multiple, [2-6](#), [4-7](#)
 Oracle Intelligent Agent uses the, [7-26](#)
 reasons for multiple in a group, [4-8](#)
 troubleshooting configuration problems, [7-34](#)

virtual servers, [1-2](#)
 configuration, [2-6](#)
 configuration problems, [7-34](#)
 definition, [2-6](#)
 discovering, [7-25](#)
 listener failure and, [7-34](#)

virtual servers (*continued*)
 multiple, [4-7](#)
 multiple network names, [2-6](#)
 network configuration for, [7-16](#)
 updates to tnsnames.ora files during
 configuration of, [7-16](#)

W

wide area network, [1-7](#)

Windows cluster
 See clusters

Windows Failover Cluster
 cluster software, [1-1](#)
 Is Alive polling, [2-5](#)
 Pending Timeout value, [7-32](#)

Windows operating system groups
 ORA_DBA, [7-11](#)
 ORA_sid_DBA, [7-11](#)

wizards, [1-4](#)
 input, [4-2](#)

workloads
 setting Pending Timeout to accommodate
 heavy, [7-32](#)