



Resource Failover Using Oracle Fail-safe 2.0.5 and Microsoft Cluster Server

TechNote

**First Edition (March 1998)
Part Number 298219-001
Compaq Computer Corporation**

Notice

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL. THIS INFORMATION IS PROVIDED "AS IS" AND COMPAQ COMPUTER CORPORATION DISCLAIMS ANY WARRANTIES, EXPRESS, IMPLIED OR STATUTORY AND EXPRESSLY DISCLAIMS THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR PARTICULAR PURPOSE, GOOD TITLE AND AGAINST INFRINGEMENT.

This publication contains information protected by copyright. No part of this publication may be photocopied or reproduced in any form without prior written consent from Compaq Computer Corporation.

© 1998 Compaq Computer Corporation.
All rights reserved. Printed in the U.S.A.

The software described in this guide is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of the agreement.

Compaq, Deskpro, Compaq Insight Manager, ProLiant, ROMPaq, SmartStart, NetFlex, registered United States Patent and Trademark Office.

Netelligent is a trademark and/or service mark of Compaq Computer Corporation.

Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Resource Failover Using Oracle Failsafe 2.0.5 and Microsoft Cluster Server

First Edition (March 1998)
Part Number 298219-001

Contents

Chapter 1

About This TechNote

Introduction.....	1-1
Expected Audience	1-1

Chapter 2

Overview of a Failover

What is a Resource Failover?.....	2-1
What are the Uses of a Resource Failover?.....	2-1
What are the Advantages of a Resource Failover?	2-2
What are the Disadvantages of a Resource Failover?	2-2

Chapter 3

Model Resource Failover Environment

Introduction.....	3-1
Model Environment	3-3
Hardware Requirements.....	3-5
Server Configuration for the ProLiant Clusters	3-5
Disk Storage.....	3-5
Server Configuration for the Standalone Server / Client	3-5
Software Requirements	3-6
Server Configuration for the ProLiant Clusters	3-6
Server Configuration for the Standalone Server / Client	3-6

Chapter 4

Installation of Microsoft Cluster Server

Chapter 5

Preliminary Tasks

Pre-installation checklist	5-1
Editing the Windows NT Hosts files	5-1

Chapter 6

Oracle Server 7.3 Installation

Installing Oracle Workgroup Server on both Servers	6-1
Installing Oracle Client Products in the Workstation.....	6-3

Chapter 7

Oracle Failsafe 2.0.5 Installation

Updating the Required Support Files to 7.3.3.3.2.....	7-1
Installing Oracle Failsafe Server	7-2
Installing Oracle Failsafe Manager	7-3

Chapter 8

Oracle Enterprise Manager

Configuring SQL*Net to use the Enterprise Manager	8-1
Building the Repository and Discovering New / Existing Services.....	8-2

Chapter 9

Oracle Failsafe Manager

Launching the Failsafe Manager	9-1
Creating Failsafe Groups	9-3
Standalone Databases.....	9-5
Creating Sample Standalone Databases	9-6
Adding a Standalone Database to a Failsafe Group.....	9-8
Verify each Failsafe Group	9-11

Chapter 10

Using the New Repository

Creating a New Repository and Discovering the Failsafe Groups.....	10-1
--	------

Chapter 11
Partitioned Data and Partitioned Workload

Chapter 12
Testing the Failsafe Cluster

Chapter 13
Client and Server Considerations

A Re-connect is Still Needed	13-1
Data Loss and Data Integrity	13-3
Maintaining Data Integrity - Reliable Transactions	13-3

Appendix
Hardware Configuration

Server 1	A-1
Server 2	A-2

Glossary

Index

Chapter 1

About This TechNote

Introduction

The purpose of this TechNote is to provide our customers a practical and effective way of implementing a successful Compaq Cluster resource failover configuration for an Oracle database. This TechNote shows in particular how to setup Microsoft Cluster Server 1.0 and Oracle Failsafe Server 2.0.5 to deliver an Oracle database clustering solution.

Expected Audience

This TechNote is intended for anyone installing, configuring and administering a cluster running:

- Microsoft Windows NT Enterprise Server 4.0 w/ Service Pack 3
- Microsoft Cluster Server 1.0
- Oracle Workgroup Server 7.3.3 for NT or Oracle Server 7.3.4 for NT
- Oracle Failsafe Server 2.0.5

This TechNote contains technical material. It requires that the reader have access to manuals, knowledge or familiarity of the following:

- Installing and setting up Compaq Servers
- Installing and administering Microsoft Windows NT server
- Installing the Oracle Server and Oracle Client products
- Using the Oracle Enterprise Manager
- Setting up SQL*Net Connections

Chapter 2

Overview of a Failover

What is a Resource Failover?

Resource failover is a cluster strategy that allows you to obtain high levels of server availability. Although this can be achieved by a various number of ways, this TechNote will use the Microsoft Cluster Server 1.0 and Oracle Failsafe Server 2.0.5 to accomplish this.

Resource Failover clustering takes advantage of cluster-aware resource(s) optimized for Microsoft Cluster Server. Independent Software Vendors, such as Oracle, provide these cluster-aware resources with their product so that their products can run efficiently with Microsoft Cluster Server. These cluster-aware resources are optimized for the Microsoft Cluster Server.

A simpler failover model is the service failover, which takes advantage of generic services instead of cluster-aware resources. In a service failover configuration, a service is defined to Microsoft Cluster Server as a resource. The Failover Manager ensures that the service is running on only one member of the cluster at a given time. This configuration uses services that are not cluster-aware, but by the use of a - Generic Service - resource type, we can implement a service failover even without a cluster-aware resource.

What are the Uses of a Resource Failover?

Resource failover can be used both for planned and unplanned service outages. A planned outage occurs when one of the nodes in the cluster is taken offline for upgrades or maintenance. An unplanned outage occurs when one of the cluster nodes or the software on that node fails to function correctly. During both kinds of outages, a resource failover cluster strategy can move the services running on one of the cluster nodes to another cluster node almost transparently from the client systems.

What are the Advantages of a Resource Failover?

Resource failover provides the following benefits:

- a highly available computing environment
- allows you to take advantage of each individual node by running its primary function and at the same time serving as a back-up of the other server, and
- enables a system to have fault tolerance and fault recovery features

A service outage can be very costly. By reducing the probability of a service outage, resource failover minimizes the losses in almost any industry.

What are the Disadvantages of a Resource Failover?

During a failover, the performance of the surviving node may degrade when it takes the load of the failed node.

Chapter 3

Model Resource Failover Environment

Introduction

The resource failover model is configured using two Compaq ProLiant servers running Windows NT Enterprise Server 4.0 and Microsoft Cluster Server 1.0. Also installed on both servers was the following software:

- Oracle Workgroup Server 7.3.3 (you can also use Oracle Server 7.3.4 for NT)
- Updated Required Support Files 7.3.3.3.2, and
- Failsafe 2.0.5 Server

The goal of this model is to be able to move any of the Oracle resources running from one server to the other server. This provides the clients a highly available database server in the event of a server failure or maintenance. This resource failover model also shows (1) partitioned workload and (2) partitioned data, providing a more effective and efficient use of the cluster. The cluster server manages the server that runs the database services at any given time. The data files, control files, redo log files for each instance will be located in a shared drive where only one of the servers can use it at any given time.

It is important to remember that even if resources are shared by two servers, these resources (such as disks, database service, and IP address) are not actually shared at the same time. These resources are actually “owned” by only one server at a time, but any server can release that ownership so that it can be hosted on the other server.

3-2 Model Resource Failover Environment

Everything discussed above describes the clustering architecture that Microsoft Cluster Server uses – the “shared-nothing” cluster architecture. Oracle Failsafe 2.0.5 uses Microsoft Cluster Server and Oracle Workgroup Server 7.3.3 to implement clusters for an Oracle database. In a pure shared-nothing architecture, files are partitioned among the instances running on the nodes of a loosely coupled system, which you will use for your cluster. Each instance or node “owns” a subset of the shared data, and all access to this subset of the shared data is performed exclusively by the instance or node that owns it. This approach is simple and elegant. It employs a software architecture that directly maps to the hardware system and has the potential to provide scalable performance on loosely coupled systems.

The key function of the Oracle Failsafe Server includes the following:

- optimized failover response time
- specify proper dependencies for database and cluster resources
- verify proper configuration of Oracle Failsafe databases
- fully configure SQL*Net for each Oracle Failsafe database
- monitor Oracle Failsafe database to ensure they are running correctly
- coordinate complex tasks such as moving a standalone database into a Failsafe group

In addition to being able to failover a resource, you can also perform data partitioning, workload partitioning and static load balancing. Data partitioning and workload partitioning allows each node to serve independent workloads and databases. Thus allowing each node to perform independent tasks. As a result of workload partitioning, you can perform static load balancing, which improves the overall performance of the cluster by moving services from a heavily used node to a lightly used node.

This TechNote shows you how to build a database cluster that delivers workload partitioning, data partitioning, and static load balancing using Microsoft Cluster Server 1.0 and Oracle Failsafe Server 2.0.5.

Model Environment

Please keep the following tables accessible since we will use all of these in setting up the Cluster.

**Table 3-1
System and Cluster Name**

Name of Node 1	OFS2-NODE1
IP Address of Node 1	110.6.120.101 (subnet mask 255.255.0.0)
Interconnect IP of Node 1	110.4.120.101 (subnet mask 255.255.255.0)
Name of Node 2	OFS2-NODE2
IP Address of Node 2	110.6.120.102 (subnet mask 255.255.0.0)
Interconnect IP of Node 2	110.4.120.102 (subnet mask 255.255.255.0)
Name of Client	OFS2-CLIENT
IP Address of Client	110.6.120.103 (subnet mask 255.255.0.0)
Name of Cluster	OFS2-CLUSTER
IP Address of Cluster	110.6.120.110 (subnet mask 255.255.0.0)

**Table 3-2
Application Names**

Database Server	ORACLE v 7.3.3 Workgroup Server
Cluster Server / Manager	ORACLE Failsafe 2.0.5 with Required Support Files 7.3.3.3.2

3-4 Model Resource Failover Environment

Table 3-3
Cluster Groups and Resources

Resource Group Name	REPOSITORY
Disk E:	Physical Disk
IP Address	110.6.120.111 (subnet mask 255.255.0.0)
Network Name	OFS2-DB1
Instance ID	OFS1
Instance Name	OFS1db
Service ID	REPOSITORY.world
Service Name	Repository
Resource Group Name	CUSTOMER
Disk E:	Physical Disk
IP Address	110.6.120.112 (subnet mask 255.255.0.0)
Network Name	OFS2-DB2
Instance ID	OFS2
Instance Name	OFS2db
Service ID	CUSTOMER.world
Service Name	Customer
Resource Group Name	SALES
Disk F:	Physical Disk
IP Address	110.6.120.113 (subnet mask 255.255.0.0)
Network Name	OFS2-DB3
Instance ID	OFS3
Instance Name	OFS3db
Service ID	SALES.world
Service Name	Sales

Hardware Requirements

Important: For this cluster, a standalone server is used as the temporary server to hold the local repository and at the same time the client. Using a standalone server to hold the repository is not necessary to configure a Failsafe cluster if you already have an existing repository.

Server Configuration for the ProLiant Clusters

Two Compaq ProLiant 2500 servers, each with:

- 2 processors / 128 Mbytes RAM
- SMART-2/P Array Controller
- Compaq Netelligent 10/100 TX PCI UTP Network Interface Controller
- Compaq Fibre Channel Host Controller
- Compaq SCSI hard disk drives

Disk Storage

Compaq ProLiant Fiber Channel Disk Storage System F2 with:

- Compaq ProLiant Fibre Channel Array Controller
- Compaq SCSI hard disk drives (configured with 3 logical volumes)

Server Configuration for the Standalone Server / Client

Compaq ProLiant 850 Server with:

- 1 processor / 64 Mbytes RAM
- Compaq SCSI hard disk drives

Software Requirements

Server Configuration for the ProLiant Clusters

- Compaq SmartStart and Support Software Release 3.4 or later
- System Configuration Utility 2.38 or later
- Array Configuration Utility 1.20 or later
- Microsoft Windows NT Enterprise Server 4.0
- Microsoft Windows NT Service Pack 3
- Compaq NT Support Software Utility 2.03 or later
- Microsoft Cluster Server 1.0
- Oracle Workgroup Server 7.3.3 for NT
- Oracle Required Support Files 7.3.3.3.2
- Oracle Failsafe Server 2.0.5

Server Configuration for the Standalone Server / Client

- Compaq SmartStart and Support Software Release 3.4 or later
 - System Configuration Utility 2.38 or later
 - Microsoft Windows NT Server 4.0
 - Microsoft Windows NT Service Pack 3
 - Compaq NT Support Software Utility 2.03 or later
 - Oracle Workgroup Server 7.3.3 for NT
 - Oracle Required Support Files 7.3.3.3.2
-

Chapter 4

Installation of Microsoft Cluster Server

Microsoft Cluster Server should be installed on both servers before installing Oracle Failsafe. Detailed installation instructions are not provided in this document. Refer to *Introduction to Compaq ProLiant Clusters* for complete instructions. The following table provides the necessary information to complete the installation of Microsoft Cluster Server on both servers:

Table 4-1
Microsoft Cluster Server Installation Specifics

Domain Name	OFS2-DOMAIN
Cluster Name	OFS2-CLUSTER
Cluster IP address	110.6.120.110 (subnet mask 255.255.0.0)
Shared Disks	Disk E:, F:,G:
Cluster Admin User / Password	Administrator / <password>
Location of Cluster Executables	Default (C:\WinNT\System32)
Quorum Disk	Disk E:

Keep this information and everything from the model environment section ready. You will use this in installing the Microsoft Cluster Server.

Chapter 5

Preliminary Tasks

Pre-installation checklist

Before moving forward with the Oracle Failsafe installation, the following tasks should already be completed.

- Hardware configuration for each server
- Array Configuration for the Fiber Channel shared storage
- Windows NT Enterprise Server 4.0 is installed
- Service Pack 3 for Windows NT 4.0 is applied
- Compaq NT SSD 2.03 or later (NT Support Software Disk) is applied
- Fiber Channel Storage driver is installed and both servers can access the shared disks.
- Shared disks are configured and have the same drive letters on both nodes.
- Microsoft Cluster Server is installed and running.

Editing the Windows NT Hosts files

Note: Throughout this paper, the standalone server OFS2-CLIENT, is referred to and used as a client to the cluster nodes.

In addition to the cluster IP address 110.2.30.3 with a cluster name OFS2-CLUSTER that was used during the installation of the cluster, each instance of oracle needs an IP address and a network name. It is important to remember that the IP addresses of the clusters should be static. A DHCP server cannot provide the cluster IP addresses. However, the server IP addresses can be static or provided by a DHCP server.

5-2 Preliminary Tasks

For this test environment, static IP addresses are used for the cluster, servers and workstation. Add the IP addresses of the cluster and the servers to the HOSTS file of each server. Use the virtual server IP addresses below. Add the following IP addresses in **C:\Winnt\System32\Drivers\Etc\Hosts** file of OFS2-NODE1, OFS2-NODE2 and OFS2-CLIENT:

```
110.6.120.101OFS2-NODE1# first node
110.6.120.102OFS2-NODE2# second node

110.6.120.110OFS2-CLUSTER# cluster

110.6.120.111OFS2-DB1# repository database
110.6.120.112OFS2-DB2# customer database
110.6.120.113OFS2-DB3# sales database
```

The first two IP addresses are the static IP addresses of the nodes that you assigned during Windows NT server installation. These are the IP addresses where the client as well as the two servers will communicate over the LAN. The third IP address is for the cluster. The remaining IP addresses are used by each resource group that will be created later.

Please make sure that the **HOSTS** file is saved without any file extension. You should replicate this file to OFS2-CLIENT, OFS2-NODE1 and OFS2-NODE2.

After creating the **HOSTS** file, make sure that you can PING both cluster nodes from your client, and each node can ping the other node. Try to PING the cluster nodes from your client:

```
C:\> ping OFS2-NODE1
Pinging OFS2-NODE1 [110.6.120.101] with 32 bytes of data:
Reply from 110.6.120.101: bytes=32 time<10ms TTL=128
Reply from 110.6.120.101: bytes=32 time<10ms TTL=128
Reply from 110.6.120.101: bytes=32 time<10ms TTL=128
Reply from 110.6.120.101: bytes=32 time<10ms TTL=128
C:\> ping OFS2-NODE2
Pinging OFS2-NODE2 [110.6.120.102] with 32 bytes of data:
Reply from 110.6.120.102: bytes=32 time<10ms TTL=128
Reply from 110.6.120.102: bytes=32 time<10ms TTL=128
Reply from 110.6.120.102: bytes=32 time<10ms TTL=128
Reply from 110.6.120.102: bytes=32 time<10ms TTL=128
```

Repeat the same commands from every node on your cluster, including the workstation. It is important that the IP addresses that the PING command returns are the IP addresses specified for client connection.

Chapter 7

Oracle Failsafe 2.0.5 Installation

Updating the Required Support Files to 7.3.3.3.2

The Required Support Files 7.3.3.3.2 is necessary to run Failsafe 2.0.5.

1. From the Control Panel → Services, **stop all Oracle-related Services**.
2. Insert the Failsafe CD into the server, and click **No** when prompted to install Failsafe. Run **D:\73332\Setup.exe** from the Start → Run menu.
3. Select the language in which you want to run Oracle.
4. Enter your Company name and a **local** home directory, such as **C:\ORANT**.
5. A message to remind you to shutdown all Oracle services appears. Click **Yes** to continue.
6. Oracle Installer analyzes file dependencies and installs all necessary files in your home directory **C:\ORANT**. This takes a while depending on the speed of your CD-ROM drive.
7. Please read the message after the installation. Review the Release Notes in the Failsafe CD. Click **OK** to finish the install.

Install the Required Support Files 7.3.3.3.2 to all machines that will run either Failsafe Server or Failsafe Manager.

Installing Oracle Failsafe Server

Important: Install Microsoft Cluster Server before installing the Failsafe Server. The Failsafe Server should be installed on both nodes of the cluster, while the Failsafe Manager should be installed on the client. For this cluster, install the Failsafe Manager on all the machines.

Follow these steps to install the Failsafe Server:

1. Insert the Failsafe CD on the server, and click **No** when prompted to install Failsafe. Run **D:\Win32\Install\Setup.exe** from the Start → Run menu.
 2. Select the language in which you want to run Oracle.
 3. Enter your Company name and a **local** home directory, such as **C:\ORANT**.
 4. The Software Asset Manager appears. Select **Failsafe Manager 2.0.5** and **Failsafe Server 2.0.5**. Click **Install**.
 5. Select the Failsafe Documentation that you wish to install.
 6. A message displays saying an upgrade to the Enterprise Manager will be needed. Click **OK** to continue.
 7. Oracle Installer analyzes file dependencies and installs all necessary files in your home directory **C:\ORANT**. This takes a while depending on the speed of your CD-ROM drive.
 8. Please read the message after the installation. Review the Release Notes in the Failsafe CD. Click **OK** to finish the install.
 9. **Restart** the system.
-

Installing Oracle Failsafe Manager

The Oracle client products and the Required Support Files 7.3.3.3.2 should have already been installed before installing the Failsafe Manager.

1. Insert the Failsafe CD into the server, and click **No** when prompted to install Failsafe. Run **D:\Win32\Install\Setup.exe** from the Start → Run menu.
2. Select the language in which you want to run Oracle.
3. Enter your Company name and a **local** home directory, such as **C:\ORANT**.
4. The Software Asset Manager appears. Select **Failsafe Manager 2.0.5** then click **Install**.
5. Select the Failsafe Documentation that you wish to install.
6. A message displays saying an upgrade to the Enterprise Manager is needed. Click **OK** to continue.
7. Oracle Installer analyzes file dependencies and installs all necessary files in your home directory **C:\ORANT**. This takes a while depending on the speed of your CD-ROM drive.
8. Please read the message after the installation. Review the Release Notes in the Failsafe CD. Click **OK** to finish the install.
9. **Restart** the system.

Chapter 8

Oracle Enterprise Manager

Configuring SQL*Net to use the Enterprise Manager

Important: Configure SQL*Net on both cluster nodes and the standalone server. For this cluster you will use a temporary repository **TEMP-REPO** located outside the cluster, which will only be used to launch Enterprise Manager. **You can use your existing repository if you already have one**. Later on, you will build a clustered repository in one of the shared disk of the cluster.

The new Enterprise Manager requires the new **Start → Programs → Oracle for Windows NT → Net8 Easy Config** program. However, you can still use the **SQL*Net Easy Config** for any of your existing programs.

Follow these steps to configure SQL*Net for the temporary repository:

1. Start the **Start → Programs → Oracle for Windows NT → Net8 Easy Config** program to Launch the Oracle Service Name wizard.
2. Select **Add New Service** and type in the New Service Name: **TEMP-REPO**. Click **Next**.
3. Select **TCP/IP** as your networking protocol to connect to the database. Click **Next**.
4. Type the Hostname: **OFS2-CLIENT** and leave the default port number. Click **Next**.
5. Type the System ID (SID) for your database. Use **ORCL**. Click **Next**.
6. Click **Test Service** and login using user **SYSTEM** and password **MANAGER**.
7. Make sure the test is successful, otherwise check the following:
 - a. All entered information is correct
 - b. OracleListener service from the Control Panel → Services icon is started
8. Click **Finish** to complete the setup.

Building the Repository and Discovering New / Existing Services

Building the repository is only done once, every time you launch the Enterprise Manager for the first time using a new service name that does not have a repository.

1. Before proceeding, check from your Control Panel → Services, that the **Oracle Intelligent Agents are started for both cluster nodes**.
 2. Start the **Start → Programs → Oracle Enterprise → Enterprise Manager** program.
 3. Login as follows:
 - Username: SYSTEM
 - Password: MANAGER
 - Service: TEMP-REPO (or your existing repository)
 - Connect as: NORMAL
 4. A message appears saying that the Repository Manager, Enterprise Manager and Software Manager sub-components are not found. This creates all 3 sub-components. Click **OK** to continue.
 5. This operation may take a while. Once the operation is complete, the Enterprise Manager launches the Discover New Services Wizard.
 6. Add the 2 cluster nodes OFS2-NODE1 and OFS2-NODE2 to the list of nodes that you want services to be discovered.
 7. Select to retrieve service information **Immediately after Wizard completion** and choose **Yes** to update your local network configuration file.
 8. The wizard shows you the list of nodes that you want services to be discovered. Click **Finish** to confirm your choices.
 9. A progress window shows you the list of nodes as they are discovered. Click **Close** after all nodes are discovered.
-

Your Enterprise Manager should be able to show both nodes and all existing databases and services in your nodes, such as your starter databases. These starter databases are not used in this cluster.

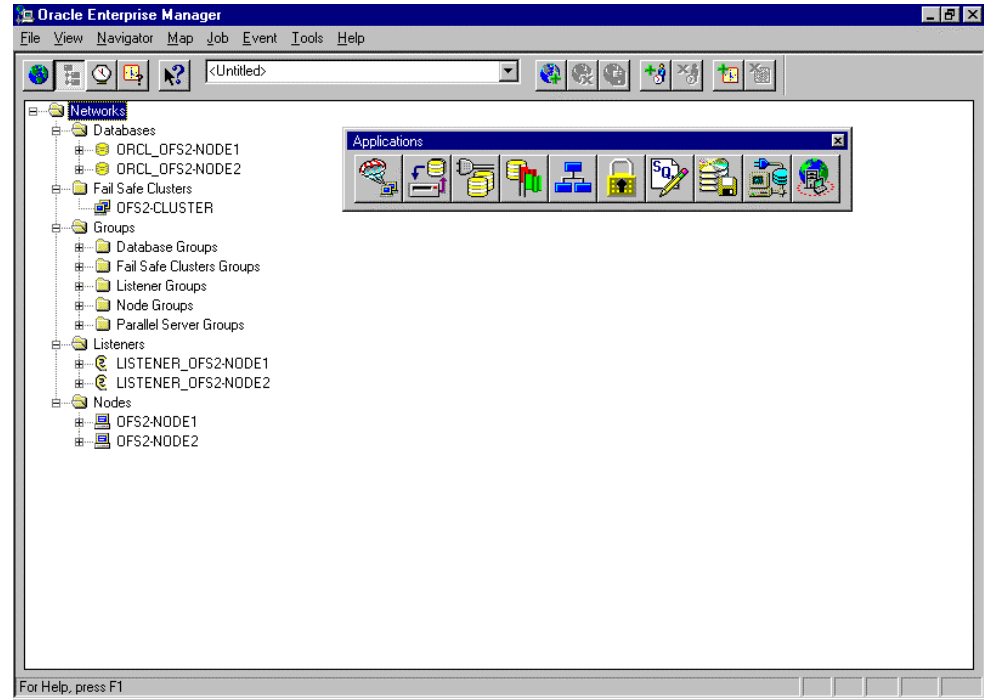


Figure 8-1. Oracle Enterprise Manager Window

Chapter 9

Oracle Failsafe Manager

Launching the Failsafe Manager

The Failsafe Manager can be launched from your desktop or the Enterprise Manager. Click on **Tools** → **Failsafe Manager** from the Enterprise Manager.

1. Login to Failsafe Manager using the cluster administrator user and password. Use the following to login:

Username: *Administrator*
Password: *<password>*
Cluster Alias: *OFS2-CLUSTER*
Domain: *OFS2-DOMAIN*

9-2 Oracle Failsafe Manager

2. This is your Failsafe Manager:

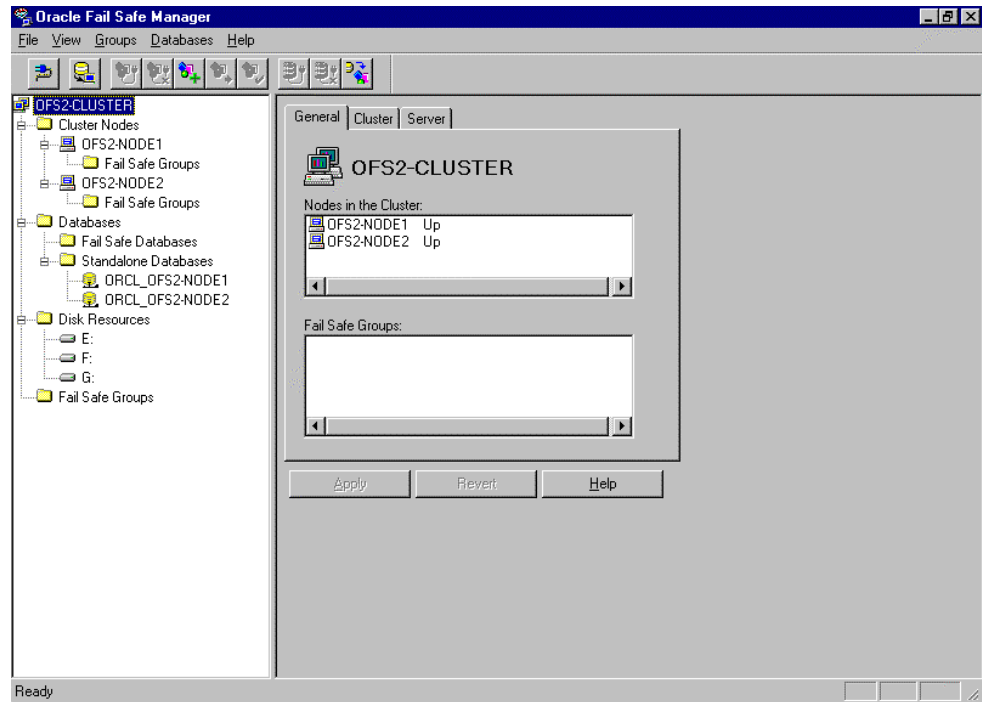


Figure 9-1. Oracle Failsafe Manager Window

All database-related cluster administration should be done in the Failsafe Manager. You can also launch the Failsafe Manager from the **Start → Programs → Oracle Enterprise** menu. You can also read the Failsafe Documentation and go to the Failsafe Tutorial from the same menu.

Creating Failsafe Groups

Refer to the following tables to create all Failsafe groups.

Table 9-1
Failsafe Group 1 - REPOSITORY

Group Name	REPOSITORY
Network to Use	Client Connection
Network Name	OFS2-DB1
IP Address	110.6.120.111
Subnet Mask	255.255.0.0
Failover Period	6 hours
Failover Threshold	10 attempts
Failback Policy	Prevent Failback; No Failback
Preferred Node	None

Table 9-2
Failsafe Group 2 – CUSTOMER

Group Name	CUSTOMER
Network to Use	Client Connection
Network Name	OFS2-DB2
IP Address	110.6.120.112
Subnet Mask	255.255.0.0
Failover Period	6 hours
Failover Threshold	10 attempts
Failback Policy	Failback Immediately
Preferred Node	OFS2-NODE1

Table 9-3
Failsafe Group 3 – SALES

Group Name	SALES
Network to Use	Client Connection
Network Name	OFS2-DB3
IP Address	110.6.120.113
Subnet Mask	255.255.0.0
Failover Period	6 hours
Failover Threshold	10 attempts
Failback Policy	Failback Immediately
Preferred Node	OFS2-NODE2

To create a failsafe group, select **Group** → **Create** from the menu.

1. Specify the name of the group, for example **REPOSITORY**. Click **Next**.
2. Select **Networks accessible by clients**. Enter the following items:
 - Network to use: *Client Connection*
 - Network Name: *OFS2-DB1*
 - IP Address: *110.6.120.111*
 - Subnet Mask: *255.255.0.0*
3. Specify the Failover Period and Failover Threshold. Click **Next**.
4. Select whether to Failback or Prevent a Failback. Click **Next**.
5. If you want a Failback, select a Preferred Node. Click **Next**.
6. A summary screen shows all your selections. Click **OK** to confirm.

Repeat these steps for all Failsafe groups. From the cluster administrator, you can also see the groups and resources that were created.

Standalone Databases

As an exercise, you will create your own standalone databases to the above failsafe groups that were just created. In production environments, it is also helpful to test on sample databases to plan a move of a production database from a standalone server to a failsafe cluster.

If you have a properly configured standalone database that is accessible from your client, and the data files are located in the external storage that you are planning to share with another node, the Enterprise Manager and the Failsafe Manager are able to see them during Discovery. By moving standalone databases to the Failsafe groups, you are configuring the standalone database to become a failsafe database. The Failsafe Manager automatically configures both nodes accordingly.

Creating Sample Standalone Databases

Refer to the following tables to create all sample databases.

Table 9-4
Sample Database 1 – REPOSITORY

Service ID	REPOSITORY
Disk Resource	Disk E:
User Name	Internal
Password	Internal

Table 9-5
Sample Database 2 – CUSTOMER

Service ID	CUSTOMER
Disk Resource	Disk F:
User Name	Internal
Password	Internal

Table 9-6
Sample Database 1 – SALES

Service ID	SALES
Disk Resource	Disk G:
User Name	Internal
Password	Internal

To create sample standalone databases:

1. Click the Database → Create Sample Database from the menu.
2. Enter the information from each table accordingly on each sample database you want to create
3. Click **OK** when the sample database is created.

Adding a Standalone Database to a Failsafe Group

Refer to the following tables to add each standalone database to a failsafe group.

Table 9-7
Add REPOSITORY database to REPOSITORY Group

Group Name	REPOSITORY
Service ID	REPOSITORY.world
Instance Name	OFS1
Database Name	OFS1DB
Parameter File	E:\Ofsdb\Ofs1\Param\InitOFS1.ora
User Name	Internal
Password	Internal

Table 9-8
Add CUSTOMER database to CUSTOMER Group

Group Name	CUSTOMER
Service ID	CUSTOMER.world
Instance Name	OFS2
Database Name	OFS2DB
Parameter File	F:\Ofsdb\Ofs2\Param\InitOFS2.ora
User Name	Internal
Password	Internal

Table 9-9
Add SALES database to SALES Group

Group Name	SALES
Service ID	SALES.world
Instance Name	OFS3
Database Name	OFS3DB
Parameter File	G:\Ofsdb\Ofs3\Param\unitOFS3.ora
User Name	Internal
Password	Internal

Follow these steps to add each standalone database to a failsafe group:

1. Right-click on the standalone database you want to add to a failsafe group and select **Add to Group**
2. **Select the correct group.** Inspect the default values making sure they are correct. Enter the user name and password. For this cluster use the user **INTERNAL** and password **INTERNAL**.
3. Please wait until the operation is complete.
4. Once completed, add an agent to the failsafe group by **selecting the disk resource** then click **OK** to continue.
5. A message appears confirming that the standalone database was added to the failsafe group. Click **OK** to continue or **HELP** to view more details on how to update your client connections to the newly created services.

Repeat these steps to all the standalone databases that you want to add to a failsafe group.

9-10 Oracle Failsafe Manager

Your Failsafe Manager should look like this:

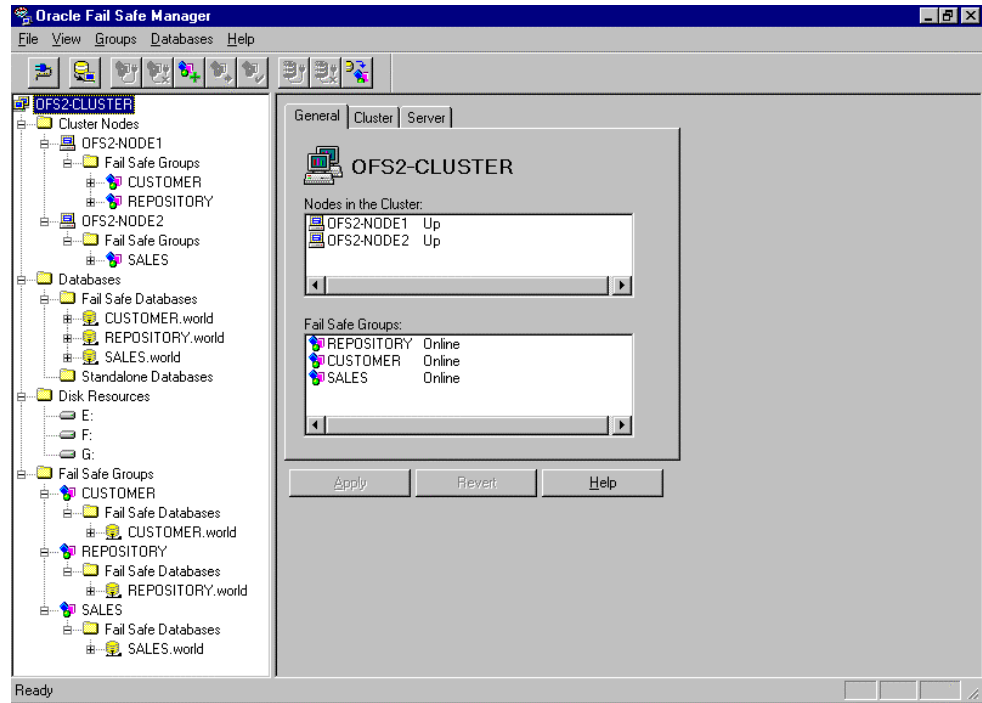


Figure 9-2. Oracle Failsafe Manager Window

Verify each Failsafe Group

It is important to test the following:

- the client can see each failsafe group through the network, and
- the client can connect to the listener that is running on each failsafe group

You should see a similar screen when you execute the following test from your client:

```
C:\> ping OFS2-DB1
```

```
Pinging OFS2-DB1 [110.6.120.111] with 32 bytes of data:  
Reply from 110.6.120.111: bytes=32 time<10ms TTL=128  
Reply from 110.6.120.111: bytes=32 time<10ms TTL=128  
Reply from 110.6.120.111: bytes=32 time<10ms TTL=128  
Reply from 110.6.120.111: bytes=32 time<10ms TTL=128  
C:\> tnsping repository
```

```
TNS Ping Utility for 32-bit Windows: Version 2.3.3.0.0  
Copyright (c) Oracle Corporation 1995. All rights  
reserved.  
Attempting to contact  
(ADDRESS=(COMMUNITY=tcp.world)(PROTOCOL=TCP)(Host=OFS2-  
DB1)(Port=1521))  
OK (30 msec)
```

Repeat these tests to all Failsafe groups and make sure your client can access each virtual server and each Failsafe database.

Chapter 10

Using the New Repository

Creating a New Repository and Discovering the Failsafe Groups

Now that you created a new **clustered repository**, you will start using this repository by following the same steps that you used for your temporary repository.

From the client, follow these steps to use the new repository:

1. Exit the Failsafe Manager, then exit the Enterprise Manager.
2. Select the **Start → Programs → Oracle Enterprise → Enterprise Manager** program.
3. Login as follows:
 - Username: SYSTEM
 - Password: MANAGER
 - Service: REPOSITORY
 - Connect as: NORMAL
4. A message appears telling that the Repository Manager, Enterprise Manager and Software Manager sub-components are not found. This creates all 3 sub-components. Click **OK** to continue.
5. This operation may take a while. Once the operation is complete, the Enterprise Manager launches the **Discover New Services Wizard**.
6. Add the 3 failsafe groups that were created: OFS2-DB1, OFS2-DB2 and OFS2-DB3 to the list of nodes that you want services to be discovered.
7. Select to retrieve service information **Immediately after Wizard completion** and choose **Yes** to update your local network configuration file.

10-2 Using the New Repository

8. The wizard shows you the list of nodes that you want services to be discovered. Click **Finish** to confirm your choices.
9. A progress window shows you the list of nodes as they are discovered. Click **Close** after all nodes are discovered.

Your Enterprise Manager should be able to show all your Failsafe groups and databases.

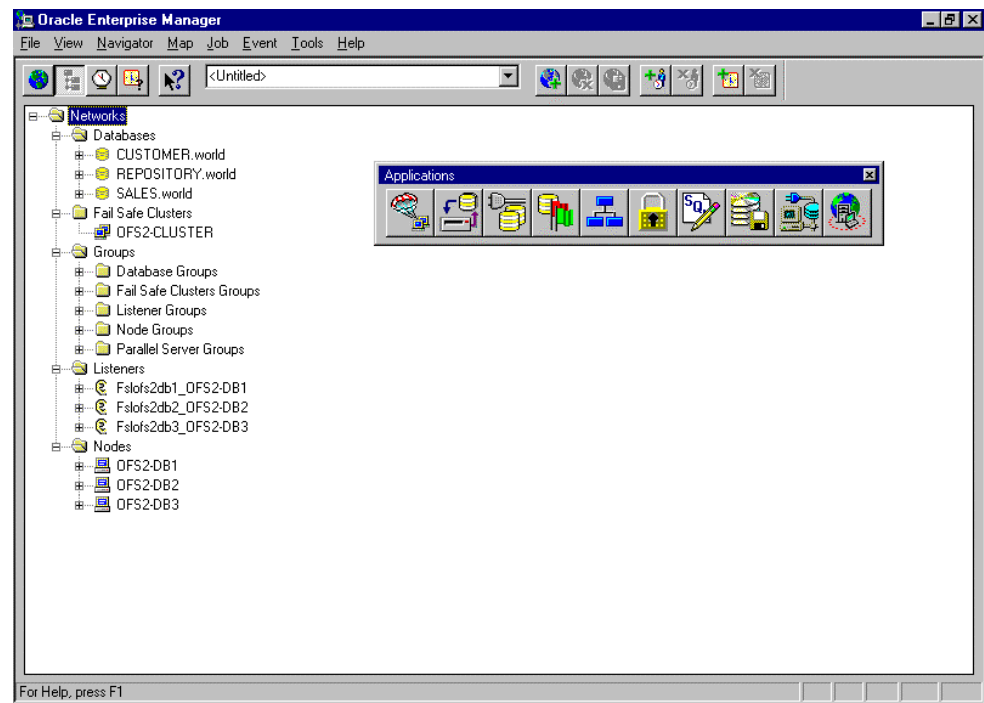


Figure 10-1. Oracle Enterprise Manager Window

Chapter 11

Partitioned Data and Partitioned Workload

View the screen shown below:

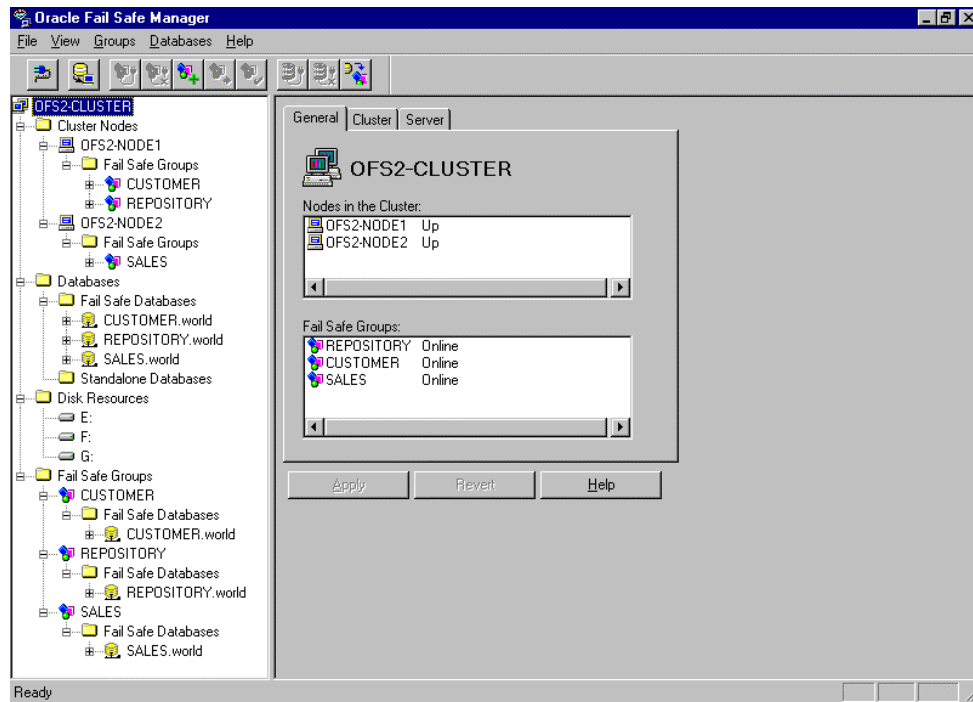


Figure 11-1. Oracle Failsafe Manager Window

Each node is running different databases. Database REPOSITORY.world and CUSTOMER.world are running on OFS2-NODE1, while database SALES.world is running on OFS2-NODE2.

This partitioning of services running on individual failsafe groups can be described as **data partitioning**. At anytime you can move any of these databases on any node and the client can still connect to the same network name and IP address, and connect to the same instance using the same connect string. This is possible since the database instance and the data files are defined and contained on a single failsafe group.

11-2 *Partitioned Data and Partitioned Workload*

Running these failsafe groups on separate nodes can be described as **workload partitioning**. At anytime you can distribute the workload and move any of the failsafe groups to any node. By partitioning the workload, you are able to independently move each service from one node to another. As a result, you are actually performing a **static load balancing** between the 2 cluster nodes.

Data partitioning and Workload partitioning allows each node to serve independent workloads and databases. Thus, allowing each node to perform independent tasks. As a result of workload partitioning, you can perform static load balancing, which improves the overall performance of the cluster by moving services from a heavily used node to a lightly used node.

Chapter 12

Testing the Failsafe Cluster

From your client workstation:

1. Launch the SQL*Plus application by clicking on the Program Icon and login to your database using the following command:

```
C:\> PLUS33 SYSTEM/MANAGER@CUSTOMER
```

2. Try to connect to the same database using the same connect string, as shown below:

```
SQL*Plus: Release 3.3.2.0.2  
Copyright (c) Oracle Corporation 1979, 1994. All rights reserved.
```

```
Connected to:  
Oracle7 Server Release 7.3.2.2.0 - Production Release  
With the distributed, replication and parallel query options  
PL/SQL Release 2.3.2.2.0 - Production
```

```
SQL> connect system/manager@customer  
Connected.  
SQL>
```

3. Verify that your client can connect to the database using the above connect string. Currently, OFS2-NODE1 is running the database services.
4. From the Failsafe Manager, move the failsafe group from OFS2-NODE1 to OFS2-NODE2 to simulate a failover.

12-2 Testing the Failsafe Cluster

5. After the failover, try to connect to the database using the same connect string as shown below:

```
SQL*Plus: Release 3.3.2.0.2  
Copyright (c) Oracle Corporation 1979, 1994. All rights  
reserved.
```

```
Connected to:  
Oracle7 Server Release 7.3.2.2.0 - Production Release  
With the distributed, replication and parallel query  
options  
PL/SQL Release 2.3.2.2.0 - Production
```

```
SQL> connect system/manager@customer  
Connected.  
SQL> connect system/manager@customer  
ERROR: ORA-12571: TNS:packet writer failure
```

```
Connected.  
SQL>
```

Since the client is connected to the cluster, as configured in your SQL*Net, even with a failover of the database services, the client can re-connect to the database using the same database alias.

Chapter 13

Client and Server Considerations

A Re-connect is Still Needed

This TechNote does not provide a way to remove the need to re-connect to the database. A client application that checks if there is a connection to the database prior to executing any requests from the database will work very well. A simple algorithm is shown below to illustrate this:

```
// FUNCTION TO CHECK IF THERE IS A CONNECTION
// RE-CONNECT IF DISCONNECTED

FUNCTION CHECK_CONNECTION(CONNECT_ATTEMPTS) : BOOLEAN
BEGIN
  FOR I = 1 TO CONNECT_ATTEMPTS DO
  BEGIN
    IF (CONNECTED) THEN
      RETURN (TRUE)
    ELSE
      CONNECT TO DATABASE
    END IF
  END
  RETURN (FALSE)
END

// MAIN PROGRAM

BEGIN
  IF CHECK_CONNECTION(NUMBER_OF_ATTEMPTS) = TRUE THEN
    ::
    READ THE DATABASE
    ::
  END IF
END
```

The algorithm shown above is just an example of automating a re-connect. It shows a good way to request services from the database, whether the client wants to read or write to the database, and whether the client is connected to a single server or a cluster. This kind of algorithm can be implemented in more elegant ways, such as the use of exception handlers.

13-2 *Client and Server Considerations*

A classic example to avoid an explicit re-connect of a client is the use of a web browser that is connected to a web server, which is also using an Oracle database server.



Data Loss and Data Integrity

This Technote is intended to work with the existing techniques that resolve data loss and data integrity used with standalone databases, and at the same time, minimize the downtime of the system. It can also provide a means of scheduling maintenance to the servers without depriving the clients the access to the system.

Data loss and data integrity should be resolved using the techniques presented by Microsoft and Oracle. You would have to implement these techniques similar to how you would implement them if you were connected to a single server.

Maintaining Data Integrity - Reliable Transactions

Maintaining reliable transactions is a must in database applications, regardless of whether it is a clustered database or not. Creating smart clients that will automatically re-connect to the database server on a clustered environment can make a failure transparent, but not data corruption and incomplete transactions. Microsoft Cluster Server is not a solution for minimizing incomplete transactions.

There are a few alternatives:

- Build smart client and server applications that will try to preserve the states of the transaction being submitted by the client to the server, and keeping the results of the transaction being processed from the server to the client.
- Build you own transaction server that will work with your database server and clients to handle your transactions.
- Use Microsoft Transaction Server to handle transaction bundling and preserving states of your transactions. Allowing you to build thin and simple client and server applications that are easier to manage.

In a clustered solution, the techniques that you use to keep data integrity on a standalone database server will still be applied.

13-4 *Client and Server Considerations*

A classic example is the credit-debit transaction happening in a bank transfer. The transfer from one account to another account should be contained in a single transaction. Otherwise, a successful credit and a failed debit will create money that did not come anywhere. This scenario can happen in standalone database servers and can still happen even if you have a highly available database cluster. The only major advantage that you can get from a clustered database is high availability, static load balancing, and a potentially scalable environment in the future. It does not, however, solve or avoid any problems caused by data corruption and incomplete transactions.

Appendix

Hardware Configuration

Server 1

Compaq ProLiant 2500, 2P P5/200, 128-MB memory

Table A-1
Server Configuration Information

Slot Number	Description	IRQ
Embedded NIC	Netelligent 10/100 TX (Client Network)	5
Embedded Disk Ctrl	C875 Ultra Fast Wide SCSI Controller	14
Slot 1	Netelligent 10/100 TX (Interconnect Network)	9
Slot 2	Fibre Channel Host Adapter	11

Table A-2
Disk Configuration Information

Disk Controller	Controller Order	SCSI ID or Logical Drv	Disk Configuration	Raid Level	Drive Letter
Embedded Disk Ctrl	1st	0	2.1-GB	none	C
	2nd	5	CD-ROM	n/a	D
Fibre Host Ctrl	3rd	0	2 x 2.1-GB	1	E
		1	2 x 2.1-GB	1	F
		2	2 x 2.1-GB	1	G

A-2 Hardware Configuration

Server 2

Compaq ProLiant 2500, 2P P5/200, 128-MB memory

Table A-3
Server Configuration Information

Slot Number	Description	IRQ
Embedded NIC	Netelligent 10/100 TX (Client Network)	5
Embedded Disk Ctrl	C875 Ultra Fast Wide SCSI Controller	14
Slot 1	Netelligent 10/100 TX (Interconnect Network)	9
Slot 2	Fibre Channel Host Adapter	11

Table A-4
Disk Configuration Information

Disk Controller	Controller Order	SCSI ID or Logical Drv	Disk Configuration	Raid Level	Drive Letter
Embedded Disk Ctrl	1st	0	2.1-GB	none	C
	2nd	5	CD-ROM	n/a	D
Fibre Host Ctrl	3rd	0	2 x 2.1-GB	1	E
		1	2 x 2.1-GB	1	F
		2	2 x 2.1-GB	1	G

Glossary

Availability	The ability of a system to appear to the users as if it was operating continuously, even though failures have occurred.
Cluster	<p>A cluster is a set of independent computer systems working together as a single system. The cluster software directs client requests to one of the systems in the cluster based on resource accessibility, server load, and application requirements. When a system in the cluster fails the cluster software will recover and disperse the work from the failed system amongst the remaining systems in the cluster. When the overall load exceeds the capabilities of the systems in the cluster, additional systems may be added to the cluster. The failure of a system in the cluster will not affect the other systems and in certain cases the client applications should be unaware of the failure, except for some loss of service for a short period of time.</p> <p>Given the above definition, it would be hard to distinguish a cluster from a general distributed system; therefore, clusters provide the following additional capabilities:</p> <p>Every system in the cluster must have full connectivity and communications with all other systems in the cluster.</p> <p>Systems in the cluster know all the other systems in the cluster; they join in order to form a cluster. The cluster systems also know when systems join or leave the cluster, via a connection manager</p> <p>Systems in the cluster know what resources are available on all other members of the cluster.</p> <p>There is a virtual circuit, provided by the connection manager that provides reliable delivery of messages between systems in the cluster</p>
Connect Descriptor	A specially formatted description of the destination for a network connection. Connect descriptors are constructed using a set of keywords and values mapped to service names.

Failover	The time it takes a resource to failover from one system to another. For example, if a resource is to be polled at a 1-second interval, then we might decide that after two poll intervals, that a resource is dead and that it should be restarted or failed over. Also, let's assume that the local restart count for the resource is 4, then we would require 4 of these timeouts before the resource is failed over to another system. Given a 1-second interval, we would have a period of 2 seconds times the 4 retries - or a total of 8 seconds before the resource failed over to another system.
Listener	An executable program that enables an Oracle7 server to accept connections from client machines over SQL*Net.
Member	A system which has joined a cluster
Node	<i>See</i> system.
Oracle Instance	Every time a database is started, a system global area (SGA) is allocated and Oracle background processes are started. The system global area is an area of memory used for database information shared by the database users. The combination of the background processes and memory buffers is called an Oracle instance.
Quorum	<p>A voting mechanism maintained by the connection manager. A simple quorum voting mechanism is used to indicate when more than half of the available votes are present in the cluster. Note, that a member (or many members) in the cluster can have more than one vote. In addition, a resource available within the cluster can have a vote. This latter ability allows two systems cluster to function in the presence of a communication or system failure, by having the system that continues functioning and reserves the resource to obtain a quorum.</p> <p>There are also change quorums and operational quorums, which are described in quorum.doc.</p>

Quorum resource	A quorum resource is any resource that has writable persistent storage for maintaining cluster database change log records.
Reliability	The dictionary defines this as meaning dependable or trustworthy. Our intention is to make the cluster, as a whole, reliable, in that clients can depend on the services provided by the cluster.
Resource	A physical or logical entity that is managed (see resource manager below). They include, but are not limited to: disks, tapes, printers, CD-ROMs, IP address, etc.
Server	The layer of software running on a cluster node which performs service registration or advertisement. Any software that uses this layer will inherit its properties.
Server application	The actual application software running on a cluster node, regardless of whether it does service registration.
Service	A data set or operation set exported by application servers to their clients. This term is distinguished from a Windows NT Service.
Service Name	A short, convenient name mapped to a network address contained in a TNS connect descriptor. Users need only know the appropriate service name to make a TNS connection.
Shared disk	A cluster organization that requires all systems of the cluster to have access to shared Disks.
Shared nothing	A cluster organization in which disks are 'owned' by a single system of the cluster and requests for disk access is directed to the system that 'owns' the disk.
Shared resource	A cluster organization in which some resources are accessible to all systems in the cluster.

G-4 *Glossary*

SID	Oracle Server ID, this is not the same as the SID of Windows NT.
Single system image	A cluster that is 'logically' viewed by clients as being a single system. Services are provided by the cluster that is available to clients, without distinction as to which system in the cluster is actually providing the service. The cluster is also managed as a single system, with the ability to administer all systems in the cluster from a single system within the cluster.
Startup Wait Time	This parameter sets the number of seconds that the listener sleeps before it responds to the first listener control STATUS command. This feature assures that a listener with slow communication lines has adequate time to start up before responding to a STATUS request. The default of 0 means that the listener responds to the first STATUS command without any delay. Any positive numeric value is acceptable.
SQL*Net	The Oracle client/server communication software that offers transparent operation to Oracle tools or databases over any type of network protocol and operating system.
System	A complete computer system capable of operating independently (either a single processor system or an SMP system), for example, a PC.
TNS	Transparent Network Substrate (TNS) is the Oracle networking technology that provides a single application interface to all industry-standard networking protocols.

Index

A

- adding a standalone database to a failsafe group, 9-8
- advantages of a resource failover, 2-2
- Array Configuration Utility, 3-6

B

- building the repository, 8-2

C

- client and server considerations, 13-1
- client configuration, 3-6
- Compaq
 - Fibre Channel Host Controller, 3-5
 - Netelligent 10/100 TX PCI UTP Network Interface Controller, 3-5
 - NT Support Software Utility, 3-6
 - ProLiant 2500 servers, 3-5
 - ProLiant 850 Server, 3-5
 - ProLiant Fiber Channel Disk Storage System F2, 3-5
 - ProLiant Fibre Channel Array Controller, 3-5
 - ProLiant servers, 3-1
 - SmartStart, 3-6
- Configuring SQL*Net to use the Enterprise Manager, 8-1
- Creating a New Repository, 10-1
- Creating Failsafe Groups, 9-3

D

- data integrity, 13-3
- data loss, 13-3
- data partitioning, 11-2
- disadvantages of a resource failover, 2-2
- discovering the Failsafe Groups, 10-1
- disk storage, 3-5

E

- editing the Windows NT Hosts files, 5-1
- expected audience, 1-1

F

- Failsafe 2.0.5 Server, 3-1

H

- hardware requirements, 3-5

I

Installation

- Microsoft Cluster Server, 4-1
- Oracle Client Products in the Workstation, 6-3
- Oracle Failsafe Manager, 7-3
- Oracle Failsafe Server, 7-2
- Oracle Workgroup Server, 6-1

L

- launching the Failsafe Manager, 9-1

Index -2

M

- Microsoft Cluster Server 1.0, 3-1
- Microsoft Cluster Server, installation of 4-1
- Microsoft Windows NT Enterprise Server 4.0, 3-6
- Microsoft Windows NT Service Pack 3, 3-6
- Model Environment, 3-3

O

- Oracle
 - Client Products in the Workstation, installing 6-3
 - Enterprise Manager, 8-1
 - Failsafe 2.0.5 Installation, 7-1
 - Failsafe Manager, 9-1
 - Failsafe Manager, installing 7-3
 - Failsafe Server 2.0.5, 3-6
 - Failsafe Server, installing 7-2
 - Failsafe Server, key functions, 3-2
 - Required Support Files 7.3.3.3.2, 3-6
 - Server 7.3 installation, 6-1
 - Workgroup Server 7.3.3 for NT, 3-6
 - Workgroup Server 7.3.3, 3-1
 - Workgroup Server, installing 6-1
- overview of a failover, 2-1

P

- partitioned data, 11-1
- partitioned workload, 11-1
- PING the cluster nodes, 5-3
- planned outage, 2-1
- pre-installation checklist, 5-1

R

- re-connect, 13-1
- reliable transactions, 13-3
- resource failover, 2-1
- resource failover, advantages of a 2-2
- resource failover, disadvantages of, 2-2
- resources, shared, 3-1

S

- server configuration, 3-5, 3-6
- service failover, 2-1
- Smart 2/P Array Controller, 3-5
- software requirements, 3-6
- standalone databases, 9-5
- standalone databases, creating, 9-5
- static load balancing, 11-2
- System Configuration Utility, 3-6

T

- testing the Failsafe Cluster, 12-1

U

- unplanned outage, 2-1
- updating the Required Support Files to 7.3.3.3.2, 7-1
- using the New Repository, 10-1

V

verify each Failsafe Group, 9-11

W

Windows NT Enterprise Server 4.0,
3-1

Windows NT Hosts files, editing 5-1
workload partitioning, 11-2