




High Availability for Oracle 9i Using Double-Take



High Availability for Oracle 9i Using Double-Take published November 2002

NSI and Double-Take are registered trademarks of Network Specialists, Inc. All other products are trademarks of their respective companies. © 1996–2002 NSI Software

Double-Take Support for Application Failover

Double-Take's file system replication process is application independent and replicates any file system changes (including permissions and attributes) written to NTFS, FAT or FAT32 file systems by any application or process, subject to specific exceptions called out in the *User's Guide* or *readme.txt* file. Maintaining point-in-time consistent file system replicas and providing server monitoring and automatic or manual failover of the server name and IP address are the primary functions of the Double-Take software and we offer support to qualified customers should these functions fail to operate in accordance with our published documentation, regardless of what application or process is manipulating the data.

NSI Software may provide application notes and other documents that provide implementation guidelines on how to use Double-Take functions and replicas to manually or automatically failover or recover many popular third party applications and a general process to accomplish failover or recovery of many other third party applications. While these steps are believed to be accurate for the specific configuration, Double-Take version, and application versions originally tested, due to the number of possible configurations and variables, NSI Software can only test selected combinations and may provide only limited support for the operation and configuration of third party applications or the behavior of those applications before, during, or after failover, in its discretion. In cases where NSI Software has no direct access to or experience with a particular application or configuration, NSI Software support may also be limited to only the actual replication of the file system data and failover (name and IP address) of the server.

For assistance in validating, implementing or troubleshooting these or other possible configurations with third party applications, NSI Software and its partners may offer professional services on a fee basis to apply best practices for assisting with third party applications to recover automatically or manually using replicated data.

This, and any other, application note is provided solely for the convenience of our customers and is not intended to bind NSI Software to any obligation.

Table of Contents

Introduction	1
Requirements	1
Protecting Your Oracle Data	2
0.1 Install Software on the Source	2
0.2 Install and Configure Software on the Target	2
0.3 Configure Double-Take and Begin Mirroring and Replication	3
0.4 Configure failover and begin failure monitoring	7
Monitoring Failover	10
Restoring Your Oracle Data	12

Introduction

Oracle is a scalable, reliable, flexible, and high-performance relational database management system for many server-based operating systems. NSI Software's Double-Take provides real-time enterprise data protection and replication. Double-Take can be used to provide high availability for Oracle.

This document describes the steps necessary to configure Double-Take 4.x to provide high availability for Windows 2000 servers running Oracle 9i. These procedures allow a secondary server to assume the identity and role of a failed primary Oracle server while maintaining the availability of Oracle services with minimal disruption or data loss.

To complete these instructions, you will install Oracle and Double-Take, and configure Double-Take for replication and failover. Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications including Double-Take and Oracle.

NOTE: Double-Take allows you to configure one target to monitor and failover for one or more source machines. In a one-to-one configuration, you will want to replicate your Oracle data to the same location on the target so that failover is automatic. In a many-to-one configuration, the data will need to be replicated to a unique location and then renamed to the corresponding Oracle directory on the source before failover occurs.

This application note focuses on a single Oracle server being replicated to a single target.

Requirements

- ◆ Two servers that meet one of the following operating system requirements:
 - ◆ Microsoft Windows NT 4.0 with Service Pack 4 or higher
 - ◆ Microsoft Windows 2000

NOTE: The two servers should both be running the same operating system. Although cross-platform mirroring and replication are available, NSI Software recommends that the two servers be the same platform for effective failover and fallback.

- ◆ Both servers must be connected to the same physical network
- ◆ It is recommended that both source and target servers be standalone servers. (You may experience problems with promotion and demotion during failover if either of the machines are Primary or Backup Domain Controllers.)
- ◆ Two licensed copies of Oracle 9i
- ◆ Two licensed copies of Double-Take 4.x

Protecting Your Oracle Data

Install Software on the Source

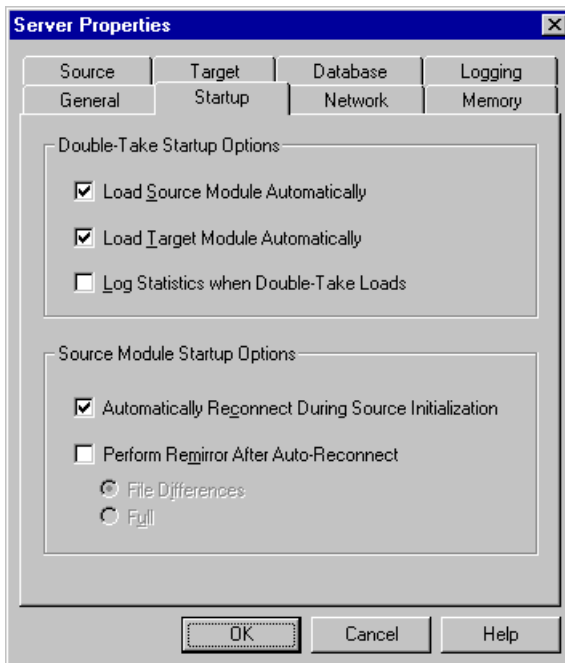
1. Install Oracle on the source, if it is not already installed.
2. Install Double-Take on the source using the installation defaults. See the Double-Take *Getting Started* guide for details.

Install and Configure Software on the Target

1. Install Double-Take on the target using the installation defaults. See the Double-Take *Getting Started* guide for details.
2. Install Oracle on the target using the same installation options used on the source machine.
3. Set the Oracle service to manual startup so that the failover and failback scripts that you will be creating will control the stopping and starting of the Oracle service
 - a. In **Control Panel, Administrative Tools, Services**, double-click the Oracle service.
 - b. Click the **General** tab.
 - c. Select **Manual** from the **Startup type** drop-down list and click **OK**.

Configure Double-Take and Begin Mirroring and Replication

1. Select **Start, Programs, Double-Take, Management Console**.
2. Double-click your source machine to log on.
3. If you are using Double-Take version 4.1 or earlier, you will need to disable auto-remirror on auto-reconnect so that the source does not remirror files after failback. In version 4.2 and later, the source automatically recognizes that a restore is required and will not remirror. If you are using 4.1 or earlier, complete steps a-c below. If you are using 4.2, you can continue with the next numbered step.
 - a. Right-click the source machine and select **Properties**.
 - b. Select the Startup tab.



- c. By default, **Perform Remirror After Auto-Reconnect** will be selected. Disable this option so that the source does not remirror files after failback. Click **OK** to continue.

NOTE: If you disable this option and a auto-disconnect occurs, you will need to remirror manually after the connection is reestablished.

4. Right-click your source machine and select **New, Replication Set** and enter the desired name for the replication set.

5. Select the Oracle data you wish to protect. The table below will help you verify that you have selected all pertinent Oracle files necessary to enable high availability on your target machine.

- ◆ **File Name**—This column lists each file name you need to select for replication. These files are assigned names during your initial Oracle installation and configuration causing some of these files to be unique to your environment. For example, if you used the default Oracle installation, your instance parameter file will be `init.ora` and located in the `\oracle\admin\instance_name\pfile` folder in the destination drive. If you supplied an instance name of `prod`, your `init.ora` file will be located in the `\oracle\admin\prod\pfile` folder by default. Files with unique names will be identified in the table with a wildcard asterisk (*).

Additionally, some of the files are identified with a number sign (#) placeholder in the file name. This placeholder identifies redundant files and/or multiple files which may exist because of tablespace configurations. For example, you have `control01.dbf`, `control02.dbf`, `undotbs01.dbf`, and `undotbs02.dbf`.

- ◆ **File Description**—This column gives a brief description of each file.
- ◆ **Default Location and Table to Query**—The **Default Location** column identifies the directory where a specific file is stored if you choose the default installation directory. If you did not use the default directory and do not know the location of a specific file, query the table identified in the **Table to Query** column to determine the location. For detailed information on querying tables, see your Oracle reference guide. If you do not know the location of files that are not associated with a table, locate those files using the Windows Find feature.

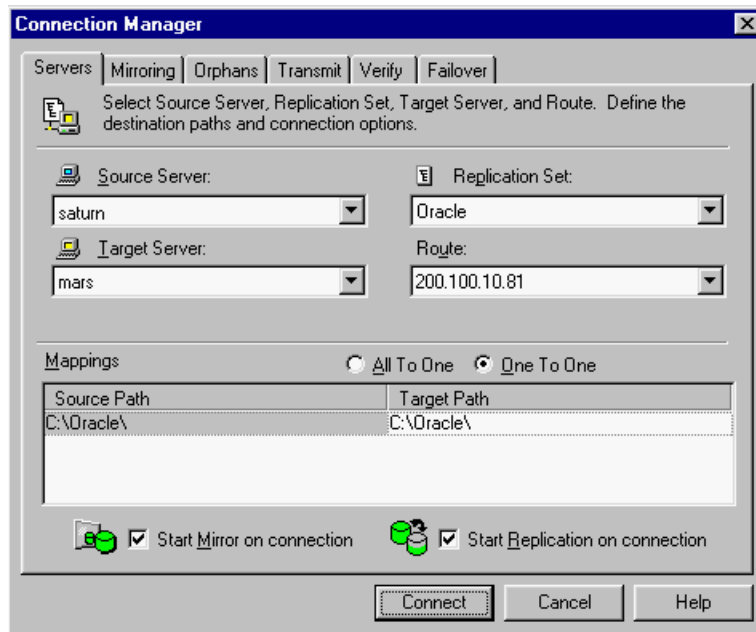
File Name	File Description	Default Location	Table to Query
<code>init.ora</code>	Parameter File	<code>\oracle\admin\instance_name\pfile</code>	N/A
<code>control#.ctl</code>	Control File	<code>\oracle\oradata\instance_name</code>	V\$CONTROLFILE
<code>redo#.log</code>	Online Redo Log	<code>\oracle\oradata\instance_name</code>	V\$LOGFILE
<code>system*.dbf</code>	System Tablespace	<code>\oracle\oradata\instance_name</code>	V_\$DATAFILE
<code>indx#.dbf</code>	Index Tablespace	<code>\oracle\oradata\instance_name</code>	V_\$DATAFILE
<code>tools#.dbf</code>	Tools Tablespace	<code>\oracle\oradata\instance_name</code>	V_\$DATAFILE
<code>undotbs#.dbf</code>	Rollback Tablespace	<code>\oracle\oradata\instance_name</code>	V_\$DATAFILE
<code>temp#.dbf</code>	Temporary Tablespace	<code>\oracle\oradata\instance_name</code>	V_\$DATAFILE

File Name	File Description	Default Location	Table to Query
users#.dbf	Data Tablespace	\oracle\oradata\ instance_name	V_\$DATAFILE
drsys#.dbf	Text-related Schema Objects	\oracle\oradata\ instance_name	V_\$DATAFILE
cwmlite#.dbf	OLAP Tablespace	\oracle\oradata\ instance_name	V_\$DATAFILE
odm#.dbf	ODM and ODM_MTR schema objects	\oracle\oradata\ instance_name	V_\$DATAFILE
xdb#.dbf	XML data and metadata	\oracle\oradata\ instance_name	V_\$DATAFILE
listener.ora	Service Configuration File	\oracle\ora90\ network\admin	N/A
tnsnames.ora	Client Configuration File	\oracle\ora90\ network\admin	N/A
pwd*.ora	Password Information File	\oracle\ora90\database	N/A

NOTE: The Listener service will not start on the target if `listener.ora` is replicated from the source, but the target is not in a failed over state. This file contains configuration information that is not valid on the target unless it is in a failed over state. If you need to start Oracle on the target server without failing over, do not include `listener.ora` in the replication set.

- After you have selected all of the Oracle data you wish to protect, right-click the replication set name and select **Save** to save the replication set.

7. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.



8. The **Source Server**, **Target Server**, **Replication Set**, and **Route** fields will automatically be populated. If you have multiple IP addresses on your target, verify the **Route** field is set to the correct network path. (For detailed information on connecting a source and target, reference the *Double-Take User's Guide*.)
9. Select the **One To One** mapping so that the replication set data is transmitted to the same directory structure on the target.
10. Click **Connect** to start the mirror and replication processes.

Configure failover and begin failure monitoring

- I. If a failure occurs, you will want to have the Oracle services start on the target machine automatically. To do this, create a batch file called `POSTOVER.BAT` using the sample batch file below. Save the batch file to the same directory where your Double-Take files are installed.

NOTE: The following sample batch files are available on the NSI Software web site at www.nsisoftware.com/download/orcl9iscrip.exe.

POSTOVER.BAT

```
rem ## Sample Oracle post-failover script
rem ## This script makes the following two assumptions:
rem ## 1. You have created a directory c:\oracle_fo
rem ## 2. The dbx_start.sql file referenced in this script is located
rem ## in the oracle failover directory c:\oracle_fo
rem ## Note: %a is a constant recognized in the command line mode and should NOT
rem ## be substituted.

rem ## The next two lines establish the working directories. Specify the location
rem ## of your Oracle program files in the first line. Do not leave any spaces
rem ## around the equal (=) sign.
set ORACLE_HOME=C:\ORACLE\ORA90
set ORACLE_FO=C:\ORACLE_FO

rem ## The following line CREATES the Oracle instance(s).
rem ## Substitute your oracle instance names for instance_name1 and instance_name2. Multiple instances
rem ## are separated by a space.
for %a IN (instance_name1 instance_name2) DO %ORACLE_HOME%\bin\oradim -new -sid %a -srvc oracleservice%a
-startmode auto -pfile %ORACLE_HOME%\admin\instance_name\init.ora

rem ## The following line STARTS the Oracle instance(s).
rem ## Substitute your oracle instance names for instance_name1 and instance_name2. Multiple instances
rem ## are separated by a space.
for %a IN (instance_name1 instance_name2) DO NET START ORACLESERVICE%a

rem ## The following line starts other Oracle services. The services specified
rem ## are started on an as-needed basis depending on the Oracle operations
rem ## you are running. The Oracle HTTP Server and the Oracle Agent are started
rem ## in the following example. Add additional services as needed separating
rem ## the service names by a space. The outcome is redirected to a log file in
rem ## the failover directory. Check this log file for errors.
for %a IN (OracleOraHomeHTTPServer OracleOraHomeAgent) DO NET START %a >>%ORACLE_FO%\start_services.log

rem ## The following lines mount and open the databases. Note that the same
rem ## command is repeated for each database. Also note that separate files
rem ## (DB1_Start.sql and DB2_Start.sql) are called for the different databases
rem ## so that the correct path to the parameter file (init.ora) is used.
rem ## The outcome is redirected to a log file in the failover directory. Check
rem ## this log file for errors. Substitute your instance names for instance_name1
rem ## and instance_name2.
set ORACLE_SID=instance_name1
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db1_start.sql>%ORACLE_FO%\start_%ORACLE_SID%.log
set ORACLE_SID=instance_name2
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db2_start.sql>%ORACLE_FO%\start_%ORACLE_SID%.log

rem ## The following line starts the Oracle Listener service. The outcome is
rem ## redirected to the failover directory where a log file is created.
rem ## Check this log file for errors.
NET START OracleOraHomeTNSListener >>%ORACLE_FO%\start_services.log
```

The `DBX_START.SQL` script used in the post-failover script is included on the following page. You may have multiple files, `DB1_START.SQL`, `DB2_START.SQL`, etc., if you have multiple databases.

DBx_START.SQL

```
# Sample db_start.sql script called during the Double-Take post-failover sequence.
# The following line identifies the Oracle instance to start.
# Replace instance_name with the name of the instance.
startup pfile=c:\oracle\admin\instance_name\pfile\init.ora
exit
```

NOTE: After failover is complete, users will need to log back into the Oracle database. You may want to devise an automatic process to notify the users of this information and include it at the end of the post-failover script.

2. After a failure is resolved, you will be ready to bring your source back online. At this time, you will want to stop the Oracle services on the target automatically. To do this, create a batch file called `PREBACK.BAT` using the sample batch file below. Save the batch file to the same directory where your Double-Take files are installed.

PREBACK.BAT

```
rem ## Sample Oracle pre-failback script

rem ## This script makes the following two assumptions:
rem ## 1. You have created a directory c:\oracle_fo
rem ## 2. The db_stop.sql file referenced in this script is located
rem ## in the oracle failover directory c:\oracle_fo
rem ## Note: %%a is a constant recognized in the command line mode and should NOT
rem ## be substituted.

rem ## The next two lines establish the working directories. Specify the location
rem ## of your Oracle program files in the first line. Do not leave any spaces
rem ## around the equal (=) sign.
set ORACLE_HOME=C:\ORACLE\ORA90
set ORACLE_FO=C:\ORACLE_FO

rem ## The following lines stops the Oracle databases.
rem ## Substitute your oracle instance names for instance_name1 and
rem ## instance_name2.
rem ## If you only have one Oracle database, remark out the
rem ## third and fourth lines of this section.
set ORACLE_SID=instance_name1
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db_stop.sql >%ORACLE_FO%\stop_%ORACLE_SID%.log
set ORACLE_SID=instance_name2
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db_stop.sql >%ORACLE_FO%\stop_%ORACLE_SID%.log

rem ## The following lines delete each instance service.
rem ## Substitute the name of the Oracle instance to delete for instance_name1
rem ## and instance_name2. Multiple instances are separated by a space.
FOR %%a IN (instance_name1 instance_name2) DO %ORACLE_HOME%\bin\oradim -delete -sid %%a

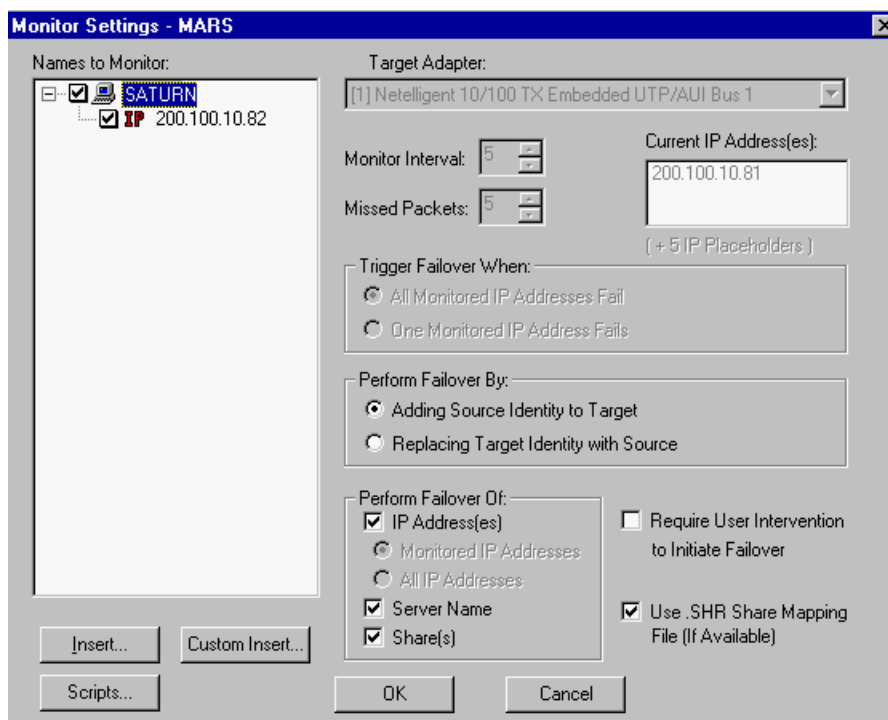
rem ## The following lines stop other non-instance specific Oracle services.
rem ## The value in parenthesis are the services to stop.
FOR %%a IN (OracleOraHomeTNSListener OracleOraHomeHTTPServer OracleOraHomeAgent) DO net stop %%a
>>%ORACLE_FO%\stop_service.log
```

The `DB_STOP.SQL` script used in the pre-failback script is included below.

DB_STOP.SQL

```
# Sample db_stop.sql script called during the Double-Take pre-failback script.
# Shutdown the database immediately.
shutdown immediate
exit
```

3. Select **Start, Programs, Double-Take, Failover Control Center**.
4. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, and click **OK**.
5. To add a monitor for the selected target, click **Add Monitor**. Type the name of the source machine and click **OK**. The Monitor Settings window will open.
6. In the Monitor Settings window, mark the IP address that is going to failover and verify that **Adding Source Identity to Target** is selected.



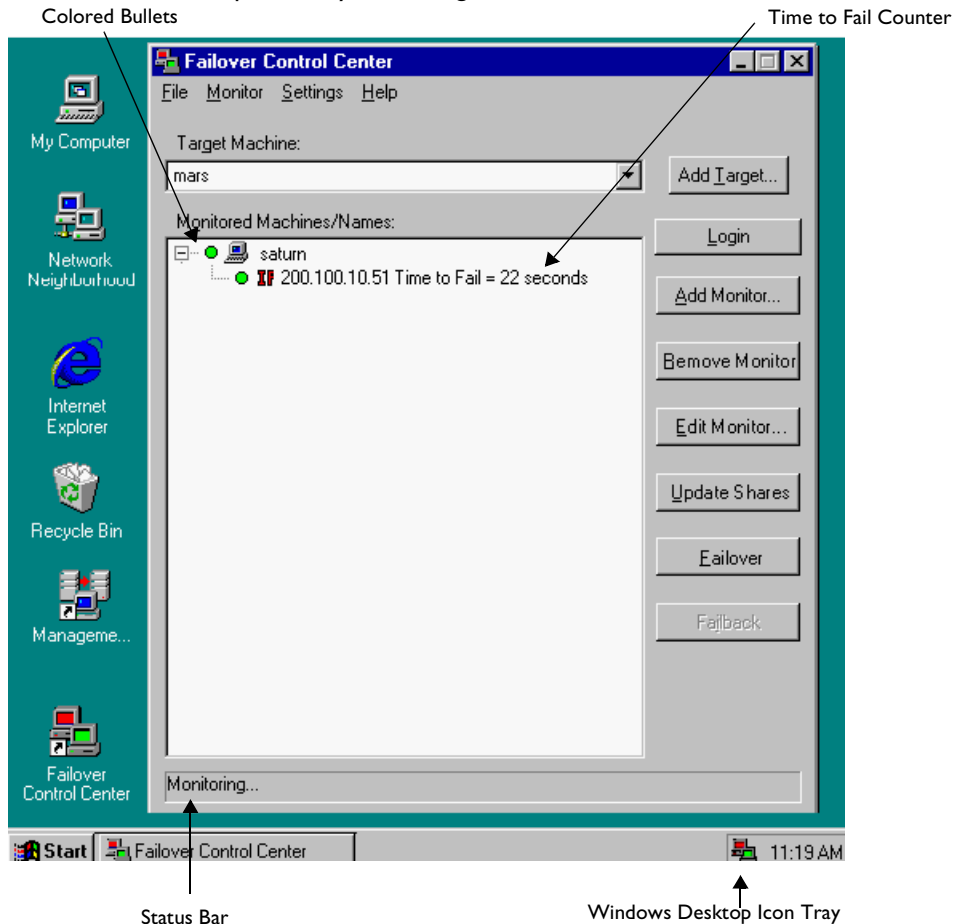
7. Click **Scripts** and specify the scripts that were created previously.
8. Click **OK** to go back to the Monitor Settings dialog box.
9. Click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source.

Monitoring Failover

Now that replication and failover monitoring are configured and started, you will need to know if and when there is a problem. Since it can be essential to quickly know the status of your machines, Double-Take offers various methods for monitoring the status of failover. When the Failover Control Center is running, you will see four visual indicators:

- ◆ The Failover Control Center Time to Fail counter
- ◆ The Failover Control Center status bar located at the bottom of the window
- ◆ The Failover Control Center colored bullets to the left of each IP address and source machine
- ◆ The Windows desktop icon tray containing a failover icon



NOTE: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

The following table identifies how the visual indicators change as the status of failover changes.

	Time to Fail Countdown	Status Bar	Colored Bullets	Desktop Icon Tray
Source is Online	The Time to Fail counter is counting down and resetting each time a heartbeat is received from the source machine.	The status bar indicates that the target machine is monitoring the source machine.	The bullets are green. ^a	The Windows desktop icon tray contains a failover icon with red and green computers.
Source Fails and Failover is Initiated	The Time to Fail countdown value is 0.	The status bar displays the source machine and IP address currently being assumed by the target.	The bullets are red.	The Windows desktop icon tray contains a failover icon with red and green computers.
Failover is Complete	The Time to Fail counter is replaced with the "Failed Over" message.	The status bar indicates that monitoring has continued.	The bullets are red.	The Windows desktop icon tray contains a failover icon with a red computer.

a. When the **Time to Fail** value has decreased by 25% of the entire timeout period, the bullet changes from green to yellow, indicating that the target has not received a response from the source. The yellow bullet is a caution signal. If a response from the source is received, the countdown resets and the bullets change back to green. If the countdown reaches zero without the target receiving a response from the source, failover begins.

Once failover is complete, any clients logging into the server will be automatically directed to the target. Depending on the type of client software being used, clients with an existing connection may have to be restarted in order to establish a connection with the target.

NOTE: For additional detailed information on failover and other monitoring tools, see the *Double-Take User's Guide*.

Restoring Your Oracle Data

If your source experiences a failure, such as a power, network, or disk failure, your target machine will stand in for the source while you resolve the source machine issues. During the source machine downtime, data is updated on the target machine. When your source machine is ready to come back online, the data is no longer current and must be updated with the new data on the target machine.

1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.

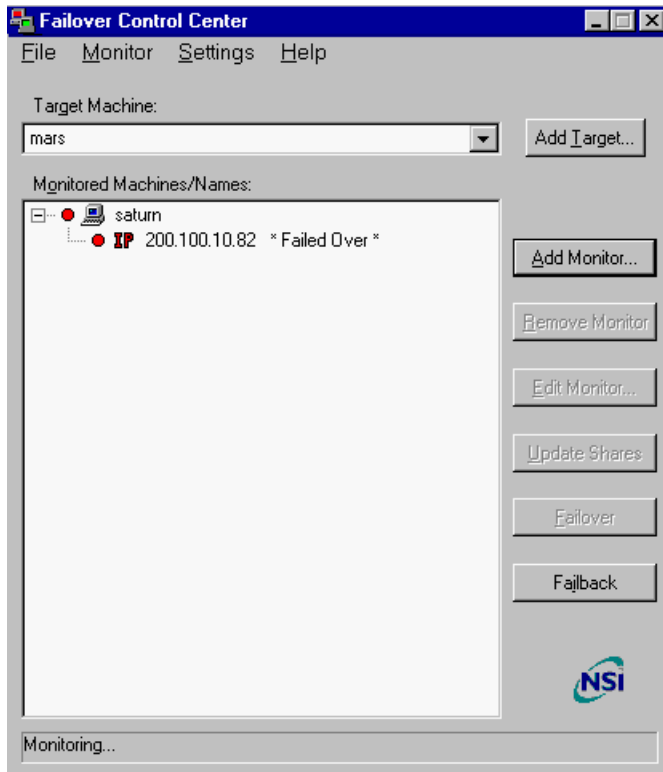
NOTE: If you must rebuild your hard drive, continue with step 3. If you do not need to rebuild your hard drive, continue with step 7.

3. Install Windows 2000 or NT. Since your source machine is not connected to the network, go ahead and use the source's original name and IP address.
4. Install Double-Take using the installation defaults.

NOTE: Verify that the **transactional applications** option is selected on the Double-Take Optimizations screen since Oracle is a transactional database application. See the Double-Take *Getting Started* guide for further details.

5. Install Oracle using the installation settings previously used.
6. Set the Oracle service to manual startup.
 - a. In **Control Panel, Administrative Tools, Services**, double-click the Oracle service.
 - b. Click the **General** tab.
 - c. Select **Manual** from the **Startup type** drop-down list and click **OK**.
7. **Verify that Oracle is not running on the source.** The Oracle services must not be running at this time. Depending on the type of failure, your services may be set to manual startup but could still be running. **Stop your Oracle services and set them to manual.**
8. Select **Start, Programs, Double-Take, Failover Control Center**.
9. Select the target machine that is currently standing in for the failed source.

-
10. Select the failed source and click **Failback**.



The pre-failback script entered during the failover configuration stops the Oracle services on the target so that no additional changes can be made.

11. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
12. Connect the source machine to the network.
13. After the source is back online, select whether or not you want to continue monitoring this source machine (**Continue** or **Stop**).

-
14. To begin the restoration process, open the Double-Take Management Console and select **Tools, Restoration Manager**.

NOTE: You can also run the Double-Take DTCL automated restoration script, which can be found in the Double-Take *User's Guide*, to complete the remaining steps in this section.

15. Complete the appropriate fields as described below.
- ◆ **Original Source**—The name of the source machine where the data original resided.
 - ◆ **Restore From**—The name of the target machine that contains the replicated data.
 - ◆ **Replication Set**—The name of the replication set to be restored.
 - ◆ **Restore To**—The name of the machine where you the data will be restored. This may or may not be the same as the original source machine.
16. Identify the correct drive mappings for the data and any other restoration options necessary. For detailed information on the restoration options, see the Double-Take *User's Guide*.
17. Verify that the selections you have made are correct and click **Restore**. The restoration procedure time will vary depending on the amount of data that you have to restore.
18. After the restoration is complete, start the Oracle services on the source machine.
19. Reestablish the Double-Take Oracle replication set connection.

At this time, your data is restored back to your source machine, the source machine is again the primary Oracle server, and, if you selected to continue failover monitoring, the target is available to stand in for the source in the event of a failure.