Guidelines for using Oracle® 10g with Double-Take® Linux

Guidelines for using Oracle 10g with Double-Take for Linux

Revision 1.0.0 published November 2007

Double-Take, GeoCluster, and NSI are registered trademarks of Double-Take Software, Inc. Balance, Double-Take for Virtual Systems, Double-Take for Virtual Servers, and Double-Take ShadowCaster are trademarks of Double-Take Software, Inc. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the United States and other countries. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Microsoft, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective companies.

© 2007 Double-Take Software, Inc. All rights reserved.

Double-Take® Support for Application Failover

The Double-Take® file system replication process is application independent and replicates any file system changes (including permissions and attributes) written to any Linux file system (ext2 and ext3) by any application or process, subject to specific exceptions called out in the *User's Guide* or readme 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 Double-Take, and Double-Take Software offers 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.

Double-Take 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, Double-Take 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 Double-Take Software has no direct access to or experience with a particular application or configuration, Double-Take 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, Double-Take 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 Double-Take Software to any obligation. Although we try to provide quality information, Double-Take Software makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained in this document.

Table of Contents

1. Document Overview

This document is a Double-Take® application note. An application note provides guidelines on the use of Double-Take in a specific environment.

This document contains:

- **Document Overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Solution Overview**—Explains how the application works with Double-Take and describes the considerations that you must weigh when implementing your Double-Take solution. Review this section to make sure that you understand the theory involved with using Double-Take and your application. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.
- **Sample Implementation**—Describes a specific example of how to use Double-Take for this solution. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

1.1. Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Double-Take solution is to be deployed. You many need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of:

- Double-Take
- Oracle®

1.2. Expectations

Application notes are intended to provide a framework for configuring a Double-Take solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guidelines as well as an example configuration procedure that has been tested for a specific environment.

This document assumes that you are comfortable working with your operating system, Double-Take, and Oracle.

1.3. Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application, and Double-Take. This application note does not provide step-by-step instructions for using standard operating system, application, and Double-Take functionality.

The following documents contain additional information that you may need while setting up this solution:

- Double-Take *User's Guide* or online documentation
- Reference guides or documentation for Oracle

1.4. Getting help

Double-Take Software has application notes that describe how to configure Double-Take with a variety of popular third-party applications. These application notes are available on the Application Notes page of the Double-Take Software support web site (`http://support.doubletake.com`).

For help using Double-Take, refer to the Double-Take online manual or online help.

Double-Take Software offers professional services on a fee basis to assist you in identifying the best practices for implementing a solution in your environment. Visit `http://www.doubletake.com/what-we-offer/services/` for more information.

2. Solution Overview

Oracle is a scalable, reliable, flexible, and high-performance relational database management system for many server-based operating systems. 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 to provide high availability for Linux® servers running Oracle 10g. 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.

2.1. Requirements

- Two servers running Red Hat® Enterprise Linux® version 4.2 through 4.5

  NOTE: The two servers should both be running the same operating system.

- Both servers must be connected to the same physical network
- Two licensed copies of Oracle 10g (the number of licenses required may vary based on your individual Oracle license agreement)
- Two licensed copies of Double-Take for Linux

If you are using Oracle 10g for Linux servers (64Bit), the following minimum versions are required:

- Two servers running Red Hat Enterprise Linux 4.2 through 4.5 (64-bit)

  NOTE: The two servers should both be running the same operating system.

- Both servers must be connected to the same physical network
- Two licensed copies of Double-Take for Linux
- Two licensed copies of Oracle 10g for Linux (x64) (the number of licenses required may vary based on your individual Oracle license agreement)
- The system hardware must support 64-bit applications

2.2. Permissions

The source and target machines must share the same user accounts that have the same GUID. Double-Take Software strongly recommends that customers use unified network accounts in order to facilitate transferring the database from the source to the target.

2.3. Modifying the sample script files

Sample script files are available at `http://support.doubletake.com`.

After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Double-Take is installed.

The sample script files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

For Oracle, you must verify that the sample scripts include the correct version number in the directory path for the home directory.

3. Sample Implementation

This section describes an example of how to configure Double-Take and Oracle. Use these procedures as a guideline for creating your own implementation.

Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

3.1. 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.

3.2. Install and configure software on the target

1. Install Oracle on the target using the same installation options used on the source machine.
2. Stop the Oracle service.
   a. In the application bar, select **Applications, System Settings, Server Settings, Services**.
   b. Select the Oracle service, then select **Stop**.
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.

   Create an `/etc/init.d/` Oracle script to start and stop the Oracle service manually. An example `init.d` script is included in ***Configure failover and begin failure monitoring***. Once configured properly in a Linux server, the Oracle service can be started or stopped by executing the command

   `# /etc/init.d/oracle start / stop`.
4. The command `/etc/init.d/oracle start/stop` can be used inside a post-failover and pre-failback script to start and stop Oracle services on the target server.
5. Install Double-Take on the target using the installation defaults. See the Double-Take *Getting Started* guide for details.

3.3. Configure and begin mirroring and replication

1. From the machine where the Double-Take Management Console in installed, select **Start, Programs, Double-Take, Linux Management Console**.
2. Double-click your source machine to log on.
3. Right-click your source machine and select **New**, **Replication Set** and enter the desired name for the replication set.
4. Select the Oracle data you wish to protect. The following table 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/` directory in the destination drive. If you supplied an instance name of `orcl`, your `init.ora` file will be located in the `/oracle/admin/orcl/pfile/ directory` 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 may 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 Linux **Search for files** feature.

File Name	File Description	Default Location	Table to Query
`init.ora.xxxxx`	Initialization Parameter File	`/oracle/product/10.2.1/admin/`	N/A
`control#.ctl`	Control File	`/oracle/oradata/instance_name/`	V$CONTROLFILE
`redo#.log`	Online Redo Log	`/oracle/oradata/instance_name/`	V$LOGFILE
`system*.dbf`	System Tablespace	`/oracle/oradata/instance_name/`	V_$DATAFILE
`undotbs*.dbf`	Undo Tablespace	`/oracle/oradata/instance_name`	V_$DATAFILE
`sysaux*.dbf`	System Auxiliary Tablespace	`/oracle/oradata/instance_name/`	V_$DATAFILE
`users#.dbf`	Users Tablespace	`/oracle/oradata/instance_name/`	V_$DATAFILE
`example*.dbf`	Example Tablespace	`/oracle/oradata/instance_name/`	V_$DATAFILE
`temp*.dbf`	Temporary Tablespace	`/oracle/oradata/instance_name/`	V_$DATAFILE
`listener.ora`	Listener File	`/$ORACLE_HOME/network/admin`	N/A
`tnsnames.ora`	Service Configuration File	`/$ORACLE_HOME/network/admin`	N/A
`pwd*.ora`	Password Information File	`/$ORACLE_HOME/dbs/`	N/A
`spfile*.ora`	Server Parameter File	`/$ORACLE_HOME/dbs/`	N/A

NOTE: There is a unique numeric identifier at the end of the `init.ora` file (xxxxx) that uniquely identifies initialization files when there is more than one installed.

Besides the initialization parameter file (`init.ora`), which is stored in a text format, administrators can choose to maintain the initialization parameters in a binary server parameter file (`spfile.ora`). This is installed by default.

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.

5. 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.
6. Right-click on the source and select **Properties**.
7. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
8. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.
9. 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, see the Double-Take *User's Guide*.)
10. Select **One to One** to map the replication set data from the source to an identical volume/directory structure on the target.
11. Click **Connect** to start the mirror and replication processes.

NOTE: If you start Oracle and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. If the updated data on the target is needed, restore the data from the target to the source.

3.4. Configure failover and begin failure monitoring

1. If a failure occurs, you will want to have the Oracle services start on the target machine automatically. To do this, create a shell script file called `POST_FAILOVER.SH` and save the file to the `/usr/bin/` directory where the Double-Take commands are installed. The `POST_FAILOVER.SH` script should include the following contents

```
*******************************************************
#!/bin/bash
/etc/init.d/oracle start
*******************************************************
```

2. The following example shows the `/etc/init.d/oracle` script and SQL script to start and stop Oracle services.

Sample_Init.d

```
############***** BEGINING OF INIT.D SCRIPT FOR ORACLE SERVICE ***** #############

#!/bin/bash
ORACLE_USER="oracle"
##### Path of installed oracle home directory in the server. #####
ORACLE_HOME="/u01/app/oracle/product/10.2.1/db_1/"
##### Set user name. #####
SYS_USER="sys"
##### Set password of the username. ######
SYS_PASS="password"
case "$1" in
start)
echo "*** Starting Oracle *** "
su - $ORACLE_USER -c "$ORACLE_HOME/bin/lsnrctl start"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/ora_start"
###### here "ora_start" is a shell script created inside $ORACLE_HOME/bin directory. This script contains
###### the following sqlplus command :--
###### ----> $sqlplus "sys/password as sysdba" @/home/oracle/start.sql
###### Again /home/oracle/start.sql is the path of a sql script "start.sql" used for starting the database.
###### Following are the contents of sql script.
###### ----> startup ; exit ;
;;
stop)
echo "*** Stopping Oracle *** "
su - $ORACLE_USER -c "$ORACLE_HOME/bin/lsnrctl stop"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/ora_shut"
###### here "ora_shut" is a shell script created inside $ORACLE_HOME/bin directory. This script contains
###### the following sqlplus command :--
###### ----> $sqlplus "sys/password as sysdba" @/home/oracle/shut.sql
###### Again /home/oracle/shut.sql is the path of a sql script "shut.sql" used for shutting down the
###### database. Following are the contents of sql script..
###### ----> shutdown immediate ; exit ;
;;
isqlstart)
echo "***** Starting Oracle iSQL Plus ****"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/isqlplusctl start"
echo "***** Note : You can access service at url: http://$(hostname):5560/isqlplus"
;;
isqlstop)
echo "***** Stopping Oracle iSQL Plus ****"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/isqlplusctl stop"
;;
emstart)
echo "***** Starting Oracle Enterprise Manager 10g Database Control ****"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/emctl start dbconsole"
echo "***** Note : You can access service at url: http://$(hostname):1158/em"
;;
emstop)
echo "***** Stopping Oracle Enterprise Manager 10g Database Control ****"
su - $ORACLE_USER -c "$ORACLE_HOME/bin/emctl stop dbconsole"
;;
*)
echo $"Usage : $0 {start|stop|isqlstart|isqlstop|emstart|emstop}"
exit 1
esac
exit 0
############# ***** END OF INIT.D SCRIPT FOR ORACLE SERVICE ***** #############
```

```
*****************************************************************************************************
NOTE :--- Copy this script to the /etc/init.d/ directory of the linux server, and name it as "oracle".
Create a soft link (Run level 3)
$ cd /etc/rc.d/rc3.d/
$ ln -s /etc/init.d/oracle S100oracle
$ ln -s /etc/init.d/oracle K100oracle
Alternatively use chkconfig command to add script:
$ chkconfig --add oracle
$ chkconfig --list oracle
Please note that you can just use script directly, To start Oracle type command:
$ /etc/init.d/oracle start
To stop Oracle type command:
$ /etc/init.d/oracle stop
*****************************************************************************************************
NOTE ---> Create the following scripts at the directories as mentioned in the INIT.D script, if you are using
this script for oracle service.
ora_start ---> $ORACLE_HOME/bin/
ora_shut ---> $ORACLE_HOME/bin/
start.sql ---> /home/oracle/
shut.sql ---> /home/oracle/

*****************************************************************************************************
```

3. 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 shell script file called `PREFAILBACK.SH` and save the script file to the `/usr/bin/` directory where the Double-Take commands are installed. The `PREFAILBACK.SH` script should include the following:

```
****************************************************************
#!/bin/bash
/etc/init.d/oracle stop
****************************************************************
```

4. Select **Start**, **Programs**, **Double-Take**, **Failover Control Center**.
5. 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**.
6. 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.
7. In the Monitor Settings window, mark the IP address that is going to failover.
8. Click **Scripts** and specify the scripts that were created previously, using `POST_FAILOVER.SH` for the target post-failover script and `PREFAILBACK.SH` for the target pre-failback script.

NOTE: The scripts are processed using the same account running the Double-Take service.

9. Click **OK** to go back to the Monitor Settings dialog box.
10. 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. For details on monitoring failover, see the Double-Take *User's Guide*.

NOTE: In the event of a source machine failure and the target machine stands in, users that were connected during the time of the failure will be required to relogin to the database.

3.5. 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 6.

3. Install your operating system. 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.
5. Install Oracle using the installation settings previously used.
6. On the source, stop the Oracle service so that all Oracle files are closed on the source and the Double-Take target can restore the changes back to the source.
7. Select **Start, Programs, Double-Take, Failover Control Center**.
8. Select the target machine that is currently standing in for the failed source.
9. 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.
10. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
11. Stop the Oracle service on the source.
12. Delete the `.db` and `.sts` files from the `/var/lib/DT` directory.
13. Restart the Double-Take service (`#service DT restart`).
14. Connect the source machine to the network.
15. After the source is back online, select whether or not you want to continue monitoring this source machine (**Continue** or **Stop**).

NOTE: Verify that the Double-Take connection on the source has been disconnected. If it is not, right-click the connection in the Double-Take Management Console and select **Disconnect**.

16. Since the replication sets were deleted, you will need to re-create those replication sets.
17. 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.

18. Complete the appropriate fields.
    - **Original Source**—The name of the source machine where the data originally 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 the data will be restored. This may or may not be the same as the original source machine.
19. 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*.
20. 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.
21. After the restoration is complete, start the Oracle services on the source machine.
22. 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.