



Hitachi Dynamic Replicator VX Solution for Oracle (Solaris)

Copyright © 2009 Hitachi, Ltd., Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements. All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.


Hitachi is a registered trademark of Hitachi, Ltd. in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd. in the United States and other countries.

All other trademarks, service marks, and company names in this document are properties of their respective owners.

Contents

1	Introduction to the Solution	5
2	Protecting Oracle.....	6
2.1	Discovering Oracle volumes.....	6
2.2	Protecting Oracle Volumes	7
2.3	Discovering replication details.....	10
2.4	Consistency	12
3	Failover.....	15
3.1	Planned Failover through CX UI.....	15
3.2	Planned Failover through CLI.....	18
3.3	Unplanned Failover through CX UI	19
3.4	Unplanned failover through CLI	20
4	Failback.....	21
4.1	Oracle Discovery	22
4.2	Reverse replication.....	23
4.3	Planned failover.....	23
5	Protecting Oracle database running on Solaris Zones	24
5.1	Introduction to the solution	24
5.2	Protect	25
5.2.1	Prepare	25
5.2.2	Discover Solaris Zones	25
5.2.3	Replicate Solaris zones / Database.....	26
5.2.4	Consistency tags.....	26
5.3	Recovery	28
5.3.1	Rollback all target volumes	28
5.3.2	Recovering zones	29
5.3.3	Recovering Oracle database	31
6	Trouble shooting	32

Conventions

Keywords, command buttons and other such fields are enclosed in “ ” while being bold (for example, to denote  “Next” is used).

Inputs for commands and Variables are shown in *Italics*

File names and paths are shown in **bold**

Commands are shown in **Courier new font**

Optional keywords and arguments are enclosed within [].



Notes:

Contain suggestions or tips.



Caution:

Contains critical information

Purpose of the document

This document will attempt to uncover Oracle Database support for UNIX-based platforms.

Supported OS & Database configurations

- Solaris 8,9, and 10 on sparc machines
- Solaris 10 on x86 64 bit machines.
- Oracle database 10g and oracle database 11g.

Before you begin

- VX agent installed on production and DR servers
- FX agents are installed on production and DR servers
- CX server is up and running
- All agents are pointed to the same CX server with valid licenses
- The database volume should be mounted on the same mount point for both source and target host (for example, “/u01/db1”).
- Oracle is installed on both source and target servers and oracle is running on source production server.
- If the Oracle database was not created through dbca (database configuration assistant) then add an entry (e.g., SID,ORACLE HOME in “/var/opt/oracle/oratab” on target host).

1 Introduction to the Solution

This document is divided into three major sections

[Protect](#)

This section explains how to protect oracle and the steps involved. Oracle is protected in three steps

The first step in protecting Oracle is to perform Oracle Discovery. This is done through the command line of the production machine (source host). Discovery is performed whenever a new data base or a database volume is added to oracle. Discovery lists all the volumes that are to be protected.

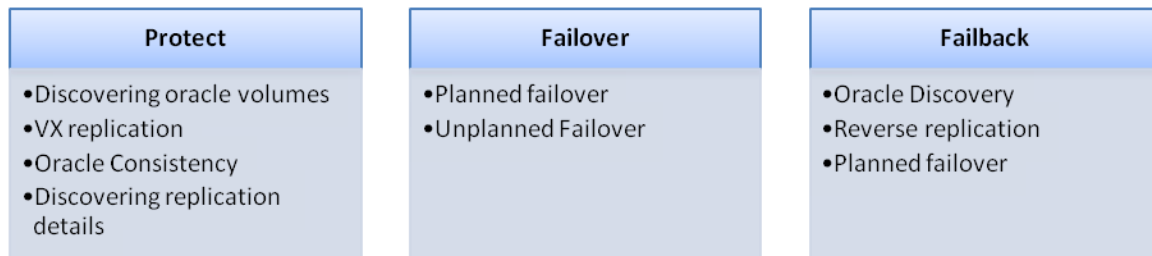


Figure 1

VX replication

All the discovered volumes are replicated to the DR Oracle server

Oracle Consistency

Consistency tags are issued on the production volume at regular intervals

Discovering replication details

All information required to perform an unplanned failover is collected from the source on a daily basis and backed up to the target host through a single FX job.

[Failover](#)

A Failover is performed when there is an outage on the production server. If the outage is expected then a planned failover is performed. For unexpected failover an unplanned failover is performed. Each of them can be performed through CLI or CX UI.

[Failback](#)

A failback is performed once the production server has recovered from the outage and is ready to resume its operations. To perform a failback a reverse replication is set from DR oracle Server to Production server and then a planned failover is performed.

2 Protecting Oracle

Oracle is protected in three steps, the first of them being Oracle Discovery followed by VX replication and Oracle Consistency.

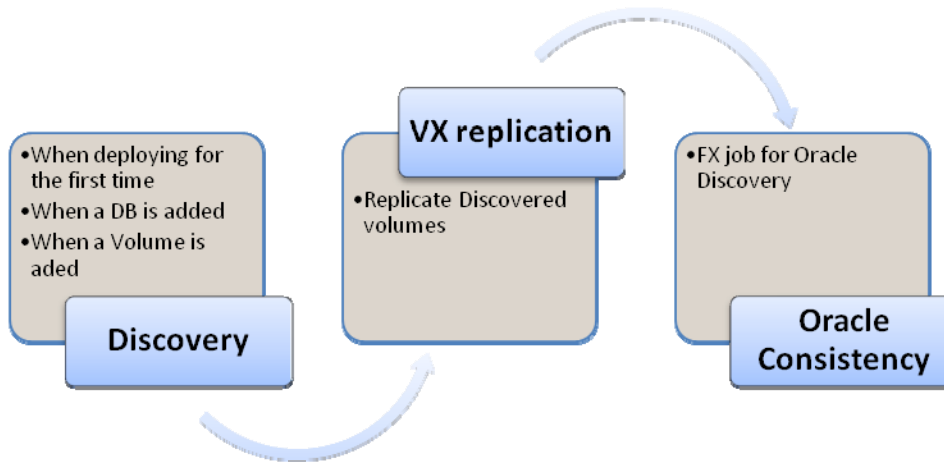


Figure 2

2.1 Discovering Oracle volumes

Enterprise databases like oracle usually span over multiple volumes. To get a list of volumes being used by oracle, access the console of the source host and navigate to the VX agent installation folder and change the directory to “**scripts**”.

Then run the “**application.sh**” script, the syntax is:

```
./application.sh -a oracle -h <ORACLE_HOME> -c <ORACLE_SID> -u <ORACLE_OS_USER>
```

The below picture shows the output for oracle discovery

```
[root@oral scripts]# ./application.sh -a oracle -h /orahome/product/10.2.0/db_1/ -c hyd -u oracle
Oracle_home = /orahome/product/10.2.0/db_1/
Oracle_sid = hyd
ORACLE_USER = oracle
Verifying Database.....
hyd using following filesystems...
/oraarch → Volume to be replicated
Putting the Database hyd in begin backup mode....
Database altered.
Tag:oracle_20080611203125
Generating VACP tag for the below file systems...
Issuing the Consistency tag for the below file systems...
```

Figure 3

Where

ORACLE_home: Oracle home directory

ORACLE_SID : Oracle instance name

ORACLE_OS_USER: Oracle user name

2.2 Protecting Oracle Volumes

All discovered volumes are replicated to the DR server.

Step 1. Setup a replication pair for the volumes shown in the above step. Since there are two volumes that oracle is using, we replicate both the volumes to the remote site. Login to the CX UI, click on Volume protection and select the source volume and click on “**Start Replication**”

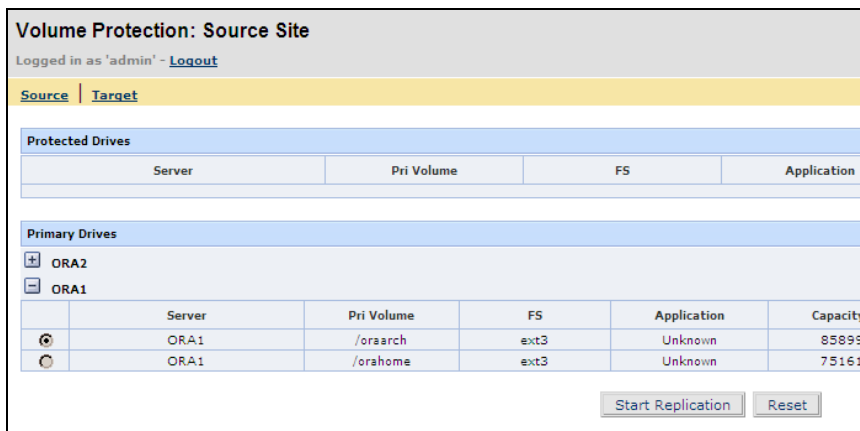


Figure 4: Selecting first source volume

Step 2. Select the target volume and scroll down to set the “**Replication Options**”

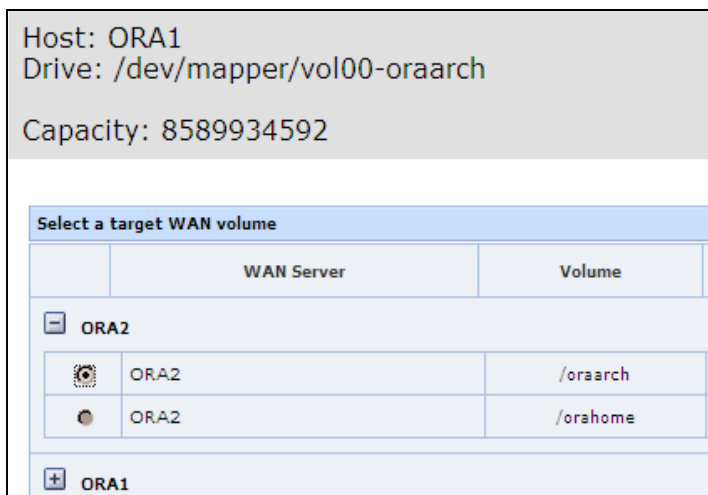


Figure 5: Selecting target volume

Step 3. All the process servers pointed to the CX server are listed here, you may choose a desired process server which will handle all the offload activities specific to this replication pairs. For better scalability you may point more process servers to the CX server. Select the process server and scroll down to set the “**Replication Options**”

Process Server	
<div> <div>Select Process Server</div> <div> ITGTR5U2-64.inmage.in(10.0.164.73) w2k3-32PPS(10.0.227.227) SR5U2-64.inmage.in(10.0.248.73) </div> </div>	<div>Number of Pair Configured</div> <div>2</div>

Figure 6

Step 4. Encryption and compression are optional. “**CDP retention option**” however needs to be enabled then click on “**Submit**”

Replication Options	
<input type="checkbox"/>	Secure transport from Source to InMage CX
<input type="checkbox"/>	Secure transport from InMage CX to destination
<input checked="" type="checkbox"/>	Use fast resync instead of off-load resync algorithm
Use compression: CX Based Compression <small>(Overrides existing 1-N replication pairs)</small>	
Add to volume consistency group: New Volume Group	
CDP Retention	
<input checked="" type="checkbox"/>	Enable CDP Retention option
Automatic Resync Options	
<input type="checkbox"/>	Start between hours 18 : 00 and 6 : 00 after waiting 30 minutes. (All times are local to CX)
<div> <input type="button" value="Submit"/> <input type="button" value="Cancel"/> <input type="button" value="Reset"/> </div>	

Figure 7: Enable CDP retention option

Step 5. Define the retention policy in the next screen and click on “**Submit**”.

Volume Protection: Retention Options
 Logged in as 'admin' - [Logout](#)

Pair Details			
Server	Pri Volume	Remote Server	Volume
ORA1	/dev/mapper/vol00-oraarch	ORA2	/dev/mapper/vol00-oraarch

Retention Logging Policy

Retention Policy	Roll-backward		
Retention Log Size	0.00 (MB)	Current Retention Log Size	0.00 (MB)
Unused Space	256.00 (MB)		
Retain changes upto	<input type="text" value="256"/> <input type="text" value="MB"/> (Cannot be less then 256 MB)		
Retain changes upto the (time)	<input type="text"/> (Days) <input type="text"/> (hrs.)		
On insufficient disk space	<input checked="" type="radio"/> Purge older retention logs <input type="radio"/> Pause differentials		
Log data directory	<input type="text" value="/orahome"/> (Eg:- K:\log_data) /orahome drive is suggested for storing rollback log files.		

Configure Threshold for Alerts

Alert when disk space utilization reaches	<input type="text" value="80"/> %
---	-----------------------------------

Figure 8: Define retention policy

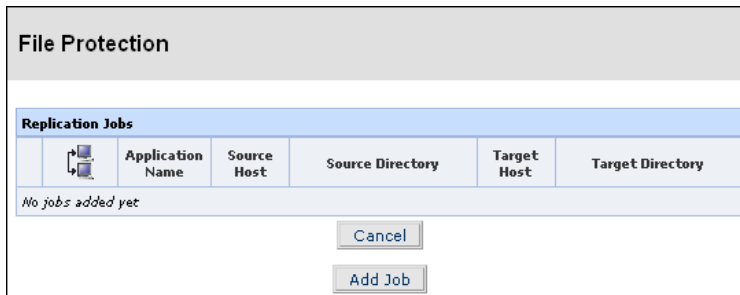
Step 6. Repeat the same process for the other volume and wait until both the replication pairs enter “**Differential Sync**”

Protection Status										
Logged in as 'admin' - Logout										Server Time: Jun-9-2008 16:54:40
Volume Protection Status										
Server	Volume	Group	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)	Resync progress	RPO	Status	Resync Required	View Details
ORA1->ORA2	/oraarch->/dev/mapper/vol00-oraarch	Volume /dev/mapper/vol00-oraarch	0	3.63	0.02	N/A	6.35 minutes	Differential Sync	NO	

Figure 9:

2.3 Discovering replication details

Step 1. Once the replication pair reaches “**Differential Sync**”, setup a single FX job to discover replication details on a daily basis. The objective of this is to ensure that all the required details are present on the target host to perform an unplanned failover when required. While remaining on the CX UI, click on “**File Protection**” and then click on “**New Job Group Wizard**”. You should see the File Protection screen as shown below.



File Protection

Replication Jobs					
	Application Name	Source Host	Source Directory	Target Host	Target Directory
No jobs added yet					

Figure 10

Step 2. Enter the “**Application Name**” and “**Job Description**” then select the source and target hosts. Select the FX template as “**Oracle (Unix/Linux) Discovery**” and click on “**Next**” to continue



File Protection Wizard: Replication Pair

Replication Hosts

Application Name:

Job Description:

Source		Destination	
Host	Host	Host	Host
<input checked="" type="radio"/> ORA1 [Unix]	<input checked="" type="radio"/> ORA1 [Unix]	<input type="radio"/> ORA1 [Unix]	<input type="radio"/> ORA1 [Unix]
<input type="radio"/> ORA2 [Unix]	<input type="radio"/> ORA2 [Unix]	<input checked="" type="radio"/> ORA2 [Unix]	<input checked="" type="radio"/> ORA2 [Unix]

Directory:

Figure 11

Step 3. You should now see the FX “**Job Options**” screen, scroll down to fill up the “**pre execution script pathname**”. Most of the script is filled up, you will need to enter the oracle user name, oracle home and oracle SID as indicated.

```
"/usr/local/InMage/Vx/scripts/application.sh -a oracle -discover -u <ORACLE_USER_NAME> -h
<ORACLE_HOME> -c <ORACLE_SID>"
```

Similarly enter the oracle SID for the “**post execution script pathname (destination)**”, click on “**Finish**” to continue.

```
"/usr/local/InMage/Vx/scripts/oracletargetdiscovery.sh <ORACLE_SID>"
```

Figure 12

Step 4. The next screen opens up, this FX job is set to run once a day and you may change this by clicking on the “**Set Schedule**” command button. Click on “**Finish**” to save the job

Figure 13



Notes:

We recommend that you run this job manually as soon as you configure it.

2.4 Consistency

Issuing consistency tags on the production volume is vital to a successful recovery. Under unexpected outages, the DR server will be recovered to a latest consistent point. Issuing a consistency tag regularly minimizes data loss.

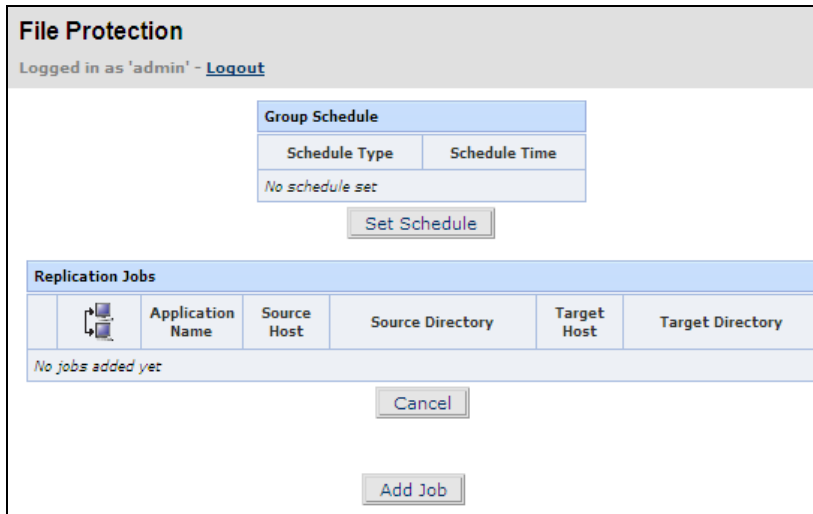
Step 1. Once the above replication pair enters “**Differential Sync**”, proceed to setup a consistency job. Open the CX UI, and then click on “**File Protection-> New Job Group Wizard**”



The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as 'admin' - Logout'. Below this is a section titled 'Applications' with a filter dropdown set to 'Application Name' and 'Application Per Page' set to 5. The main area displays 'No Replication Applications To Display'. At the bottom, there are two buttons: 'New Job Group Wizard' and 'Manage Templates'.

Figure 14

Step 2. The next screen opens up, click on “**Add Job**” to continue



The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as 'admin' - Logout'. Below this is a section titled 'Group Schedule' with a table for 'Schedule Type' and 'Schedule Time'. The table is empty, and the text 'No schedule set' is displayed. Below the table is a 'Set Schedule' button. Below this is a section titled 'Replication Jobs' with a table for 'Application Name', 'Source Host', 'Source Directory', 'Target Host', and 'Target Directory'. The table is empty, and the text 'No jobs added yet' is displayed. Below the table are 'Cancel' and 'Add Job' buttons.

Figure 15

Step 3. The FX replication pair screen opens up as shown below. Enter the “**Application name**”, “**Job description**” then select the production Oracle server for “**source**” and DR Oracle server for “**Destination**”. From the drop down menu, select the FX template as “**Oracle (UNIX) Consistency**”. Then click on “**Next**”

File Protection Wizard: Replication Pair
 Logged in as 'admin' - [Logout](#)

Replication Hosts

Application Name:

Job Description:

Source		Destination	
	Host		Host
<input checked="" type="radio"/>	ORA1 [Unix]	<input type="radio"/>	ORA1 [Unix]
<input type="radio"/>	ORA2 [Unix]	<input checked="" type="radio"/>	ORA2 [Unix]
Directory		Directory	
<input type="text"/>		<input type="text"/>	

Figure 16

Step 4. FX job Options screen opens up, scroll down to “**Miscellaneous Options**” to fill up the pre script

Send RPO alert if minutes passed

Send E-mail alert if minutes passed without job progress

Pre execution script pathname

Post execution script pathname

Pre execution script pathname (destination)

Post execution script pathname (destination)

Catch All job modifier for power users only

Figure 17

Syntax:

```
/usr/local/InMage/Vx/scripts/application.sh -a oracle -h <ORACLE_HOME>
-c <ORACLE_SID> -u <ORACLE_OS_USER>
```

Then click on “**Finish**” to proceed to the next step

Step 5. You will be returned to one of the previous screens with all the values filled up, by default the consistency job is scheduled to run every fifteen minutes. To change the frequency of execution, click on **“Set Schedule”**. Click on **“Finish”** to save the job.

File Protection



Logged in as 'admin' - [Logout](#)

Group Schedule

Schedule Type	Schedule Time
Run Every	15 Minutes

Set Schedule

Replication Jobs

		Application Name	Source Host	Source Directory	Target Host	Target Directory
Run order 1						
		Oracle	ORA1	/usr/local/InMage/Vx/failover_data/	ORA2	/usr/local/InMage/Vx/failover_data/

Details

Remove

Cancel

Add Job

Finish

Figure 18

3 Failover

3.1 Planned Failover through CX UI

A planned failover is preferred either to test the setup or when an outage is expected (for example, a mock drill).

Step 1. A planned failover can be performed either through CX UI or CLI. To perform a planned failover through CX UI, setup a single FX job with the pre defined template.

Open the CX UI, click on “**File Protection -> New Job Group Wizard**”,

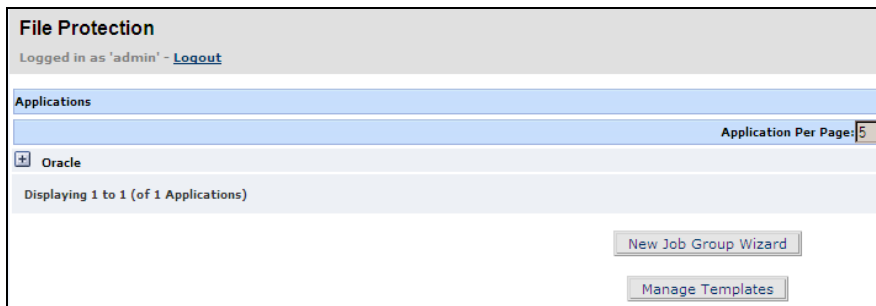


Figure 19

Step 2. The next screen opens up, click on “**Add Job**” to proceed

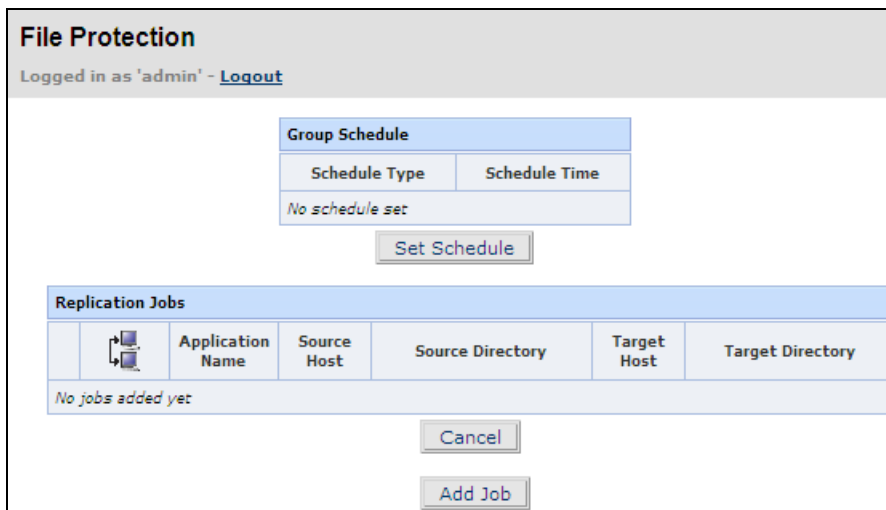


Figure 20

Step 3. Enter the “Application Name”, “Job Description” and then select the production Oracle server for “Source” and DR Oracle server for “Destination”. Select the FX template as “Oracle(UNIX) Planned Failover”. Click on “Next” to proceed.

The screenshot shows the 'File Protection Wizard: Replication Pair' window. At the top, it says 'Logged in as 'admin' - Logout'. Below this is a section titled 'Replication Hosts'. It contains two main columns: 'Source' and 'Destination'. Each column has a 'Host' section with a table of available hosts and a 'Directory' section with a text input field. In the 'Source' column, 'ORA1 [Unix]' is selected. In the 'Destination' column, 'ORA2 [Unix]' is selected. Below these columns is a dropdown menu showing 'Oracle(UNIX) Planned Failover'. At the bottom are 'Next ->' and 'Cancel' buttons.

Figure 21

Step 4. The FX job Options page opens up, scroll down to “Miscellaneous Options” to edit the source pre script and target post script

The screenshot shows the 'FX job Options' page, specifically the 'Miscellaneous Options' section. It contains several fields: 'Send E-mail alert if' (set to 5 minutes), 'Pre execution script pathname' (set to /usr/local/InMage/Vx/scripts/application.sh -a oracle -h /orahome), 'Post execution script pathname' (empty), 'Pre execution script pathname (destination)' (empty), 'Post execution script pathname (destination)' (set to /product/10.2.0/db_1 -c hyd -u oracle -t My_Tag-planned failover), and 'Catch All job modifier' (empty). At the bottom are '<- Back', 'Finish ->', and 'Cancel' buttons.

Figure 22

Source pre script

```
/usr/local/InMage/Vx/scripts/application.sh -a oracle -h <ORACLE_HOME>
-c <ORACLE_SID> -u <ORACLE_OS_USER> -t <name of a tag> -planned
```

Target post script

```
/usr/local/InMage/Vx/scripts/application.sh -a oracle -h <ORACLE_HOME>
-u <ORACLE_OS_USER> -t <name of the same tag as given above> -planned -failover
```

Then click on “Finish”

Step 5. The next screen opens up, the job is set to **“Run On Demand”**. You may choose to change this if desired. Click on **“Finish”** to save the job.

File Protection
Logged in as 'admin' - [Logout](#)

Group Schedule

Schedule Type	Schedule Time
Once At	On Demand

[Set Schedule](#)

Replication Jobs

	Application Name	Source Host	Source Directory	Target Host	Target Directory
Run order 1					
	Oracle	ORA1	/usr/local/InMage/Vx/failover_data/	ORA2	/usr/local/InMage/Vx/failover_data/

[Details](#) [Remove](#) [Cancel](#)

[Add Job](#)

[Finish](#)

Figure 23

Step 6. To start the job, click on **“File Protection”** then select the job and click on **“Start”**. The job completes indicating that the planned failover is complete.

File Protection
Logged in as 'admin' - [Logout](#)

Applications Filter:

Application Per Page:

Job Description	Status	Source Host	Source Directory	Target Host	Target Directory
planned failover...	Not started...	ORA1	/usr/local/InMage/Vx/failover_data/	ORA2	/usr/local/InMage/Vx/failover_data/

Results 1-1 of 1 << < 1 > >>

[Stop](#) [Start](#) [Details](#) [Remove](#)

Displaying 1 to 1 (of 1 Applications)

[New Job Group Wizard](#)

[Manage Templates](#)

Figure 24



Notes:

Ensure that you stop Consistency and Discovery jobs after failover, else they will fail

3.2 Planned Failover through CLI

Step 1. Access the source console and then issue the command

```
./application.sh -a oracle -h <ORACLE_HOME> -c <ORACLE_SID> -u <ORACLE_OS_USER> -t <name of the tag> -planned
```

Step 2. The result of this command shows another command that is to be executed on the target to complete the failover

```
[root@ora1 scripts]# ./application.sh -a oracle -h /orahome/product/10.2.0/db_1/
-c hyd -u oracle -t "My_Tag" -planned
Oracle_home = /orahome/product/10.2.0/db_1/
Oracle_sid = hyd
ORACLE_USER = oracle
Verifying Database.....
hyd using following filesystems...

/oraarch
```

Figure 25:

```
#####
###
application.sh -a oracle -h /orahome/product/10.2.0/db_1/ -c hyd -u oracle -t My_
Tag -planned -failover
#####
#####
```

Figure 26

Step 3. Copy the command shown in the previous step and issue it on the target console, this completes planned failover through console

```
[root@ora2 scripts]# ./application.sh -a oracle -h /orahome/product/10.2.0/db_1/
-c hyd -u oracle -t My_Tag -planned -failover

/oraarch /dev/mapper/vol100-oraarch
Checking the Consistency Tag Status on target.....
Stopping Vx agent....
InMage VX Agent daemon is running...
Stopping InMage VX Agent daemon ...
All VX-related processes are successfully terminated!
```

Figure 27:

3.3 Unplanned Failover through CX UI

Unplanned failover is performed if an outage is unexpected or if the production oracle server is down. To perform an unplanned failover, at least one consistency tag is required on the DR Oracle server.

Step 1. Similar to setting up a planned failover through CX UI, setup a single FX job for unplanned failover. Enter the “**Application Name**”, “**Job Description**” and then select the production Oracle server for Source and DR Oracle server for Destination, then select the FX template as “**Oracle (UNIX) Unplanned Failover**”. Click on “**Next**” to continue

The screenshot shows the 'File Protection Wizard: Replication Pair' window. At the top, it says 'Logged in as 'admin' - Logout'. Below this is a section titled 'Replication Hosts'. It contains two main columns: 'Source' and 'Destination'. Under 'Source', there is a table with two rows: 'ORA1 [Unix]' and 'ORA2 [Unix]', each with a radio button. Under 'Destination', there is a similar table with 'ORA1 [Unix]' and 'ORA2 [Unix]' and radio buttons. Below these tables are fields for 'Directory' and a dropdown menu currently showing 'Oracle(UNIX) Unplanned Failover'. At the bottom are 'Next ->' and 'Cancel' buttons.

Figure 28

Step 2. The FX job Options screen opens up, scroll down to “**Miscellaneous Options**” to edit the target post script as marked in the picture below.

The screenshot shows the 'FX job Options' screen. It has several input fields: 'Pre execution script pathname', 'Post execution script pathname', 'Pre execution script pathname (destination)', 'Post execution script pathname (destination)', and 'Catch All job modifier'. The 'Post execution script pathname (destination)' field is highlighted with a black box and contains the text 'ome/product/10.2.0/db_1 -c hyd -u oracle -unplanned failover'. At the bottom are '<- Back', 'Finish ->', and 'Cancel' buttons.

Figure 29

```
/usr/local/InMage/Vx/scripts/application.sh -a oracle -h  
<ORACLE_HOME> -c <ORACLE_SID> -u <ORACLE_OS_USER> -unplanned -  
failover
```

Step 3. The next screen opens up, click on “**Finish**” to save the job. And Start the job to perform an unplanned failover through CX UI.

The screenshot shows a web interface titled "File Protection". At the top, it says "Logged in as 'admin' - [Logout](#)". Below this is a "Group Schedule" section with a table for "Schedule Type" and "Schedule Time". The table has two rows: "Once At" and "On Demand". Below the table is a "Set Schedule" button. The main section is "Replication Jobs", which contains a table with columns: "Application Name", "Source Host", "Source Directory", "Target Host", and "Target Directory". Below the table, it says "Run order 1". There is one row in the table with the following values: "Oracle", "ORA1", "/usr/local/InMage/Vx/failover_data/", "ORA2", and "/usr/local/InMage/Vx/failover_data/". Below the table are buttons for "Details", "Remove", and "Cancel". At the bottom of the interface are buttons for "Add Job" and "Finish".

Group Schedule	
Schedule Type	Schedule Time
Once At	On Demand

[Set Schedule](#)

Replication Jobs					
	Application Name	Source Host	Source Directory	Target Host	Target Directory
Run order 1	Oracle	ORA1	/usr/local/InMage/Vx/failover_data/	ORA2	/usr/local/InMage/Vx/failover_data/

[Details](#) [Remove](#) [Cancel](#)

[Add Job](#)

[Finish](#)

Figure 30

3.4 Unplanned failover through CLI

To perform an unplanned failover through CLI, access the DR Oracle server's console. Then navigate to the “**scripts**” directory under VX agent installation directory to issue the following command

```
./application.sh -a oracle -h <ORACLE_HOME> -c <ORACLE_SID> -u
<ORACLE_OS_USER> -unplanned -failover
```

```
[root@ora2 scripts]# ./application.sh -a oracle -h /orahome/product/10.2.0/db_1/
-c hyd -u oracle -unplanned -failover

/oraarch /dev/mapper/vol100-oraarch /dev/mapper/vol100-oraarch
Checking the Consistency Tag Status on target.....
Stopping Vx agent....
InMage VX Agent daemon is running...
Stopping InMage VX Agent daemon ...
```

Figure 31

4 Failback

A failback is performed in three steps

- Discovery
- Reverse replication and
- Planned failover




Figure 32

To perform a failback, Oracle Discovery command is executed on the DR server. The DR server will now act as the source for the replication pair. Once the discovery is complete, proceed to setup reverse replication. i.e. from DR server to production server.

Once the production server is back online, all the data changes occurred on the DR server needs to be updated on the production server. To achieve this, set a reverse replication from DR server to the production server.

Once the replication is reaches “**Differential Sync**”, proceed to perform a planned failover as described in [Planned Failover through CX UI](#) on page 15.

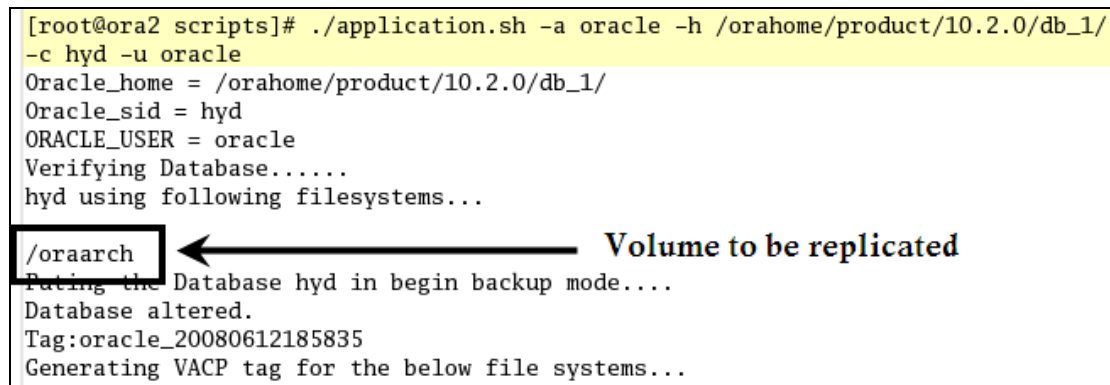
4.1 Oracle Discovery

To get a list of volumes being used by oracle, access the console of the original target host and navigate to the scripts folder under the VX agent installation folder.

Then run the “**application.sh**” script, the syntax is:

```
./application.sh -a oracle -h <ORACLE_HOME> -c <ORACLE_SID> -u  
<ORACLE_OS_USER>
```

The below picture shows the output for oracle discovery



```
[root@ora2 scripts]# ./application.sh -a oracle -h /orahome/product/10.2.0/db_1/  
-c hyd -u oracle  
Oracle_home = /orahome/product/10.2.0/db_1/  
Oracle_sid = hyd  
ORACLE_USER = oracle  
Verifying Database.....  
hyd using following filesystems...  
/oraarch  
Putting the Database hyd in begin backup mode....  
Database altered.  
Tag:oracle_20080612185835  
Generating VACP tag for the below file systems...
```

Figure 33

4.2 Reverse replication

Set a reverse VX replication. i.e. from DR Oracle server to production Oracle server to update all the production oracle server with all the data changes occurred during the outage.

Protection Status

Logged in as 'admin' - [Logout](#)

Server Time: Jun-10-2008 09:13:17

Volume Protection Status

Server	Volume	Group	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)	Resync progress	RPO	Status	Resync Required	View Details
ORA2->ORA1	/dev/mapper/vol00-oraarch -> /oraarch	Volume /oraarch	0	0	0	0 %	0	Resyncing (Step 1)	YES	

Figure 34

4.3 Planned failover

To perform a failback through the CX UI, setup an FX job.

Select the DR Oracle server as “**source**” and the production Oracle server as the “**Destination**”. Then select the FX template as “**Oracle(UNIX) Planned Failover**”. The rest of the process is similar to a planned failover as described in the section [Planned Failover through CX UI](#) on page 15.

File Protection Wizard: Replication Pair
Logged in as 'admin' - [Logout](#)

Replication Hosts

Application Name:

Job Description:

Source		Destination	
	Host		Host
<input type="radio"/>	ORA1 [Unix]	<input checked="" type="radio"/>	ORA1 [Unix]
<input checked="" type="radio"/>	ORA2 [Unix]	<input type="radio"/>	ORA2 [Unix]
Directory		Directory	
<input type="text"/>		<input type="text"/>	

Oracle(UNIX) Planned Failover

Figure 35

Once the FX job is complete, a failback is completed. You will have to set up the replication pair again to protect the production Oracle server

5 Protecting Oracle database running on Solaris Zones

5.1 Introduction to the solution

This section of the document describes steps involved in protecting and recovering Oracle running on Solaris Zones. The protect section contains four sequential steps.

[Prepare](#): The VX install location is exported to the Solaris zones. By doing this the scripts on the VX install location are available on each of the zones. These scripts are used just before issuing the vacp consistency tag.

[Discover](#): This step is performed to identify the file systems that are to be replicated.

[Replicate](#): The discovered file systems, zone data are replicated through the CX UI.

[Consistency](#): Vacp tags are issued on the production volumes at regular intervals through a single FX job.

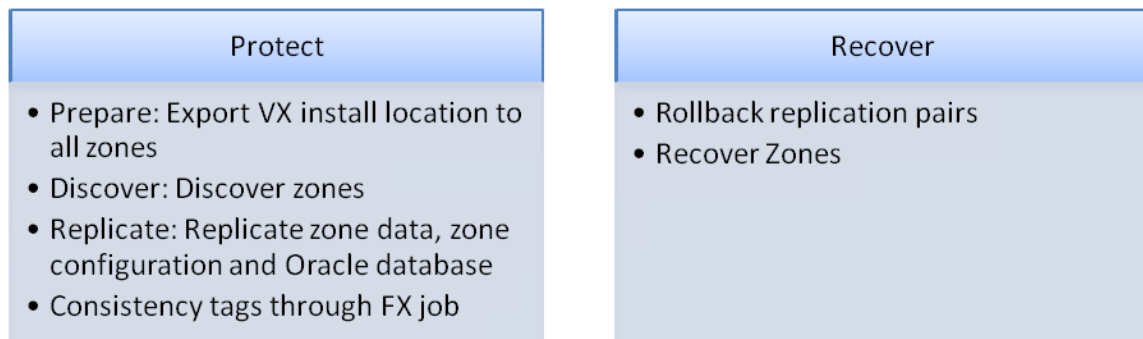


Figure 36

[Recover](#)

Recovery is achieved in two steps as shown in the picture above. You will first need to rollback the replicated volumes to a desired consistent point, preferably the latest consistency point. The zone data is now present on the DR server (target machine). You may active the zone through the command line as explained in the [Recovering zones](#) section on page 29

5.2 Protect

5.2.1 Prepare

Login to the production server's base operating system and use the “`zoneadm list -cv`” command to get the list of zones. Then navigate to the “`<VX install path>/scripts/zones`” to run the “`inimage_zones_config.sh`” script. You should be prompted to reboot the zones, hit Y to reboot the zones. This will export the VX agent install path to each of the zones.

```
[A-bash-3.00# zoneadm list -cv
ID NAME          STATUS  PATH                                BRAND  IP
0 global          running /                                     native shared
5 ora1-zone       running /ZONE/zones/ora1-zone             native shared
6 ora2-zone       running /ZONE/zones/ora2-zone             native shared
-bash-3.00# pwd
/usr/local/InMage/Vx/scripts/zones
-bash-3.00# ./inimage_zones_config.sh
This script reboots all the zones. Do you want to proceed Y/N
y
/usr/local/InMage/Vx exported to ora1-zone zone...
Rebooting ora1-zone
zoneadm: zone 'ora1-zone': WARNING: e1000g0:1: no matching subnet found in netmask
/usr/local/InMage/Vx exported to ora2-zone zone...
Rebooting ora2-zone
zoneadm: zone 'ora2-zone': WARNING: e1000g0:2: no matching subnet found in netmask
-bash-3.00#
```

Figure 37

5.2.2 Discover Solaris Zones

There can be multiple zones running on a physical Solaris machine. A typical zone consists of directories, file systems, partitions or a combination of them exported directly to the zones. You can find out the resources used a zone through the following command

#zonecfg -z <ZONE-NAME> export

```
-bash-3.00# zoneadm list -cv
ID NAME          STATUS  PATH                                BRAND  IP
0 global          running /                                     native shared
2 ora2-zone       running /ZONE/zones/ora2-zone             native shared
3 ora1-zone       running /ZONE/zones/ora1-zone             native shared
-bash-3.00# zonecfg -z ora2-zone export
create -b
set zonepath=/ZONE/zones/ora2-zone
set autoboot=true
set pool=pool_default
set ip-type=shared
add inherit-pkg-dir
set dir=/lib
end
add inherit-pkg-dir
set dir=/platform
end
add inherit-pkg-dir
set dir=/sbin
end
add inherit-pkg-dir
set dir=/usr
```

Figure 38

5.2.3 Replicate Solaris zones / Database

Switch to the CX UI and replicate the source volumes with CDP retention enabled. The replication pair starts from “Resync step 1” and then moved to “Resync step 2” before reaching “Differential Sync”. Move to the next section to set a consistency job.

Protection Status									
<div> <div>25 Pairs</div> <div>Page 1 of 1</div> </div>									
Volume Protection Status									
Server	Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPO	Status	Resync Required
				On CX-PS	On Target				
QSO1064SR20317S->QSO1064SR20318T	/ZONE -> /dev/dsk/c1t2d0s0	0	0	0	0	N/A	0.35 minutes	Differential Sync	NO
QSO1064SR20317S->QSO1064SR20318T	/Oracle1 -> /dev/dsk/c1t2d0s4	0	0	0	0	N/A	0.42 minutes	Differential Sync	NO

Figure 39

5.2.4 Consistency tags

Switch to the “File Protection” screen to set a single FX job to automate issuing consistency tags on all the production volumes. Click “File Protection -> New Job Group Wizard ->Add Job”. Enter the name of the application, job description and Source and Destination directory as “<vx installation directory>/failover_data”. Click “Next” to continue.

File Protection Wizard: Replication Pair

Replication Hosts

Application Name: ZoneConsistency

Job Description: for ZoneConsistency

Source

Destination

Host

Host

Directory

Directory

Template not selected

Next ->

Cancel

Figure 40

The Job Options screen appears, scroll down to Miscellaneous Options to enter the pre-script as `"/usr/local/InMage/Vx/scripts/zones/zonesconsistency.sh"`

Miscellaneous Options

☐

Create temporary files in:

☒

Do not cross filesystem boundaries

☒

IO timeout in seconds:

9600

☒

Specify port number to use:

874

☐

Limit bandwidth to

0

 KB/s

☐

Push data from source to target (implies lower CPU load on source)

☒

Pull data from source to target

CPU throttle (source)

0

Send RPO alert if

0

 minutes passed

Send E-mail alert if

5

 minutes passed without job progress

Pre execution script pathname:

/usr/local/InMage/Vx/scripts/zones/zonesconsistency.sh

Post execution script pathname:

Pre execution script pathname (destination):

Post execution script pathname (destination):

Catch All job modifier: for power users only

< - Back

Submit

Cancel

Figure 41

Click **“Submit”** to continue to the next screen. Schedule the job to run every half hour

File Protection Status											
Filter	Job Description	Application	Status	Source Host	Source Directory	Target Host	Target Directory	Scheduled Type	GID	JID	Job Instance
<div>Submit</div> <div>Clear</div>	<div></div>	<div>Select</div>	<div>Select</div>						<div>Select</div>	<div>Select</div>	<div>Select</div>
<div></div>	for ZoneConsiste...	ZoneConsistency	Not started...	QSO1064SR20317S	/usr/local/InMage/Vu/failover_data/	QSO1064SR20318T	/usr/local/InMage/Vu/failover_data/	Run Every	5	5	206
<div></div>	for ZoneConsiste...	ZoneConsistency	Completed	QSO1064SR20317S	/usr/local/InMage/Vu/failover_data/	QSO1064SR20318T	/usr/local/InMage/Vu/failover_data/	Run Every	4	4	204
<div></div>	for ZoneConsiste...	ZoneConsistency	Completed	QSO1064SR20317S	/usr/local/InMage/Vu/failover_data/	QSO1064SR20318T	/usr/local/InMage/Vu/failover_data/	Run Every	4	4	203

Figure 42

5.3 Recovery

5.3.1 Rollback all target volumes

Access the recovery page and select all the replication pairs associated with Zones and Oracle database then click “Rollback”

Volume Recovery : Recovery Snapshots

[Recovery Snapshots](#)
[Scheduled Snapshots](#)

Replication Pair Details					
<input type="checkbox"/>	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input checked="" type="checkbox"/>	QSO1064SR20317S	/ZONE	QSO1064SR20318T	/dev/dsk/c1t2d0s0	7
<input checked="" type="checkbox"/>	QSO1064SR20317S	/Oracle1	QSO1064SR20318T	/dev/dsk/c1t2d0s4	8

Figure 43

You will need to enter the mount point for each of the replication pair as shown in the picture below. Select the “Using Application consistency & Event based” option and scroll down to set further options.

Volume Recovery: Rollback Options (Multiple Pairs)

Pair Details					
Server	Primary Volume	Remote Server	Target Volume	Replication Pool	Mount Point
QSO1064SR20317S	/dev/dsk/c1t2d0s0	QSO1064SR20318T	/dev/dsk/c1t2d0s0	7	<input type="text" value="/ZONES"/>
QSO1064SR20317S	/dev/dsk/c1t2d0s4	QSO1064SR20318T	/dev/dsk/c1t2d0s4	8	<input type="text" value="/Oracle1"/>

Recovery Options		
Logs Available	From 2009/08/03 16:51:30:915 (GMT)	To 2009/08/04 11:26:53:108 (GMT)
Recovery Based On		
<input type="radio"/> Using Time <input checked="" type="radio"/> Using Application consistency & Event based		

Figure 44

Select the desired consistency tag from the list, preferably the latest tag and click “Submit”

Search Result

	Accuracy	Timestamp	Application	Tag Name
<input checked="" type="radio"/>	Exact	2009/8/4 11:21:23:12	File System	FileSystem4a7819b2
<input type="radio"/>	Exact	2009/8/4 8:4:24:145	File System	FileSystem4a77eb88
<input type="radio"/>	Exact	2009/8/4 7:34:19:864	File System	FileSystem4a77e47b
<input type="radio"/>	Exact	2009/8/4 7:4:19:437	File System	FileSystem4a77dd72

<< < 1 2 3 4 > >>

Recovery Points Accuracy:
 Exact
Approximate
Not guaranteed

Figure 45

You will be returned back to the recovery screen where you can monitor the progress of the rollback. Once the target volumes are rolled back proceed to the next step to activate the zones on the DR server.

Target Drive Rollback Status								
	Host	Rollback Drive	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message
	QSO1064SR20318T	/dev/dsk/c1t2d0s0	Complete	100%	2009/8/4 11:21:23:12	-	Tag Based Tag FileSystem4a7819b2 Accuracy	-
	QSO1064SR20318T	/dev/dsk/c1t2d0s4	Complete	100%	2009/8/4 11:21:23:12	-	Tag Based Tag FileSystem4a7819b2 Accuracy	-

[Release Drive](#)

Recovery Points Accuracy: - Exact - Approximate - Not guaranteed

Figure 46

5.3.2 Recovering zones

Step 1. All the zone data is available on the rolled back target volume. To activate these zones, you will first need to create the create zones using the replciated zone configuration files. You may use the following command to create the zones.

Zonecfg -z <name of the zone>

Observe that the newly created zones are listed with status as “configured”

```
-bash-3.00# zonecfg -z ora1-zone -f /ZONE/ora1-zone.config
-bash-3.00# zonecfg -z ora2-zone -f /ZONE/ora2-zone.config
-bash-3.00# zoneadm list -cv
ID NAME STATUS PATH BRAND IP
0 global running / native shared
- ora1-zone configured /ZONE/zones/ora1-zone native shared
- ora2-zone configured /ZONE/zones/ora2-zone native shared
```

Figure 47

Step 2. Once the zones are created, run the following command to attach the newly created zones on the DR server.

Zoneadm -z <name of the zone> attach

```
-bash-3.00# zoneadm -z ora1-zone attach
zoneadm: zone 'ora1-zone': WARNING: pools facility not active; zone will not be bound
zoneadm: zone 'ora1-zone': WARNING: Pools facility not active; zone will not be bound
zoneadm: zone 'ora1-zone': WARNING: e1000g0:1: no matching subnet found in netmasks (
of 255.0.0.0.
-bash-3.00# zoneadm -z ora2-zone attach
zoneadm: zone 'ora2-zone': WARNING: pools facility not active; zone will not be bound
zoneadm: zone 'ora2-zone': WARNING: Pools facility not active; zone will not be bound
zoneadm: zone 'ora2-zone': WARNING: e1000g0:1: no matching subnet found in netmasks (
of 255.0.0.0.
-bash-3.00# zoneadm list -cv
ID NAME STATUS PATH BRAND IP
0 global running / native shared
- ora1-zone installed /ZONE/zones/ora1-zone native shared
- ora2-zone installed /ZONE/zones/ora2-zone native shared
```

Figure 48

Step 3. Repeat the same step for all the zones created. The status of the zones should now change from “configured” to “installed”.

Step 4. Now that the zones are installed, you will need to start the zones

zoneadm -z <name of the zone> boot

```
-bash-3.00# zoneadm -z ora1-zone boot
zoneadm: zone 'ora1-zone': WARNING: pools facility not active; zone wi
zoneadm: zone 'ora1-zone': WARNING: Pools facility not active; zone wi
zoneadm: zone 'ora1-zone': WARNING: e1000g0:1: no matching subnet foun
of 255.0.0.0.
-bash-3.00# zoneadm -z ora2-zone boot
zoneadm: zone 'ora2-zone': WARNING: pools facility not active; zone wi
zoneadm: zone 'ora2-zone': WARNING: Pools facility not active; zone wi
zoneadm: zone 'ora2-zone': WARNING: e1000g0:2: no matching subnet foun
of 255.0.0.0.
-bash-3.00#
```

Figure 49

Step 5. You may see the list of zones and their status by using the following command

zoneadm list -cv

```
-bash-3.00# zoneadm list -cv
ID NAME          STATUS    PATH                                BRAND  IP
 0 global         running   /                                  native shared
 7 ora1-zone      running   /ZONE/zones/ora1-zone             native shared
 8 ora2-zone      running   /ZONE/zones/ora2-zone             native shared
-bash-3.00#
```

Figure 50

5.3.3 Recovering Oracle database

You will need to login to each of the nodes to execute the following commands to recover Oracle

Use the below command to login to the zone.

zlogin <name of the zone>

login to the oracle and the database should be mounted without any errors.

```
-bash-3.00$ sqlplus '/as sysdba'

SQL*Plus: Release 10.2.0.2.0 - Production on Thu Aug 6 04:02:04 2009

Copyright (c) 1982, 2005, Oracle. All Rights Reserved.

Connected to an idle instance.

SQL> startup;
ORACLE instance started.

Total System Global Area  452984832 bytes
Fixed Size                  1280756 bytes
Variable Size               134219020 bytes
Database Buffers            314572800 bytes
Redo Buffers                 2912256 bytes
Database mounted.

Database opened.
SQL> SQL> SQL>
SQL> quit
```

Figure 51

6 Trouble shooting

If an error occurs when mounting the database there may be problems with the control-files or data files, or with resources required to open these files.

Follow the steps below to recover the database.

Step 1. The location of the control files are specified in the init.ora file. Try mounting using each control file in turn.

eg: "Shutdown abort",
edit the init.ora to refer to ONE of the control files only,
"startup nomount",
"alter database mount"
Repeat for each control file to see if any control file works.

Step 2. Recover the database using the following command

```
SQL> recover the database until cancel
```

Point to the full path of the redo log file when asked for the path of the next archive log. Provide the path of the each of the redo file until sequence# matches and recovery continues.

Step 3. Open the database

Once the recovery is complete, open the database

```
SQL> alter database open resetlogs;
```


Hitachi Data Systems

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

