

# HP OpenView OmniBack II 3.5

For HP-UX, Windows NT and Windows 2000

Technical Whitepaper

## **Performing Oracle 8i (OPS) Split Mirror Backups Using OmniBack II**

---

An HP-UX only Paper

# **Table of Contents**

1. Introduction .....	7
2. Architecture.....	8
3. OmniBack II .....	10
3.1. Prerequisites.....	10
4. Configuring the XP256 Integration .....	11
4.1. Configuring Business Copy .....	11
4.2. Inquiry.ksh Output .....	13
4.3. Configuration Files.....	14
4.4. Starting the RAID Manager Instances.....	16
4.5. Managing Disk Pairs.....	17
4.5.1. Create Pair .....	17
4.5.2. Split Pair .....	17
4.5.3. Resynchronize Pair .....	17
4.5.4. Restore Pair .....	18
4.5.5. Delete Pair.....	18
4.6. Other Available Commands .....	18
4.7. Preparing the Application Server.....	19
4.8. Preparing the Backup Server .....	19
5. Setting up the RMAN Catalog Database.....	21
5.1. Installing Oracle.....	21
5.2. Creating the tablespaces.....	22
5.3. Creating the User RMAN.....	22
5.4. Setting up the Recovery Catalog.....	23
5.5. Setting up OPS for RMAN .....	23
5.6. Setting up Net8.....	24
5.7. Registering the OPS Database .....	26
6. OmniBack II Oracle Integration.....	27
6.1. Setting up Remote Login .....	27
6.2. Filesystem Structure on the Backup Server .....	27
6.3. Linking OmniBack and Oracle.....	28
6.4. Setting up a User.....	28
6.5. Configuring the Oracle 8 Client on the Application Servers .....	29
6.6. Configuring the Oracle 8 Client on the Backup Server .....	30
7. Backing up Oracle Parallel Server.....	32

7.1.	Backing up the Database .....	32
7.2.	Backing Up the Archived Logs .....	37
7.2.1.	Backup the Archived Logs by Each Node.....	37
7.2.2.	Backup Archived Logs from one Node.....	40
7.2.2.1.	Shared Archived Log Directories .....	40
7.2.2.2.	Backup Archived Log Files using Connect Command in RMAN .....	41
7.2.2.3.	Copying Archived Logs to one Node .....	42
8.	Restore and Recovery .....	43
8.1.	Restore to the Application Server.....	43
8.2.	Restore to the Backup Server .....	44
9.	Troubleshooting .....	45

## **Table of Figures**

Figure 1 - System Layout.....	8
Figure 2 - Example Output from the Application Server .....	11
Figure 3 -Example Output from Backup Server.....	12
Figure 4 - <code>inquiry.ksh</code> output from the Application Server .....	13
Figure 5 - <code>inquiry.ksh</code> output from the Backup server .....	14
Figure 6 - Configuration file for the Application Server .....	14
Figure 7 -Business Copy XP Configuration file for the Backup Server.....	15
Figure 8 - Environment Variables for HORCM .....	16
Figure 9 - Example Listener Configuration File .....	24
Figure 10 - Example of the file <code>tnsnames.ora</code> .....	25
Figure 11 - Retrieving the list of files. ....	32
Figure 12 - Resolving a File .....	33
Figure 13 - Mapping the Physical Device File to LDEVs .....	33
Figure 14 - Putting the Database in Backup Mode .....	34
Figure 15 - Splitting the Disks.....	34
Figure 16 - Putting the Database in Normal Operation.....	35
Figure 17- Example Backup Specification .....	36
Figure 18 - Backup Specification for Archivelogs produced by 'erwin1' .....	39
Figure 19 - Backup Archived Logs with Shared Directories .....	40
Figure 20 - Backing up Archived Logs from one Node using the <code>connect</code> clause .....	41

# **1. Introduction**

This paper provides a guideline how to configure an OmniBack II split mirror backup environment on HP-UX, using RMAN, Oracle 8i Parallel Server and the high-end disk array HP SureStore E XP256.

It describes the steps necessary for the installation, configuration and administration of OmniBack II in this kind of environments. It also discusses the different scenarios for backups and restores in OPS environments. Performance is not a focus of the paper as it only describes the concepts of this type of backups.

It will start with an overview about architecture of the example environment, which is used throughout this paper.

Then it covers the installation of OmniBack in such an environment. It will explain the different components that need to be distributed throughout the environment.

The next chapter will explain the necessary steps that need to be done in order to enable the zero downtime backup for the HP SureStore E XP256 disk array. It will start with commands that allow you to discover the aArray. Then it will show example outputs on which example configuration files will be explained. It will furthermore provide examples on the most common commands of the Raid Manager. The preparation of the application Servers and the backup Server will be covered as well.

The next chapter explains the steps, that are necessary to setup the Recovery Manager (RMAN) catalog database. It will show how to install Oracle on the backup server, how to create the database scheme, the Net8 Configuration steps and how to register the production database.

How to integrate Oracle and OmniBack will be explained in chapter 6. It will describe how to setup remote login to the production database, how to configure the application and the backup servers and more.

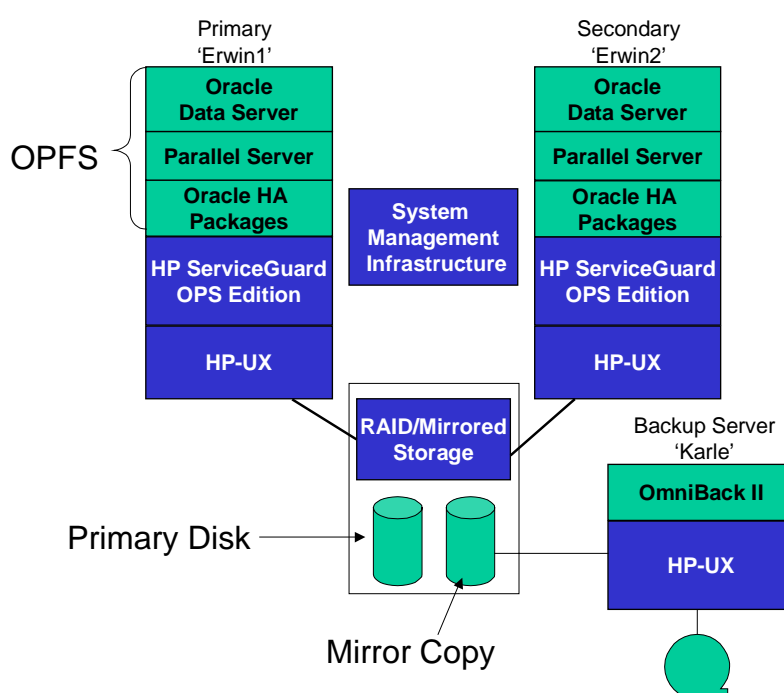
Chapter 7 will explain in detail how the Oracle Parallel Server database will be backed up. It will focus on different methods for backup of the archived log files.

The next chapter focuses on restore and recovery. It will provide different possibilities for those tasks.

At the end of this paper, troubleshooting will be covered. It will explain different error messages, their causes and the possible solution for the problems.

## 2. Architecture

The following chapter will explain the setup the environment. It will explain the different hard- and software components used for this integration. It is based on the Oracle Parallel Fail Safe (OPFS) solution jointly developed by HP and Oracle. OPFS is built upon Oracle Parallel Server running in a two-node primary/secondary configuration. In this configuration, all connections to the database are through the primary node called 'erwin1' throughout the paper. The secondary node, 'erwin2' serves as a backup, ready to provide services should an outage at the primary occur. Figure 1 - System Layout, shows high level view of the environment in which this integration is installed.



**Figure 1 - System Layout**

On the HP 9000 platform, Oracle Parallel Fail Safe combines Oracle's best high availability features with HP's flagship ServiceGuard clustering technology to create a premier high availability (HA) solution, offering the highest possible levels of availability for critical business and internet applications. Oracle and HP both recognize that outages are inevitable. With Oracle Parallel Fail Safe running on an HP ServiceGuard cluster, you can recover very quickly from outages and minimize the disruption to end-user applications, thus maximizing application availability. Oracle Parallel Fail Safe takes advantage of HP's clustering technology to rapidly detect and recover from failures. By tightly integrating HP ServiceGuard OPS Edition with Oracle Parallel Server, Oracle and HP offer a highly robust and resilient, yet very simple availability solution.

Then there is the backup server, who is responsible for backing up the bulk data from disk to tape using the split off mirror copy within the XP256 disk array. Bulk data means the shared database files. The archived log files, which are written by each of the database servers independently, will be backed up either over the network or backup devices that are connected locally on the database servers. It is recommended to use Fiber Channel technology for the connection of the different servers to the library for performance and fault tolerance reasons.

The recovery catalog database, which will be covered later in the paper, can be installed anywhere in the environment. However, this paper will explain the setup of the recovery catalog on the backup server.

## **3. OmniBack II**

This chapter will explain the steps necessary to enable HP OpenView OmniBack II 3.5 for the XP 256 Split Mirror backup integration. It assumes that you already have a regular HP OpenView OmniBack II 3.5 installation. It will not cover how to setup an OmniBack II cell. Please read the HP-OpenView OmniBack Installation and Licensing Guide for more information.

### **3.1. Prerequisites**

In order to enable HP OpenView OmniBack II for the split mirror backup, the HP SureStore E Agent (SSEA) functionality, a license (B7023AA) must be installed. Furthermore you need to distribute the SureStore E Agent together with the On-Line Agent for Oracle8 to all nodes within your environment.

This integration supports only the OmniBack II Windows user interface. Therefore the standard OmniBack II Windows user interface must be installed within the OmniBack II cell on a Windows NT system running the Windows NT Workstation version 4.0 with Service Pack 3.0.

The following HP SureStore E XP 256 components are required for this integration.

- DKC microcode version should be 52-42-51 or newer.
- RAID Manager XP version A.01.02.03 (HP-UX) or A.01.03.01 (Windows NT) installed on both the application system and the backup system.

Business Copy XP or Continuous Access XP microcode and license installed.

The `inquiry256.ksh` (HP-UX) must be installed. Contact the Hewlett-Packard support center to obtain these utilities.

Please refer to the HP OpenView OmniBack II HP SureStore E Disk Array XP 256 Configuration Guide, for more information.



## 4. Configuring the XP256 Integration

It is assumed that the disks are already visible to all, the database- and the backup servers. It is also assumed that the Oracle production database is already installed and running on the production servers.

This chapter will explain how to setup the business copy configuration for one application server and the backup server. It will not describe the configuration files for the second application server. Furthermore it will explain the different configuration tasks that are necessary on each server in order to use HP-OpenView OmniBack II's split mirror facility.

It is assumed that the reader of the paper is already familiar with the XP256 terminology, since it will not explain all the different expressions in detail.

### 4.1. Configuring Business Copy

In order to find the primary volumes on the application server issue the following command on the HP-UX command prompt:

```
$ ioscan -fnC disk
```

The output of the command will be similar to this one:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	5	8/8.0.0	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d0		/dev/rdisk/c0t0d0	
disk	6	8/8.0.1	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d1		/dev/rdisk/c0t0d1	
disk	7	8/8.0.2	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d2		/dev/rdisk/c0t0d2	
disk	8	8/8.0.3	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d3		/dev/rdisk/c0t0d3	
disk	9	8/8.0.4	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d4		/dev/rdisk/c0t0d4	
disk	15	8/8.0.5	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d5		/dev/rdisk/c0t0d5	
disk	16	8/8.0.6	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c0t0d6		/dev/rdisk/c0t0d6	
disk	18	8/8.1.0	sdisk	CLAIMED	DEVICE	HP OPEN-9-CM
			/dev/dsk/c0t1d0		/dev/rdisk/c0t1d0	
disk	37	8/12.0.0	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c1t0d0		/dev/rdisk/c1t0d0	
disk	38	8/12.0.1	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c1t0d1		/dev/rdisk/c1t0d1	
disk	39	8/12.0.2	sdisk	CLAIMED	DEVICE	HP OPEN-9

Figure 2 - Example Output from the Application Server

The highlighted disks are used in the example configuration files that will be explained later in this chapter. The turquoise disks are the physical disks used on the application system as the so-called primary volumes (P-Vol). The pink marked disk is the command device, identified by the entry '-CM'. This disk is used for the communication with the disk array. The Raid Manager sends the `ioctl` commands to this device. One command device can be used by multiple Raid Managers.

The same command issued on the backup server will result in the following example output:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	28	8/4.0.0	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d0		/dev/rdisk/c5t0d0	
disk	29	8/4.0.1	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d1		/dev/rdisk/c5t0d1	
disk	30	8/4.0.2	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d2		/dev/rdisk/c5t0d2	
disk	31	8/4.0.3	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d3		/dev/rdisk/c5t0d3	
disk	32	8/4.0.4	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d4		/dev/rdisk/c5t0d4	
disk	33	8/4.0.5	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d5		/dev/rdisk/c5t0d5	
disk	34	8/4.0.6	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d6		/dev/rdisk/c5t0d6	
disk	35	8/4.0.7	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t0d7		/dev/rdisk/c5t0d7	
disk	36	8/4.1.0	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t1d0		/dev/rdisk/c5t1d0	
disk	37	8/4.1.1	sdisk	CLAIMED	DEVICE	HP OPEN-9
			/dev/dsk/c5t1d1		/dev/rdisk/c5t1d1	
disk	38	8/4.1.2	sdisk	CLAIMED	DEVICE	HP OPEN-9-CM
			/dev/dsk/c5t1d2		/dev/rdisk/c5t1d2	

**Figure 3 -Example Output from Backup Server**

The highlighted entries are the disks that will be used for the example configuration files that follow. The disks marked green are the physical disks the backup server is connected to. They are the mirror copies for the P-Vols on the backup server, the so-called secondary volumes (S-Vol). The entry marked yellow is the command device, used on the backup server.

## 4.2. Inquiry.ksh Output

Use the output from the command `inquiry256.ksh` as well to find more information, you need for the creation of the configuration files. The following shows an example output of this command from the application server. The command is usually installed in the directory `/usr/contrib/bin`.

```
$ inquiry256.ksh
```

Device File	---	Port	Target	Lun	CU:Ldev	Serial#	Type	Size
/dev/rdisk/c0t0d0	:	CL1A	0	0	0:00	00035390	OPEN-9	7211520 Kbytes
...								
/dev/rdisk/c0t0d5	:	CL1A	0	5	0:05	00035390	OPEN-9	7211520 Kbytes
/dev/rdisk/c0t1d0	:	CL1A	1	0	0:0d	00035390	OPEN-9-CM	7211520 Kbytes

**Figure 4 - inquiry.ksh output from the Application Server**

Again the highlighted entries correspond to the ones that are used later for the configuration files. The colors used here correspond also to the ones used in Figure 2 - Example Output from the Application Server. The following information can be read from this output.

- **Device File:**  
This identifies the HP-UX device file under which the server accesses this disk
- **Port:**  
The port is a logical naming for the interface cards within the XP256 Disk array. CL1 refers to slot P-S and A refers to one of the four SCSI ports per slot. In this case it is the first one. It could also be the case that it shows CL2, which refers to slot V-Z. Please refer to the XP256 hardware manual for more information.
- **Target and LUN:**  
This refers to the SCSI Target- and LUN ID.
- **CU:Ldev**  
This identifies a device uniquely within the disk array. The first parameter specifies the control unit and the second one the device number. 256 Ldevs are grouped onto one control unit. The XP has four control units numbered from 0 to 3. The number CU:Ldev displayed in decimal can be used later for commands like `'pairdisplay'`. It can range from 0 to 1023.
- **Serial Number:**  
This identifies the Disk Array uniquely
- **Type and size:**  
This shows the type and size of disks.

This is an example output of the `inquiry.ksh` command on the backup server. The disks, used in the examples are again highlighted in the corresponding colors.

Device File	---	Port	Target	Lun	CU:Ldev	Serial#	Type	Size
...								
/dev/rdisk/c5t0d6	:	CL2B	0	6	0:06	00035390	OPEN-9	7211520 Kbytes
/dev/rdisk/c5t0d7	:	CL2B	0	7	1:08	00035390	OPEN-9	7211520 Kbytes
/dev/rdisk/c5t1d0	:	CL2B	1	0	1:09	00035390	OPEN-9	7211520 Kbytes
...								
/dev/rdisk/c5t1d2	:	CL2B	1	2	0:0d	00035390	OPEN-9-CM	7211520 Kbytes

**Figure 5 - inquiry.ksh output from the Backup server**

### 4.3. Configuration Files

Use this output to create the following configuration file for the application servers. Here only the configuration file for 'erwin1' is shown. In order to be able to perform the split on 'erwin2' as well, you need to create a file that is adopted to this server. Please read on for explanation of the different parameters.

```
# Business Copy XP P-VOL Configuration File
# Application Server: erwin1.hp.com (122.122.122.122)
# Filename: bc1.conf
# Split Units: su1, su2
# Backup Server: karle.hp.com (122.122.122.123)

HORCM_MON
#ip_address      service          poll(10ms)      timeout(10ms)
122.122.122.122  bc1              1000            3000

HORCM_CMD
#dev_name        dev_name          dev_name
/dev/rdisk/c0t1d0

HORCM_DEV
#dev_group  dev_name      port#      TargetID      LU#
su1         vg02_ops     CL1-A      0              0
su2         vg04_arch    CL1-A      0              5

HORCM_INST
#dev_group  ip_address      service
su1         122.122.122.123  bc1
su2         122.122.122.123  bc1
```

**Figure 6 - Configuration file for the Application Server**

This is the corresponding configuration file on the Backup Server.

```
# Business Copy XP S-VOL Configuration File
# Backup Server: karle.hp.com (122.122.122.123)
# Filename: bcl_svol.conf
# Split Unit: su1, su2
# Application Server: erwin1.hp.com (122.122.122.122)

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
122.122.122.123  bc1          1000            3000

HORCM_CMD
#dev_name        dev_name        dev_name
/dev/rdsk/c5t1d2

HORCM_DEV
#dev_group      dev_name        port#           TargetID        LU#
su1             vg02_ops        CL2-B           0                7
su2             vg04_arch        CL2-B           1                0

HORCM_INST
#dev_group      ip_address      service
su1             122.122.122.122  bc1
su2             122.122.122.122  bc1
```

**Figure 7 -Business Copy XP Configuration file for the Backup Server**

The first parameter `HORCM_Mon` describes the IP address, port number and paired volume error monitoring interval of the local server. It is used by the HORC Manager for monitoring and communicating.

The second parameter describes the special file (raw device) that is the HORC control device. The control device can be seen from the `ioscan` output by the tag `OPEN-9-CM`.

The `HORCM_DEV` parameter describes the addresses in the RAID of the physical volumes corresponding to the paired logical volume names. It specifies a group name to which the disks belong. This group name must be unique and must be the same on the remote server. The device name is used to describe a name for the paired logical volume composing the group.

The `port#` is used to describe a port number of the physical volume in the RAID connected with `dev_name`. The port numbers, which can be specified for the RAID, are `CL1-A` to `CL1-R` or `CL2-A` to `CL2-R`.

The parameter `TargetID` is used to describe a SCSI target ID of the physical volume on the port, specified in `port#`.

The parameter `LU#` describes the logical unit number of the physical volume with the number specified in `TargetID`.

The `HORCM_INST` parameter is used to describe the network address of the remote server, which contains the group for each group name. It is used to change the status of the paired volume on the remote server. The parameter `dev_group` is a group name described in the `dev_group` of `HORC_DEV`. The parameter `ip_address` describes the network address of the remote server. The parameter `service` describes the port name assigned to the HORCM communication path. This must be registered in the `/etc/services` file. A port number can also be specified.

#### 4.4. Starting the RAID Manager Instances

Use the following shell script to set environment variables before you start the HARC Manager.

```
#!/bin/ksh
export HORCC_MRCF=1
export HORCMINST=1
export HORCC_LOG=/opt/HORCM/bc/horcclog1
export HORCM_CONF=/opt/HORCM/bc/bc1_svol.conf
export HORCM_LOG=/opt/HORCM/bc/horcmlog1
```

**Figure 8 - Environment Variables for HORCM**

The variable `HORCC_MRCF` describes the functionality you want to perform. Specify '1' for Business Copy XP and '0' for Continuous Access XP. `HORCMINST` specifies the instance for the RAID Manager, which is used to communicate with the disk array.

The parameter `HORCM_CONF` must be set if you don't use the default path for the configuration files, which is `$HORCM_HOME/etc`.

Use the following command to start the Instance after you ran the configuration script:

```
$HORCM_HOME/usr/bin/horcmstart.sh
```

## 4.5. Managing Disk Pairs

The following operations can be performed on disk pairs. This chapter will be based on Business Copy XP but can also be applied to Continuous Access XP.

### 4.5.1. Create Pair

To create the initial Business Copy XP pair you have to pick LDEVs on the application server and on the backup server. The pair can be created using the following command:

```
paircreate -g sul -d vg02_ops -v1
```

The parameter `sul` and `vg02_ops` point to the device in the RAID manager configuration file. If you start this command on the application server the device `vg02_ops` will become the primary volume. The corresponding device on the backup server will become the secondary volume.

Monitor the status of the pair using the following command:

```
pairdisplay -g sul -l
```

This command produces the following output:

Group	PairVol(L/R)	(Port#,TID,LU-M)	Seq#	LDEV#.P/S	Status	Seq#	P-LDEV#	M
sul	vg02_ops(L)	(CL2-B , 0, 7-0)	35390	264..S-VOL	PAIR	-----	0	-

You see that the paired volume identified by the name `vg02_ops` is in status `PAIR` and that it is in the group `sul`.

### 4.5.2. Split Pair

To split the pair the devices must be in status `pair`. Use the `pairdisplay` to check this before you try to split the disks.

```
pairsplit -g sul -d vg02_ops
```

### 4.5.3. Resynchronize Pair

To resynchronize the pair, that means to setup the mirror again, you issue the following command:

```
pairresync -g sul -d vg02_ops
```

This will only transfer the changes since the split from the P-Vol to the S-Vol. It will not be a complete copy of the data.

#### 4.5.4. Restore Pair

To restore the data from the primary volume to the secondary volume, that means to do a restore, issue the following command.

```
pairrestore -g sul -d vg02_ops
```

#### 4.5.5. Delete Pair

If you want to destroy the pair issue the following command:

```
pairsplit -g sul -d vg02_ops -S
```

The `-s` specifies the mode to be simplex, that means no pair. This command returns the secondary volume to the standard pool of disks again.

### 4.6. Other Available Commands

The following commands can be found in `/HORCM/usr/bin`. Please refer to the Raid Manager Manuals for more information on the commands.

- `horcctl`
- `horcmshutdown.sh`
- `horctakeover`
- `inqraid`
- `paircurchk`
- `pairevtwait`
- `pairmon`
- `pairvolchk`
- `raidar`
- `raidqry`
- `raidscan`

Please see the manuals for the Raid Manager for more information on these commands.

You will get help from the command itself with the parameter `'-h'`.



## 4.7. Preparing the Application Server

The volume group configuration of the application server needs to be exported and imported on the backup server. Make sure that you use the same volume groups on all application servers. This way, you only need to do the steps once. For this integration, this is the case anyway, therefore only one example will be shown. Make sure to deactivate the volume group on the application server first. You can use the following command:

```
vgchange -a n /dev/vg02
```

Use the following command to export the volume group configuration on the application server.

```
vgexport -p -m vg02 /dev/vg02
```

This will create a map file with the name `vg02`, which should be copied to the backup server. The `-p` option will only do a preview and not the actual export. This is mandatory, since otherwise the volume group information will be deleted.

Now split the disks using the `pairsplit` command.

```
pairsplit -g sul -d vg02_ops
```

Do that for all device groups that you created in the configuration file. This will make the target volumes accessible to the backup server.

## 4.8. Preparing the Backup Server

This chapter describes the steps that need to be done to configure the backup server for the split mirror backup. The configuration of the volume groups must be the same as on the application server. This can be done by importing the volume group configuration into the backup server. Transfer the mapfile to a directory on the backup server using HP-UX commands like `ftp`.

First create the volume group directory on the backup server:

```
mkdir /dev/vg02
```

Now create a special file for the volume group,

```
mknod /dev/vg02/group c 64 0xYY0000
```

where `YY` stands for the minor number of the volume group.

After this step you can import the volume group configuration for `vg02`.

```
vgimport -m $MAPFILELOCATION/vg02 -v /dev/vg02 <pv_path>
```

Example:

- `vgimport -m /tmp/vg02 -v /dev/vg02 /dev/dsk/c5t0d7`

The `<pvpath>` points to the disks that have been identified as S-Vol on the backup server.

Repeat the steps in chapter 4.7 and 4.8 for each volume group you want to backup. The database on the system is stored on raw devices therefore there is no need to create the mount-points on the backup server. If you have a database that is based on a filesystem, you need to create the exact same mount-points as on the application system.

## **5. Setting up the RMAN Catalog Database**

This chapter covers the installation and configuration of the Recovery Manager (RMAN) Catalog. In this example, the Catalog Database is installed on the Backup Server. It is also possible to install it anywhere in the backup environment. It will only give an overview about the necessary steps from. Please refer to the official Oracle Documentation if you need more specific descriptions. The installation of an Oracle database on a server requires approximately 1GB of disk space. Make sure that is available on the backup server. Please check also other system requirements in your Oracle Documentation.

### **5.1. Installing Oracle**

Create a directory for the Oracle Software. In this example the following directory will be used:

```
/opt/oracle
```

---

NOTE: If possible you should use the same directory as on the application server. This makes succeeding integration tasks much easier.

---

Create the Unix user groups 'dba' and 'oinstall' on the backup server using SAM.

Create the user 'oracle' on the backup server using SAM.

Switch to the user 'oracle' using the command:

```
su - oracle.
```

Set the file creation mask of the user 'oracle' to '022'.

Add the following entries to the `.profile` file of the oracle user:

```
ORACLE_BASE=/opt/oracle/orasrc
ORACLE_HOME=$ORACLE_BASE/product/8.1.5
export ORACLE_BASE
export ORACLE_HOME
```

Mount the Oracle Installation CD-ROM as the user root (see Oracle Installation Guide).

Switch to the user 'oracle'.

Run `./runInstaller` from the Oracle Installation CD-ROM and follow the installation procedure. This includes specifying the global database name, the SID, the directory for

the database files and more. In this example the SID of the database is 'rmancat'. The directory for the database files is \$ORACLE\_HOME/dbs

## 5.2. Creating the tablespaces

After the installation is completed, you need to create the rman catalog. The following command can be used to create the tablespace RCVCAT:

```
SVRMGR> CREATE TABLESPACE RCVCAT
2> DATAFILE 'rcvcat.dbf' SIZE 50M
   DEFAULT STORAGE ( INITIAL 50K NEXT 1000K MINEXTENTS 1
   MAXEXTENTS 121
   PCTINCREASE 0)
ONLINE;
```

Please refer to your Oracle Manual for more information about the different options. The amount of space, the recovery catalog uses can be determined using the following table:

SYSTEM tablespace	50 Megabytes
TEMP tablespace	5 Megabytes
Rollback tablespace	5 Megabytes
Recovery catalog tablespace (RCVCAT)	10 Megabytes per year
Controlfiles	2*5 Megabytes
Online redo log groups	6*1 Megabyte

**Table 1 - Datafile Size for Recovery Catalog**

## 5.3. Creating the User RMAN

Export the SID of the database and start the Server Manager. Then on the Server Manager command prompt enter the following commands, which create the user `rman` with the password `rman` with its privileges:

```
SVRMGR> create user rman identified by rman
2> temporary tablespace temp
3> default tablespace rcvcat quota unlimited on rcvcat;
4> grant recovery_catalog_owner to rman;
5> grant connect, resource to rman;
```

## 5.4. Setting up the Recovery Catalog

Then you need to create the recovery catalog using the following commands that create the catalog database in the tablespace 'RCVCAT':

```
$ rman
RMAN> connect catalog rman/rman
RMAN> create catalog tablespace "RCVCAT"
```

It should the say:

```
RMAN-06431: recovery catalog created
```

## 5.5. Setting up OPS for RMAN

This chapter describes the necessary changes that need to be done on the OPS servers in order to be able to backup the database using RMAN.

Edit LOG\_ARCHIVE\_FORMAT parameter in the init<SID>.ora using %s to include the log sequence number as part of the filename and %t to include the thread number. Use capital letters (%S and %T) to pad the filename to the left with zeroes. For example, enter:

```
LOG_ARCHIVE_FORMAT = arch_%t_%s.arc
```

The destination for the archived log files should be set as well. Use the following parameter to do so:

```
LOG_ARCHIVE_DEST = "<DEFAULT_PATH>/<DEST_NODEx>/"
```

The variable <DEST\_NODEx> refers to the local archived log directory in each application server, where the 'x' should refer to the thread number. Refer to chapter 7.2 for more information on this parameter.

## 5.6. Setting up Net8

In order to setup Net8 please modify the following files on the backup server:

```
$ORACLE_HOME/network/admin/listener.ora
```

```
$ORACLE_HOME/network/admin/tnsnames.ora
```

This is an example listener file for the backup server.

```
Listener=
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=TCP)(SERVER=karle.hp.com)(PORT=1521))
    (ADDRESS=(PROTOCOL=ipc)(KEY=PNPKEY))
  )
  SID_LIST_LISTENER=
    (SID_LIST=
      (SID_DESC=
        (SID_NAME=rmancat)
        (ORACLE_HOME=/home/oracle/orasrc/product/8.1.5)
      )
    )
)

STARTUP_WAIT_TIME_listener = 0
CONNECT_TIMEOUT_listener = 10
```

**Figure 9 - Example Listener Configuration File**

Please refer to the official Oracle Manuals for more information on the different parameters.

This an example of the file `tnsnames.ora` on the backup server.

```
rmancat =
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP)(SERVER = karle.hp.com)(PORT = 1521))
  (CONNECT_DATA = (SERVICE_NAME = rmancat))
)
op1=
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP)(SERVER = erwin1.hp.com)(PORT = 1521))
  (CONNECT_DATA = (SERVICE_NAME = op))
)
op2=
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP)(SERVER = erwin2.hp.com)(PORT = 1521))
  (CONNECT_DATA = (SERVICE_NAME = op))
)
op=
(DESCRIPTION =
  (ADDRESS_LIST =
    (LOAD_BALANCE = OFF)
    (FAILOVER = ON)
    (ADDRESS = (PROTOCOL = TCP)(SERVER = erwin1.hp.com)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(SERVER = erwin2.hp.com)(PORT = 1521))
  )
  (CONNECT_DATA = (SERVICE_NAME = op)
    (FAILOVER_MODE =
      (TYPE = SELECT)
      (METHOD = BASIC)
      (BACKUP = op2)
    )
  )
)
```

**Figure 10 - Example of the file `tnsnames.ora`**

Modify the '`tnsnames.ora`' file on the application systems as well. That means you should make sure that the catalog database can be accessed from each system that will be backed up and that the system, on which the Recovery Catalog is installed can also reach the databases. The entry 'op' is configured as Transparent Application Failover. It has connect failover from 'erwin1' to 'erwin2', and TAF failover to 'op2'. The (BACKUP=...) clause is used as without this the actual connect strings tried on TAF failover can cause confusing results such as attempting to reconnect back to the failed node.

Please refer to the official Oracle Manuals for more information on the different parameters.

## 5.7. Registering the OPS Database

If the network is set up correctly, the recovery and the OPS database are up and running the production database needs to be registered in the recovery catalog. This can be done using the following command:

```
$ rman catalog rman/rman@rmancat target system/manager@op  
RMAN> register database;
```

If this command fails you might have not setup remote login to the target database. In this case refer to chapter 6.1 for more information on how to setup remote login.



## **6. OmniBack II Oracle Integration**

This chapter describes the steps, necessary to configure the Oracle 8i / OmniBack Integration. First you have to configure each of the Oracle Database Servers. This needs to be done in order to be able to backup the Archived Log files from each server, since in an OPS configuration each node, that is part of the cluster, writes it's own Archived Log files. Then you configure the Backup system. This is necessary for the Split Mirror backup integration.

### **6.1. Setting up Remote Login**

In order to allow remote login to the database on the application servers, a password file needs to be created. Modify the following entry in the

`$ORACLE_HOME/dbs/init<SID>.ora` file:

```
Remote_login_passwordfile=shared
```

to

```
Remote_login_passwordfile=exclusive
```

and issue the following command on the application servers after you shut down the database:

```
$ orapwd file=$ORACLE_HOME/dbs/orapw<SID> password=pwd \  
  entries=max_users
```

After you created the password file, you need to grant `sysdba` and `sysoper` to the user, that you want to use for logging in remotely with OmniBack II. For example if the user, that you want to use is `system` you specify:

```
SVRMGR> grant sysdba, sysoper to system;
```

Now you need to copy the file to the `$ORACLE_HOME/dbs/` directory on the backup server. Do that on all the nodes, that are part of OPS and that write archived log files that you are going to backup.

### **6.2. Filesystem Structure on the Backup Server**

Make sure that the Oracle8 software on the Application Server and the Backup Server have the same directory structure. In addition the following files need to be copied from the Application system to the backup system:

`init<SID>.ora` This file has to be copied to the same location on the Backup Server. If this is not possible, you should create a link to this file. In addition the entry `parallel_server=true` needs to be set to `parallel_server=false`.

Copy also the directory `$ORACLE_BASE/admin/<SID>` to the Backup Server.

Please make sure that the permissions are set correctly and that the files are owned by user 'oracle' on the application server. Please refer to the HP OpenView OmniBack II Unix Integration Guide for more information. If you cannot have the same directory structure you can use HP-UX links to create it.

### 6.3. Linking OmniBack and Oracle

Linking the OmniBack shared library with Oracle is not necessary anymore with Oracle 8i. In order to link the shared library with Oracle, shutdown the Oracle Databases you want to backup and issue the following command on the application system.

```
ln -s /opt/omni/lib/libob2oracle8_64bit.sl \  
$ORACLE_HOME/lib/libobk.sl
```

### 6.4. Setting up a User

In order to start an Oracle8 backup session a user needs an operating system logon on the system, where an Oracle8 Server is running.

In addition, this operating system user has to be registered in the Oracle8 database and identified by Oracle8 through the operating system identification.

This means, that the Oracle8 Server does not request connection information from an application started under such user account, but only checks whether the operating system user is registered in the database.

Further on, this user is allowed to backup or restore an Oracle8 database. In order to start a backup of an Oracle8 database using OmniBack II, this user has to become the owner of the OmniBack II backup specification. As the owner of the backup specification, the user has to be added to either the OmniBack II `admin` or `operator` user group.

You can identify this user by running the following command on the

Oracle8 Server system:

```
$ ps -ef | grep ora_pmon_<ORACLE_SID>
```

## 6.5. Configuring the Oracle 8 Client on the Application Servers

In order to be able to backup the archived log files on the application servers, it is necessary to configure the Oracle 8 client on both nodes that produce the archived log files. Furthermore in order to backup the database files from the shared disks using the failover alias, it is necessary to configure a separate client for that as well.

The paper contains an example of the `tnsnames.ora` file in chapter 5.6. This file exists on each of the servers involved in the backup process. It has three aliases that point to the same target database. The first alias to the production database in the file is called 'op1'. This will be used to backup the local archived log files on node 'erwin1'. The second one is 'op2' and will be used for the same purpose on node 'erwin2'. The last entry will be used to backup the database files using the failover alias.

Execute the following command to configure an Oracle8 Server using the OmniBack II command line interface. Make sure you run this command on the Application Server as the user root for all the different database aliases that are configured in the `tnsnames.ora` file:

```
/opt/omni/lbin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> \  
<Target_Database_Login>  
<Recovery_Catalog_Login>
```

The variables are defined as follows:

<ORACLE\_SID> The name of the Oracle8 Server instance alias.

<ORACLE\_HOME> The home directory of the Oracle8 Server instance.

<Target\_Database\_Login> The format of the login information is  
<user\_name>/<password>@<service> It is used to connect to the Target Database via Recovery Manager. The service name must be the alias.

<Recovery\_Catalog\_Login> The format of the login information is  
<user\_name>/<password>@<service>. It is used to connect to the Recovery Catalog Database via Recovery Manager.

In the example environment above the command for the configuration of the application server 'erwin1', would look as follows:

```
/opt/omni/lbin/util_oracle8.exe -CONFIG op1 \  
/OraSrc/app/oracle/product/8.1.5 \  
system/manager@op1 rman/rman@rmancat
```

Repeat these steps for 'op2' and 'op' accordingly.

You can also use the OmniBack II Windows NT User to configure the integration on the Application Server. Please check the HP OpenView OmniBack II Unix Integration manual for more information. Also refer to this manual if you want to know in detail, what this command does.

After you configured the integration on the application server, it is recommended to create an Oracle backup specification in order to test the configuration.

## 6.6. Configuring the Oracle 8 Client on the Backup Server

Execute the following command to configure an Oracle8 Client using the OmniBack II command line interface. Make sure you run this command on the Backup Server as the user root:

```
/opt/omni/lbin/util_oracle8.exe -CONFIG_SMB \  
<ORACLE_SID>  
<ORACLE_HOME>  
<INITFILE_LOCATION>  
<CONTROL_FILE_LOCATION>  
<TARGET_DB_CONNECTION_STRING>  
<RECOVERY_DB_CONNECTION_STRING>
```

The variables are defined as follows:

<ORACLE\_SID> The name of the Oracle8 Server instance on the Application Server.

<ORACLE\_HOME> The home directory of the Oracle8 Server instance on the Application Server.

<INITFILE\_LOCATION> The full pathname of the `init<SID>.ora` file on the Application Server. The default path is `$ORACLE_HOME/dbs`.

<CONTROL\_FILE\_LOCATION> The full path name of the raw file or the directory on the mirrored volumes, to which OmniBack II copies the Oracle8 control file of the Application server. This device needs to have a valid mirror that can be split during the backup process. If your database files and the control file are on raw logical volumes, you need to specify a filesystem, since OmniBack needs to copy the controlfile in place on the backup server, using operating system commands.

<Target\_Database\_Login> The format of the login information is `<user_name>/<password>@<service>`. It is used to connect to the Target Database via Recovery Manager. Use the failover alias as the service descriptor. This way you will be able to start a backup on either of the production servers.

<Recovery\_Catalog\_Login> The format of the login information is  
<user\_name>/<password>@<service>. It is used to connect to the Recovery Catalog  
Database via Recovery Manager.

In the example environment above command to configure the backup server 'karle',  
would look as follows:

```
$ /opt/omni/lbin/util_oracle8.exe -CONFIG_SMB \  
op \  
/OraSrc/app/oracle/product/8.1.5 \  
/OraSrc/app/oracle/product/8.1.5/dbs/ \  
initop.ora \  
/dev/vg02/obctl.ctl \  
system/manager@op \  
rman/rman@rmancat
```

After these steps, you are able to backup the database using the split mirror backup  
facility. Make sure that the owner of the backup specification is the Oracle User from the  
Application System. Please check the HP-OpenView OmniBack II Sure Store E Disk  
Array Configuration Guide for more information about the creation and different options of  
a backup specification.

## **7. Backing up Oracle Parallel Server**

You need to create two backup specifications to start a backup from OmniBack II. It is not the objective of this paper to describe the different options and possibilities that are available within OmniBack. Please see the *HP OpenView OmniBack II Integration Guide* for information on how to create backup specifications and the available options. Also refer to the Oracle8i Backup and Recovery Guide for further information about the available Recovery Manager commands. This chapter focuses on methods that are available to backup the database using split mirror backups. The next chapters will describe ways that are available to backup the archived log files that are created on each node.

### **7.1. Backing up the Database**

The first backup that will be described is an on-line backup of the database using the split mirror facilities. Since the focus of the Oracle Parallel Server is high availability, it will describe on-line backups of the database only. There will be no explanation of off-line backups. The following list shows the different steps, OmniBack II performs during a split mirror backup of the database.

- OmniBack II executes the `ob2rman.exe` command on the backup (R2) server. This command retrieves a list of files to back up from the Oracle8 database on the application (R1) system and starts the resolving process of these files. The `ob2rman.exe` command also copies the Oracle8 database control file to the mirrored devices (`CONTROL_FILE_LOCATION`) on the application (R1) server. The following shows an example session output from the OmniBack II Monitor:

```
[Normal] From: ob2rman.exe@karle Time: 02/21/00 07:11:36

The list of data files of the Production Database has been made.

The file list:

/dev/vg02/rsystem01.dbf
/dev/vg02/roemrep01.dbf

[Normal] From: ob2rman.exe@karle Time: 02/21/00 07:11:36

Making list of database control files of the Production Database ....

[Normal] From: ob2rman.exe@karle Time: 02/21/00 07:11:36

The list of database control files of the Production Database has been
made. The file list: /dev/vg02/rcontrol01.ctldev
```

**Figure 11 - Retrieving the list of files.**

The following shows an example output for the resolving process. This is done for each file or backup object that has been found out to be backed up. You see that each file is resolved down to the physical device file for the disk on which it resides.

```
[Normal] From: SSEA@erwin2 "" Time: 21/02/2000 7:11:32 AM

Backup object /dev/vg02/rssystem01.dbf

Was successfully resolved to the following objects:

Volume group: /dev/vg02

Disk: /dev/rdisk/c0t0d0
```

**Figure 12 - Resolving a File**

Then OmniBack finds the corresponding LDEV for each of the character device files.

```
[Normal] From: SSEA@erwin2 "" Time: 21/02/2000 7:11:33 AM

The following disks were successfully mapped from
character device path names to LDEVs:

Disk                Target Lun SEQ#      Port LDEV
-----
/dev/rdisk/c0t0d0   0    0 35390    1A  000h ( 0)
/dev/rdisk/c0t0d5   0    5 35390    1A  005h ( 5)
-----
```

**Figure 13 - Mapping the Physical Device File to LDEVs**

- The `ob2rman.exe` command then sets the Oracle8 database into backup mode, starts the procedure to split the devices on which the database is installed, and, after the split, puts the database out of the backup mode.

The following output from the OmniBack II Monitor shows how OmniBack II puts the database into backup mode. It also shows that OmniBack II creates a backup of the controlfile to the location specified during the configuration. This is necessary because OmniBack II cannot find a point in time for the controlfile. If it would use the split copy of the controlfile, there would not be the guarantee for consistency. If your database is based on raw devices managed by Oracle itself, you still need to have a filesystem, that can be split of for the destination of the controlfile backup.

```
[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:08

        Issuing ALTER SYSTEM SWITCH LOGFILE on the Production Database ...

[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:08

        Issuing ALTER TABLESPACE ... BEGIN BACKUP on the Production Database

[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:16

        Issuing ALTER DATABASE BACKUP CONTROLFILE TO ... for the Production
        Database ...

[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:19

        The Production Database is in backup mode.

        Backup Control file is at "/ArchLog/ctrlop.ctl"
```

**Figure 14 - Putting the Database in Backup Mode**

The following output shows how OmniBack does the split of the links from the P-Vol to the S-Vol.

```
[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:19

        Starting split of mirror disks with database data and backup control files.

[Normal] From: SSEA@erwin2 "" Time: 21/02/2000 7:12:12 AM

        Preparing the Application System for splitting of disks.

[Normal] From: SSEA@erwin2 "" Time: 21/02/2000 7:12:13 AM

        The Application System is prepared for splitting of disks.

[Normal] From: SSEA@karle "" Time: 21/02/2000 7:12:20 AM

        Splitting disks.

[Normal] From: SSEA@karle "" Time: 21/02/2000 7:12:21 AM

        Splitting of disks completed.

[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:21

        Split of mirror disks succeeded.
```

**Figure 15 - Splitting the Disks**



After the split of the disks has succeeded, OmniBack changes the state of the database again. It puts it out of backup mode into normal operation. If you look at the timestamps you will see, that the database has only been a small amount of time in backup mode.

```
[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:21

        Issuing ALTER TABLESPACE ... END BACKUP for the Production Database

[Normal] From: ob2rman.exe@karle "op" Time: 02/21/00 07:12:25

        The Production Database is in non-backup mode.
```

#### Figure 16 - Putting the Database in Normal Operation

- The `ob2rman.exe` command then starts the procedure to prepare the mirror disks on the backup (R2) server. This involves activating the volume groups and mounting filesystems if applicable.
- The `ob2rman.exe` command starts the RMAN backup command, `rman`, on the backup (R2) server. Then it sends the Oracle8 RMAN Backup Command Script to the standard input of the `rman` command.
- Oracle8 RMAN contacts the Oracle8 database, which contacts OmniBack II via SBT API and initiates a backup.
- RMAN reads data from the L-DEV mirror disks and sends it to OmniBack II for writing to the backup devices.
- At the very end OmniBack automatically performs a backup of the Oracle 8 Recovery Database is performed. OmniBack II is using the standard Oracle 8 export utility. Then `ob2rman.exe` backs up the file, created by the export command.

The following is an example backup specification for the database files on the application server using the split mirror facility of OmniBack II.

```

BARLIST "ops_db"
OWNER oracle dba erwin1.hp.com
GROUP "OPS"
DYNAMIC 1 5
SSE
{
}
SSE
{
    BC erwin1.hp.com karle.hp.com
    -establish
    -disable_bs
}
DEVICE "dlt7k_dev1"
{
    -sync
    -pool "OPS_DB"
}

CLIENT "op" erwin1.hp.com
{
    -exec ob2rman.exe
    -args {
        "-backup"
    }
    -input {
        "run {"
            "allocate channel 'dev_0' type 'sbt_tape'"
            "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST=ops_db)';"
            "allocate channel 'dev_1' type 'sbt_tape'"
            "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST=ops_db)';"
            "allocate channel 'dev_2' type 'sbt_tape'"
            "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST=ops_db)';"
            "allocate channel 'dev_3' type 'sbt_tape'"
            "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST=ops_db)';"
            "backup incremental level <incr_level> filesperset 1 "
            "format 'ops_db<op_%s:%t>.dbf'"
            "database "
            "include current controlfile"
            ";"
            "}"
        }
    -public
}

```

**Figure 17- Example Backup Specification**

The amount of channels you allocate should always be the sum of the concurrencies of the backup devices. For example the backup device, that was used here, has a concurrency of four. So four channels have been allocated.

---

NOTE: If you use the failover facility for the failover of the connection to the application servers, you might get a warning from OmniBack, which is similar to 'received connect from client system that is not specified in the barlist'. This message is a warning and can be ignored.

---

The backup specifications for Oracle 8 can be found in the following directory:

```
/etc/opt/omni/barlists/oracle8/
```

## 7.2. Backing Up the Archived Logs

The procedures for backing up the archived log files in OPS environments are slightly different than the ones in regular Oracle 8i. Archive logs are generally only accessible by the node on which they were created. In OPS you have different backup options:

- Have each node back up its own archive logs
- Move the archived logs to one node, and then back them up
- Write the archived logs to more than one destination so that both nodes can access all logs all the time.

The following chapters will describe the different possibilities in detail. It should be considered to mirror the archived redo logs in a way, that RMAN always has access to all threads and to avoid loss of data. Consider the situation, where a node has already written some archived logs and goes down for whatever reason. RMAN will have no chance to access the archived logs of the crashed node. Only incomplete recovery will be possible and the database needs to be opened with resetlogs option, which causes a loss of data.

### 7.2.1. Backup the Archived Logs by Each Node

The first possibility explains how you can backup the archived log files, if you don't share them between the different servers. This is the preferred solution for backing up the archived log files. Recovery Manager can automatically enable each node to back up its own archive logs. Each node can then back up the archived logs it has created.

In order to enable archiving on the nodes that are going to be backed up, you need to specify the

```
log_archive_dest = "<DEFAULT_PATH>/<DEST_NODE1>/"
```

in the `initop.ora` file on each of the database servers. If you have for example `/OraSrc/app/oracle/admin/op/` as the `<DEFAULT_PATH>` you could create a subdirectory called 'arch1' as the destination for the archived log files produced on the instance with thread one. Do the same for the other node, where you use the following directory:

```
/OraSrc/app/oracle/admin/op/arch2/
```

This way, you will be able to use the `like` clause in the RMAN script, which will only backup the archived log in the `log_archive_dest` visible to that instance.

Use the OmniBack II Windows NT user interface in order to create a backup specification for the server 'erwin1', using the SID 'op1' created during the configuration of the Application system. Do the same for the server 'erwin2'. Specify either a local or a device that is connected to the dedicated device server. Save the backup specifications.

The backup of the archived log files on each node, needs to be performed after each database backup run. You can execute post-database backup commands within OmniBack II in order to start a backup of the archived log files after the database backup or have your own methods to start the backup of the logs. Without these logs, your database backups will be useless.

Then you need to manually edit the backup specifications in order to enable them to backup only specific archived log files. The backup specifications will be stored on the cell manager system in the directory:

```
/etc/opt/omni/barlists/oracle8/
```

The following is an example backup specification for the server 'erwin1':

```
BARLIST "archivelogs_op1"
OWNER oracle dba erwin1.hp.com
GROUP "OPS"
DESCRIPTION "Backup of the archivelog files on erwin1"
DYNAMIC 1 5
DEVICE "local_dev_erwin1"
{
  -sync
}
CLIENT "op1" erwin1.hp.com
{
  -exec ob2rman.exe
  -args {
    "-backup"
  }
  -input {
    "run {"
      "allocate channel 'dev_0' type 'sbt_tape'"
      "parms ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op1,OB2BARLIST=archivelogs_op1)";"
      "allocate channel 'dev_1' type 'sbt_tape'"
      "parms ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op1,OB2BARLIST=archivelogs_op1)";"
      "allocate channel 'dev_2' type 'sbt_tape'"
      "parms ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op1,OB2BARLIST=archivelogs_op1)";"
      "backup incremental level <incr_level> filesperset 1 "
      "format 'archivelogs_op1<op1_%s:%t>.dbf'"
      "archivelog like '%/arch1/%';"
    }"
  }
}
```

**Figure 18 - Backup Specification for Archivelogs produced by 'erwin1'**

Here you see how you can make use of the 'like' clause within the backup command in the RMAN script. Here RMAN is told to look for the archived log files that contain a directory called 'arch1'. Create the same kind of backup specification for the second node and modify the script accordingly. That means you would enter like '%/arch2/%'.

The number of allocated channels within the RMAN script is determined by the sum of the concurrent datastreams the devices, specified in the backup specification, accept. If you have for example 2 tape devices with concurrency 4, the amount of channels that can be allocated would be 8.

## 7.2.2. Backup Archived Logs from one Node

If you are using one node to back up your archive logs, when Recovery Manager compiles the list of logs to be archived, it must be able to check that the archived logs exist. To do this it must be able to read the headers of all archived logs on all nodes. This chapter describes examples that can be used to perform the archived log backups from one node.

### 7.2.2.1. Shared Archived Log Directories

In this case, all nodes see all log files produced by each other. This is the simplest method of backing up the archived log files. You can use the following backup specification to backup the archived logs from either node. You can use the failover alias 'op' for that kind of backup.

```
BARLIST "all_archived_logs"
OWNER oracle dba erwin1.hp.com
GROUP "OPS"
DYNAMIC 1 5
DEVICE "local_dev_erwin1"
{
  -sync
}
CLIENT "op" erwin1.hp.com
{
  -exec ob2rman.exe
  -args {
    "-backup"
  }
  -input {
    "run {"
    "allocate channel 'dev_0' type 'sbt_tape'"
    "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST= all_archived_logs)';"
    "allocate channel 'dev_1' type 'sbt_tape'"
    "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST= all_archived_logs)';"
    "allocate channel 'dev_2' type 'sbt_tape'"
    "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST= all_archived_logs)';"
    "backup incremental level <incr_level> filesperset 1 "
    "format ' all_archived_logs <op_%s:%t>.dbf'"
    "archive log all;"
    "}"
  }
  -public
}
```

**Figure 19 - Backup Archived Logs with Shared Directories**

There is no need to modify the backup specification, created by the OmniBack II Windows NT user interface.

### 7.2.2.2. Backup Archived Log Files using Connect Command in RMAN

This is the second method that can be used to backup the archived log files from one node. This example will make use of the 'connect' clause in the 'allocate' command of an RMAN script.

---

NOTE: This feature is only available in Oracle 8.1.6. It will not be possible to use it with Oracle 8.1.5 and lower. This is because of a bug (Oracle bug # 837079) that was fixed in 8.1.6: When backing up archivelogs and the user has specified a particular channel for doing the backup, then RMAN should perform the log validation on that channel rather than the default channel because the logs may not be accessible from the default channel.

---

```
BARLIST "all_archived_logs"
OWNER oracle dba erwin1.hp.com
GROUP "OPS"
DYNAMIC 1 5
DEVICE "local_dev_erwin1"
{
  -sync
}
CLIENT "op" erwin1.hp.com
{
  -exec ob2rman.exe
  -args {
    "-backup"
  }
  -input {
    "run {"
      "allocate channel 'dev_0' type 'sbt_tape'"
      "connect system/manager@op1"
      "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST= all_archived_logs)';"
      "allocate channel 'dev_2' type 'sbt_tape'"
      "connect system/manager@op2"
      "parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=op,OB2BARLIST= all_archived_logs)';"
      "backup incremental level <incr_level> filesperset 1 "
      "format ' all_archived_logs <op_%s:%t>.dbf'"
      "archivelog like '%/arch1/% channel dev_0"
      "archivelog like '%/arch2/% channel dev_1"
      "}"
    }
  }
  -public
}
```

**Figure 20 - Backing up Archived Logs from one Node using the connect clause**

Again you need to have this backup scheduled, after the database has been backed up or you can put the backup command as a post-execution command of the database backup.

### 7.2.2.3. Copying Archived Logs to one Node

You can use operating system tools to copy the archived log files to one location. If you use a script, you can put the execution of this script as a pre-exec command before the actual backup of the archived logs is started. You need to perform the following steps before the backup:

```
alter system archive log current;  
  
rcp erwin2:/OraSrc/app/oracle/admin/op/arch2/*\  
    /OraSrc/app/oracle/admin/op/arch2
```

This will cause a log switch on the server 'erwin2', and will then copy the existing log files to the log directory on 'erwin1'. After that, you can use the same backup specification as it is explained in chapter 7.2.2.1 started on node 'erwin1'.



## **8. Restore and Recovery**

The Oracle8 database objects can be restored using the Oracle8 Recovery Manager (RMAN) only. It is not possible to initiate a restore or recovery from within the OmniBack II user interface. Use the RMAN `restore` command to restore datafiles, control files, or archived redo logs from backup sets. When you do a restore in an OPS split mirror environment, you have several different possibilities. This chapter will first describe the restore to the application server and then the restore to the backup server.

The commands you use in Recovery Manager for OPS are the same as those you use to recover single instance environments. Recovery Manager automatically restores and applies the archive logs required. By default, Recovery Manager restores archive logs to the `LOG_ARCHIVE_DEST` directory of the instances to which it connects.

If you are using multiple nodes to restore and recover, this means that the archive logs may be restored to any of the nodes performing the restore / recovery. The node that actually reads the restored logs and performs the roll forward is the target node to which the connection was initially made. You must ensure that the logs are readable from that node.

### **8.1. Restore to the Application Server**

In order to be able to do a fast restore it is recommended that you have local devices connected to the application server. This can either be a locally attached tape library or one that is shared between the multiple systems within the environment using fiber channel hardware. This paper doesn't go into detail with regards to library sharing in SAN environments. Please refer to the HP-OpenView OmniBack II Administrator's Guide for more information about device sharing.

The database on the application server must not be open, when you do the restore. Therefore use the following command to shut it down first and mount it afterwards:

```
$ svrmgrl
SVRMGR> connect internal;
SVRMGR> shutdown immediate;
SVRMGR> startup mount
```

Then you need to connect to the recovery catalog database on the backup server from any of the application server using the following RMAN commands:

---

NOTE: If you did your backups using Oracle 8.1.6 you may specify the 'set autolocate on' for the restore of the archived log files to the application server

---

```

$ rman target system/manager@op catalog rman/rman@rmancat
RMAN> run {
    allocate channel 'dev_0' type 'sbt_tape';
# This line restores the database
    restore database;
}

run {
allocate channel 'dev_0' type 'sbt_tape';
# set a new archive log destination for all logs so that it
# matches the LOG_ARCH_DEST of the host on which restore was
# started - here marked with an 'X'
set archivelog destination to '/OraSrc/app/oracle/admin/op/archX';
restore archivelog all;
recover database;
}

```

Please refer to the Oracle 8i Backup and Recovery guide for more information about restoring a whole database, tablespaces, controlfiles, etc.

---

NOTE: You must specify as many channels for parallel restore as the total device concurrency, you use for the restore. So if you use 2 DLT drives, each with a concurrency of 3, you need to allocate 6 channels.

---

## 8.2. Restore to the Backup Server

The disks where the database should reside, must be available to the backup server before you can proceed. That means you need to shutdown the database on the application server, split the disks in SIMPLEX mode, enable the disks on the backup system and then mount them. Now you can use Recovery Manager to perform a restore to the backup server. After the restore completed, you can mount the database on the backup server in order to do a recovery. After you successfully recovered the database, you need to shut it down. Then unmount the disks, de-activate the volume groups and create device pairs on the backup server, where the primary volumes are the disks on which the restored database resides. The last step involves the resynchronization of the disks from the backup server.

## **9. Troubleshooting**

This chapter describes some of the most common error messages for such kind of environment. Please refer to the official OmniBack II and Oracle manuals for more information regarding troubleshooting.

### **Message:**

```
RMAN-20242: specification does not match any archivelog in the
recovery catalog
```

Here it could be that the archives have been backed up previously with 'delete input' option. Run the command

```
RMAN> 'change archivelog like '%/archX/%' crosscheck'
```

on the RMAN command prompt of each host, where the X stands for the corresponding archiving directory.

### **Message:**

```
ORA-00439: feature not enabled: Parallel Server
DISCN
```

```
-----
```

```
ORA-01507: database not mounted
```

```
Server Manager complete.
```

```
[Normal] From: ob2rman.exe@karle "op" Time: 03/29/00 10:38:49
Shutdown database on the Backup Server ...
```

The database was mounted on the Backup Server. However, the parallel server entry in the init<SID>.ora was not changed to parallel\_server=false.

### **Message:**

```
10:13:08 : (4204) OBK-sbt: [Major] From: OB2BAR@hansel.cup.hp.com
"op" Time: 03/29/00 10:13:08
```

```
Received ABORT request from BSM (ERR: Error starting
backup/restore (BMA cannot
```

```
10:13:08 : (4197) OBK-sbt: [Major] From: OB2BAR@hansel.cup.hp.com
"op" Time: 03/29/00 10:13:08
```

```
Lost socket connection to BSM (Details unknown.)
```

Check your device configuration. This can be caused, when you manually change the backup device in the files /etc/opt/omni/barlists/oracle8/\*. Change the backup device using the GUI and save the barlist again.

**Message:**

RMAN-02002: unexpected end of input file reached

This occurs if RMAN detects a syntax error within the RMAN script. This is usually caused by changing the RMAN Script in the barlist manually.

**Message:**

RMAN-06089: archived log  
/OraSrc/app/oracle/admin/op/arch/arch\_T2\_S178.arc not found or out  
of sync with catalog

This message occurs if the archived log files are not accessible on the server, for which the backup is configured. Either the log file has been deleted or the server, on which the backup is started from, is not the one who writes the logfile. Please check the thread number in the log file and check if this node was writing the serverfile. Check if the file exists. Issue 'change archivelog like '%/archX/%' crosscheck' on the RMAN command prompt of each host, where the X stands for the corresponding archiving directory.

**Message:**

RMAN-04005: error from target database: ORA-01031: insufficient  
privileges

OmniBack II used a login string to the database, which was correct, but the user that has been specified, doesn't have the correct privileges. Check whether the specified user has sysdba and sysoper privileges. If not please check 6.1 Setting up Remote Login.

**Message:**

[Major] From: ob2rman.exe@karle "op" Time: 04/05/00 09:46:16  
Check of the OB2/Oracle8 configuration failed.  
[12:8300] The Oracle8 SQL\*Net login information to the  
database instance is not correct or database is not  
running.

If this message occurs randomly after a successful configuration, you may want to enter IP addresses instead of the fully qualified hostnames in the listener.ora and tnsnames.ora files. Otherwise, the configuration, the user, etc. may have been changed between the last backup. In this case you need to configure the environment again as described in chapter 6.

## References

---

- [1] Backup and Recovery OPS - Guidelines for Backup and Recovery Oracle 8i Parallel Server - Technical Whitepaper - Juergen Paasch, Heiko Robert, Hogan Flake (Oracle)
- [2] Setting up a Manual Online and Offline Backup and Restore Solution using HP SureStore E Business Copy XP, HP OpenView OmniBack II and Oracle RMAN - Taeko Wong (HP)
- [3] HP-OpenView OmniBack II official documentation
- [4] Oracle 8i official documentation
- [5] Oracle Parallel Fail Safe on HP ServiceGuard Clusters, Technical Whitepaper - Ken Lee (HP)

## Acknowledgements

---

Krishnadev Telikicherla, Oracle Corporation

Francisco M. Sanchez, Oracle Corporation

Harvey Skinner, HP

Taeko Wong, HP

Ken Lee, HP

Harald Burose, HP

---

Edition 001

Last updated: April 26, 2000

Hewlett-Packard, HP and OpenView  
are registered trademarks of

Hewlett-Packard Company.  
Other trademarks are the property  
of their respective owners.

Technical information in this  
Document is subject to change  
without notice.

Copyright © 1999 Hewlett-Packard Company  
Printed in USA 10/99



Storage  
Management

HP OpenView **Solutions**