



hp storage solutions

march 2003

solutions
implementation
blueprint

StorageWorks Rapid Backup for mySAP Business Suite on HSG80 on Oracle on Tru64

executive summary

Businesses deploy mission-critical SAP/Oracle complex applications to meet business requirements. The ongoing availability of these applications to users across the organization becomes hypercritical when performing backups and restores.

This solution satisfies the needs of customers, who require high-availability and high-performance computing capability on most reliable configurations. Customers require mySAP Business Suite in business applications for a 7X24 production environment.

Rapid Backup Solutions was developed by HP StorageWorks Engineering, SAP Global Solutions to complement SAP business applications designed for businesses that need to run 7X24 – SAP/Oracle databases. These integrated storage solutions solve the three major business problems associated with backup and recovery today:

- Shrinking backup window
- Complexity of backup
- The time window for rapid restore and recovery

overview

This solution is integrated with SAP/Oracle and involves streaming backup data from a mirror instead of an original disk. Because it minimizes the impact on the production system and simplifies backup management, the Rapid Backup Solution delivers quantifiable business benefits.

This solution also guarantees rapid restore and recovery of data in a reasonable time, ensuring that the database is up and running in a short period. This means that the customer mean time between failures (MTBF) is drastically reduced.

This solution uses the HSTK utility, which directly communicates to HSG80 subsystems and performs the split/mirror operations at HSG80 controller level.

Table of contents		
	executive summary	1
	overview	1
	Table of contents	2
	business problem overview	3
	hp's answer to the business problem	3
	reduce performance of database during backup	3
	hp solutions overview	4
	general recommendation	5
	SAP/Oracle backup process flow	5
	restoring lost files from BCV units	9
	storage layout recommendation	10
	Mirrorsets (RAID 1)	10
	Striped mirrorsets (RAID 1+0)	11
	detailed description, how to set up	13
	system setup and installation	13
	limitations	13
	configuring the backup system:	14
	non-disruptive backup for mySAP Business Suite solutions with different configuration	17
	description of lab setup	23
	testing environment	23
	bill of materials	23
	MSL5026 tape library notes	25
	SAN setup and configuration	25
	Host Bus Adapter configuration performance considerations	25
	storage subsystem configuration notes	26
	MA8000/EMA12000 installation	26
	rapid restore and recovery procedures	26
	Case 1: Full restore and recovery from BCV units to the production system:	26
	Case 2: Partial restore and recovery	29
	Case 3: Restoring and recovering from tape	29
	how to set up storage as backup media	31
	appendix a: examples of lab setup	36
	disk layout	36
	example of listener.ora	36
	example of tnsnames.ora	37
	example of SAPDBA profile	38
	appendix b: Brbackup, Rapid-Restore, Rapid-Recovery logs	46
	script: do_split.ksh	46
	scripts: do_resync.ksh	46
	Brbackup output Log	47
	Rapid Recovery Log	52
	Rapid-Restore Log	55
	appendix c: SAP/Oracle storage growth and impact on HSG80	59
	appendix d: SAP/ORACLE backup and recovery strategies	60
	online backup internal operation	60
	appendix e: HSTK installation overview	64
	references	65
	more information	65

business problem overview

In a mission critical SAP/Oracle 24/7 production environment, customers need to perform either online backup or offline backup to protect company data. During online backup, users may experience performance impact on the production system because the database is in backup mode for a long period. This impact is caused by Oracle generating up to 60% more log information and requiring up to 250% more CPU time for inserts while the database is in backup mode. During offline backup, the Oracle database has to be shutdown for a long period, disrupting service from the mySAP Business Suite system. Due to the service level agreement, they cannot afford to shut down the database, and the customer wants the SAP/Oracle database to be up and running on a 24/7 basis. This solution overcomes all the problems described above using HSG80 Split/Mirror technology. The HSTK utility can be integrated with SAP R/3 Brbackup to prevent impact of this kind on the production system during backup.

hp's answer to the business problem

reduce performance of database during backup

Figure 1 shows Brbackup performance and the impact on the database in different types of backup.

During offline file backup from the production system, the database is not available to users for 00:47 minutes. This therefore impacts the business. In an order entry system, for example, if users cannot place orders during this period, the customer will lose a lot of business.

During online backup from the production system, the database is in backup mode for 00:47:15 minutes. This therefore also impacts the performance of the database even though users can access the system.

During offline-split backup from the backup server, the database is not available for 00:17:58 minutes. This therefore also impacts the business due to non-availability of the database to users.

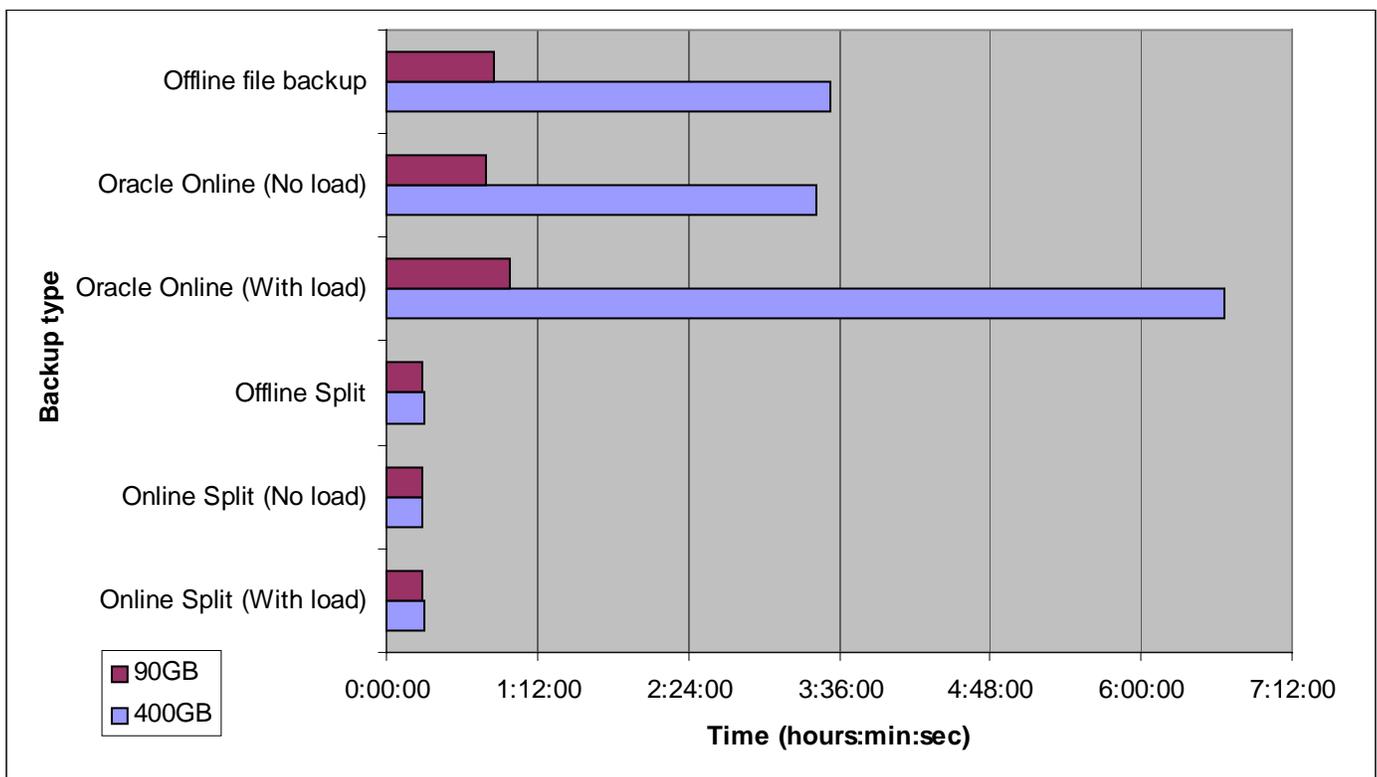


Figure 1, reduced database performance/database not available

During online-split backup from the backup server with load on production system, database performance is slightly affected but users can still use the system. The time window for the split is as short as possible. Hence, the online-split backup is considered better for mission-critical applications such as SAP.

During online-split backup from backup server without load on production system, database performance is not affected so users can use the database normally. This backup can therefore be considered as high-availability and the best option for mission-critical applications such as mySAP Business Suite in a 24/7 environment.

hp solutions overview

This solution takes advantage of the ability of storage subsystems to clone data. The process of cloning data on the disk subsystem has been integrated with SAP's BRBACKUP utility. The solution described in this paper covers Tru64 and Oracle using HP StorageWorks Modular Array.

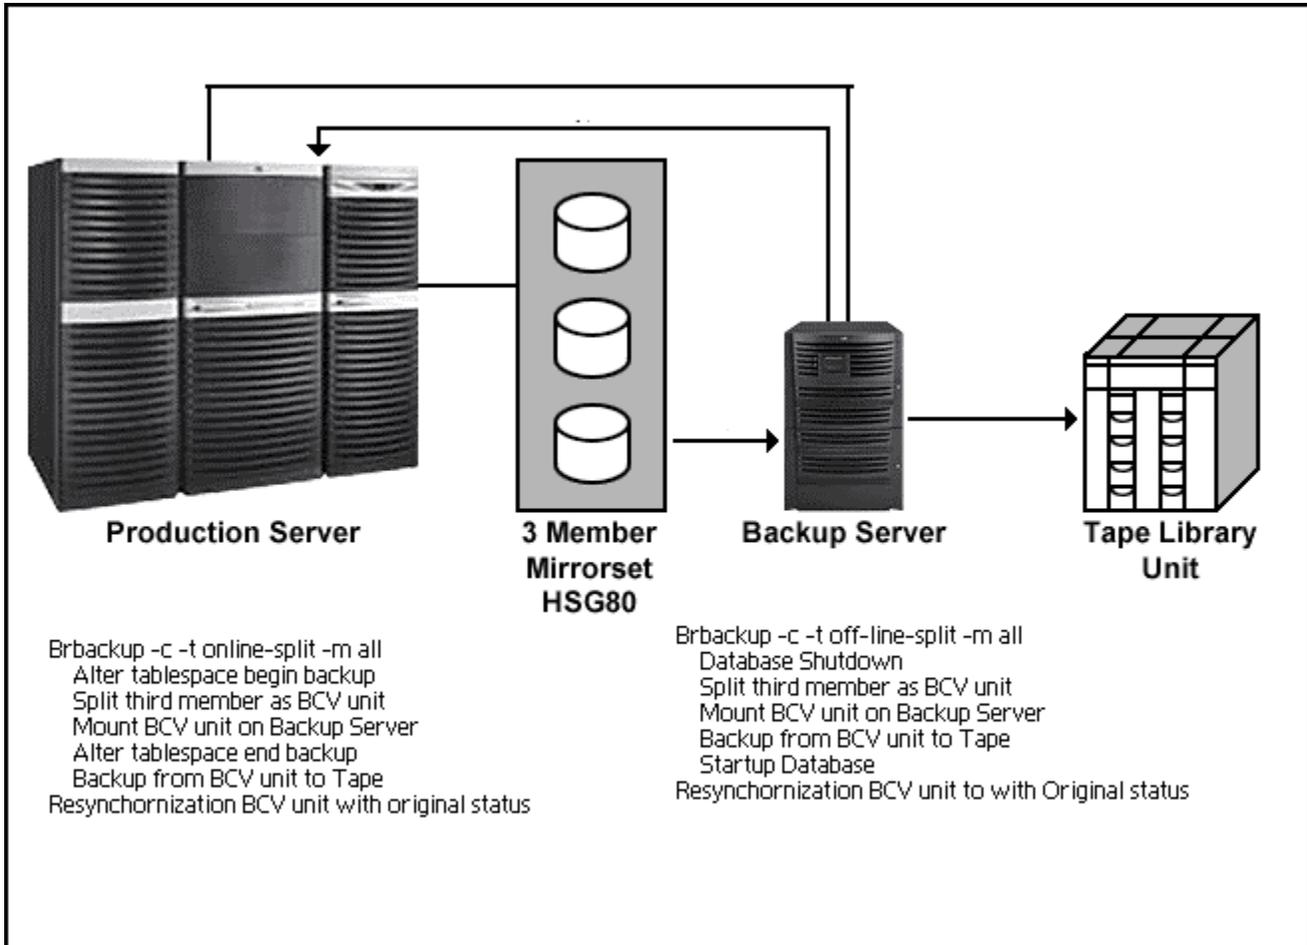


Figure 2, Split Mirror solutions overview

This solution helps customers to achieve Return On Investment (ROI) in several ways:

- Point in time backup with almost no performance impact on business applications.
- Reduced time during which the database is in backup mode for online backup or reduced time during which the database is unavailable for offline backup.
- Less time required to get the SAP system up and running after a failure where data has to be restored.
- Rapid-recovery is possible within a reasonable time in the event of a disaster affecting the production environment.

The solution relies on quickly creating a copy of the database and having a dedicated backup server backing up the data (as shown in Figure 2), thereby minimizing the impact on the production system. The database backup can be performed using backup methods supported by Brbackup (with local tape, backint, etc).

Dedicated scripts are used to communicate with the disk subsystem's HSG80 controller to perform the split operation for the 3rd member of the mirrorset and to resynchronize the mirror without affecting the production system. For online backup, these operations are carried out in the background so users will not experience any disruption of the production system. For offline backup, disruption of the SAP system is greatly reduced - downtime is measured at 15 minutes. Also, the backup operations are executed in the backup system, further reducing the load on the production system.

general recommendation

The non-disruptive backup solution has been simulated in a business-critical environment using specific test configurations. SAP R/3 is one of the most intensive applications for I/O disks in the industry. As the disk subsystem is the slowest part of an R/3 system, it is extremely important to design it for optimum performance. For an SAP R/3 system, the disk subsystem for the SAP/Oracle databases is commonly designed using RAID fault-tolerant settings.

Therefore, in our test we used RAID 1+0, striped mirror volume sets for performance and reliability. To achieve our goals and optimum throughput, we distributed the data evenly on several disks to balance the I/O of different groups of files.

We recommend that you create a three-member striped mirrorset (a striped set whose members are mirrorsets). The split-off third member is known as a Business Continuance Volume (BCV). After the BCV has been backed up to tape, it can be resynchronized immediately or it can be retained for a day depending upon the business needs.

The main advantage of keeping a BCV unit for more time is that the moment you find that you have lost a data file or want to restore a file for some reason (such as ensuring consistency) then you can restore the data file immediately as long as BCV units are available on the backup system. This way the production system can be up and running again very quickly.

From a Restore perspective, backing up data to disk and restoring from disk can be faster than similar procedures using tape. Once the files are readily available on disk, in the event of a database crash they don't have to be restored from an offline storage media, thus reducing MTTR (mean time to recover). Most of the delay occurs due to the time it takes to restore data from tape. This is one of the main reasons that we recommend maintaining the most recent online backup on disk.

We also recommend that you schedule online-backup during the low period of DML (data manipulation language) activity and preferably when there is less load on the production system, in order to avoid generating more checkpoints in the SAP/Oracle database.

For security and better performance, database data files and database online redo log files should reside on different disks.

The overall performance of the disk subsystem depends heavily on the number of disk drives installed. The more drives or spindles installed in a volume, the more I/Os per second and the greater the transfer rate the volume is able to provide.

In a typical mySAP Business Suite installation the internal structure of Oracle tables is constantly changing because of update activities. Since most SAP applications are oriented toward heavy online transaction processing, the internal structure of SAP tables typically becomes less optimal over time. Reorganization substantially improves table performance.

SAP/Oracle backup process flow

For a backup, the third disk of a three-member mirrorset is detached and mounted on the backup system. The backup process (streaming data to tape) is started on the backup system without interfering with the disks on which the productive database is running. After successful Brbackup, the third disk rejoins the productive database's mirrorset.

This integration does not create new mirrors to perform the backup. It uses the existing 3 member mirror pairs. The pair of disks is already synchronized and in paired status at the beginning of the backup.

Application systems and file systems always keep important data in the system cache. The system cache must be flushed to the disks in order to make the data consistent with the backup system before the mirrors are split and the backup is started.

online split mirror backup

An online split backup is more sophisticated than offline split backup since the database is continuously available to users, thus enabling them to continue to access their SAP system. Specifically for online split, the BRBACKUP tool in the SAPDBA tool suite alters the status of the tablespaces to backup mode. Soon after split has been completed, BRBACKUP alters the tablespaces to end backup mode. Other than a change in one entry in the SAP backup configuration file (initSID.SAP), the software configurations and BRBACKUP procedures are the same for both online and offline split mirror backups. In general, BRBACKUP performs the steps outlined in Figure 3.

The following steps take place during online split mirror backup:

- Creates a copy of the database control file and places it into `/oracle/SID/sapbackup`
- Switches log file
- Puts the tablespaces into online backup mode
- Calls an external script that initiates the split
 - flushes the file system
 - calls the script that performs the split
 - call the script that mounts the BCV on the backup server
- Changes the status of tablespaces, so they are no longer in backup mode
- Switches log file
- Executes backup on backup server to back up the content of the BCV
- Optionally resync the BCV unit with original mirrorset

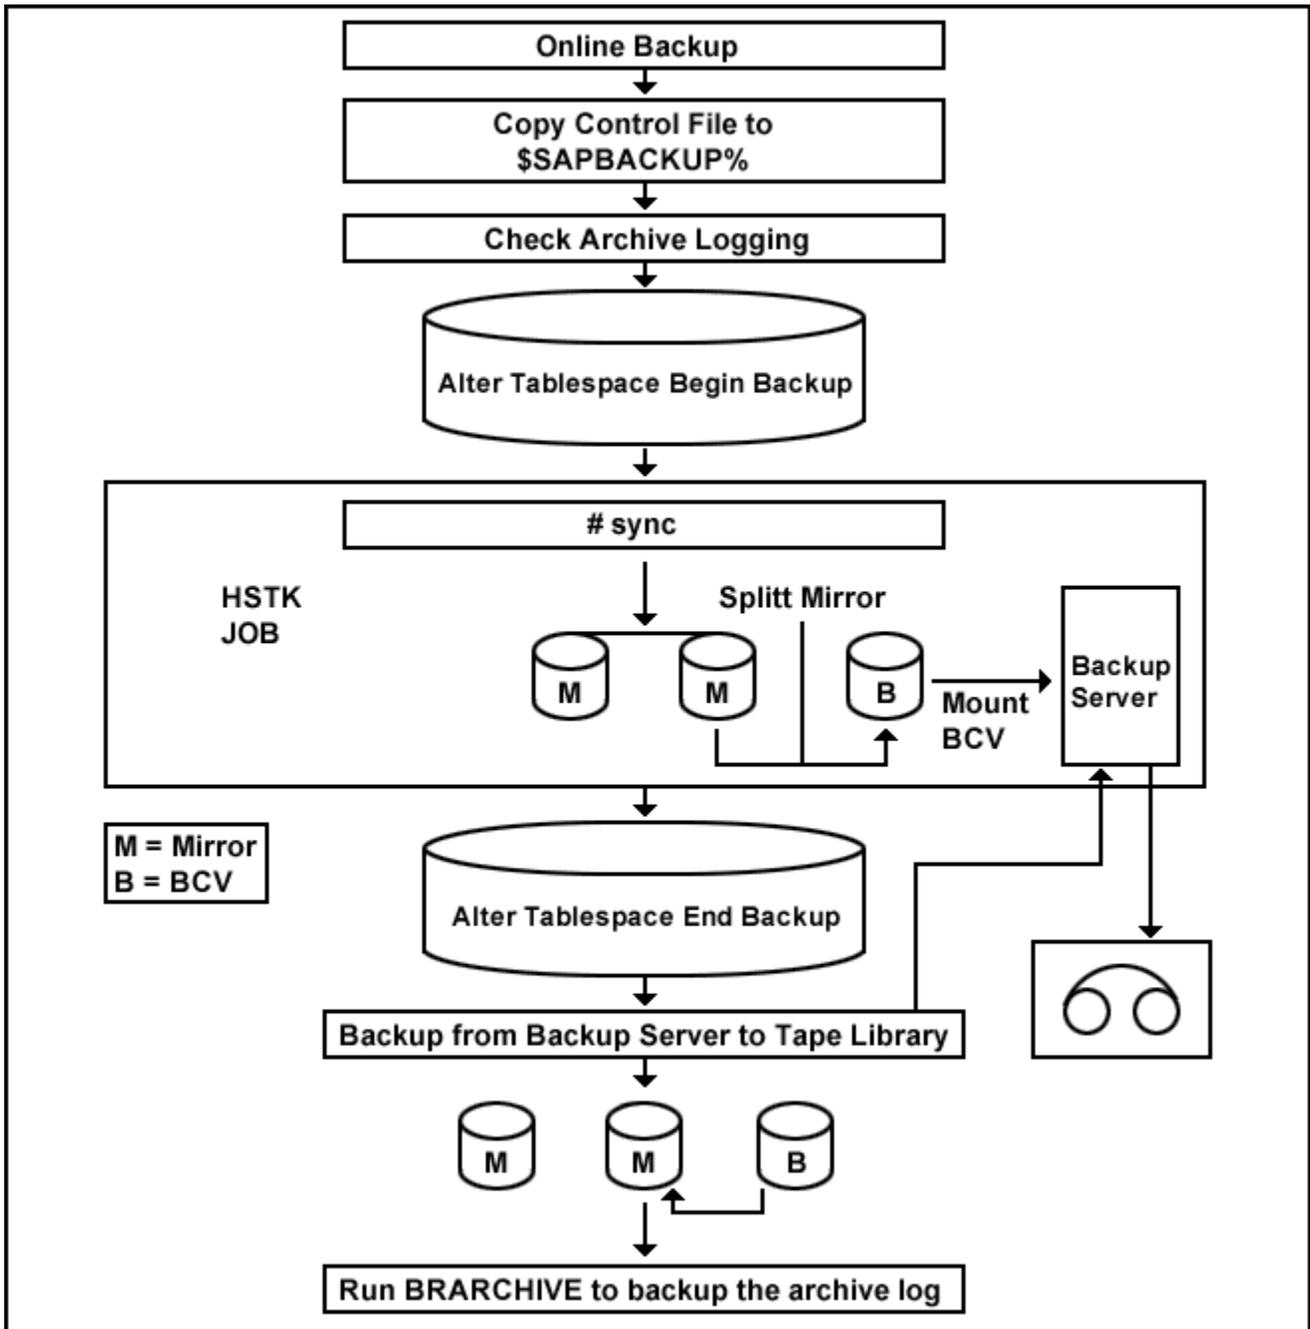


Figure 3, SAP/Oracle Online Split Mirror Backup

For point in time recovery the archive logs need to be applied in addition to recovering data from BCV units for recovery. During online split backup and after successful backup, preserve or back up archive logs to different media. This process is very important for point in time recovery.

offline split mirror backup

During SAP offline split backup, the database is shut down to perform the BCV split. Soon after the split has been completed, BRBACKUP restarts the production database, thereby making the system available to users again. The physical backup is then performed on the backup system. This is the simplest backup strategy and results in a consistent point-in-time database backup.

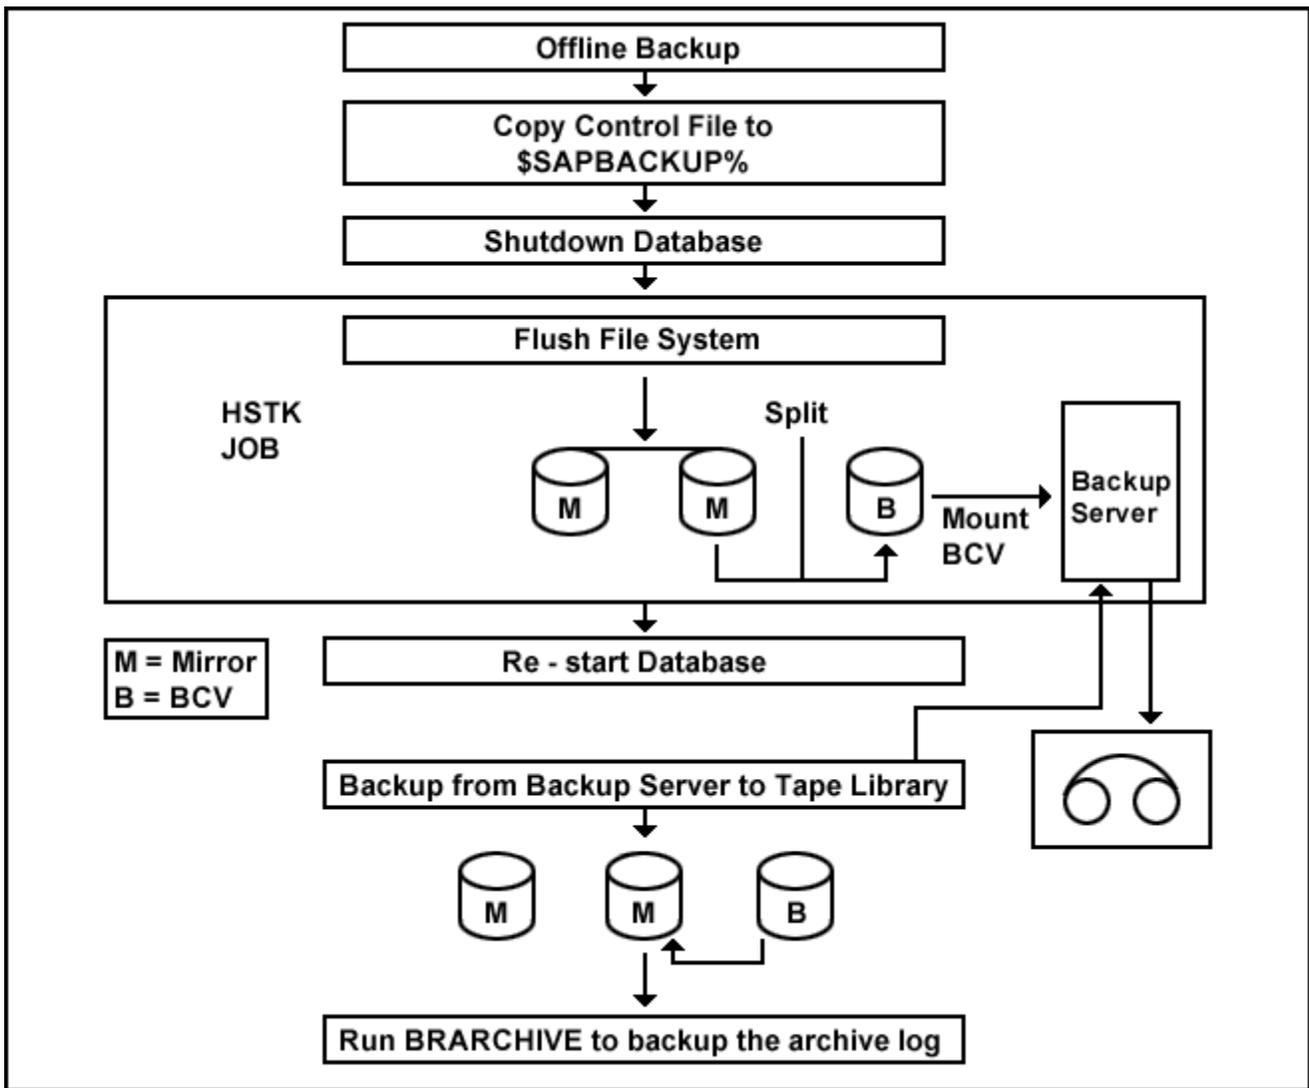


Figure 4, SAP/Oracle Offline Split Mirror Backup

The following steps takes place during offline split mirror backup:

- Creates a copy of the database control file and places it into `/oracle/SID/sapbackup`
- Shuts down the database
- Calls an external script that initiates the split
- Flushes the file system
- Calls the HSTK script that performs the spit
- Calls the HSTK script that mounts the BCV on the backup server
- External script returns control to Brbackup
 -
 -
- Restarts the database
- Executes backup on backup server from BCV unit to tape
- Optionally resync the BCV unit with original mirrorset

If you resync the BCV unit immediately after Brbackup has been successfully completed, you will not be in a position to restore the files immediately. You can therefore decide how long you want to retain the BCV units, depending upon your business needs.

restoring lost files from BCV units

SAP Brbackup offloads backup data from the production system to the backup system and then communicates with backint (backint is a type of interface between Brbackup and backup software vendor backup utility). For security reasons, most backup utility vendors record the name of the backup server in the index catalog of the files stored on tape. During restore the utility compares the name of the server the restore is initiated on with the name stored in the index catalog. Due to this security check, it is not possible to successfully restore a file to the productive server from a tape created on the backup server.

If you lose a volume on the productive database server, then restoring lost files from tape as well as from BCV units, requires manual intervention. The volume to which data must be restored has to be presented to the backup system. The backup system then restores the file(s) from tape to the volume. After a successful restore, the volume then has to be presented back to the database server. Given the number of steps involved in this process, it is generally faster and easier to just restore the data from tape directly back to the original volume, even if this involves restoring via the network.

Figure 5 shows that in the event of original LUN loss due to hardware error or corrupted data files, you can perform the recovery from the BCV unit, if the BCV unit is still available.

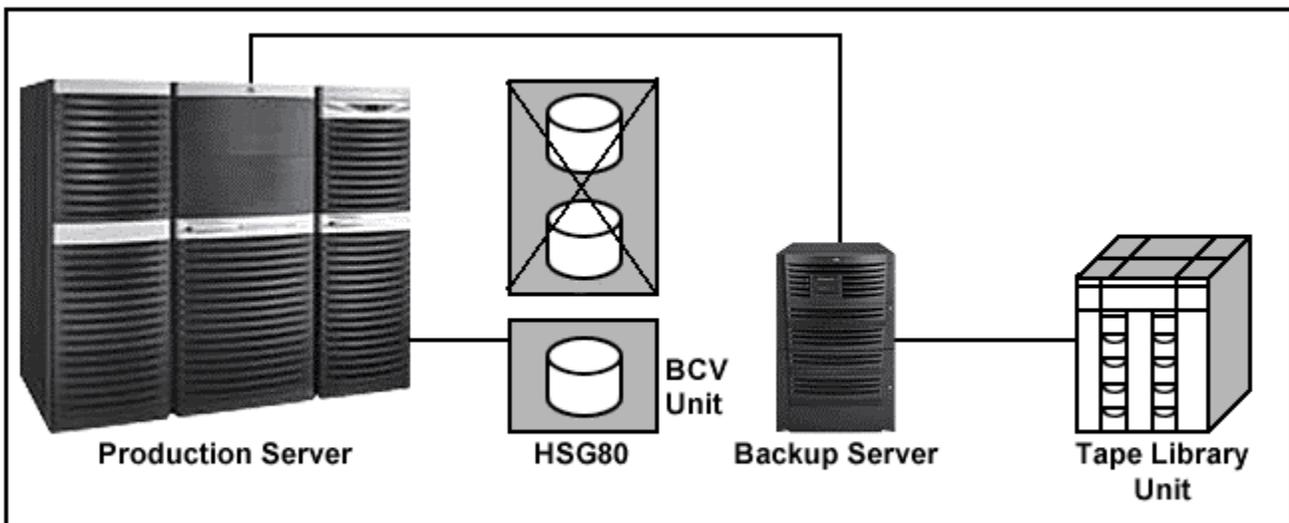


Figure 5, Rapid Restore from BCV unit

applying redo logs

Complete recovery is recovering the database up to the point in time of failure. The database must be running in archive log mode and you must have a complete set of archived redo logs. You can perform complete recovery on the datafile, tablespace or database level.

If tablespace recovery is required, then take the specific tablespace offline. Restore the necessary files from backup.

Applying the required redo logs to recover the datafiles up to the point of failure.

Time-based recovery recovers the database to a given time in the past.

Change-based recovery recovers the database to a particular SCN (system change number).

Cancel-based recovery recovers the database up through a given archived redo log. Once the final redo log is applied, the administrator enters the cancel command to halt recovery. For incomplete database recovery, use the following procedure.

```
SvrMgrl> shutdown immediate;
```

Restore control file, all datafiles, and all archive logs from backup.

```
SvrMgrl> startup mount; #(do not open database)
```

```
SvrMgrl> Alter database XXX recover database until...;
```

```
SvrMgrl> Alter database open Resetlogs;
```

Alter database open resetlogs must be used to open the database after incomplete recovery. This resets the log sequence number to one and synchronizes the database.

storage layout recommendation

When implementing non-disruptive backup solutions based on HP StorageWorks MA/EMA storage, the storage configuration requires some attention. In general, a correct configuration will use all available storage hardware resources (SCSI buses, controllers, etc.) in a balanced way. The following sections discuss some general storage layout recommendations for non-disruptive backup solutions based on MA/EMA storage.

Mirrorsets (RAID 1)

When configuring mirrorsets, one should generally adhere to the following rules:

- Each member of a mirrorset should use a separate SCSI bus to ensure good distribution of I/O
- Mirrorset members should be seated in different disk shelves for data protection reasons
- Mirrorset members should be of the same disk type

In a MA/EMA configuration using single bus disk enclosures (4214 or 4314) (Figure 6), each disk shelf is attached to a separate SCSI channel and arranged in a configuration using dual bus disk enclosures (4254 or 4354) (Figure 7). Each half of a disk shelf is attached to an SCSI channel. This means that some extra attention is needed for dual bus configurations to make sure that mirrorset members are not seated in the same disk shelf. This also means that, when the above rules are applied, a dual bus configuration has a maximum of three members in a mirrorset, whereas a mirrorset in a single bus configuration could have up to six members. Non-disruptive backup solutions based on MA/EMA storage subsystems typically require three member mirrorsets.

The example in Figure 6 shows a single bus configuration with a source mirrorset on target 0 with two members on bus 5 and 6 (DISK50000, DISK60000). When target 15 is reserved for BCV disks, there are four possibilities for such a mirrorset to allocate its third member. The example also shows how disks with the highest target numbers can be reserved as BCV disks for source mirrorsets on the lowest target numbers. Although not strictly required, it makes BCV disks more easily distinguishable from their source mirrorsets, while still being able to use all available SCSI buses and disk shelves efficiently.

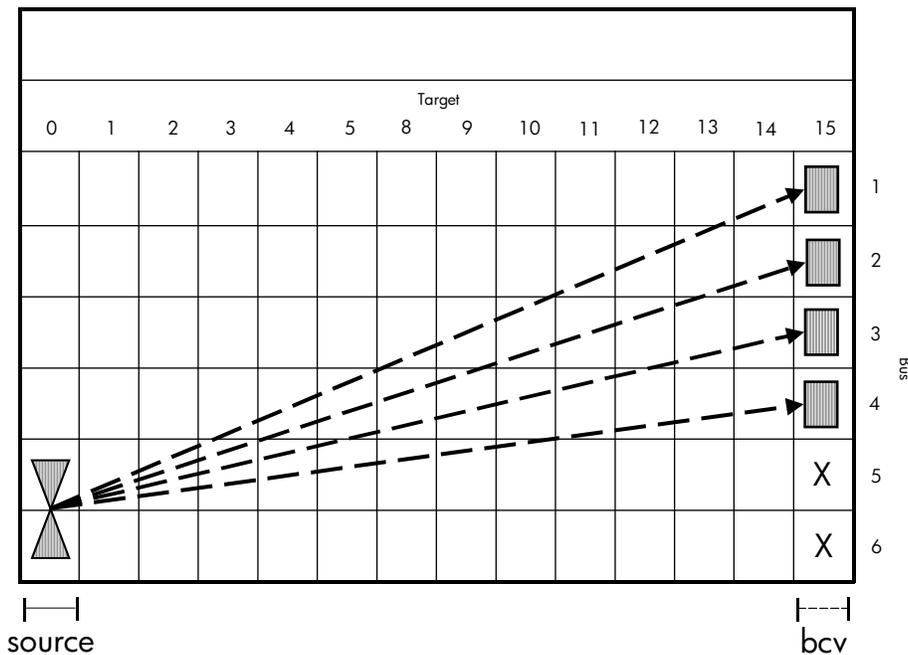


Figure 6: Four possibilities for a mirrorset to allocate a third member on target 15 (single bus configuration)

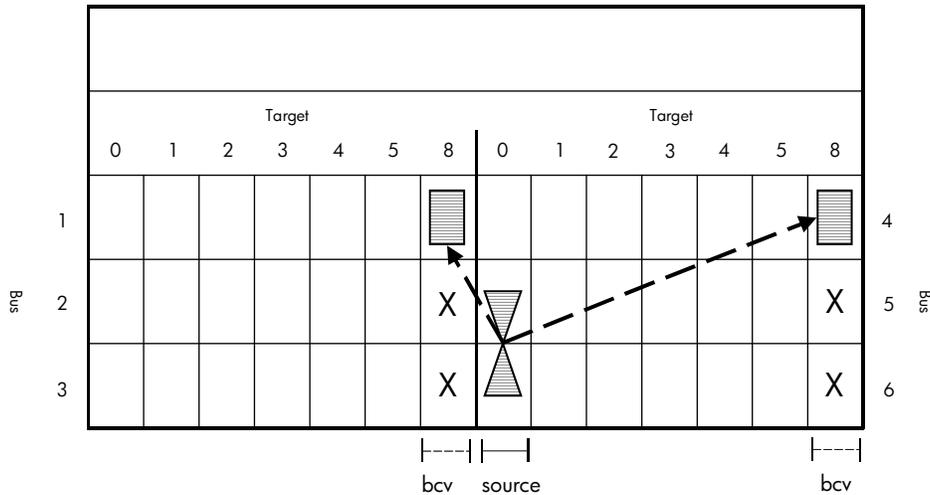


Figure 7: Two possibilities for a mirrorset to allocate a third member on target 8 (dual bus configuration)

In a dual bus configuration, one has to keep in mind that each disk shelf holds two SCSI buses. This means that two more SCSI buses need to be excluded for third member allocation due to data protection reasons. Figure 7 shows the source mirrorset from the previous example, but this time in a dual bus configuration. The source mirrorset has its members on target 0 and buses 5 and 6 (DISK50000, DISK60000). If disks with highest target number 8 are reserved as BCV disks for source mirrorsets on target 0, the figure shows two possibilities for this mirrorset to allocate its third member. Buses 2 and 3 are excluded in addition to buses 5 and 6.

Striped mirrorsets (RAID 1+0)

Multiple mirrorsets can be combined into stripesets to balance I/O across all available SCSI buses and create storage sets with higher storage capacity. When defining a stripeset of mirrorsets, one should bear in mind that the stripeset should actually engage all SCSI buses. As a correctly defined two-member mirrorset engages two SCSI buses, we recommend that a stripeset contain at least 3 two-member mirrorsets.

A second thing to bear in mind is the first and second write-behavior of mirrorsets. When a write request is issued to a mirrorset, the lowest member of the mirrorset receives the write request first. After that, the write is copied to the second member of the mirrorset. As the ultimate objective should be to balance all simultaneous I/O over all SCSI buses, it is good practice to have an equal number of even and odd mirrorsets in a stripeset and to distribute these over all available SCSI buses. The source stripeset shown in Figure 8 contains three even mirrorsets on target 0 and three odd mirrorsets on target 1, thus balancing both general I/O and first mirrorset writes over all available SCSI buses. Figure 8 also shows how the third members for the BCV in such a source stripeset can be distributed over all the available SCSI buses. This ensures that the third member mirrorset synchronization engages all SCSI buses. Furthermore, the read I/O on the BCV stripeset caused by the backup process will also be evenly distributed over all SCSI buses. In the example, all disks with target 15 are reserved as the BCV area for the stripeset containing all disks with target 0 and 1. Reserving disks with equal target numbers makes sizing easier and source disks clearly distinguishable from their BCV counterparts.

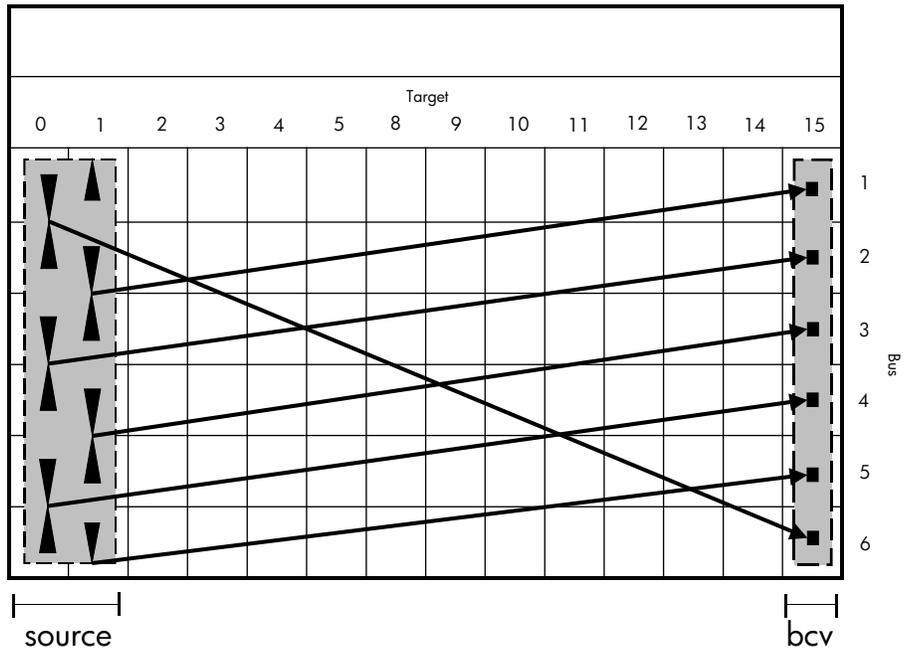


Figure 8: Third member allocation on target 15 of a striped mirrorset (single bus configuration)

Alternatively, stripesets can also be defined with the mentioned minimum of three mirrorsets, using all SCSI buses. As first mirrorset writes will not be balanced within such a stripeset itself, we recommend that you have an equal number of these stripesets containing even and odd mirrorsets. Then, you should equally balance odd and even stripesets for those stripesets that are written to simultaneously. Figure 9 shows a dual bus configuration with two stripesets. If those stripesets are written to simultaneously, an equal distribution of first mirrorset writes will also be achieved. As shown in the figure, disks with target number 8 are reserved for third member allocation of the source stripesets.

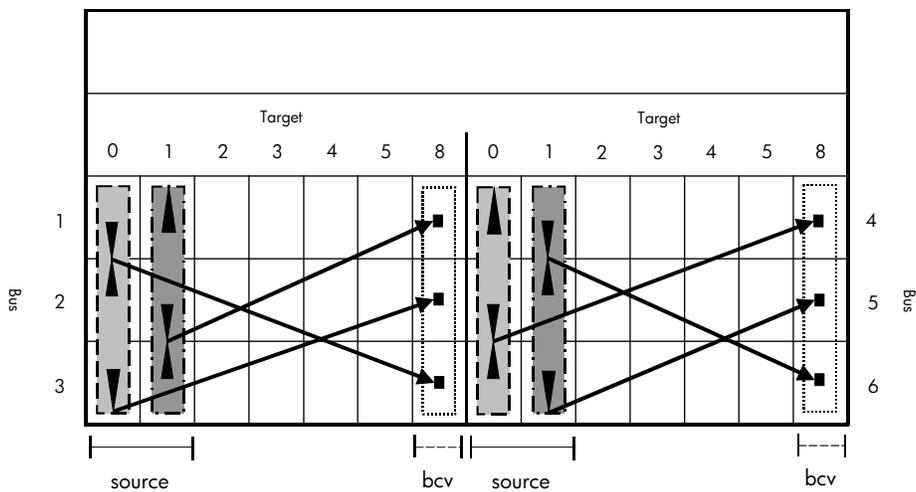


Figure 9: Third member allocation on target 15 of a striped mirrorset (dual bus configuration)

detailed description, how to set up

Please refer to OSS-Note 511972 Split Mirror Backup with HP Storage for latest updates.

system setup and installation

This solution guide assumes that you have mySAP Business Suite/Oracle installed and functioning well on the production system in a Tru64 Unix environment.

To execute Rapid Restore and Non-Disruptive Backup Solutions, the following configurations need to be made:

- Install Tru64 Unix with Advfs on the Alpha-Backup system, same configuration as production system in respect of version and patches (Refer Alpha Tru64 UNIX Installation Guide, for complete installation of Operating System.)
- On both the systems NFS (Network File System) has to be configured. (See Alpha Tru64 UNIX Installation Guide.)
- Install HSG tool kit (HSTK) on backup system. (See <http://storage.igo.cpqcorp.net/perl-scripts/V1.60/>.)
- Install CPQcmdsript utility on backup system and the file Tru64V5_cpqcmdschr_V10A.tar from www.hp.com/go/storage.
- Install hp StorageWorks Command Scripser version 1.0B (see www.hp.com/go/storage for download and installation instructions.)
- This solution is intended for configuration of HSG80 storage in RAID 1+0 with 3-member mirrorset.

limitations

This solution does not provide support for

- One or two member mirrorset, Raid 0 or Raid 5 configurations.
- Earlier versions of SAP R/3 4.6C and Tru64 Unix 5.0

configuring the production system

Step 1

The `/etc/exports` files from the production system should be configured as in the following example in order to share mount points from the production system to the backup system.

Example: EVM > SID and evmbck64 is the backup system

```
/oracle/EVM/saparch -root=0 -rw = evmbck64
/oracle/EVM/saptrace -root=0 -rw = evmbck64
/oracle/EVM/sapcheck -root=0 -rw = evmbck64
/oracle/EVM/sapreorg -root=0 -rw = evmbck64
/oracle/EVM/sapbackup -root=0 -rw = evmbck64 -rw=evmbck64 ,access=evmbck64
share -F nfs -o rw=evmbck64 -d "allow evmbck64 to access" /oracle/EVM/sapbackup
share -F nfs -o rw=evmbck64 -d "allow evmbck64 to access" /oracle/EVM/sapreorg
share -F nfs -o rw=evmbck64 -d "allow evmbck64 to access" /oracle/EVM/saptrace
share -F nfs -o rw=evmbck64 -d "allow evmbck64 to access" /oracle/EVM/saparch
share -F nfs -o rw=evmbck64 -d "allow evmbck64 to access" /oracle/EVM/sapcheck
```

Step 2

In the file `.rhosts` from the Oracle/SID directory, the backup system host name and user name should be present.

e.g

```
# cat .rhosts
EVMBCK64 root
EVMBCK64 oraevm
```

Step 3

Create an entry for the backup system in `/etc/hosts.equiv` on the database server, so that it is allowed to execute rsh commands.

Step 4

Create an Oracle password file in the database server. If your password file already exists, you can maintain the same internal (e.g. manager) password. The password should be identical for production system and backup system, in order to establish a remote connection from the backup system.

```
orapwd file=%ORACLE_HOME%/dbs/orapw password=manager entries=100
```

e.g

```
orapwd file=/oracle/EVM/817_64/dbs/orapw password=manager entries=100
```

Include the entry `remote_login_passwordfile = exclusive` in the file `initSID.ora` in the database server.

The Oracle user system must have SYSOPER rights

```
SVRMGR>grant SYSOPER to SYSTEM;
```

```
SVRMGR>grant SYSDBA to SYSTEM;
```

Step 5

On the production system issue the following command to copy the control files to sapbackup directory

```
svrmgrl>alter database backup controlfile to '/oracle/SID/sapbackup/cntrlSID.dbf';
```

configuring the backup system:

In order to accomplish successful Rapid Restore and Non-Disruptive Backup, using HSTK for SAP R/3 Brbackup tools, you need to perform the steps described in this chapter. appendix a: examples of lab setup on page 36 contains detailed example information for the lab installation.

Step 1

Create user accounts and groups on the backup system similar to those on the production system:

1. oraevm with the primary group dba.
2. evmadm with the Unix group sapsys.

e.g in our test oraevm, evmadm where evm is SID

Step 2

Copy the following files from the production system to the backup system in the respective directory and make the appropriate changes (such as hostname etc.). This is necessary in order to create an Oracle environment in the backup system that is similar to the production system. The environment variable files

```
.profile  
.Login  
.cshrc  
.dbenv_<hostname>.ksh
```

Step 3

The ORACLE_HOME directory structure must match the SAP standard installation and should be identical to that of the production system. Brbackup expects this in the backup system. The ORACLE_HOME directory is /oracle/< SID>, where SID is the three-character system Identifier for the R/3 database instance. The following directory structure should be mounted on the backup system using NFS software.

Example: The entry NFS mount point in FSTAB in the backup system

```
Evmsap64:/oracle/EVM/sapbackup /oracle/EVM/sapbackup nfs rw 1 0
```

Default directories:

Directory structure	Example
#ORACLE_HOME/sapbackup	/oracle/EVM/sapbackup
#ORACLE_HOME/saparch	/oracle/EVM/saparch
#ORACLE_HOME/sapreorg	/oracle/EVM/sapreorg
#ORACLE_HOME/sapcheck	/oracle/EVM/sapcheck
#ORACLE_HOME/saptrace	/oracle/EVM/saptrace

Step 4

The SAP tools Brbackup, Brarchive, Sapdba, and Brrestore evaluate the environment variable from the backup system. To create the environment available on the backup system please create the following symbolic links. In the following example evmsap64 is the production system in our lab test.

```
#mkdir /oracle/SID/sapbckup
#ln -fs evmsap64:/oracle/SID/sapbackup /oracle/SID/sapbckup
#ln -fs evmsap64:/oracle/SID/sapreorg /oracle/SID/sapreorg
#mount evmsap64:/oracle/SID/sapbackup /oracle/SID/sapbckup
```

Step 5

Installing Oracle with sqlnet on the backup system.

1. Please refer to the OS-Dependent Installation Unix Kernel parameter from the SAP Basis Installation on Unix: Oracle guide and ensure that the appropriate kernel parameters are changed in your backup system.
2. You need to create the following directory on the backup system.

```
/oracle/stage/817_64/Disk1/SAP
/oracle/stage/817_64/Disk2
```

3. Mount the Oracle RDBMS_1 CD, which comes from the SAP installation kit.

```
#mount -r /dev/disk/cdrom0c /cdrom
#cd /oracle/stage/817_64
#SAPCAR -xvf /cdrom/UNIX/DEC_64/OR181764.SAR
```

4. Mount the Oracle RDBMS_2 CD, which comes from the SAP installation kit.

```
#cd /oracle/stage/817_64
#SAPCAR -xvf /cdrom/UNIX/DEC_64/OR281764.SAR
```

5. Mount the Oracle RDBMS_1 CD

```
cd /oracle/stage/817_64/Disk1/SAP
./RUNINSTALLER
```

Enter the proper path name for installing directory as

```
/oracle/EVM/817_64 and the group is dba.
```

Ensure that you have successfully installed Oracle on the backup system. To avoid errors, ensure that you are using the proper x'term, enough space, the proper default Oracle environment and refer to SAP OSS notes in case of other errors.

Step 6

The Oracle net files need to be copied from the production system to the backup system in the respective directory.

```
/oracle/SID/817_64/network/admin/tnsnames.ora
/oracle/SID/817_64/network/admin/sqlnet.ora
```

Create directories for sapdata (sapdata1, sapdata2, etc) on the backup system as in the production system for mounting BCV units. Please refer to appendix a: examples of lab setup, disk layout.

Step 7

You need to copy the latest version of all brtools files from the production system (/usr/sap/SID/sys/exe/run/) to the backup system in the respective directory /usr/sap/SID/SYS/exe/run. This is required for execution of Brbackup and Brrestore from the backup system.

Step 8

Copy InitSID.sap, InitSID.ora from the production system to the backup system in /Oracle/SID/817_64/dbs.

The parameter backup_type in the profile init<SID>.sap must be set as follows:

```
backup_type = offline_split or online_split
```

The online_split option specifies an online backup of the mirror disks. The offline_split option specifies an offline backup of the mirror disks. If the Brbackup command is invoked from the command line with one of the above options, it overwrites the value set in the parameter backup_type in the init<SID>.sap profile.

```
# backup type [offline | offline_force | offline_standby | offline_split
# | offline_stop | online | online_cons | online_split]
# default: offline
backup_type = online_split
```

Step 9

The parameter `split_cmd` in the profile `init<SID>.Sap` must be set as follows:

`<split_cmd>` is a shell script called by Brbackup to split the mirror disks.

```
# split mirror disks command
# used if backup_type = offline_split | online_split
# caution: if successful, exit code 0 and no output!
# no default
# split_cmd = "<split_cmd> [$]"
split_cmd = "/oracle/EVM/backup_scripts/do_split.ksh"
```

Note: `do_split.ksh` shell script contains the execution of the HSTK job, which splits the 3rd member of the mirror, constructs the BCV and mounts on the backup system.

Step 10

The parameter `resync_cmd` in the profile `init<SID>.Sap` must be set as follows:

`<resync_cmd>` is a shell script called either by Brbackup or externally to resynchronize the mirror disks after the backup using the `<split_cmd>` script. If this parameter is not set, then the resynchronization process is not carried out under the control of Brbackup. The script `resync.ksh` executes HSTK. The HSTK mounts the BCV units, deletes the BCV units and resynchronization takes place with the original mirrorset.

```
# resync_cmd = "<resync_cmd> [$]"
resync_cmd = "/oracle/EVM/backup_scripts/do_resync.ksh"
```

Note: We would recommend keeping the BCV units until the next Brbackup execution, so that the data can be restored immediately from the BCV unit in the event of data recovery. In this case you can use SAP batch scheduling to execute the procedure `do_resync.ksh` for resynchronization of the BCV unit with the original mirrorset. This job can be scheduled a few hours before Brbackup execution. The script `do_resync.ksh` is a shell script which unmounts the mountpoints from the Backup Server and deletes the BCV unit and brings the deleted 3rd member to the mirrorset.

Step 11

To establish the link between the backup system and the production system, an instance string (SID) must be defined in the file `initSID.sap`. Set the profile parameter `primary_db` to the SID of the production server (e.g. `primary_db=EVM`). The instance string to the production database is needed in order to connect the backup system with the production host. While configuring Oracle SQLNET, the connections (alias names) are defined in the configuration file `tnsnames.ora` by the user.

```
# Oracle instance string to primary database
# [primary_db = <inst_str> | LOCAL]
# no default
# primary_db = <inst_str>
primary_db = EVM
```

Step 12

The path names and directory structure for mount points on the backup system should be the same as on the production system. The directory `ORACLE_HOME/sapbackup` has to be shared by the two nodes using NFS mount.

Step 13

The ownership of all the files must match (the same administrative users with the same user and group IDs) on the production server node and on the backup node. The Oracle instance has to be accessed from the remote node with SYSOPER authorization. The authorization SYSOPER has to be granted to the system user in the backup system. An Oracle password file `$ORACLE_HOME/dbs/orapw<SID>` has to be present containing the valid database password of the system user and the parameter `remote_login_passwordfile = exclusive` has to be set in the parameter file `init<SID>.ora`.

Example:

```
orapwfile=/oracle/EVM/817 64/dbs/orapw password=manager entries=100
```

Check the connection to the production system database server using `tnsping SID`, as specified in the file

```
/oracle/SID/817/database/initSID.ora
```

and

```
initSID.sap.
```

Test the connection to the database server

```
# svrmgr1
SVRMGR> connect system@SID
Password: *****
Connected.
SVRMGR>
```

Step 14

The 3rd party backup tool has to be installed on the backup server and configured to use the Backint interface. Please refer to the example of InitEVM.utl in appendix a: examples of lab setup.

Note: If `split_cmd/resync_cmd` is completed successfully, an exit code of 0 is returned. If the command is not successful, a return code of >0 is returned, as well as a message describing the cause of the error.

It is good practice to keep a copy of `init/SID/817_64/dbs/...` in a different node, in case recovery is required.

non-disruptive backup for mySAP Business Suite solutions with different configuration

The SAP backup/restore functions are among the most important operations in the production environment. Using the HSTK procedure `split/mirror` functions enable customers to easily back up their production database without affecting database availability. This solution allows customers to perform complete and fully automated backup of SAP data in a production environment. Due to the fact that SAP is a write-intensive system, we strongly recommend configuring RAID 1+0 in HSG80 subsystem.

small configuration

The first example is a standard SAP/Oracle implementation, using ACS 8.6P and ma8000. If the database size is 216 GB, 3 striped mirrorsets are needed: 3 x 18 disks x 18.2 GB. This example shows a small configuration with one HSG80 Array dual-redundant controller with the minimum of 512 MB cache installed. Such a small configuration can have up to 216 GB of SAP/Oracle tablespace.

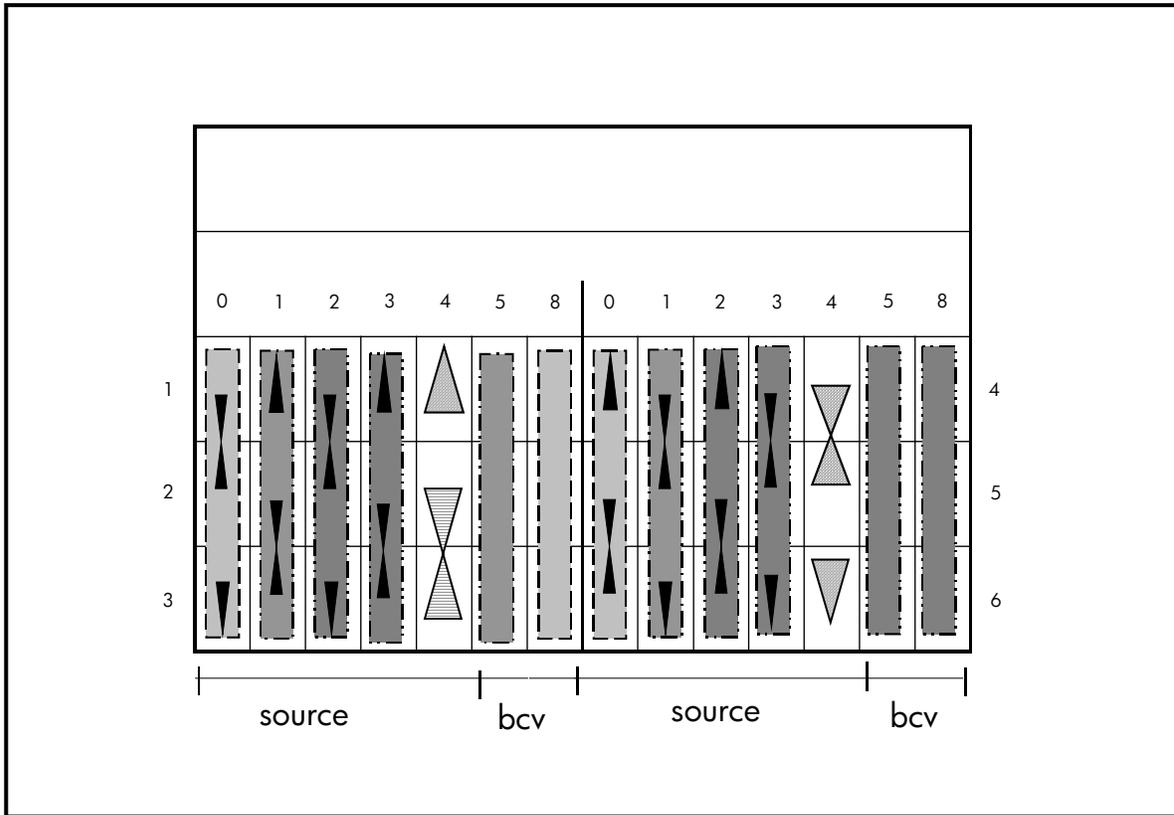


Figure 10: Small configuration on ma8000

Table 1: Small configuration on ma8000, legend

Description	Raid level	Disks	BCV?	Legend
Sapdata 1	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 2	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 3	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 4	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
OriglogA/B	R1	2 x 18.2 GB (18 GB)	Yes	
MirrlogA/B	R1	2 x 18.2 GB (18 GB)	No	
Saparch	R0	2 x 36.4 GB (36 GB)	No	

medium configuration

The second example is a standard SAP/Oracle implementation implemented using two ma8000. If the database size is 432 GB, you need 36 disks x 18.2 GB for each striped (R 0+1) three-member mirrorset, including BCV units. As shown in Table 2, you can configure SAPDATA1, SAPDATA3, SAPDATA5 and SAPDATA7. As shown in Table 3, you can configure SAPDATA2, SAPDATA4, SAPDATA6 and SAPDATA8. For each striped (R 0+1) three-member mirrorset, you need 9 disks X 18.2 GB. The example shows a medium configuration with one HSG80 array controller with the minimum of 512 MB cache installed. A medium configuration like this can have up to 432GB of SAP/Oracle tablespace.

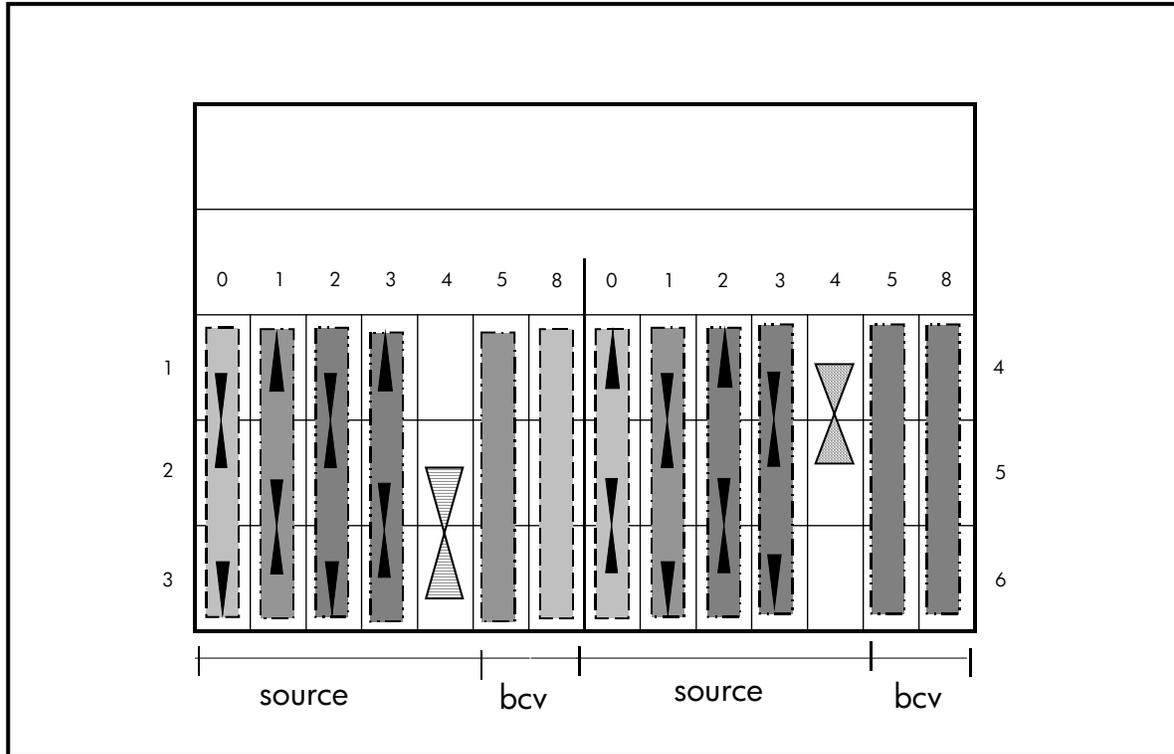


Figure 11: Medium configuration (subsystem 1 of 2, ma8000)

Table 2: Medium configuration, legend (subsystem 1 of 2, ma 8000)

Description	Raid level	Disks	BCV?	Legend
Sapdata 1	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 3	R1+0	3x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 5	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 7	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
OriglogA/B	R1	2 x 18.2 GB (18 GB)	Yes	
SAP / ORACLE Binaries	R1	2 x 18.2 GB (18. GB)	No	

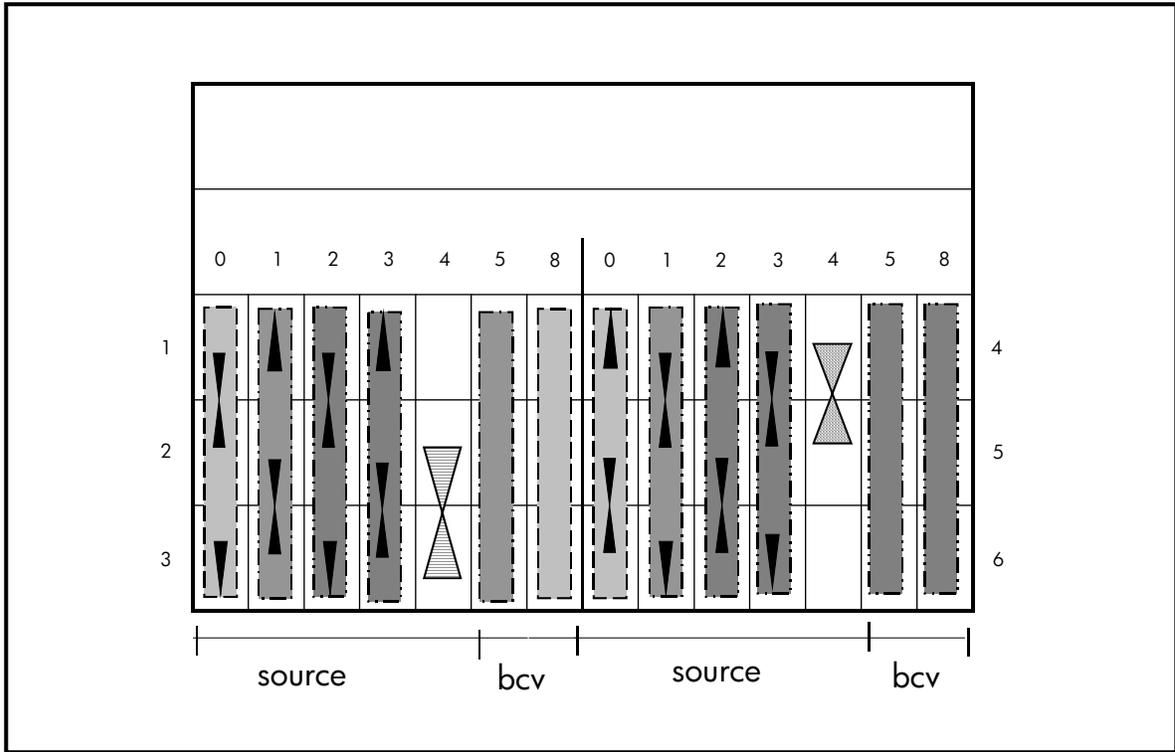


Figure 12: Medium configuration (subsystem 2 of 2, ma8000)

Table 3: Medium configuration, legend (subsystem 2 of 2, ma8000)

Description	Raid level	Disks	BCV?	Legend
Sapdata 2	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 4	R1+0	3x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 6	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
Sapdata 8	R1+0	3 x 2 x 18.2 GB (54 GB)	Yes	
MirrlogA/B	R1	2 x 18.2 GB (18 GB)	No	
Saparch	R0	2 x 36.4 GB (36 GB)	No	

large configuration (1.3 Terabyte SAP/Oracle DB)

The third example is a standard SAP/Oracle implementation, implemented using two ema12000. For a 1300 MB database, you need 18 disks per unit to create a three-member striped mirrorset, so you need 72 disks x 36 GB. Also, you need 18 disks per unit to create a three-member striped mirrorset, so you need 72 disks x 18.2 GB. This example shows a large configuration with two HSG80 array controllers with the minimum of 512 MB cache installed. Such a large configuration can have up to 1300 MB of SAP/Oracle tablespace.

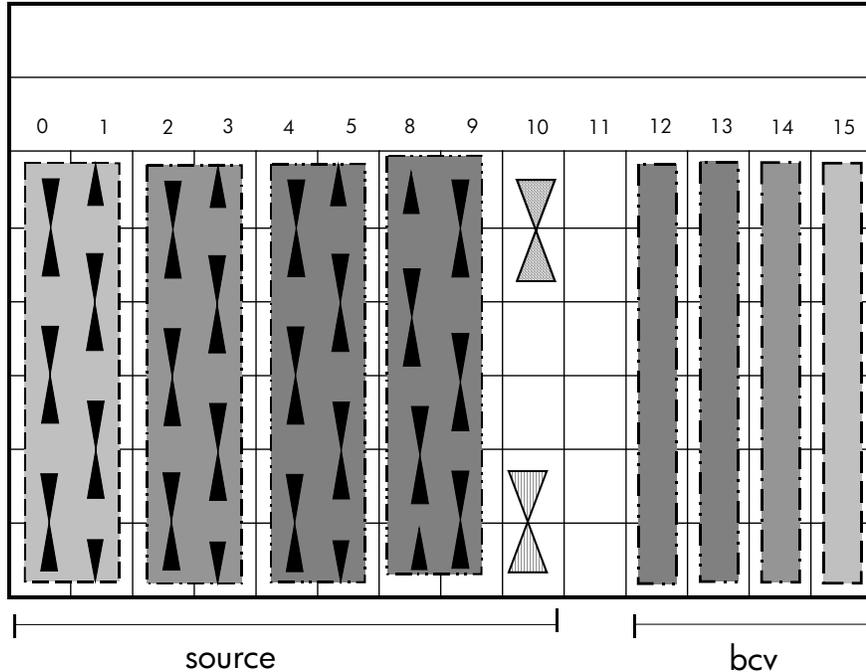


Figure 13: Large configuration (subsystem 1 of 2, ema12000)

Table 4: Large configuration, legend (subsystem 1 of 2, ema12000)

Description	Raid level	Disks	BCV?	Legend
Sapdata 1	R1+0	6 x 2 x 36.4 GB (216 GB)	Yes	
Sapdata 3	R1+0	6 x 2 x 18.2 GB (108 GB)	Yes	
Sapdata 5	R1+0	6 x 2 x 36.4 GB (216 GB)	Yes	
Sapdata 7	R1+0	6 x 2 x 18.2 GB (108 GB)	Yes	
OriglogA/B	R1	2 x 18.2 GB (18 GB)	Yes	
SAP / Oracle Binaries	R1	2 x 36.1 GB (36 GB R1)	No	

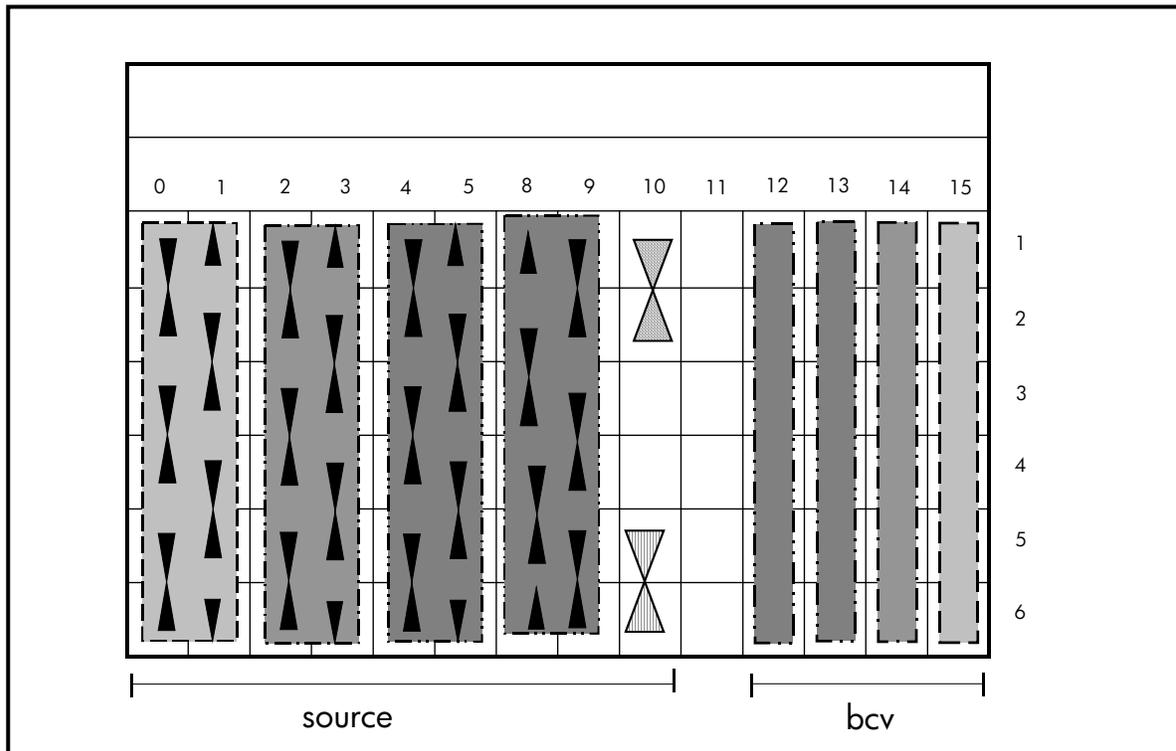


Figure 14: Large configuration (subsystem 2 of 2, ema12000)

Table 5: Large configuration, legend (subsystem 2 of 2, ema12000)

Description	Raid level	Disks	BCV?	Legend
Sapdata 2	R1+0	6 x 2 x 36.4 GB (216 GB)	Yes	
Sapdata 4	R1+0	6 x 2 x 18.2 GB (108 GB)	Yes	
Sapdata 6	R1+0	6 x 2 x 36.4 GB (216 GB)	Yes	
Sapdata 8	R1+0	6 x 2 x 18.2 GB (108 GB)	Yes	
MirrorlogA/B	R1	2 x 18.2 GB (18 GB)	Yes	
SAP Archive	R0	2 x 36.1 GB (36 GB R1)	No	

description of lab setup

testing environment

The following section describes the SAP Oracle Tru64 Unix backup/restore solutions and includes the hardware, software, and firmware components.

bill of materials

Table 6 provides a bill of materials that includes some of the products used in the SAP/Oracle non-disruptive backup/restore solutions.

Item	Item Description
1	HSG Tool Kit (HSTK)
2	ACS 8.6P Controller Software
3	Operating System Platform Software kit for Tru64 Unix 5.1A
4	EMA 12000 with 2 Chassis Four HSG80 Controllers
5	Hard Disk Drives, 18.2 GB, 10K UWSE
6	Two Alpha Server DS20(Production system and Backup system)
7	Fibre Channel HBA (switched fabric), KGPSA-CA
8	Fibre Channel Storage 2 SAN Switches, 16 port
9	FC Optical Short Wave GBIC
10	Fibre Channel Optical Cable
11	Modular Data Router
12	Super DLT 110/220 Tape Library
13	VERITAS® NetBackup/Computer Associates ARCserve 2000/Legato Networker

Table 6, Bill of Materials for SAP/Oracle Rapid Restore-Backup solution

table explanation

1. HSTK software: HSTK software is needed on the backup system. It is needed to communicate or coordinate communication with hosts and storage subsystems for splitting and resynchronization of the mirrors.
2. Controller firmware – One instance of ACS is needed for each HSG80 controller. ACS firmware is in the form of a PCMCIA card inserted into the front panel of the controller. ACS 8.6P also supports cloning and snapshots.
3. OS Platform Kit CDROM – StorageWorks Solutions Software ra8000/ma8000 Fibre Channel Solutions for Tru64 Unix installation files for the specific operating system and host involved.
4. Disk array – For our tests we used two ema12000/4214 and one ma8000/4254. The subsystems were configured with Multi-Bus failover.
5. Disks – The choice of disk type does not affect the functionality of these solutions. Two factors that need to be considered for disk type are rpm and the number of spindles. For applications such as SAP/Oracle, the access pattern is random. Therefore, it is desirable to use as many spindles as possible to increase parallel access to the volume as opposed to fewer disks. In our test we created six mirrorsets, each with three members, and then striped them together to form one logical unit.
6. Backup system – In our test we used one Alpha DS20 server as the production system and another Alpha DS20 server as the backup system. The backup system must be powerful enough to maintain performance levels during a backup/restore operation
7. Fibre Channel Host Bus Adapter (HBA) – The Tru64 Unix Backup system supports the following Fibre Channel (FC) host bus adapters (HBA): KGPSA-CA. Please refer to the Host Bus adapter driver information on the storage website: <http://www.hp.com/go/storage>.

8. Fibre Channel Interconnect – SAN interconnect switches allow for significantly better performance, whereas loops (hubs) present significantly limited performance and connectivity options. Storage arrays and tape can reside on the same-switched fabric, making a simpler, more extensible installation possible. Disk and tape cannot exist on the same loop; dual loops are required in this case.
9. FC Optical Short Wave GBIC – Gigabit Interconnect (GBIC) – GBICs are not available with the hubs or switches and must be ordered separately.
10. Fibre Channel Optical Cable – Order the length needed. Each controller comes with one five-meter cable. All other cables will have to be ordered separately.
11. Modular Data Router – HP StorageWorks modular data router (Fibre Channel Tape Controller). SCSI attached tape libraries require special Host Bus Adapters (do not use standard server HBA adapters). Fibre Channel-attached tape libraries (being SCSI devices) require a special Fibre Channel to SCSI bridge (FCTC). Requires VHDCI cables for SCSI connection. #341176.B21 for 6ft cable or #341177-B21 for 12ft cable.
12. MSL5026 Super DLT 110/220 Library –SDLT 110/220, 35/70 GB DLT tape drives.
13. Third-party backup software – Computer Associates ARCserve 2000 or VERITAS NetBackup or Legato Networker.

Server	Model	RAM
Database Server	DS20	4096 MB
Backup system	DS20	2046 MB
Application Server	Proliant 5500	1024 MB

Table 7, Hardware Configuration Overview

Tru64 Unix version 5.1A	Oracle RDBMS version 8.1.7.2.3
Advfs version 4	SAP R/3 version 4.6c
C compiler	Legato Networker 6.1
NFS	NLS

Table 8, Software configuration overview

	StorageWorks 18.2 GB 10,000 RPM Ultra3 SCSI (internal hard drives)
	StorageWorks 36.4 GB 10,000 RPM Ultra3 SCSI (internal hard drives)
StorageWorks KGPSA-BC, Fibre Channel Host Bus Adapter	StorageWorks SAN Switch 16 x 2
(2) StorageWorks HSG80 array controller ACS v.8.6P	(2) StorageWorks RAID array ma8000 Compaq DS-RZ1DD-VW 18.2 GB 10,000 RPM UltraSCSI 16-bit
StorageWorks modular data router (Fibre Channel Tape Controller)	Fibre Channel to SCSI tape controller with (2) 2-port SCSI modules
StorageWorks msl 5026 Tape Library	StorageWorks Super DLT 110/220 tape drives

Table 9, StorageWorks components

MSL5026 tape library notes

MSL5026 tape library is configured and connected to a modular data router (MDR) using an SCSI connection. The MDR is connected to an SAN switch via Fibre. Use of tape libraries in an SAN environment enables you to select the location of the tape library, allowing you to place it in location different to that of the storage to be backed up. It also provides the option to move application data directly from the HSG80 storage system to the tape storage system without using the traditional communication networks. Having the msl5026 connected to an MDR enables data to travel over the high-speed SAN as opposed to a traditional LAN.

SAN setup and configuration

In our test, we selected SAN Fabric topology to achieve the best performance. Fabric switches provide full-speed data transfer between all ports on the SAN switch interface. This is recommended because SANs provide a host of benefits, including better resource utilization, improved performance, better security, centralized management and high availability. As shown in Figure 15, the SAN Fabric solutions are scalable to allow for easy addition of switches with minimal changes to fabric, which permits you to set up a much larger enterprise and company-wide SAN. Each port on the switch provides 100 MB/second bandwidth, so performance is not affected by the addition of devices. SANs provide superior interoperability as well as greater reliability and robustness.

Host Bus Adapter configuration performance considerations

As shown in Figure 15, the configuration with no single point of failure (NSPOF) utilizes two storage subsystems and each subsystem has dual HSG80 controllers with four active controller host ports. The primary advantage of this configuration is the ability to provide redundant independent paths to the storage. This is a tested configuration that can be utilized for intensive I/O applications and very intensive bandwidth applications.

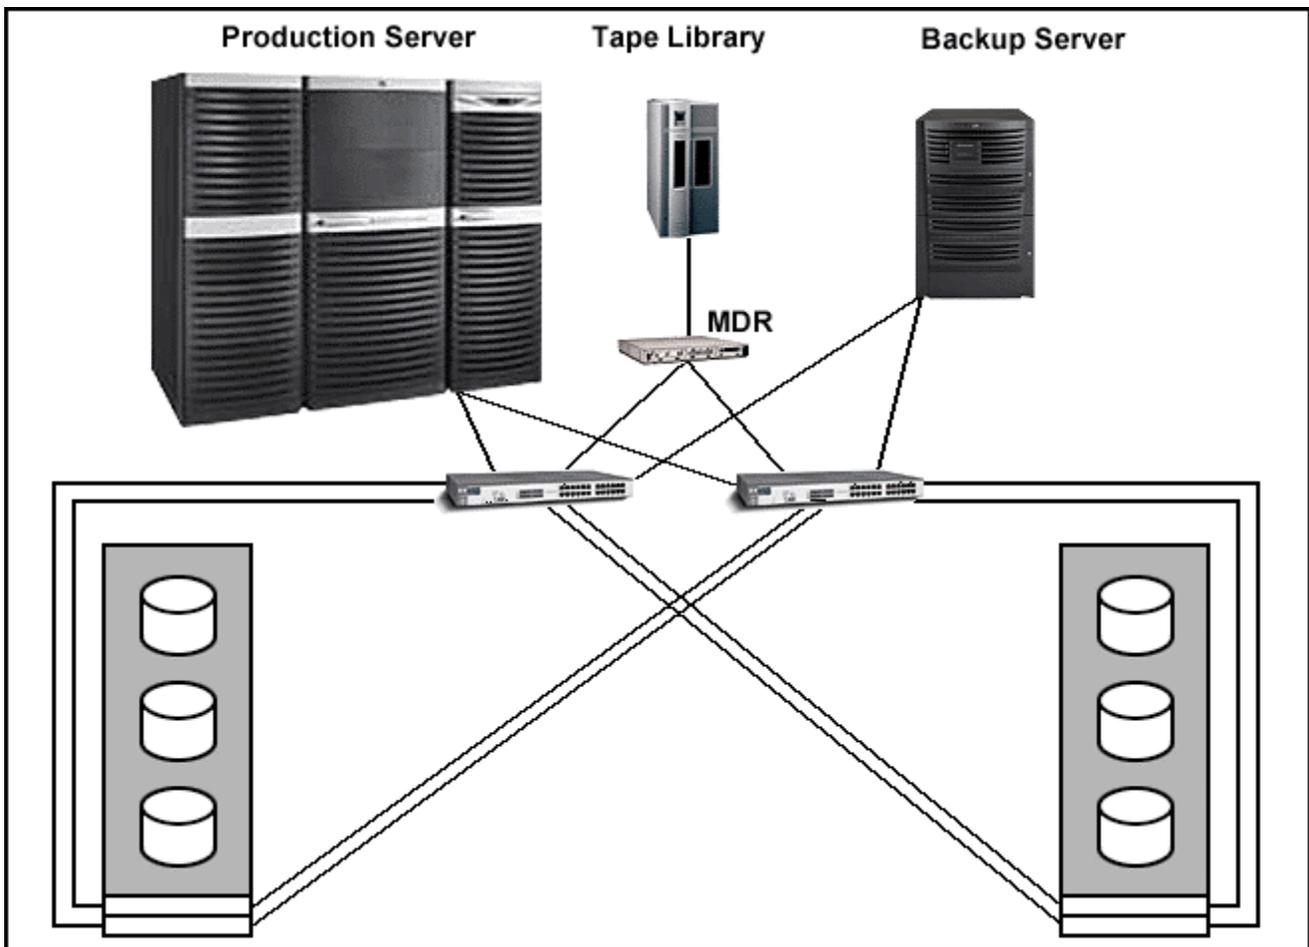


Figure 15, Fibre channel configuration

storage subsystem configuration notes

Fibre Channel Loops

Three subsystems have dual HSG80 controllers configured in multiple-bus failover mode

Created 13 LUNs (D1-D13) with 2 active controller host ports.

The Oracle log files and the Oracle archive files have to be on different disks.

The archive files are on single disk.

MA8000/EMA12000 installation

The HSG80 software (ACS 8.6-1 P) is shipped separately. The HSG80 will not function without the ACS software. To enhance availability, it is possible to install redundant HSG80 Array Controllers, where each of them is able to take over if the other controller fails. Each array controller has a separate cache module with its own external battery. Optionally, an additional external battery can be installed per cache module for redundancy purposes. The two cache modules can be configured to mirror each other, enabling one cache module to fail without bringing the storage subsystem to a halt.

Note: Refer to the Compaq StorageWorks HSG80 Array Controller ACS Version 8.6, Configuration and CLI Reference Guide for a complete description of the HSG80 configuration and multiuse failover configuration procedures.

rapid restore and recovery procedures

Note: Please refer to SAP notes 4157, 491160 and 96848 for logical and hardware errors in addition to the following explanations.

Case 1: Full restore and recovery from BCV units to the production system:

The following procedure only applies in the event that the entire SAP /Oracle database needs to be restored. For example, if D1 has been damaged due to logical error or hardware problem, we can use BCV LUN D51 to return to the original status of LUN D1. Ensure that the database has been shutdown completely on the production system before executing the procedure.

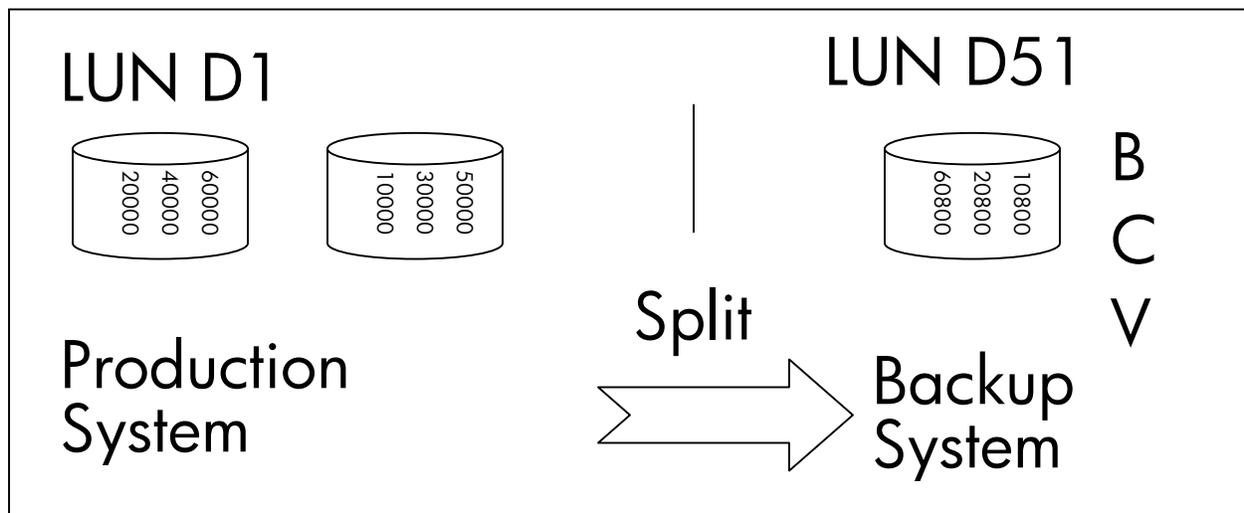


Figure 16, After split of BCV mounted on backup system

Note: The following procedure is an example pertaining to one LUN (in our case the single LUN for SAPDATA1). You need to repeat this procedure for other split LUNs to return the original status of the production system.

This example assumes LUN D1 consists of striped mirrorset 3x3 times like M1A (disk60000, disk50000, disk10800), M1B (disk40000, disk30000, disk20800), M1C (disk20000, disk10000, disk60800). The sapdata1 belongs to stripeset R1M1 (M1A, M1B, M1C). After the split of the 3rd member of the mirrorset, HSTK will reconstruct BCV LUN known as D51.

In order to recover the data from the BCV unit to the production system as shown in the Figure 17, you can use the following procedure. Using the HSG80 controller CLI console command procedure as shown in step 1, you can reconstruct the original status using the BCV unit. Depending on the storage configuration, you need to carefully repeat the following procedure. Once the procedure has been completed for required LUNs and the mount points have been mounted, then you do not need to wait until the resynchronization takes place to bring up the database. During this time you can start restoring control files and archive logs from the latest backup after Brbackup has been successfully completed. During the process of building up BCV unit, the HSTK tool creates LUN D51 with mirrorset BM1A, BM1B, BM1C. You can refer to the HSG80 Array Controller ACS configuration guide for further storage configuration details.

Step 1

Cloning BM1A, BM1B, and BM1C with other mirrorset to bring the original status.

```
HSG80>SET BM1A NOPOLICY
HSG80>SET BM1A MEMBER=3
HSG80>SET BM1A REPLACE=DISK60000
HSG80>SET BM1A REPLACE=DISK50000
HSG80>SET BM1B NOPOLICY
HSG80>SET BM1B MEMBER=3
HSG80>SET BM1B REPLACE=DISK40000
HSG80>SET BM1B REPLACE=DISK30000
HSG80>SET BM1C NOPOLICY
HSG80>SET BM1C MEMBER=3
HSG80>SET BM1C REPLACE=DISK20000
HSG80>SET BM1C REPLACE=DISK10000
HSG80>DELETE D51
HSG80>RENAME RBCV R1M1
HSG80>ADD UNIT D1 R1M1
HSG80>SET D1 DISABLE=ALL
HSG80>Set D1 Enable = SAP64A, SAP64B, SAP64C, SAP64D
```

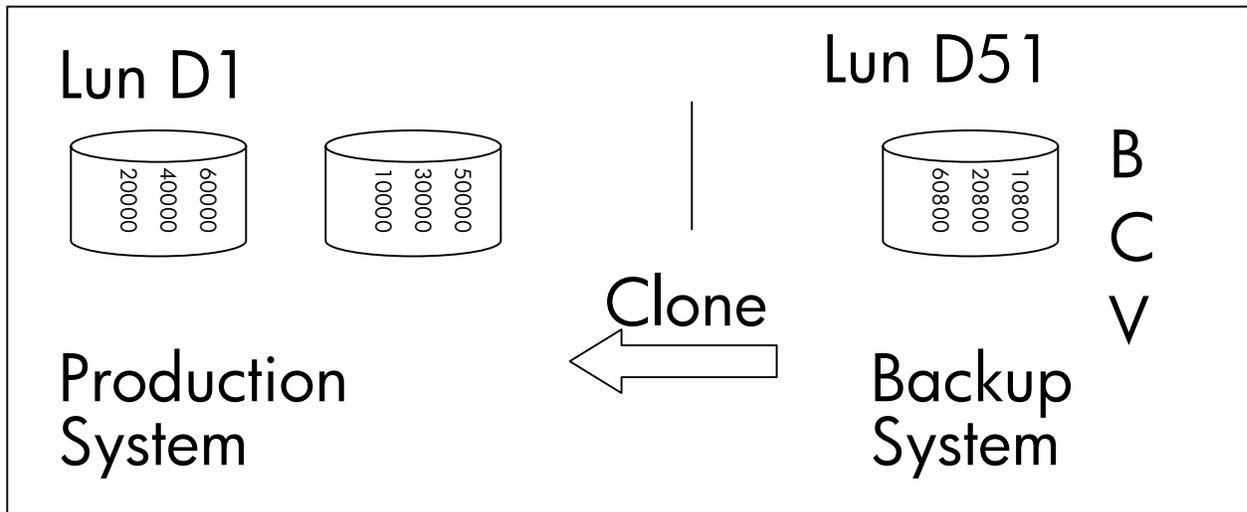


Figure 17, The BCV resynchronization takes place with the original mirrorset

The above procedure shows D1 reconstructed and enabled for the production system. The following procedure is used to link the LUN to the ADVFS filesystem from the production system.

```
# hwmgr -scan scsi
hwmgr: Scan request successfully initiated
# hwmgr -view dev
HWID: Device Name           Mfg      Model           Location
-----
 4: /dev/dmapi/dmapi
 5: /dev/scp_scsi
 6: /dev/kevm
46: /dev/disk/floppy0c           3.5in floppy   fdi0-unit-0
57: /dev/disk/cdrom0c    DEC      RRD47 (C) DEC   bus-0-targ-5-lun-0
58: /dev/disk/dsk0c     COMPAQ   BD009122C6     bus-1-targ-0-lun-0
59: /dev/disk/dsk1c     COMPAQ   BB00911CA0     bus-1-targ-1-lun-0
60: /dev/disk/dsk2c     COMPAQ   BB00911CA0     bus-1-targ-2-lun-0
61: /dev/disk/dsk3c     DEC      RZ1CD-CS (C) DEC bus-1-targ-3-lun-0
62: /dev/disk/dsk4c     DEC      RZ1CD-CS (C) DEC bus-1-targ-4-lun-0
63: /dev/disk/dsk5c     DEC      RZ1CD-CS (C) DEC bus-1-targ-5-lun-0
64: /dev/disk/dsk6c     DEC      HSG80          bus-4-targ-0-lun-1
65: /dev/disk/dsk7c     DEC      HSG80          bus-4-targ-0-lun-3
66: /dev/disk/dsk8c     DEC      HSG80          bus-4-targ-0-lun-51
#ln -fs /dev/disk/dsk8c /etc/fdmns/SAP1_S1
# # showfsets SAP1_S1
sapdata1
  Id           : 3d4f8dce.000d40d4.1.8001
  Files        : 30, SLim= 0, HLim= 0
  Blocks (512) : 92228444, SLim= 0, HLim= 0
  Quota Status : user=off group=off
  Object Safety: off
  Fragging     : on
  DMAPI       : off
#mount SAP1_S1#sapdata1 /oracle/EVM/sapdata1
```

Once the above procedure has been completed for all LUNs and the required mount points have been mounted, including SAP, Oracle binaries, then log in to the oraSID account and initiate svrmgrl prompt.

```
#svrmgrl>connect internal;
#svrmgrl>startup mount;
#svrmgrl>recover database until time.....;
(with respect to your last control file and archive logs available)
```

When Oracle requests archive logs, specify the proper location. Once all the required archive logs have been applied, a message is displayed indicating that recovery is complete. Use the following command to open the databases.

```
#svrmgrl>Alter database open resetlogs;
```

You need to ensure the database is consistent and intact. You can use the SAPDBA tool to check and verify all the database objects and other options etc.

Case 2: Partial restore and recovery

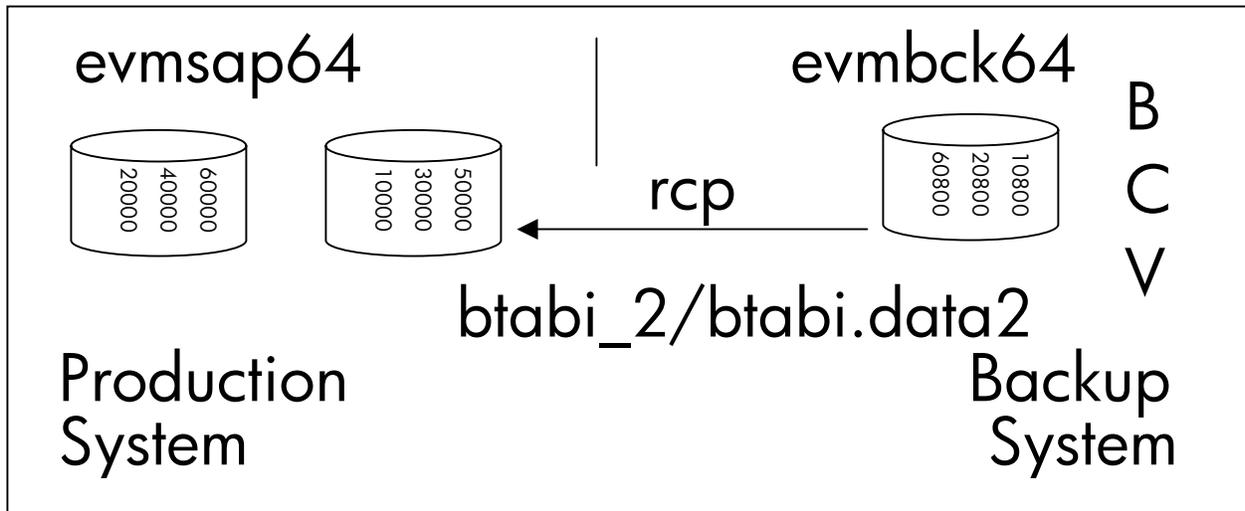


Figure 18, Partial Restore from BCV unit

If a data file gets damaged due to a logical error, a hardware error or a damaged disk, you can restore the file in a reasonable time as long as BCV unit available it is easier to restore from BCV unit to production system. The damaged data files can be copied to the original source volume. Suppose the file `oracle/EVM/sapdata5/btabi_2/btabi.data2` has got damaged on the production system and need to be restored from the BCV unit. The Figure 18 shows how the coping process takes place using the following commands. Before issuing this command please check any user accessing the particular tablespace.

```
svrmgr1> alter database datafile '/oracle/EVM/sapdata5/btabi_2/btabi.data2' offline;
#/rcp evmbck64:/ oracle/EVM/sapdata5/btabi_2/btabi.data2 evmsap64:/
oracle/EVM/sapdata5/btabi_2/btabi.data2
svrmgr1> recover database datafile 'oracle/EVM/sapdata5/btabi_2/btabi.data2';
Issue archivelog as requested and, once the recovery complete, it will display media recovery completed.
svrmgr1 > alter database datafile '/oracle/EVM/sapdata5/btabi_2/btabi.data2' online;
As shown in the above rcp command, where BCV units are on host evmbck64 and target host is evmsap64.
```

Case 3: Restoring and recovering from tape

This assumes the productive database storage is damaged and a full restore from tape must be performed. A full tape backup is available, which was made by Brbackup from the BCV unit to tape using third-party vendor backup software. Ensure your new storage has the correct firmware version (ACS 8.6P), cache memory 512 MB, and other necessary power accessories, in order to bring up the storage properly. You need to follow the following procedure before restoring the original storage configuration. This assumes that you have the original storage configuration details to further configure new storage as follows.

- Configure the storage in multibus failover mode (Refer to the HSG80 Array Controller Configuration Guide ACS 8.5)
- Create mirrorsets
- Create stripe set
- Create unit
- Rename connections to identify the backup system
- Present the unit to the backup system

Brrestore from the backup system

All the required LUNs have to be presented to the backup system. Using the oraSID account, you can issue the following command for a full restore.

```
#Brrestore -u system/password -c -m all
```

Once the above Brrestore has successfully restored, disable all the LUNs from the backup system and then present the production system. During the Brrestore, you can prepare to set up the new production system following the following steps.

operating system

Restore the operating system on the new system from the backup. If you do not have a backup then you need to install the OS and patches on the new system. The configuration should be the same as the setup for the previous production system.

SAP, Oracle and backup software

From the backup, you have to install the SAP binaries and Oracle binaries to the new system in the same ways as for the previous production system. Ensure that you have successfully installed SAP and Oracle without any errors. The following is an example for a single LUN and must be repeated for any additional LUNs.

e.g

```
HSG80>set d1 disable=all
HSG80>set d1 enable=evmsap64a, evmsap64b, evmsap64c, evmsap64c
```

Use the following command to enable the LUNs on the production system.

```
# hwmgr -scan scsi
hwmgr: Scan request successfully initiated
# hwmgr -view dev
75: /dev/disk/dsk17c      DEC      HSG80      bus-4-targ-2-lun-18
77: /dev/disk/dsk19c      DEC      HSG80      bus-4-targ-3-lun-2
78: /dev/disk/dsk20c      DEC      HSG80      bus-4-targ-3-lun-4
79: /dev/disk/dsk21c      DEC      HSG80      bus-4-targ-3-lun-6
80: /dev/disk/dsk22c      DEC      HSG80      bus-4-targ-3-lun-8
#ln -fs /dev/disk/dsk19c /etc/fdmns/SAP2_S2
root@EVMSAP64: /> showfsets SAP2_S2
sapdata2
```

```
Id          : 3d4f8f49.000de10b.1.8001
Files       :      16,  SLim=      0,  HLim=      0
Blocks (512) : 69632256, SLim=      0,  HLim=      0
Quota Status : user=off group=off
Object Safety: off
Fragging    : on
DMPAPI      : off
```

Once the above procedure is completed for all required LUNs and mount points have been mounted, log in to oraSID and invoke svrmgrl command.

```
Svrmgrl>connect internal;
Svrmgrl>recover database until time
(with respect to the latest backup of control files and archive logs)
```

When the database requests archive logs, you need to specify the correct path. After all logs have been applied successfully, Oracle Server Manager will display the message: "media recovery completed". Then issue the command

```
Svrmgrl>alter database open resetlogs;
```

Once the database is up and running again, you need to check the consistency of the database using the standard sapdba procedure.

how to set up storage as backup media

To set up remote storage as backup media you can use multi-mode Fibre Channel cables. In the primary location, you can setup your production server, storage and SAN switch. In the remote location, you can set up the backup server, remote storage and tape drives as shown in Figure 19. This ensures double protection in the event of lost files or disasters in the production environment.

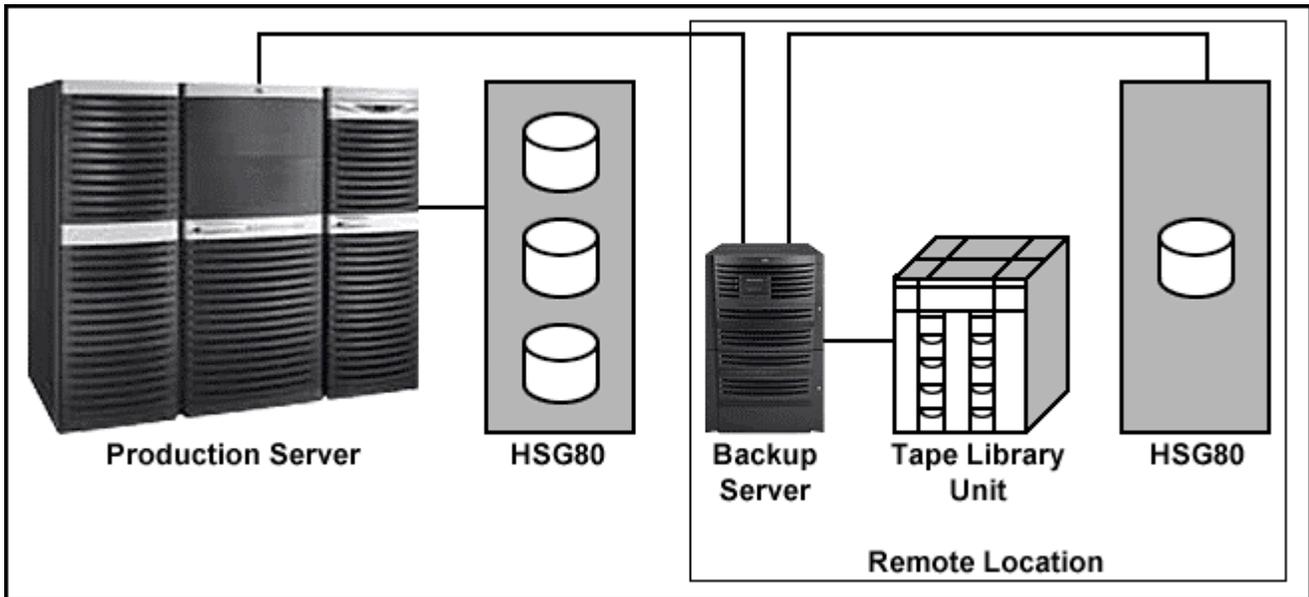


Figure 19, SAP Production environment in 7X24 remote location setup

For remote location purposes, you can choose a disk with higher capacities such as 36GB, 72GB, and 144GB in order to meet your daily backup requirements with respect to your storage capacity.

For fast backup and recovery strategies you can backup to disk and then tape.

The following example will help you to organize backup at remote storage. You can choose the best remote storage configuration to suit your environment with respect to your data capacity.

Note: If you are setting up remote storage, then consider placing the SAP archive and mirror log in the remote storage, instead of placing them at production storage. This enables you to recover the data point in time.

```
HSG80>ADD StripesetSET R1_M6 DISK60000 DISK50000 DISK40000 DISK30000 DISK20000
DISK10000
HSG80>ADD Stripeset R1_M12 DISK60100 DISK50100 DISK40100 DISK30100 DISK20100
HSG80> INIT R1_M6
HSG80> INIT R1_M12
HSG80 ADD UNIT D101 R1_M6
HSG80> ADD UNIT D102 R1_M12
#hwmgr -scan scsi      (from backup system)
#hwmgr -view dev      (from backup system)
#disklabel -rw -t ADVFS /dev/rdisk/dsk237c hsg80                (For unitD101)
#mkfdmn /dev/disk/dsk237c MONBCKUP1
#mkfset MONBCKUP1 sapdata1
#mkfset MONBCKUP1 sapdata2
#mkfset MONBCKUP1 sapdata3
#disklabel -rw -t ADVFS /dev/rdisk/dsk238c hsg80                (For unitD102)
#mkfdmn /dev/disk/dsk237c MONBCKUP2
#mkfset MONBCKUP2 sapdata4
#mkfset MONBCKUP2 sapdata5
#mkfset MONBCKUP2 sapdata6
```

The above example illustrates backup for Monday. You have to repeat the procedure in the same way for the days in the week.

Figure 20 shows that you can have one-week backup and you can rotate the backup cycle for one week. After creating the above Advfs domain and filesets, you can mount the mount points as follow.

```
#mount MONBCKUP1#sapdata1 /bkoracle/SID/sapdata1
#mount MONBCKUP1#sapdata2 /bkoracle/SID/sapdata2
#mount MONBCKUP1#sapdata3 /bkoracle/SID/sapdata3
#mount MONBCKUP2#sapdata4 /bkoracle/SID/sapdata4
#mount MONBCKUP2#sapdata5 /bkoracle/SID/sapdata5
#mount MONBCKUP2#sapdata6 /bkoracle/SID/sapdata6
```

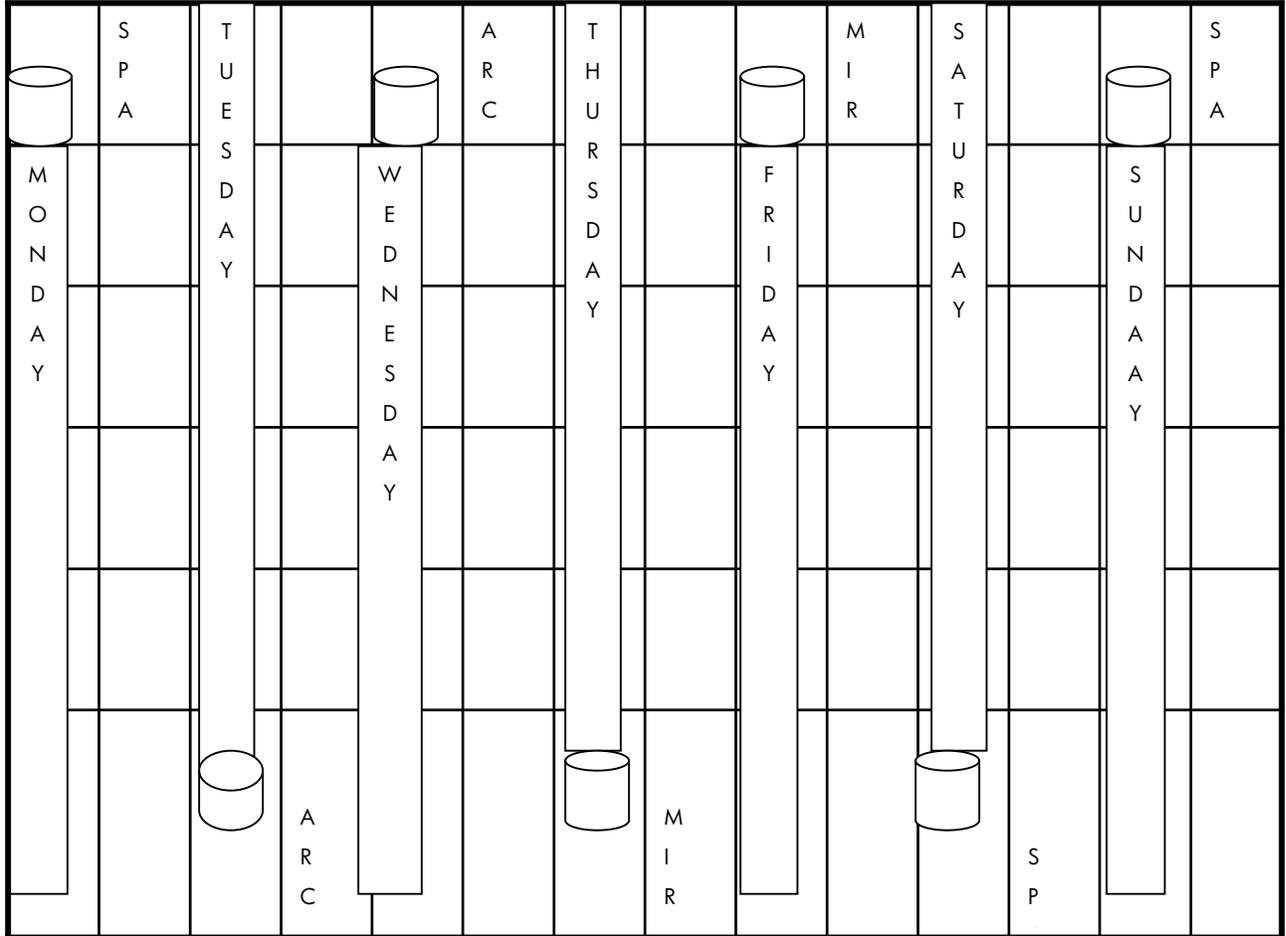


Figure 20, Remote storage backup schedule for one week

The mounting point will remain the same for other weekdays and you need to change the domains and filesets corresponding to weekdays. For example you have created domains FRIBCKUP1 and FRIBCKUP2 with necessary filesets and these filesets are to be mounted for Friday backup. Figure 20 shows an EMA 12000 and demonstrates how to plan and configure the remote storage for a weekly backup schedule.

setting up archive log and mirror log

The following example only applies if you are setting up remote storage.

```
HSG80>add mirror Archive DISK60300 DISK10500
HSG80>add LUN d103 Archive
HSG80>add mirror MLOG DISK60900 DISK11100
HSG80>add LUN D104 MLOG
#mkfdmn /dev/disk/dsk238c SAPARCH
#mkfset SAPARCH archive
```

```

#mkfdmn /dev/disk/dsk239c MIRRORLOG
#mkfset MIRRORLOG mlogA
#mkfset MIRRORLOG mlogB
#mkdir /oracle/SID/saparch
#mkdir /oracle/SID/mirrlogA
#mkdir /oracle/SID/mirrlogB
#mount SAPARCH#archive /oracle/SID/saparch
#mount MIRRORLOG#mlogA /oracle/SID/mirrlogA
#mount MIRRORLOG#mlogB /oracle/SID/mirrlogB

```

configuring the SAPDBA profile for disk-disk copy feature

In our lab test, we have tested the Brbackup disk to disk copy feature using the following specific parameters in initSID.sap profile.

Step 1

```

# backup device type
# [tape | tape_auto | tape_box | pipe | pipe_auto | pipe_box | disk
# | disk_copy | disk_standby | stage | stage_copy | stage_standby
# | util_file | util_file_online | rman_util | rman_prep]
# default: tape
backup_dev_type = disk_copy

```

Step 2

```

# new database home directory for disk_copy | disk_standby
# no default
# new_db_home = /oracle/C11
new_db_home = /bkoracle/EVM

```

As shown in Figure 21, the backup is active from the BCV unit to remote HSG80 subsystem using the Brbackup disk-to-disk copy feature. The number of backup tapes increases corresponding to the volume of data. In some business-critical environments, the administrator cannot afford to wait for a long time to restore data files. If the data file is readily available on disk, then the administrator can restore the data file as quickly as possible. The Brbackup disk-to-disk copy feature therefore benefits during the restore time and saves the productive down time.

Brbackup backs up data from BCV unit to remote HSG80 storage as shown in Figure 21. Then using third-party vendor backup software, we backed up data to tape. This implementation makes the backup and restore procedure faster than the disk to tape backup and restore procedure.

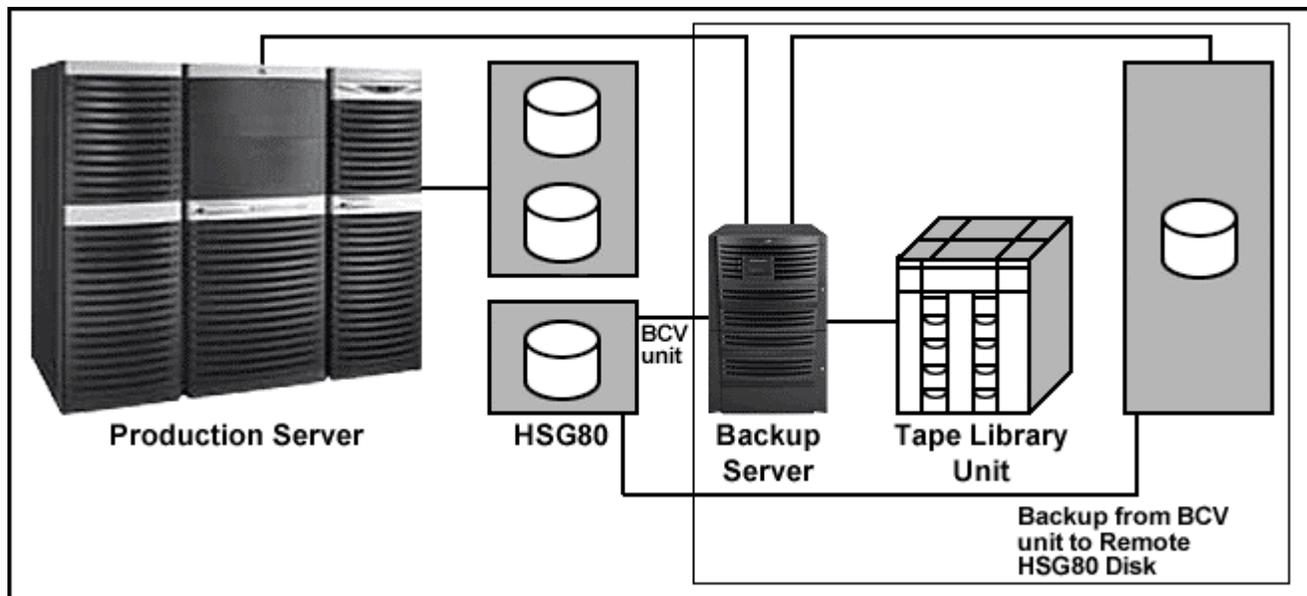


Figure 21, Brbackup Disk-to-Disk copy feature

recovery procedure from remote storage

We have performed tests for the event that the production system and production storage are completely damaged, as show in Figure 22. We have set up the new production environment in a reasonable time. The following procedures are adopted in order to perform quick-restore in a reasonable time.

Operating system backup restored on new system.

Oracle software backup restored on new system.

SAP software backup restored on new system.

Defined necessary hardware environments to the new system.

If you present the storage to a new production system then you can use the following procedure.

Use oraSID account and issue the following procedure.

```
#svrmgrl>connect internal;
```

```
#svrmgrl>startup mount;
```

```
#svrmgrl>recover database until time...;
```

If the database requests archive logs, choose the correct path. Once recovered and the message "media recovery completed" is displayed, then you can open the database.

```
#svrmgrl>alter database open resetlogs;
```

If you set up the original storage configuration and necessary mount points on a new production system, then issue brrestore command. Ensure all your source mountpoints (bkoracle/SID/sapdata1) and target mountpoints (oracle/SID/sapdata1) are present in the new production system.

```
#brrestore -u system/password -c -m all
```

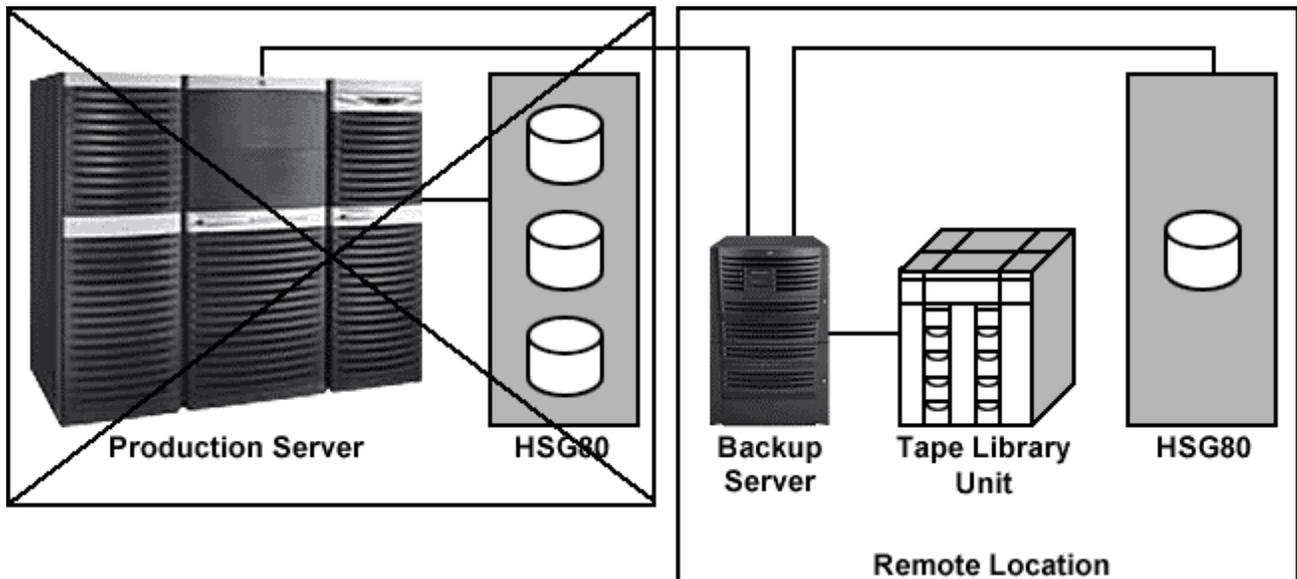


Figure 22, Disaster strikes in the production environment

If you do not have operating system backup, reinstall the same version of UNIX. Use the same production system name, TCP/IP hostname. Reinstall OS patches that existed before the disaster. Make sure that the system boots without any errors and all the network protocols are working without any errors.

To perform a successful quick-restore, the following backup procedures are implemented.

The Operating System, Oracle software and SAP binaries can be backed up once a month. This can be done as a hardware clone. Once cloning is completed, you can reduce the disk. Present this reduced disk to the backup system and mount the disk on the backup system. You can then copy the backup onto tape. This is necessary in order to estimate the total size of the backup files that would be created and to ensure that sufficient space is available on the target disks. As shown in Figure 23, the remote storage made is made available to the new production system and the system is made available to the users. If you configured the original setup as shown in Figure 19, you can use Brrestore for the entire recovery procedure.

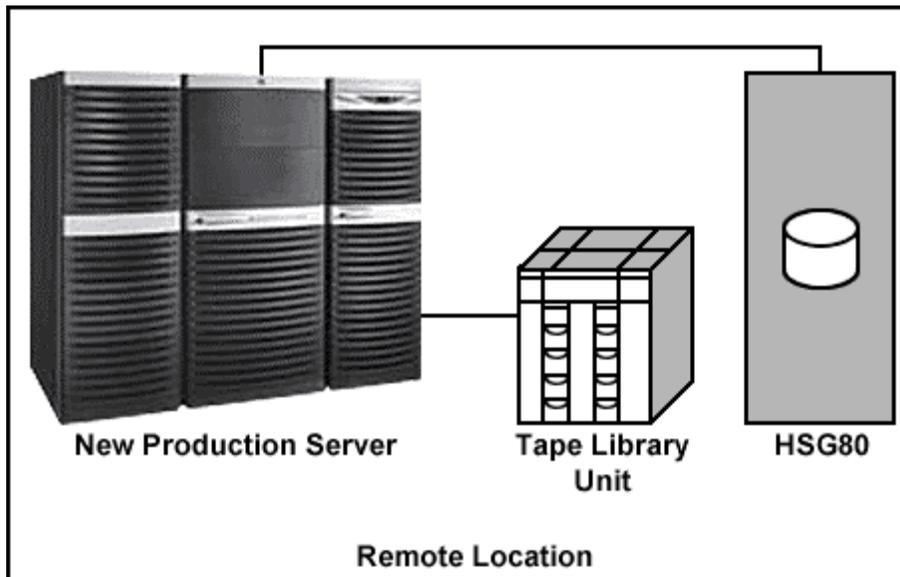


Figure 23, System up after being recovered from a remote location

Note: If you set up the new configuration as shown in Figure 19, then Brrestore is supported. Please refer to appendix b: Brbackup, Rapid-Restore, Rapid-Recovery logs. Please take maximum care before implementing this procedure. You must have thorough understanding of the restore and recovery procedure.

appendix a: examples of lab setup

disk layout

```
** Following are Production system P_mount_point category**
ORABIN#home                /oracle/EVM
SAPBIN#kernel              /usr/sap/EVM
SAPBIN#sapmnt              /sapmnt/EVM
SAPBIN#trans               /usr/sap/trans
mirrlogA#mlogA            /oracle/EVM/mirrlogA
mirrlogB#mlogB            /oracle/EVM/mirrlogB
** STD/BCV mount points on the production system db_mount_point_category**
SAP1_S1#sapdata1          /oracle/EVM/sapdata1
SAP2_S2#sapdata2          /oracle/EVM/sapdata2
SAP3_S3#sapdata3          /oracle/EVM/sapdata3
SAP4_S4#sapdata4          /oracle/EVM/sapdata4
SAP5_S5#sapdata5          /oracle/EVM/sapdata5
SAP6_S6#sapdata6          /oracle/EVM/sapdata6
SAP7_S7#sapdata7          /oracle/EVM/sapdata7
SAP8_S8#sapdata8          /oracle/EVM/sapdata8
SAP9_S9#sapdata9          /oracle/EVM/sapdata9
SAP10_S10#sapdata10       /oracle/EVM/sapdata10
SAP11_S11#sapdata11       /oracle/EVM/sapdata11
redologA#logA             /oracle/EVM/origlogA
redologB#logB             /oracle/EVM/origlogB
** Following are NFS shared mountpoint for both Production and Backup system**
EVMSAP64:/oracle/EVM/saparch      /oracle/EVM/saparch
EVMSAP64:/oracle/EVM/sapbackup    /oracle/EVM/sapbckup
EVMSAP64:/oracle/EVM/sapreorg     /oracle/EVM/sapreorg
EVMSAP64:/oracle/EVM/saptrace     /oracle/EVM/saptrace
EVMSAP64:/oracle/EVM/sapcheck     /oracle/EVM/sapcheck
** Following are after split-mirror mountpoint mounted on Backup system**

HSTK_hsg0_D71#sapdata1          /oracle/EVM/sapdata1
HSTK_hsg0_D73#sapdata3          /oracle/EVM/sapdata3
HSTK_hsg0_D75#sapdata5          /oracle/EVM/sapdata5
HSTK_hsg0_D77#logA              /oracle/EVM/origlogA
HSTK_hsg0_D84#sapdata7          /oracle/EVM/sapdata7
HSTK_hsg0_D86#sapdata9          /oracle/EVM/sapdata9
HSTK_hsg1_D72#sapdata2          /oracle/EVM/sapdata2
HSTK_hsg1_D74#sapdata4          /oracle/EVM/sapdata4
HSTK_hsg1_D76#sapdata6          /oracle/EVM/sapdata6
HSTK_hsg1_D78#logB              /oracle/EVM/origlogB
HSTK_hsg1_D85#sapdata8          /oracle/EVM/sapdata8
HSTK_hsg5_D87#sapdata10         /oracle/EVM/sapdata10
HSTK_hsg5_D88#sapdata11         /oracle/EVM/sapdata11
```

example of listener.ora

```
#####
# Filename.....: listener.ora
# Created.....: created by SAP AG, R/3 Rel. >= 4.0A
# Name.....:
# Date.....:
#####
LISTENER =
  (ADDRESS_LIST =
```

```

        (ADDRESS=
          (PROTOCOL=IPC)
          (KEY= EVM.WORLD)
        )
        (ADDRESS=
          (PROTOCOL=IPC)
          (KEY= EVM)
        )
        (ADDRESS =
          (COMMUNITY = SAP.WORLD)
          (PROTOCOL = TCP)
          (HOST = EVMSAP64)
          (PORT = 1527)
        )
      )
)
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10
TRACE_LEVEL_LISTENER = OFF
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SDU = 32768)
      (SID_NAME = EVM)
      (ORACLE_HOME = /oracle/EVM/817_64)
    )
  )
)

```

example of tnsnames.ora

```

#####
# Filename.....: tnsnames.ora
# Created.....: created by SAP AG, R/3 Rel. >= 4.0A
# Name.....:
# Date.....:
#####

```

```

EVM.WORLD=
  (DESCRIPTION =
    (SDU = 32768)
    (ADDRESS_LIST =
      (ADDRESS =
        (COMMUNITY = SAP.WORLD)
        (PROTOCOL = TCP)
        (HOST = EVMSAP64)
        (PORT = 1527)
      )
    )
  )
  (CONNECT_DATA =
    (SID = EVM)
    (GLOBAL_NAME = EVM.WORLD)
  )
)

```

example of SAPDBA profile

initEVM.sap:

```
# @(#) $Id: //bas/BIN/src/ccm/rsbr/initOSF.sap#1 $ SAP
#####
#
# SAP backup sample profile.
# The parameter syntax is the same as for init.ora parameters.
# Enclose parameter values which consist of more than one symbol in
# double quotes.
# After any symbol, parameter definition can be continued on the next
# line.
# A parameter value list should be enclosed in parentheses, the list
# items should be delimited by commas.
# There can be any number of white spaces (blanks, tabs and new lines)
# between symbols for parameter definition.
#
#####

# backup mode [all | all_data | full | incr | sap_dir | ora_dir
# | <tablespace_name> | <file_id> | <file_id1>-<file_id2>
# | <generic_path> | (<object_list>)]
# default: all
backup_mode = all

# restore mode [all | all_data | full | incr | incr_only | incr_full
# | <tablespace_name> | <file_id> | <file_id1>-<file_id2>
# | <generic_path> | (<object_list>)]
# redirection with '=' is not supported here - use option '-m' instead
# default: all
restore_mode = all

# backup type [offline | offline_force | offline_standby | offline_split
# | offline_stop | online | online_cons | online_split]
# default: offline
backup_type = online_split

# backup device type
# [tape | tape_auto | tape_box | pipe | pipe_auto | pipe_box | disk
# | disk_copy | disk_standby | stage | stage_copy | stage_standby
# | util_file | util_file_online | rman_util | rman_prep]
# default: tape
#backup_dev_type = tape
backup_dev_type = util_file
# backup root directory [<path_name> | (<path_name_list>)]
# default: $SAPDATA_HOME/sapbackup
backup_root_dir = /oracle/EVM/sapbackup

# stage root directory [<path_name> | (<path_name_list>)]
# default: value of the backup_root_dir parameter
stage_root_dir = /oracle/EVM/sapbackup

# compression flag [yes | no | hardware | only]
# default: no
compress = yes
```

```

# compress command
# first $-character is replaced by the source file name
# second $-character is replaced by the target file name
# <target_file_name> = <source_file_name>.Z
# for compress command the -c option must be set
# recommended setting for Brbackup -k only run:
# "compress -b 12 -c $ > $"
# no default
compress_cmd = "compress -c $ > $"

# uncompress command
# first $-character is replaced by the source file name
# second $-character is replaced by the target file name
# <source_file_name> = <target_file_name>.Z
# for uncompress command the -c option must be set
# no default
uncompress_cmd = "uncompress -c $ > $"

# directory for compression [<path_name> | (<path_name_list>)]
# default: value of the backup_root_dir parameter
compress_dir = /oracle/EVM/sapreorg

# brarchive function [save | second_copy | double_save | save_delete
# | second_copy_delete | double_save_delete | copy_save
# | copy_delete_save | delete_saved | delete_copied]
# default: save
archive_function = save

# directory for archive log copies to disk
# default: first value of the backup_root_dir parameter
# Oracle Parallel Server:
# archive_copy_dir = $SAPDATA_HOME/saparchglobal
archive_copy_dir = /oracle/EVM/sapbackup

# directory for archive log copies to stage
# should contain <SID> subdirectory
# default: first value of the stage_root_dir parameter
archive_stage_dir = /oracle/EVM/sapbackup

# new database home directory for disk_copy | disk_standby
# no default
# new_db_home = /oracle/C11

# stage database home directory for stage_copy | stage_standby
# default: value of the new_db_home parameter
# stage_db_home = /oracle/C11

# original database home directory for split mirror disk backup
# no default
# orig_db_home = /oracle/C11

# remote host name
# no default
# remote_host = <host_name>

# remote user name
# default: current operating system user

```

```

# remote_user = <user_name>

# tape copy command [cpio | cpio_gnu | dd | rman | rman_dd]
# default: cpio
tape_copy_cmd = cpio

# disk copy command [copy | dd | rman]
# default: copy
disk_copy_cmd = copy

# stage copy command [rcp | ftp]
# default: rcp
stage_copy_cmd = rcp

# flags for cpio output command
# default: -ovB
cpio_flags = -ovB

# flags for cpio input command
# default: -iuvB
cpio_in_flags = -iuvB

# flags for cpio command for copy of directories to disk
# default: -pdcu
cpio_disk_flags = -pdcu

# flags for dd output command
# default: "obs=16k"
# recommended setting:
# Unix: "obs=nk bs=nk", example: "obs=64k bs=64k"
# NT: "bs=nk", example: "bs=64k"
dd_flags = "obs=64k bs=64k"

# flags for dd input command
# default: "ibs=16k"
# recommended setting:
# Unix: "ibs=nk bs=nk", example: "ibs=64k bs=64k"
# NT: "bs=nk", example: "bs=64k"
dd_in_flags = "ibs=64k bs=64k"

# number of members in RMAN savesets [ 1 | 2 | 3 | 4 | tsp | all ]
# default: 1
saveset_members = 1

# additional parameters for RMAN
# rman_channels and rman_filesperset are only used when rman_util
# rman_channels defines the number of parallel sbt channel allocations
# rman_filesperset = 0 means:
# one file per saveset - for non-incremental backups
# all files in one saveset - for incremental backups
# the others have the same meaning as for native RMAN
# rman_channels = 1
# rman_filesperset = 0
# rman_kbytes = 0
# rman_readrate = 0
# rman_maxopenfiles = 0
# rman_setsize = 0

```

```

# additional parameters for RMAN version 8.1
# the parameters have the same meaning as for native RMAN
# rman_diskratio = 0
# rman_pool = 0
# rman_duplex = 0 | 1 | 2 | 3 | 4
# rman_proxy = no | yes | only

# remote copy-out command (backup_dev_type = pipe)
# $-character is replaced by current device address
# no default
copy_out_cmd = "dd ibs=8k obs=64k of=$"

# remote copy-in command (backup_dev_type = pipe)
# $-character is replaced by current device address
# no default
copy_in_cmd = "dd ibs=64k obs=8k if=$"

# rewind command
# $-character is replaced by current device address
# no default
# operating system dependent, examples:
# HP-UX:  "mt -t $ rew"
# OSF1:  "mt -f $ rew"
# AIX:   "tctl -f $ rewind"
# SINIX: "mt -f $ rew"
# SUN:   "mt -f $ rew"
# NT:    "mt -f $ rewind"
rewind = "mt -f $ rew"

# rewind and set offline command
# $-character is replaced by current device address
# default: value of the rewind parameter
# operating system dependent, examples:
# HP-UX:  "mt -t $ offl"
# OSF1:  "mt -f $ offline"
# AIX:   "tctl -f $ offline"
# SINIX: "mt -f $ offline"
# SUN:   "mt -f $ offline"
# NT:    "mt -f $ offline"
rewind_offline = "mt -f $ offline"

# tape positioning command
# first $-character is replaced by current device address
# second $-character is replaced by number of files to be skipped
# no default
# operating system dependent, examples:
# HP-UX:  "mt -t $ fsf $"
# OSF1:  "mt -f $ fsf $"
# AIX:   "tctl -f $ fsf $"
# SINIX: "mt -f $ fsf $"
# SUN:   "mt -f $ fsf $"
# NT:    "mt -f $ fsf $"
tape_pos_cmd = "mt -f $ fsf $"

# mount backup volume command in auto loader / juke box
# used if backup_dev_type = tape_box | pipe_box
# caution: if successful, exit code 0 and no output!

```

```

# no default
# mount_cmd = "<mount_cmd> $ $ $ [$]"

# dismount backup volume command in auto loader / juke box
# used if backup_dev_type = tape_box | pipe_box
# caution: if successful, exit code 0 and no output!
# no default
# dismount_cmd = "<dismount_cmd> $ $ [$]"

# split mirror disks command
# used if backup_type = offline_split | online_split
# caution: if successful, exit code 0 and no output!
# no default
# split_cmd = "<split_cmd> [$]"
  split_cmd = "/oracle/EVM/backup_scripts/do_split.ksh"
#   split_cmd = "/oracle/EVM/backup_scripts/new_dummy.ksh"
# resynchronize mirror disks command
# used if backup_type = offline_split | online_split
# caution: if successful, exit code 0 and no output!
# no default
# resync_cmd = "<resync_cmd> [$]"
  resync_cmd = "/oracle/EVM/backup_scripts/resync.ksh"
# volume size in KB = K, MB = M or GB = G (backup device dependent)
# default: 1200M
# recommended values for tape devices without hardware compression:
# 60 m    4 mm    DAT DDS-1 tape:    1200M
# 90 m    4 mm    DAT DDS-1 tape:    1800M
# 120 m   4 mm    DAT DDS-2 tape:    3800M
# 125 m   4 mm    DAT DDS-3 tape:    11000M
# 112 m   8 mm    Video tape:        2000M
# 112 m   8 mm    high density:      4500M
# DLT 2000    10/20 GB:    10000M
# DLT 2000XT  15/30 GB:    15000M
# DLT 4000    20/40 GB:    20000M
# DLT 7000    35/70 GB:    35000M
# recommended values for tape devices with hardware compression:
# 60 m    4 mm    DAT DDS-1 tape:    1000M
# 90 m    4 mm    DAT DDS-1 tape:    1600M
# 120 m   4 mm    DAT DDS-2 tape:    3600M
# 125 m   4 mm    DAT DDS-3 tape:    10000M
# 112 m   8 mm    Video tape:        1800M
# 112 m   8 mm    high density:      4300M
# DLT 2000    10/20 GB:    9000M
# DLT 2000XT  15/30 GB:    14000M
# DLT 4000    20/40 GB:    18000M
# DLT 7000    35/70 GB:    30000M
tape_size = 110000M

# volume size in KB = K, MB = M or GB = G used by brarchive
# default: value of the tape_size parameter
# tape_size_arch = 1200M

# level of parallel execution
# default: 0 - set to number of backup devices
exec_parallel = 0

# address of backup device without rewind

```

```

# [<dev_address> | (<dev_address_list>)]
# no default
# operating system dependent, examples:
# HP-UX: /dev/rmt/0mn
# OSF1: /dev/nrmt0h
# AIX: /dev/rmt0.1
# SINIX: /dev/ios0/rstape005n
# SUN: /dev/rmt/0mn
# NT: /dev/nmt0
tape_address = /dev/nrmt9h
# tape_address = /dev/ntape/tape15
# address of backup device without rewind used by brarchive
# default: value of the tape_address parameter
# operating system dependent
tape_address_arch = /dev/nrmt10h

# address of backup device with rewind
# [<dev_address> | (<dev_address_list>)]
# no default
# operating system dependent, examples:
# HP-UX: /dev/rmt/0m
# OSF1: /dev/rmt0h
# AIX: /dev/rmt0
# SINIX: /dev/ios0/rstape005
# SUN: /dev/rmt/0m
# NT: /dev/mt0
# tape_address_rew = /dev/ntape/tape11
# tape_address_rew = /dev/rmt1h
#: address of backup device with rewind used by brarchive
# default: value of the tape_address_rew parameter
# operating system dependent
# tape_address_rew_arch = /dev/rmt0h

# address of backup device with control for mount/dismount command
# [<dev_address> | (<dev_address_list>)]
# default: value of the tape_address_rew parameter
# operating system dependent
# tape_address_ctl = /dev/...
# tape_address_ctl = /dev/ntape/tape15
# address of backup device with control for mount/dismount command
# used by brarchive
# default: value of the tape_address_rew_arch parameter
# operating system dependent
# tape_address_ctl_arch = /dev/...
# tape_address_ctl_arch = /dev/ntape/tape15
# volumes for brarchive
# [<volume_name> | (<volume_name_list>) | SCRATCH]
# no default
#volume_archive = (EVMA01, EVMA02, EVMA03, EVMA04, EVMA05,
# EVMA06, EVMA07, EVMA08, EVMA09, EVMA10,
# EVMA11, EVMA12, EVMA13, EVMA14, EVMA15,
# EVMA16, EVMA17, EVMA18, EVMA19, EVMA20,
# EVMA21, EVMA22, EVMA23, EVMA24, EVMA25,
# EVMA26, EVMA27, EVMA28, EVMA29, EVMA30)
volume_archive = (PR0204,FBP448,FBP404,FBP406,FBP409,FBP414,
FBP417,FBP413,FBP405)
volume_backup = (PR0204,FBP448,FBP404,FBP406,FBP409,FBP414,

```

```

                FBP417,FBP413,FBP405)
# volumes for Brbackup
# [<volume_name> | (<volume_name_list>) | SCRATCH]
# no default
#volume_backup = (EVMB01, EVMB02, EVMB03, EVMB04, EVMB05,
#                 EVMB06, EVMB07, EVMB08, EVMB09, EVMB10,
#                 EVMB11, EVMB12, EVMB13, EVMB14, EVMB15,
#                 EVMB16, EVMB17, EVMB18, EVMB19, EVMB20,
#                 EVMB21, EVMB22, EVMB23, EVMB24, EVMB25,
#                 EVMB26, EVMB27, EVMB28, EVMB29, EVMB30)

# expiration period for backup volumes in days
# default: 30
expir_period = 30

# recommended usages of backup volumes
# default: 100
tape_use_count = 100

# backup utility parameter file
# default: no parameter file
# util_par_file = /usr/opensv/netbackup/bin/initEVM.utl
# util_par_file = /oracle/EVM/817_64/dbs/initEVM.utl
util_par_file = ?/dbs/initEVM.utl
# mount/dismount command parameter file
# default: no parameter file
mount_par_file = initEVM.mnt

# Oracle instance string to the primary database
# [primary_db = <inst_str> | LOCAL]
# no default
# primary_db = <inst_str>
primary_db = EVM
# description of parallel instances for Oracle Parallel Server
# parallel_instances = <instance_desc> | (<instance_desc_list>)
# <instance_desc_list> -> <instance_desc>[,<instance_desc>...]
# <instance_desc>      -> <Oracle_sid>:<Oracle_home>@<inst_str>
# <Oracle_sid>         -> Oracle system id for parallel instance
# <Oracle_home>        -> Oracle home for parallel instance
# <inst_str>           -> Oracle instance string to parallel instance
# Do not include the local instance in the parameter definition!
# default: no parallel instances
#
# example for initC11.sap:
# parallel_instances = (C11_002:/oracle/C11@C11_002,
#                       C11_003:/oracle/C11@C11_003)
#
# example for initC11_002.sap:
# parallel_instances = (C11:/oracle/C11@C11,
#                       C11_003:/oracle/C11@C11_003)

```

example of initEVM.utl

```

#
# This is the sample init.utl file.
#

```

```

# This file contains settings for the NetWorker SAP Interface.
#
# Note: Comments have a "#" is the first column.
#       Options name are always lowercase.
#
# Set to "yes" if you want client-side software compression.
# No compression unless set to "yes".
#
compress = yes
# Number of savesets to use for backups. Must be large enough to
# prevent files from starting an a "continued" saveset segment.
#
savesets = 40
# Number of simultaneous savesets/savestreams to send to server.
# If not set, defaults to 8. Must be between 1 to 32. Be sure
# server is configured to support at least as many streams.
#
parallelism = 8
# Uncomment to set the group to use for saves. No default.
#
group = SAPGroup
# Uncomment to set the media pool to use for saves. If not specified,
# the Default pool is used.
#
pool = SAPData
# class SAP_Backup
# schedule default
# Uncomment to set the explicit browse and retention policy to use on
# saved data (overrides the automatic policy set for the client). Must
# be in getdate format.
#
expiration = 2 weeks
# Notifications backint activity. See documentation for details.
#
# notify_start = mailx -s 'backint start' root
# notify_done = mailx -s 'backint done' root
# notify_success = mailx -s 'backint success' root
# notify_error = mailx -s 'backint error' root
# Set the NetWorker server hostname to use for backups and restores.
#
server = EVMBCK64
client = EVMBCK64
# dump more information from save and recover commands into log file.
#
verbose = yes
# Set raw partition directory other than /dev and /devices. Any files below
# these two directories are considered to be raw disk partitions. If you
# have raw partition under other directory, you must use this option.
#
# raw_dir = /oracle/data

```

appendix b: Brbackup, Rapid-Restore, Rapid-Recovery logs

script: do_split.ksh

```
#!/usr/bin/ksh
alias mydate='date +"%d-%h-%Y %H:%M:%S"'
#
export CLONE_HOME="/usr/local/hstk/scripts"
#export CLONE_DEBUG=1
LOG="/oracle/EVM/backup_scripts/do_split.log"
SYNC=/sbin/sync
COMMAND_1="$CLONE_HOME/sh/hsgcontrol.sh EVM SPLIT"
COMMAND_2="$CLONE_HOME/sh/hsgcontrol.sh EVM MOUNT"
#
mydate > $LOG
echo "#INFO: `mydate` Syncing filesystem cache" | tee -a $LOG
$SYNC >> $LOG 2>&1
$SYNC >> $LOG 2>&1
$SYNC >> $LOG 2>&1
echo "#INFO: `mydate` starting BCV SPLIT" | tee -a $LOG
if rsh evmbck64 -l root $COMMAND_1 >> $LOG 2>&1
then
    echo "#INFO: `mydate` $COMMAND_1 succeeded" | tee -a $LOG
    echo "#INFO: `mydate` starting BCV MOUNT" | tee -a $LOG
    if rsh evmbck64 -l root $COMMAND_2 >> $LOG 2>&1
    then
        echo "#INFO: `mydate` $COMMAND_2 succeeded" | tee -a $LOG
    else
        echo "#ERROR: `mydate` $COMMAND_2 did not succeed!" | tee -a $LOG
        exit 2
    fi
else
    echo "#ERROR: `mydate` $COMMAND_1 did not succeed!" | tee -a $LOG
    exit 1
fi
#
mydate >> $LOG
exit 0
```

scripts: do_resync.ksh

```
#!/usr/bin/ksh
alias mydate='date +"%d-%h-%Y %H:%M:%S"'
#
export CLONE_HOME="/usr/local/hstk/scripts"
#export CLONE_DEBUG=1
LOG="/oracle/EVM/backup_scripts/do_resync.log"
COMMAND_1="$CLONE_HOME/sh/hsgcontrol.sh EVM UNMOUNT"
COMMAND_2="$CLONE_HOME/sh/hsgcontrol.sh EVM DELETE"
COMMAND_3="$CLONE_HOME/sh/hsgcontrol.sh EVM MERGE_NOCHECK"
#
mydate > $LOG
echo "#INFO: `mydate` starting BCV UNMOUNT" | tee -a $LOG
if rsh evmbck64 -l root $COMMAND_1 >> $LOG 2>&1
then
    echo "#INFO: `mydate` $COMMAND_1 succeeded" | tee -a $LOG
    echo "#INFO: `mydate` starting BCV DELETE" | tee -a $LOG
```

```

if rsh evmbck64 -l root $COMMAND_2 >> $LOG 2>&1
then
    echo "#INFO: `mydate` $COMMAND_2 succeeded" | tee -a $LOG
    echo "#INFO: `mydate` starting BCV MERGE" | tee -a $LOG
    if rsh evmbck64 -l root $COMMAND_3 >> $LOG 2>&1
    then
        echo "#INFO: `mydate` $COMMAND_3 succeeded" | tee -a $LOG
    else
        echo "#ERROR: `mydate` $COMMAND_3 did not succeed!" | tee -a $LOG
        exit 3
    fi
else
    echo "#ERROR: `mydate` $COMMAND_2 did not succeed!" | tee -a $LOG
    exit 2
fi
else
    echo "#ERROR: `mydate` $COMMAND_1 did not succeed!" | tee -a $LOG
    exit 1
fi
#
mydate >> $LOG
exit 0

```

Brbackup output Log

```

BR051I BRBACKUP 6.10 (40)
BR055I Start of database backup: bdjpbne.and 2002-12-22 13.39.06

BR351I Restoring /oracle/EVM/sapbackup/cntrlEVM.dbf
BR355I from /oracle/EVM/sapbackup/cntrlEVM.dbf ...

BR319I Control file copied to: /oracle/EVM/sapbackup/cntrlEVM.dbf 14049280

BR428W File /bkoracle/EVM/sapdata9/btabd_1/btabd.data1 will be overwritten
BR428W File /bkoracle/EVM/sapdata9/btabd_2/btabd.data2 will be overwritten
:
:
BR428W File /bkoracle/EVM/sapdata1/system_1/system.data1 will be overwritten
BR428W File /bkoracle/EVM/sapdata1/system_2/system.data2 will be overwritten
BR428W File /bkoracle/EVM/sapdata1/cntrl/cntrlEVM.dbf will be overwritten
BR428W File /bkoracle/EVM/origlogA/cntrl/cntrlEVM.dbf will be overwritten
BR428W File /bkoracle/EVM/saparch/cntrl/cntrlEVM.dbf will be overwritten

BR101I Parameters

oracle_sid           EVM
oracle_home          /oracle/EVM/817_64
oracle_profile        /oracle/EVM/817_64/dbs/initEVM.ora
                     /oracle/EVM/817_64/dbs/initora.addon
                     /oracle/EVM/817_64/dbs/init_817.ora
sapdata_home         /oracle/EVM
sap_profile           /oracle/EVM/817_64/dbs/initEVM.sap
backup_mode          ALL
backup_type          online_split
backup_dev_type       disk_copy
new_db_home          /bkoracle/EVM

```

```

compress                no
disk_copy_cmd           copy
cpio_disk_flags        -pdcu
exec_parallel          0
primary_db             EVM
split_cmd              /oracle/EVM/backup_scripts/do_split.ksh
resync_cmd             NULL
system_info            oraevm/oraevm EVMBCK64 OSF1 V5.1 1885 alpha
oracle_info            8.1.7.2.0 8192 811 4612002
sap_info               46C SAPR3 EVM
make_info              alphaosf OCI_805_64 Jun 10 2002
command_line           brbackup -u system/***** -c -m all

```

BR116I ARCHIVE LOG LIST before backup for database instance EVM/PRIMARY

```

Database log mode      Archive Mode
Automatic archival    Enabled
Archive destination   /oracle/EVM/saparch/EVMarch
Archive format        %t_%s.dbf
Oldest online log sequence 807
Next log sequence to archive 810
Current log sequence  810          SCN: 4611692
Database block size   8192          Thread: 1
Oracle version        8.1.7.2.0

```

BR280I Time stamp 2002-12-22 13.39.32

BR315I 'Alter tablespace PSAPBTABD begin backup' successful

BR280I Time stamp 2002-12-22 13.39.33

BR315I 'Alter tablespace PSAPBTABI begin backup' successful

:
:

BR280I Time stamp 2002-12-22 13.39.38

BR315I 'Alter tablespace PSAPUSER1D begin backup' successful

BR280I Time stamp 2002-12-22 13.39.38

BR315I 'Alter tablespace PSAPUSER1I begin backup' successful

BR280I Time stamp 2002-12-22 13.39.38

BR315I 'Alter tablespace SYSTEM begin backup' successful

BR280I Time stamp 2002-12-22 13.39.38

BR295I Splitting mirror disks...

BR278I Command output of '/oracle/EVM/backup_scripts/do_split.ksh':

#INFO: 22-Dec-2002 13:39:38 Syncing filesystem cache

#INFO: 22-Dec-2002 13:39:39 starting BCV SPLIT

#INFO: 22-Dec-2002 13:47:57 /usr/local/hstk/scripts/sh/hsgcontrol.sh EVM SPLIT
succeeded

#INFO: 22-Dec-2002 13:47:57 starting BCV MOUNT

#INFO: 22-Dec-2002 13:56:32 /usr/local/hstk/scripts/sh/hsgcontrol.sh EVM MOUNT
succeeded

BR280I Time stamp 2002-12-22 13.56.32

BR296I Split of mirror disks successful

BR280I Time stamp 2002-12-22 13.56.32

BR317I 'Alter tablespace PSAPBTABD end backup' successful

```

BR280I Time stamp 2002-12-22 13.56.32
BR317I 'Alter tablespace PSAPBTABI end backup' successful
:
:
BR280I Time stamp 2002-12-22 13.56.33
BR317I 'Alter tablespace PSAPUSER1D end backup' successful

BR280I Time stamp 2002-12-22 13.56.33
BR317I 'Alter tablespace PSAPUSER1I end backup' successful

BR280I Time stamp 2002-12-22 13.56.33
BR317I 'Alter tablespace SYSTEM end backup' successful

BR340I Switching to next online redo log file for database instance EVM ...
BR321I Switch to next online redo log file for database instance EVM successful

```

BR117I ARCHIVE LOG LIST after backup for database instance EVM/PRIMARY

```

Database log mode           Archive Mode
Automatic archival          Enabled
Archive destination         /oracle/EVM/saparch/EVMarch
Archive format               %t_%s.dbf
Oldest online log sequence  875
Next log sequence to archive 878
Current log sequence        878          SCN: 4630135
Database block size         8192          Thread: 1
Oracle version              8.1.7.2.0

```

BR118I Tablespaces and data files

```

PSAPBTABD      ONLINE*    ONLINE    /oracle/EVM/sapdata9/btabd_1/btabd.data1
629153792      7 -1932786979 NOLINK    FILE
PSAPBTABD      ONLINE*    ONLINE    /oracle/EVM/sapdata9/btabd_2/btabd.data2
4294967296     44 -1932786979 NOLINK    FILE
PSAPBTABI      ONLINE     ONLINE    /oracle/EVM/sapdata5/btabi_1/btabi.data1
367009792      6 -1641351819 NOLINK    FILE
PSAPBTABI      ONLINE     ONLINE    /oracle/EVM/sapdata5/btabi_2/btabi.data2
4292608000     45 -1641351819 NOLINK    FILE
PSAPBTABI      ONLINE     ONLINE    /oracle/EVM/sapdata5/btabi_3/btabi.data3
4096000000     50 -1641351819 NOLINK    FILE
:
:
PSAPUSER1D     ONLINE*    ONLINE    /oracle/EVM/sapdata1/user1d_1/user1d.data1
5251072        28 -1068759987 NOLINK    FILE
PSAPUSER1I     ONLINE     ONLINE    /oracle/EVM/sapdata1/user1i_1/user1i.data1
5251072        29 -1068759987 NOLINK    FILE
SYSTEM         ONLINE*    SYSTEM    /oracle/EVM/sapdata1/system_1/system.data1
314580992      1 -1068759987 NOLINK    FILE
SYSTEM         ONLINE*    SYSTEM    /oracle/EVM/sapdata1/system_2/system.data2
2048000000     48 -1068759987 NOLINK    FILE

```

BR119I Redo log files

```

/oracle/EVM/origlogA/log_g11m1.dbf          20972032  11  -1473067494  INUSE
NOLINK  FILE

```

/oracle/EVM/mirrlogA/log_g11m2.dbf	20972032	11	-271487206	INUSE
NOLINK FILE				
/oracle/EVM/origlogB/log_g12m1.dbf	20972032	12	-697279307	INUSE
NOLINK FILE				
/oracle/EVM/mirrlogB/log_g12m2.dbf	20972032	12	-1794560333	INUSE
NOLINK FILE				
/oracle/EVM/origlogA/log_g13m1.dbf	20972032	13	-1473067494	INUSE
NOLINK FILE				
/oracle/EVM/mirrlogA/log_g13m2.dbf	20972032	13	-271487206	INUSE
NOLINK FILE				
/oracle/EVM/origlogB/log_g14m1.dbf	20972032	14	-697279307	INUSE
NOLINK FILE				
/oracle/EVM/mirrlogB/log_g14m2.dbf	20972032	14	-1794560333	INUSE
NOLINK FILE				

BR120I Control files

/oracle/EVM/sapdata1/cntrl/cntrlEVM.dbf	14049280	0	-1068759987	NOLINK	FILE
/oracle/EVM/origlogA/cntrl/cntrlEVM.dbf	14049280	0	-1473067494	NOLINK	FILE
/oracle/EVM/saparch/cntrl/cntrlEVM.dbf	14049280	0	1	NOLINK	FILE

BR280I Time stamp 2002-12-22 13.56.33

BR057I Backup of database: EVM

BR058I BRBACKUP action ID: bdjpbne

BR059I BRBACKUP function ID: and

BR110I Backup mode: ALL

BR077I Database files for backup:

/oracle/EVM/sapdata1/cntrl/cntrlEVM.dbf

/oracle/EVM/origlogA/cntrl/cntrlEVM.dbf

/oracle/EVM/saparch/cntrl/cntrlEVM.dbf

BR061I 60 files found for backup, total size 91611.227 MB

BR143I Backup type: online_split

BR112I Files will not be compressed

BR130I Backup device type: disk_copy

BR094I All files will be copied to the new database home directory: /bkoracle/EVM

BR140I Number of parallel copy processes: 13

BR126I Unattended mode active - no operator confirmation required

BR202I Saving init_ora

BR203I to /oracle/EVM/sapbackup/EVM ...

BR202I Saving /oracle/EVM/817_64/dbs/initEVM.sap

BR203I to /oracle/EVM/sapbackup/EVM ...

BR280I Time stamp 2002-12-22 13.56.34

BR198I Profiles saved successfully

BR202I Saving /oracle/EVM/sapdata6/poold_1/poold.data1

BR203I to /bkoracle/EVM/sapdata6/poold_1/poold.data1 ...

:

:

BR202I Saving /oracle/EVM/sapdata6/pooli_2/pooli.data2

BR203I to /bkoracle/EVM/sapdata6/pooli_2/pooli.data2 ...

BR202I Saving /oracle/EVM/sapdata1/roll_2/roll2.data1

BR203I to /bkoracle/EVM/sapdata1/roll_2/roll2.data1 ...

```

#FILE..... /oracle/EVM/sapdata1/roll_2/roll2.data1
#SAVED.... /bkoracle/EVM/sapdata1/roll_2/roll2.data1  #13/1

BR280I Time stamp 2002-12-22 14.04.29
BR063I 1 of 60 files processed - 2000.008 MB of 91611.227 MB done
BR204I Percentage done: 2.18%, estimated end time: 19:59
BR001I * _____

BR202I Saving /oracle/EVM/sapdata1/system_2/system.data2
BR203I to /bkoracle/EVM/sapdata1/system_2/system.data2 ...

#FILE..... /oracle/EVM/sapdata2/es46cd_2/es46cd.data2
#SAVED.... /bkoracle/EVM/sapdata2/es46cd_2/es46cd.data2  #8/1

BR280I Time stamp 2002-12-22 14.05.51
BR063I 2 of 60 files processed - 3800.016 MB of 91611.227 MB done
BR204I Percentage done: 4.15%, estimated end time: 17:40
BR001I ** _____

BR202I Saving /oracle/EVM/sapdata2/es46cd_1/es46cd.data1
BR203I to /bkoracle/EVM/sapdata2/es46cd_1/es46cd.data1 ...
:
:

BR202I Saving /oracle/EVM/sapdata8/el46cd_1/el46cd.data1
BR203I to /bkoracle/EVM/sapdata8/el46cd_1/el46cd.data1 ...

#FILE..... /oracle/EVM/sapdata8/el46cd_1/el46cd.data1
#SAVED.... /bkoracle/EVM/sapdata8/el46cd_1/el46cd.data1  #13/27

BR280I Time stamp 2002-12-22 14.43.11
BR063I 58 of 60 files processed - 91584.430 MB of 91611.227 MB done
BR204I Percentage done: 99.97%, estimated end time: 14:43
BR001I *****

BR202I Saving /oracle/EVM/sapbackup/cntrlEVM.dbf
BR203I to /bkoracle/EVM/origlogA/cntrl/cntrlEVM.dbf ...

#FILE..... /oracle/EVM/origlogA/cntrl/cntrlEVM.dbf
#SAVED.... /bkoracle/EVM/origlogA/cntrl/cntrlEVM.dbf  #13/28

BR280I Time stamp 2002-12-22 14.43.12
BR063I 59 of 60 files processed - 91597.828 MB of 91611.227 MB done
BR204I Percentage done: 99.99%, estimated end time: 14:43
BR001I *****

BR202I Saving /oracle/EVM/sapbackup/cntrlEVM.dbf
BR203I to /bkoracle/EVM/saparch/cntrl/cntrlEVM.dbf ...

#FILE..... /oracle/EVM/saparch/cntrl/cntrlEVM.dbf
#SAVED.... /bkoracle/EVM/saparch/cntrl/cntrlEVM.dbf  #13/29

BR280I Time stamp 2002-12-22 14.43.12
BR063I 60 of 60 files processed - 91611.227 MB of 91611.227 MB done
BR204I Percentage done: 100.00%, estimated end time: 14:43
BR001I *****

```

```

BR202I Saving reorg_log
BR203I to /oracle/EVM/sapbackup/EVM ...

BR202I Saving /oracle/EVM/sapbackup/bdjpbne.and
BR203I to /oracle/EVM/sapbackup/EVM ...

BR202I Saving /oracle/EVM/sapbackup/backEVM.log
BR203I to /oracle/EVM/sapbackup/EVM ...

BR056I End of database backup: bdjpbne.and 2002-12-22 14.43.15
BR280I Time stamp 2002-12-22 14.43.17
BR053I BRBACKUP terminated successfully with warnings

```

Rapid Recovery Log

```

# ./splmnt.ksh
# df
Filesystem          512-blocks      Used    Available Capacity  Mounted on
root_domain#root    786432          669050      93904      88%     /
/proc                0                0            0      100%    /proc
usr_domain#usr       4300800         2003432     1866592     52%     /usr
usr_domain#var       4300800         396352      1866592     18%     /var
ORABIN#home          35556384        193272      30114112     1%     /oracle/EVM
ORABIN#805_64        35556384         22210      30114112     1%     /oracle/805_64
ORABIN#817_64        35556384        4100858     30114112    12%     /oracle/EVM/817_64
ORABIN#stage         35556384        1096730     30114112     4%     /oracle/stage/817_64
SAPBIN#kernel        35551376         322624     26484464     2%     /usr/sap/EVM
SAPBIN#sapmnt        35551376         848742     26484464     4%     /sapmnt/EVM
SAPBIN#trans         35551376        7876530     26484464    23%     /usr/sap/trans
BCPBCK#rlogb         426676656       163952     188845648     1%     /oracle/EVM/origlogB
BCPBCK#rloga         426676656       191408     188845648     1%     /oracle/EVM/origlogA
BCPBCK#data1         426676656       39222310    188845648    18%     /oracle/EVM/sapdata1
BCPBCK#data2         426676656       9782512     188845648     5%     /oracle/EVM/sapdata2
BCPBCK#data3         426676656      10806592     188845648     6%     /oracle/EVM/sapdata3
BCPBCK#data4         426676656      23612992     188845648    12%     /oracle/EVM/sapdata4
BCPBCK#data5         426676656      23353744     188845648    12%     /oracle/EVM/sapdata5
BCPBCK#data6         426676656      27824192     188845648    13%     /oracle/EVM/sapdata6
BCPBCK#data7         426676656       3126480     188845648     2%     /oracle/EVM/sapdata7
BCPBCK#data8         426676656       1638464     188845648     1%     /oracle/EVM/sapdata8
BCPBCK#data9         426676656       9617488     188845648     5%     /oracle/EVM/sapdata9
BCPBCK#data10        426676656      32768288     188845648    15%     /oracle/EVM/sapdata10
BCPBCK#data11        426676656      16384160     188845648     8%     /oracle/EVM/sapdata11
BCPBCK#saptrace      426676656         236      188845648     1%     /oracle/EVM/saptrace
BCPBCK#sapcheck      426676656         32      188845648     1%     /oracle/EVM/sapcheck
BCPBCK#reorg         426676656         142      188845648     1%     /oracle/EVM/sapreorg
BCPBCK#arch          426676656      39127244     188845648    18%     /oracle/EVM/saparch
BCPBCK#mirra         426676656        81984      188845648     1%     /oracle/EVM/mirrlogA
BCPBCK#mirrb         426676656        81984      188845648     1%     /oracle/EVM/mirrlogB
#
EVMSAP64:oraevm> date
Sun Dec 22 20:44:46 CET 2002
EVMSAP64:oraevm> svrmgrl

```

Oracle Server Manager Release 3.1.7.0.0 - Production

Copyright (c) 1997, 1999, Oracle Corporation. All Rights Reserved.

Oracle8i Enterprise Edition Release 8.1.7.2.0 - Production
With the Partitioning option
JServer Release 8.1.7.2.0 - Production

```
SVRMGR> connect internal
Connected.
SVRMGR> startup
ORACLE instance started.
Total System Global Area          1064354788 bytes
Fixed Size                        103396 bytes
Variable Size                     451153920 bytes
Database Buffers                  612032512 bytes
Redo Buffers                      1064960 bytes
ORA-00214: controlfile '/oracle/EVM/sapdata1/cntrl/cntrlEVM.dbf' version 86392
inconsistent with file '/oracle/E
VM/saparch/cntrl/cntrlEVM.dbf' version 84983
SVRMGR
```

```
bSVRMGR> recover database until cancel using backup controlfile;
ORA-00279: change 4614181 generated at 12/22/2002 13:46:36 needed for thread 1
ORA-00289: suggestion : /oracle/EVM/saparch/EVMarch1_819.dbf
ORA-00280: change 4614181 for thread 1 is in sequence #819
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

```
Log applied.
ORA-00279: change 4614427 generated at 12/22/2002 13:46:49 needed for thread 1
ORA-00289: suggestion : /oracle/EVM/saparch/EVMarch1_820.dbf
ORA-00280: change 4614427 for thread 1 is in sequence #820
ORA-00278: log file '/oracle/EVM/saparch/EVMarch1_819.dbf' no longer needed for this
recovery
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

```
Log applied.
ORA-00279: change 4614674 generated at 12/22/2002 13:47:02 needed for thread 1
ORA-00289: suggestion : /oracle/EVM/saparch/EVMarch1_821.dbf
ORA-00280: change 4614674 for thread 1 is in sequence #821
ORA-00278: log file '/oracle/EVM/saparch/EVMarch1_820.dbf' no longer needed for this
recovery
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

```
Log applied.
ORA-00279: change 4614922 generated at 12/22/2002 13:47:15 needed for thread 1
ORA-00289: suggestion : /oracle/EVM/saparch/EVMarch1_822.dbf
ORA-00280: change 4614922 for thread 1 is in sequence #822
ORA-00278: log file '/oracle/EVM/saparch/EVMarch1_821.dbf' no longer needed for this
recovery
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

```
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
ORA-00279: change 4674702 generated at 12/22/2002 15:01:15 needed for thread 1
ORA-00289: suggestion : /oracle/EVM/saparch/EVMarch1_1064.dbf
ORA-00280: change 4674702 for thread 1 is in sequence #1064
ORA-00278: log file '/oracle/EVM/saparch/EVMarch1_1063.dbf' no longer needed for this
recovery
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
SVRMGR> alter database open resetlogs;
Statement processed.
```

```
SVRMGR> select * from v$instance;
INSTANCE_N INSTANCE_NAME      HOST_NAME
VERSION          S
TARTUP_T STATUS  PAR THREAD#      ARCHIVE LOG_SWITCH_ LOGINS      SHU DATABASE_STATUS
INSTANCE_ROLE
```

```
-----
-----
-----
1 EVM          EVMSAP64
8.1.7.2.0      2
2-DEC-02 OPEN  NO          1 STARTED      ALLOWED      NO  ACTIVE
PRIMARY_INSTANCE
1 row selected.
```

```
SVRMGR> select * from dba_data_files;
```

```
FILE_NAME
FILE_ID      TABLESPACE_NAME
            BYTES      BLOCKS      STATUS      RELATIVE_F AUT MAXBYTES      MAXBLOCKS
INCREMENT_  USER_BYTES  USER_BLOCK
```

```
-----
-----
-----
/oracle/EVM/sapdata3/es46ci_3/es46ci.data3
56 PSAPES46CI
1023991808    124999 AVAILABLE    56 NO          0          0
0 1023983616    124998
/oracle/EVM/sapdata2/es46cd_3/es46cd.data3
55 PSAPES46CD
1023991808    124999 AVAILABLE    55 NO          0          0
0 1023983616    124998
/oracle/EVM/sapdata4/ddici_2/ddici.data2
54 PSAPDDICI
1023991808    124999 AVAILABLE    54 NO          0          0
0 1023983616    124998
/oracle/EVM/sapdata7/ddicd_2/ddicd.data2
53 PSAPDDICD
2097152000    256000 AVAILABLE    33 NO          0          0
0 2097143808    255999
/oracle/EVM/sapdata10/ssqjd_2/ssqjd.data2
32 PSAPSSQJD
2097152000    256000 AVAILABLE    32 YES 3.4360E+10    4194302
1 2097143808    255999
/oracle/EVM/sapdata10/ssqjd_1/ssqjd.data1
31 PSAPSSQJD
/oracle/EVM/sapdata1/loadi_1/loadi.data1
18 PSAPLOADI
```

```
Sun Dec 22 21:12:50 CET 2002
EVMSAP64:evmadm> startsap
```

```
Starting SAP-Collector Daemon
```

```
-----
21:12:55 22.12.2002
Starting collector (create new process)
saposcol on host EVMSAP64 started
```

```
Checking SAP R/3 EVM Database
```

Database is running

Starting SAP R/3 Instance

Startup-Log is written to /usr/users/evmadm/startsap_EVMSAP64_00.log
Instance on host EVMSAP64 started

EVMSAP64:evmadm>

Rapid-Restore Log

BR401I BRRESTORE 6.10 (40)
BR169I Value 'online_split' of parameter/option 'backup_type/-t' ignored for
'brrestore' - 'online' assumed
BR169I Value 'disk_copy' of parameter/option 'backup_dev_type/-d' ignored for
'brrestore' - 'disk' assumed
BR405I Start of file restore: rdjpfiau.rsb 2002-12-23 08.14.20

BR101I Parameters

oracle_sid	EVM
oracle_home	/oracle/EVM/817_64
oracle_profile	/oracle/EVM/817_64/dbs/initEVM.ora /oracle/EVM/817_64/dbs/initora.addon /oracle/EVM/817_64/dbs/init_817.ora
sapdata_home	/oracle/EVM
sap_profile	/oracle/EVM/817_64/dbs/initEVM.sap
restore_mode	ALL
backup_dev_type	disk
compress	no
disk_copy_cmd	copy
exec_parallel	0
system_info	oraevm/oraevm EVMBCK64 OSF1 V5.1 1885 alpha
make_info	alphaosf OCI_805_64 Jun 10 2002
command_line	brrestore -u system/***** -c -m all

BR428W File /oracle/EVM/sapdata9/btabd_1/btabd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata9/btabd_2/btabd.data2 will be overwritten
BR428W File /oracle/EVM/sapdata5/btabi_1/btabi.data1 will be overwritten
BR428W File /oracle/EVM/sapdata5/btabi_2/btabi.data2 will be overwritten
BR428W File /oracle/EVM/sapdata5/btabi_3/btabi.data3 will be overwritten
BR428W File /oracle/EVM/sapdata1/clud_1/clud.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/clui_1/clui.data1 will be overwritten
BR428W File /oracle/EVM/sapdata7/ddicd_1/ddicd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata7/ddicd_2/ddicd.data2 will be overwritten
BR428W File /oracle/EVM/sapdata4/ddici_1/ddici.data1 will be overwritten
BR428W File /oracle/EVM/sapdata4/ddici_2/ddici.data2 will be overwritten
BR428W File /oracle/EVM/sapdata1/docud_1/docud.data1 will be overwritten
BR428W File /oracle/EVM/sapdata6/docui_1/docui.data1 will be overwritten
BR428W File /oracle/EVM/sapdata8/el46cd_1/el46cd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/el46ci_1/el46ci.data1 will be overwritten
BR428W File /oracle/EVM/sapdata2/es46cd_1/es46cd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata2/es46cd_2/es46cd.data2 will be overwritten
BR428W File /oracle/EVM/sapdata2/es46cd_3/es46cd.data3 will be overwritten
BR428W File /oracle/EVM/sapdata3/es46ci_1/es46ci.data1 will be overwritten
BR428W File /oracle/EVM/sapdata3/es46ci_2/es46ci.data2 will be overwritten
BR428W File /oracle/EVM/sapdata3/es46ci_3/es46ci.data3 will be overwritten

BR428W File /oracle/EVM/sapdata1/loadd_1/loadd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/loadi_1/loadi.data1 will be overwritten
BR428W File /oracle/EVM/sapdata6/poold_1/poold.data1 will be overwritten
BR428W File /oracle/EVM/sapdata6/poold_2/poold.data2 will be overwritten
BR428W File /oracle/EVM/sapdata6/pooli_1/pooli.data1 will be overwritten
BR428W File /oracle/EVM/sapdata6/pooli_2/pooli.data2 will be overwritten
BR428W File /oracle/EVM/sapdata1/protd_1/protd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata6/proti_1/proti.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/roll_1/roll.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/roll_1/roll.data3 will be overwritten
BR428W File /oracle/EVM/sapdata1/roll_2/roll2.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/sourced_1/sourced.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/sourced_2/sourced.data2 will be overwritten
BR428W File /oracle/EVM/sapdata4/sourcei_1/sourcei.data1 will be overwritten
BR428W File /oracle/EVM/sapdata4/sourcei_2/sourcei.data2 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_1/ssqjd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_2/ssqjd.data2 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_3/ssqjd.data3 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_4/ssqjd.data4 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_5/ssqjd.data5 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_6/ssqjd.data6 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_7/ssqjd.data7 will be overwritten
BR428W File /oracle/EVM/sapdata10/ssqjd_8/ssqjd.data8 will be overwritten
BR428W File /oracle/EVM/sapdata11/ssqji_1/ssqji.data1 will be overwritten
BR428W File /oracle/EVM/sapdata11/ssqji_2/ssqji.data2 will be overwritten
BR428W File /oracle/EVM/sapdata11/ssqji_3/ssqji.data3 will be overwritten
BR428W File /oracle/EVM/sapdata11/ssqji_4/ssqji.data4 will be overwritten
BR428W File /oracle/EVM/sapdata4/stabd_1/stabd.data1 will be overwritten
BR428W File /oracle/EVM/sapdata5/stabi_1/stabi.data1 will be overwritten
BR428W File /oracle/EVM/sapdata5/stabi_2/stabi.data2 will be overwritten
BR428W File /oracle/EVM/sapdata1/temp_1/temp.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/temp_1/temp.data2 will be overwritten
BR428W File /oracle/EVM/sapdata1/userld_1/userld.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/userli_1/userli.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/system_1/system.data1 will be overwritten
BR428W File /oracle/EVM/sapdata1/system_2/system.data2 will be overwritten
BR280I Time stamp 2002-12-23 08.14.22
BR256I Please enter 'cont' to continue, 'stop' to cancel the program:
BR257I Your reply: 'cont'
BR280I Time stamp 2002-12-23 08.14.32
BR259I Program execution will be continued...

BR456I Probably the database must be recovered due to restore from online backup

BR280I Time stamp 2002-12-23 08.14.32
BR407I Restore of database: EVM
BR408I BRRESTORE action ID: rdjpfiau
BR409I BRRESTORE function ID: rsb
BR449I Restore mode: ALL
BR419I Files will be restored from backup: bdjpbne.and 2002-12-22 13.39.06
BR416I 57 files found to restore, total size 91571.031 MB
BR424I Files will not be decompressed
BR421I Backup device type for restore: disk
BR420I Files will be restored from directory: /bkoracle/EVM
BR140I Number of parallel copy processes: 13
BR280I Time stamp 2002-12-23 08.14.32
BR256I Please enter 'cont' to continue, 'stop' to cancel the program:

BR257I Your reply: 'cont'
BR280I Time stamp 2002-12-23 08.14.36
BR259I Program execution will be continued...
BR126I Unattended mode active - no operator confirmation required

BR351I Restoring /oracle/EVM/sapdata6/poold_1/poold.data1
BR355I from /bkoracle/EVM/sapdata6/poold_1/poold.data1 ...

BR351I Restoring /oracle/EVM/sapdata11/ssqji_2/ssqji.data2
BR355I from /bkoracle/EVM/sapdata11/ssqji_2/ssqji.data2 ...

BR351I Restoring /oracle/EVM/sapdata10/ssqjd_3/ssqjd.data3
BR355I from /bkoracle/EVM/sapdata10/ssqjd_3/ssqjd.data3 ...

BR351I Restoring /oracle/EVM/sapdata6/pooli_2/pooli.data2
BR355I from /bkoracle/EVM/sapdata6/pooli_2/pooli.data2 ...

BR351I Restoring /oracle/EVM/sapdata1/roll_2/roll2.data1
BR355I from /bkoracle/EVM/sapdata1/roll_2/roll2.data1 ...

BR351I Restoring /oracle/EVM/sapdata11/ssqji_4/ssqji.data4
BR355I from /bkoracle/EVM/sapdata11/ssqji_4/ssqji.data4 ...

BR351I Restoring /oracle/EVM/sapdata10/ssqjd_8/ssqjd.data8
BR355I from /bkoracle/EVM/sapdata10/ssqjd_8/ssqjd.data8 ...

BR351I Restoring /oracle/EVM/sapdata1/temp_1/temp.data2
BR355I from /bkoracle/EVM/sapdata1/temp_1/temp.data2 ...

BR351I Restoring /oracle/EVM/sapdata4/stabd_1/stabd.data1
BR355I from /bkoracle/EVM/sapdata4/stabd_1/stabd.data1 ...

BR351I Restoring /oracle/EVM/sapdata10/ssqjd_4/ssqjd.data4
BR355I from /bkoracle/EVM/sapdata10/ssqjd_4/ssqjd.data4 ...

BR351I Restoring /oracle/EVM/sapdata9/btabd_2/btabd.data2
BR355I from /bkoracle/EVM/sapdata9/btabd_2/btabd.data2 ...

BR351I Restoring /oracle/EVM/sapdata2/es46cd_2/es46cd.data2
BR355I from /bkoracle/EVM/sapdata2/es46cd_2/es46cd.data2 ...

BR351I Restoring /oracle/EVM/sapdata5/btabi_2/btabi.data2
BR355I from /bkoracle/EVM/sapdata5/btabi_2/btabi.data2 ...

#FILE..... /oracle/EVM/sapdata2/es46cd_2/es46cd.data2
#RESTORED. /bkoracle/EVM/sapdata2/es46cd_2/es46cd.data2 #8/1

BR280I Time stamp 2002-12-23 08.30.56
BR418I 1 of 57 files restored - 1800.008 MB of 91571.031 MB done
BR204I Percentage done: 1.97%, estimated end time: 22:05
BR001I * _____

BR351I Restoring /oracle/EVM/sapdata2/es46cd_1/es46cd.data1
BR355I from /bkoracle/EVM/sapdata2/es46cd_1/es46cd.data1 ...

#FILE..... /oracle/EVM/sapdata1/roll_2/roll2.data1
#RESTORED. /bkoracle/EVM/sapdata1/roll_2/roll2.data1 #13/1

```
BR280I Time stamp 2002-12-23 08.33.23
BR418I 2 of 57 files restored - 3800.016 MB of 91571.031 MB done
BR204I Percentage done: 4.15%, estimated end time: 15:47
BR001I ** _____
:
:
#FILE..... /oracle/EVM/sapdata1/roll_1/roll.data3
#RESTORED. /bkoracle/EVM/sapdata1/roll_1/roll.data3 #4/2

BR280I Time stamp 2002-12-23 09.15.47
BR418I 56 of 57 files restored - 87664.781 MB of 91571.031 MB done
BR204I Percentage done: 95.73%, estimated end time: 9:18
BR001I *****_____

#FILE..... /oracle/EVM/sapdata6/poold_2/poold.data2
#RESTORED. /bkoracle/EVM/sapdata6/poold_2/poold.data2 #1/2

BR280I Time stamp 2002-12-23 09.15.56
BR418I 57 of 57 files restored - 91571.031 MB of 91571.031 MB done
BR204I Percentage done: 100.00%, estimated end time: 9:15
BR001I *****

BR406I End of file restore: rdjpfiau.rsb 2002-12-23 09.15.56
BR280I Time stamp 2002-12-23 09.15.56
BR403I BRRESTORE terminated successfully with warnings
```

appendix c: SAP/Oracle storage growth and impact on HSG80

It is very important to identify the SAP high-volume and high-activity tables and move them into isolated tablespaces. By keeping high-volume tables in separate tablespace(s) and placing the tablespaces on several data files on separate mount points, the administrator can manipulate the data files in the tablespace to minimize I/O contention on the disks.

Because of the high-speed, transaction-oriented nature of most SAP applications, the administrator can see high activity in specific areas of each Oracle/SAP tablespace. This is especially important for the SAP transaction tablespaces PSAPSTABD and PSAPBTABD, where it is common to see the bulk of the I/O activity.

HSTK supports RAID 1 or RAID 1+0. With HSG80 robust RAID 1+0 technology and its use of striping, hot spots can be reduced and I/O performance be improved through parallel access to all members of the array. This can be implemented after performing the following tasks for the SAP system:

- Identify the “hot” and/or large tables and move them into separate tablespaces with multiple data files in the tablespace.
- Review the I/O patterns of these tablespaces.
- Load-balance the I/O for these tablespaces by placing pieces of the tablespaces on several disks.
- Distribute data evenly across many spindles.
- Design the application efficiently with best qualified SQL statements.

SAP/Oracle stores tables and indexes in individual data blocks. When new storage space is required for a table or an index, one or more contiguous data blocks of a data file are allocated to form an extent. Each table and index is assigned to a tablespace, which consists of one or more data files at the operating system level. All table and index data is stored in the data files of tablespace.

SAP/Oracle has many storage parameters FREELISTS, FREELIST GROUPS, PCTFREE and PCTUSED that impact the growth of SAP/Oracle data objects. During installation of an SAP R/3 system, while creating an SAP table or index, the default storage parameters – INITIAL, NEXT and MAXEXTENT values are used. The first extent (INITIAL EXTENT) should be large enough for the expected table or index size. If an extent of a data object becomes full during an insert or update operation, the Oracle storage management system attempts to allocate another extent in the tablespace.

A suitably high value for PCTUSED will effectively reuse space on data blocks. So the PCTUSED relatively full blocks are placed on the freelist. Hence these blocks will be able to accept only a few rows before becoming full again, leading to more I/O.

The presence of chained rows in a table means that PCTFREE is too low or that DB_BLOCKSIZE is too small. To avoid this, a suitably high value needs to be set for PCTFREE and DB_BLOCKSIZE. In most cases within SAP, RAW and LONG RAW columns make huge rows that exceed the maximum block size for Oracle, making chained rows unavoidable. The SAP tables D010L and TST03 are good examples of this chaining problem.

SAP/Oracle tablespace reorganization overcomes the problem of fragmentation to optimize the amount of HSG80 storage space. Extents are merged together to reduce the number of extents in the database and some data files are merged together to reduce the number of data files in the database.

Always perform online reorganization for small tables and export/import for large tables during periods of low Data Manipulation Language (DML) activity (late Saturday nights, for instance). Also ensure rollback segments are big enough for read-consistent export. Additional HSG80 storage space is required to perform a re-organization in the directory \ORACLE\<SAPSID>\sapreorg.

appendix d: SAP/ORACLE backup and recovery strategies

Backup operations constitute some of the most important maintenance tasks that an administrator performs for a 24/7 database. Backups are key for SAP/Database. In a 24/7 scenario, however, they assume even more importance because during an outage, a robust backup strategy allows fast and safe recovery.

SAP/Oracle's backup capabilities include both logical and physical backups. A logical backup of the database involves reading a set of database records and writing them to a file. These records are read independent of their physical location. In Oracle the Export utility is used to perform this type of database backup. To recover using the file generated from an export, Oracle's import utility is used.

SAP R/3's kernel executables – BRBACKUP and BRRESTORE can be used to manage database backups and log file backups. During a database backup, the data files, the online redo log files, profiles, and control file are backed up.

online backup internal operation

The SAP/R3 production system remains available during online backup. Oracle automatically writes to the next online REDO log file in a round robin fashion while the online backup is in progress. In addition to the I/O during backups, online backups impose an additional performance penalty, excessive redo being written. Excessive redo is generated during online backups while the tablespaces are in backup mode (via the ALTER TABLESPACE ... BEIGN BACKUP command). In order to minimize the amount of redo generated, to reduce the adverse impact on performance, and to increase overall logical consistency (where all internal database objects are relatively current, with respect to each other's System Change Number (SCN)), backups need to be scheduled during relatively low periods of Data Manipulation Language (DML) activity. Once the checkpoint SCN is written, the data-file header will not be updated with subsequent checkpoints until the ALTER TABLESPACE ... END BACKUP command is issued.

During a "fast" checkpoint, DBWR continuously scans for dirty buffers and writes them to disk as opposed to during a regular checkpoint, where DBWR stops scanning if the number of buffers scanned is equivalent to the DB_CHECKPOINT_WRITE_BATCH initialization parameter or more than a thousand buffers are scanned and none among those are dirty. A regular checkpoint allows better balance among all Oracle processes in sharing system resources (such as CPU cycles). The check pointing process that runs during the execution of the ALTER TABLESPACE ... BEGIN BACKUP command flushes all the dirty buffers that belong to the data files in the tablespaces and ensures that only blocks that are changed during the online backup are written to the redo log file.

backup objects for SAP R/3

SAP on Tru64 UNIX integrates the Oracle backup/recovery capabilities of the following SAP Tools with the backup and recovery management.

- The program BRBACKUP brings database servers online or offline, checks the status of SAP files, and places database tablespaces into BACKUP mode to guarantee their data consistency. BRBACKUP provides online or offline backup of control files, data files, and online redo log files. It also keeps a profile and log of each backup.
- The program BRARCHIVE archives Oracle off-line redo log files by communicating with the backup software for SAP on Tru64 Unix backint interface. These files are copied by Oracle in its archiving directory. BRARCHIVE ensures that duplicates of these logs are available and that original files are not deleted prematurely. This command also keeps a profile and log of each archive.
- The program BRRESTORE recovers database data files, control files, and on-line redo log files through backup software for SAP on Tru64 Unix backint interface, BRRESTORE ensures that sufficient space is available prior to restoring these files, and removes files that will be overwritten during the recovery. This command also provides a query mode.

Figure 24 illustrates the components and architecture for the SAP/Oracle backup solutions.

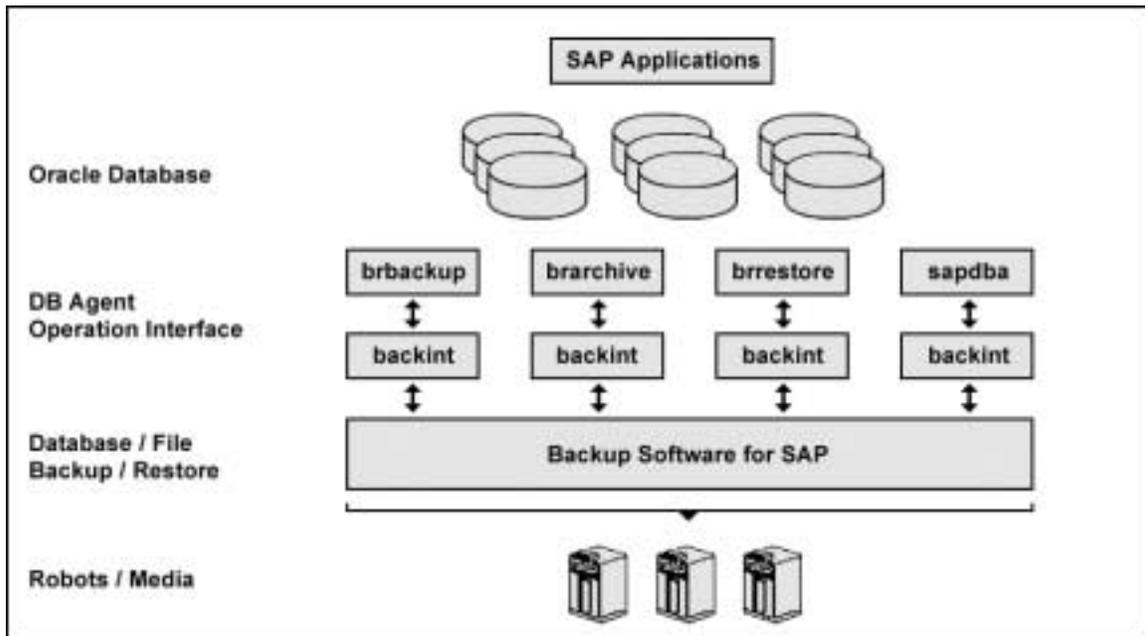


Figure 24, SAP/Oracle Backup overview

SAP Tools act as database agents, responsible for all database-related tasks. These database agents (BRBACKUP, BRARCHIVE and BRRESTORE) communicate with backup software through the backup agent software for SAP on Tru64 Unix backint interface. The SAPDBA component of SAP Tools also accesses the backup catalog used by software in order to determine the status of the valid backups.

Brbackup using Legato Networker

Figure 25 shows an example of Brbackup, Legato Networker sequence operation. NetWorker client-server technology uses a remote procedure call (RPC) network protocol to back up data. The NetWorker server software consists of several server-side services and programs that oversee backup and restore processes. The NetWorker client software consists of client-side services and user interface programs. These services and programs maintain:

- Client configuration files.
- An online client file index.
- An online media database.

NetWorker calls upon several services and programs when a request for a backup is received. These services coordinate the tasks associated with a backup or restore, and record information about what was backed up and the media containing the backed-up data.

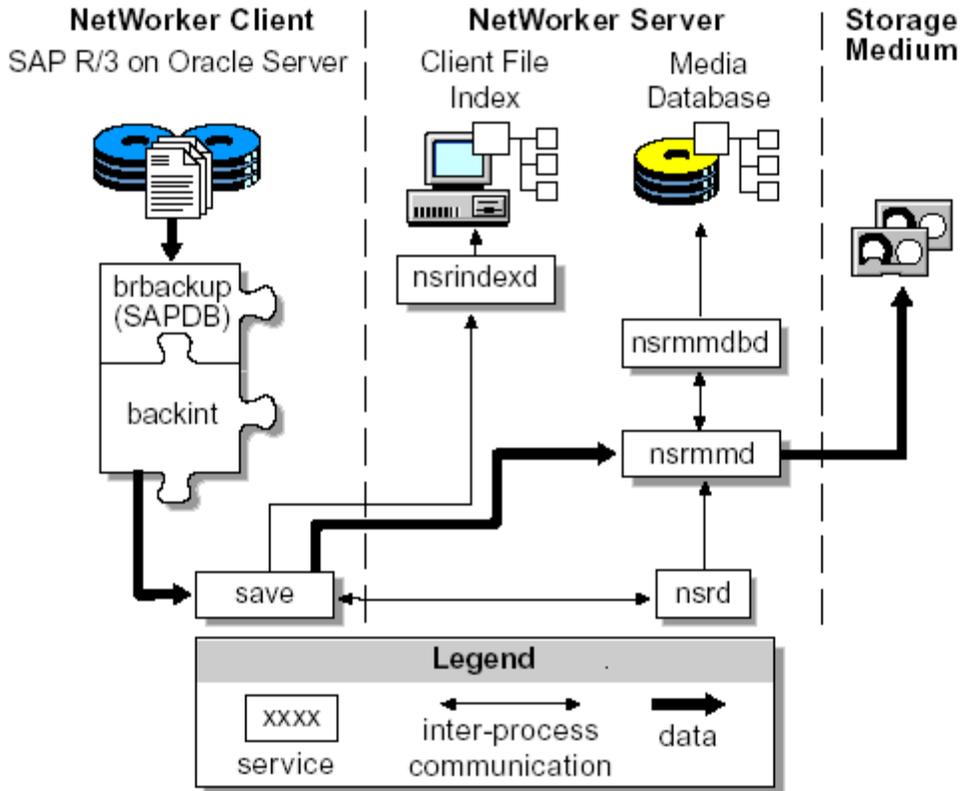


Figure 25, Brbackup using Legato Networker

For a backup, BRBACKUP calls the Networker for SAP on Tru64 Unix NT backint interface. From this point the operation is similar to a user-directed backup.

Brestore using Legato Networker

When NetWorker receives a request from a client to restore data, the server's **nsrd** service contacts the server's media database service, **nsrmmdbd**, to determine which media contains the requested saveset. Once the save set's media location is obtained, **nsrmmdbd** issues a mount request. The media is positioned at the beginning of the save set, and the save set is passed to **nsrmmdbd**. The media service then forwards the save set to the client's **recover** program, which restores the data to the client's file system.

When the server's media database service, **nsrmmdbd**, is unable to locate the required volumes or when there are no enabled devices on the NetWorker server, the restore request fails. Figure 26, shows how the services and programs managed by the NetWorker client and server interact while restoring data to a NetWorker client.

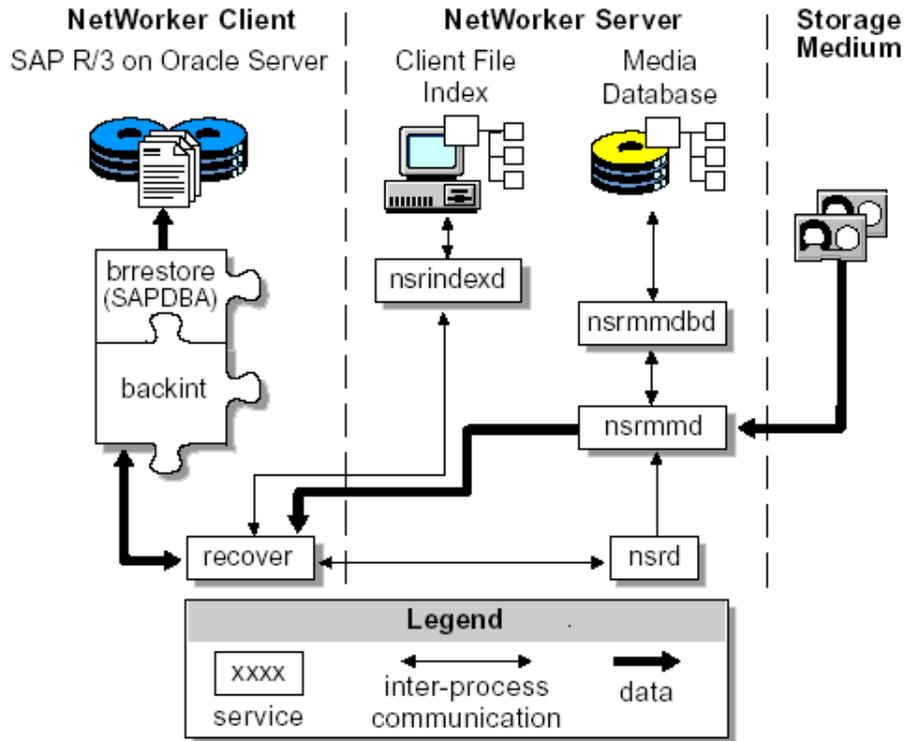


Figure 26, Brestore using Legato Networker

appendix e: HSTK installation overview

As of V1.60 of the "HSG Scripting Tool Kit", an installation utility is provided. This utility performs a number of steps required to run the "HSG Scripting Tool Kit", for you.

After unpacking the compressed tar file (the *.tar's file) into a directory, you will see the following:

```
# ls -l
total 38
-rw-r--r--  1 root    other    17176 Oct 18 13:09 install.pl
drwxr-xr-x  4 root    other     512 Oct 10 10:45 modules
drwxr-xr-x  9 root    other     512 Oct 18 13:11 scripts
```

Simply type in "perl install.pl" and you'll be asked some self-explanatory questions. The program will do the following:

Create the required directory tree underneath the target directory. If the "HSG Scripting Tool Kit" has been installed before, it will use that directory tree. It will probe for the CLONE_HOME environment variable.

1. Find out the Perl Include path.
2. Copy the Perl modules to its final destination.
3. Copy all other files to their final destination. Note that the following files are NOT copied:
 - ./config/application.act
 - ./bat/set_env.bat
 - ./sh/set_env.sh
 - ./vms/set_env.com
4. Copy the HP StorageWorks Command Scripter executable if it can be found. On Unix platforms, it will create a symbolic link from bin/cmdscript to the directory where HP StorageWorks Command Scripter is installed.
5. It will generate the set_env file for the Operating System of the server you're currently using.
6. Install the Telnet.pm Perl module (preparation for out-of-band communication).

Update any already existing "application action list" files in the \$CLONE_HOME/config directory. It will only update those entries that relate to "HP StorageWorks Data Replication Manager"

The following table is an example of the "install.pl" program.

```
# perl install.pl

Hi there !
Thanks for your interest in the HSG Scripting Tool Kit V1.60

This configuration utility installs the files for you in /usr/local/hstk
Do you want to change this directory "y/n" or type "q" to quit :n
Creating Directory /usr/local/hstk ..... OK
Determining Perl Include Path ... OK
Copying Perl Modules ..... OK
Copying Files
.....
..... OK
Finding StorageWorks Command Scripter ... OK
Installing Perl module Telnet.pm ... OK
Writing environment file ... OK
Updating existing act files .... OK
Installation of HSG Scripting Tool Kit successfully ended
```

A logfile called install.log is generated which lists all files copied, as well as the protection set on these files.

references

Oracle 24 x 7 Tips and Techniques by Venkat S. Devraj, Oracle Press
Oracle SAP Administration by Donald K. Burleson O'Reilly press
Legato Networker for SAP System Administrator Guide
Oracle 8 DBA handbook by Kevin Loney Oracle Press
SAP OSS Note 511972 Split Mirror Backup with HP Storage

more information

For more information on HP's solution for SAP, please refer to:

www.hp.com/go/SAP

www.hp.com/go/storage

notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL HP BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE, OR OTHER DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for HP products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solutions. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

©2003 Hewlett-Packard Company
3/2003 16VF-0802A-WWEN
1st edition