

March 2002

Part Number:
16DL-0302A-WWEN
Edition 1

Prepared by:
Global SAP Solutions and
Enterprise Storage Group
Compaq Computer Corporation

Contents

Introduction..... 3
Testing Environment..... 6
**SAN Setup and
Configuration 9**
**EVM – SAP Network
Operation..... 14**
Solution Planning and Sizing .. 25
**SAP/EVM Backup Process
Flow 36**
**Online and Offline Backup
Process..... 37**
Restoring Lost Files 44
**Appendix A: EVM-SAP
Backup Scripts..... 45**
**Appendix B: Oracle
Architecture on Windows
2000 for SAP 51**
**Appendix C: SAP/Oracle
Storage Growth and Impact
on HSG80 54**
**Appendix D: Data Flow from
Oracle to SAP..... 56**
**Appendix E: SAP /ORACLE
Backup and Recovery
Strategies 57**
**Appendix F: HSG80 Volume
Layout for SAP R/3-ORACLE... 60**
References 63

StorageWorks Rapid Restore and Non-Disruptive Backup for mySAP.com (Windows 2000 / Oracle)

Abstract: This document contains information on how to implement non-disruptive backup and rapid recovery for mySAP.com systems using StorageWorks Modular Array disk subsystems. This version of the white paper describes the solution for Microsoft® Windows 2000 and Oracle®.

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 COMPAQ 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 Compaq 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 solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

Compaq, NonStop, Deskpro, Compaq Insight Manager, Systempro, Systempro/LT, ProLiant, ROMPaq, QVision, SmartStart, NetFlex, QuickFind, PaqFax, and Prosignia are registered with the United States Patent and Trademark Office.

ActiveAnswers, Netelligent, Systempro/XL, SoftPaq, Fastart, QuickBlank, QuickLock are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

The following are trademarks or registered trademarks of SAP AG; ABAP/4, InterSAP, RIVA, R/2, R/3, R/3 Retail, SAP (Word), SAPaccess, SAPfile, SAPfind, SAPmail, SAPoffice, SAPscript, SAPtime, SAPtronic, SAP-EDI, SAP EarlyWatch, SAP ArchiveLink, SAP Business Workflow, and ALE/WEB. The SAP logo and all other SAP products, services, logos, or brand names included herein are also trademarks or registered trademarks of SAP AG.

Intel, Pentium and Xeon are trademarks and/or registered trademarks of Intel Corporation.

Oracle is a registered trademark of Oracle Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©2002 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

StorageWorks Rapid Restore and Non-Disruptive Backup for mySAP.com (Windows 2000 / Oracle)
White Paper prepared by Global SAP Solutions

Edition 1 (March 2002)

Overview

This solutions guide provides information on configuring Enterprise Volume Manager (EVM) as used in StorageWorks Rapid Restore and Non-Disruptive Backup for mySAP.com (Windows 2000/Oracle).

Compaq Links

For further information on configuring the Enterprise Volume Manager

- EVM Installation Guide
- EVM Network Administration Guide
- Online Help/User Guide (accessible via the Web browser interface) for Windows 2000
- Compaq SANworks Management Appliance Patch Installation Instructions.
- Compaq SANworks Enterprise Volume Manager Version 2.0 (Multi-OS) OSM/MA Supplement.

SAP Links

- Split-Mirror Disk Backup for Oracle
- Database Administration (Oracle) with SAPDBA in SAP Service
- Refer the SAP OSS-Note 511972 “Oracle/W2K split mirror backup with Compaq storage”

Introduction

Enterprise Volume Manager (EVM) Version 2.0 (Multi-OS) is a Web-enabled application that allows you to establish a network of computers that can participate in automated storage unit replication jobs using Compaq® StorageWorks Fibre Channel subsystems. EVM networks and their associated StorageWorks subsystems can be managed remotely from a central location using a Web browser.

EVM provides two methods of copying storage units: clone and snapshot. All of the underlying StorageWorks Console Line Interface (CLI) commands that are required for these methods are automated and incorporated into EVM.

Through an easy-to-use Web browser interface, EVM provides the means to create, manage, and run storage replication jobs on these subsystems. EVM jobs provide the ability to create, run, validate, monitor, review, and delete jobs. EVM Resources provides graphical displays of the available resources on the EVM network.

EVM provides a simple, yet powerful set of features that allow you to design storage unit replication jobs that can be seamlessly combined with external applications. For example, you can create:

- EVM database jobs that create nearline backup units on the subsystems.
- EVM database backup jobs that copy online databases and mount them on backup servers for automated tape backups.
- EVM database jobs that copy databases and mount them on other servers for data mining or testing.

The SAP® utilities BRBACKUP, BRARCHIVE, and BRRESTORE have to be configured before they can be used. To do this, you need to set the appropriate parameters in the initialization profile `init<SID>.sap` to use clone or snapshot jobs created by EVM for splitting, mounting and resynchronizing the Business Continuance Volume (BCV). The SAP backup (BRBACKUP) process is started on the backup host; the backup load is therefore moved from the SAP database production server to the backup server and thus does not affect the performance of the SAP production server. The BRBACKUP process provides all the information to the production database server. So this allows the user to use SAPGUI Computer Center Management System (CCMS) to monitor the online backup on the production server. The EVM command line (EVMCL) allows you to execute EVM jobs (clone, snapshot) with different combinations of parameters like `execute`, `validate`, `detail status` and `undo`. EVMCL returns with an error level, which can be used to control and monitor the EVM snapshot or clone job.

SAP does not provide the same administration and backup tools for non-Oracle databases. If other databases e.g., MS SQL are used, then a different EVM split mirror solution is implemented.

This document is designed to be used as a guideline when organizing your Oracle database backup and administration strategy.

Background

As part of the backup strategy, SAP recommends a daily backup of the entire mySAP.com production databases. However, this may be infeasible when the database is several hundred gigabytes (GB) to over one terabyte (TB) in size due to following concerns.

Server Performance: Backing up a large database causes production server performance problems as the backup process consumes database resources. In particular, CPU time, I/O bus, system bus, and disk controllers become saturated. As a result the performance of the entire mySAP.com system decreases unpredictably.

Server Availability: Normal backup operations should be performed when there is little or no activity on the mySAP.com database system (the so-called 'time window', usually at night). The timeframe to complete the backup easily becomes too short in the case of large mySAP.com databases.

Network Performance: Backing up a large mySAP.com database from the production host via network is not an option since it may lead to further loss of performance.

Compaq's Enterprise Volume Manager addresses these issues with Business Continuance Volumes (BCVs). BCVs mirror the volume where the production data files reside. They can also be split off and mounted on the backup host, and BCV volumes can resynchronize as required without loss of database performance.

Instead of backing up the production database on the production server, BRBACKUP is started on a backup server and it backs up the database files of the split mirror disks. The backup load is thereby moved from the production server to the backup server and from the production spindles to the cloned spindles and therefore does not affect the performance of the mySAP.com system during the backup process. BRBACKUP can be used to control the splitting and later synchronization of the disks.

This solution can also easily be used to create SAP database system "clones" for data mining and extraction operations.

General Recommendations

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, if not the most, CPU, memory, and disk I/O intensive application in the industry. Because the disk subsystem is the slowest part of an R/3 system it is extremely important to design it for optimal performance. For an SAP R/3 server, the disk subsystem for the SAP/Oracle databases is commonly designed using RAID fault-tolerant settings.

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

We recommend that you create a three-member striped mirror set (a striped set whose members are mirror sets). 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 kept for a day depending upon your business needs. This way, the data can be restored rapidly.

SANworks Enterprise Volume Manager can be used to create BCV clones of hardware-based, mirrored (RAID 1) or stripe-mirrored (RAID 1+0) volumes.

For security and performance results, database data files and database online redo log files should reside on different disks. As the online redo logs are written synchronously, they produce the most I/O activity of all database files.

The overall performance of the disk subsystem depends heavily on the number of disk drives installed. The more drives are installed in a volume (RAID 1+0, RAID 5), the more I/Os per second and the higher the transfer rate the volume is able to provide.

The performance of the disk subsystem can benefit from having highly sequential accessed data implemented on the fastest rotating drives. In an average mySAP.com system, 40-50 % of access to the database consists of sequential table scans.

We strongly recommend that you use the default chunk size value for your particular controller configuration. Defaults are chosen to produce optimal performance for a wide variety of I/O loads. The strip size default for HSG controllers is 256 blocks.

One can compute the total number of required I/Os for a given RAID level and subsequently determine the number of drives necessary to provide the required number of I/Os.

RAID Level	I/O Per Transaction
RAID 0	(Reads + writes) *transaction/sec
RAID (1 + 0)	(Reads+2*writes)*transaction/sec
RAID 5	(Reads+4*writes)*transaction/sec

Table 1, I/O Per Transaction

Cache Policies

In a mySAP.com system, cache or buffers exist on several different levels, where the lowest level is on disk/disk controller level and the highest is the R/3 buffers in the application server.

From a mySAP.com system point of view, the most important caches are the application caches, file system caches, controller caches, and device caches.

Modern disk controllers like Compaq's HSG80 are able to detect multiple threads or streams of sequential access. Therefore, it is important to enable them to read ahead from several areas on the disk. HSG80 caching technologies usually consist of high-speed memory that minimizes the queues in a system. At the system level, caching is used in the CPUs to reduce access to main memory, which helps with memory contention in multiprocessor systems.

At the controller and disk level, another level of cache is used to increase data transfer to and from disk media. Disk caching affects performance in the following areas:

- It reduces disk access (with cache hits).
- It reduces the negative effects of the RAID overhead.
- It assists in disk I/O request sorting and queuing.

Read cache from the HSG80 controller is beneficial in SAP R/3 application environments. It improves performance during sequential read operations.

Write-back caching from the HSG80 controller enables the application to post the write request in the controller cache, which immediately responds with a completion status. The data is written to the disk drives later. Write cache is beneficial in SAP R/3 environments, where the I/O profile includes random write requests. These write requests can be posted in the cache, increasing overall system performance.

SANworks SecurePath Notes

The software component SANworks SecurePath installed on the host in addition to the above-mentioned hardware enables Windows 2000 to fail over to the other path in the event of a hardware failure in a host bus adapter, switch, array controller or fiber cable. This ensures continuous access to the storage. In addition to providing a redundant access path to the storage, the additional hardware facilitates load balancing over the two access paths from server to storage.

Testing Environment

The following section describes the SAP Oracle Windows 2000 backup/restore solution and includes the hardware, software, and firmware components. Special configuration rules are explained where necessary.

Bill of Materials

Table 2 provides a bill of materials that includes some of the products used in the SAP/Oracle non-disruptive backup/restore solution. The part numbers in the templates were valid at the time of testing. You can choose the products best suited to your needs, including products not listed here.

Item	Item Description	Base SKU
1	EVM 2.0 Software for Window 2000-Intel	160094-B21
2	ACS 8.6P Controller Software	128697-B21
3	Operating System Platform Software kit for NT/Windows2000 (SWCC)	380551-001
4	EMA 12000 with 2 Chassis Four HSG80 Controllers	380620-001/002
5	Hard Disk Drives, 18.2 GB, 10K UWSE	380589-B21
6*	ProLiant NT Backup Server (PL6500 – typical)	153554-001
7	Fibre Channel HBA (switched fabric), KGPSA-BC	380574-001
8	Fibre Channel Storage 2 SAN Switches, 16 port	380591-B21
9	FC Optical Short Wave GBIC	158222-B21
10	Fibre Channel Optical Cable	234457-B22
11	Modular Data Router	163082-B21
12	Super DLT 110/220 Tape Library	349350-B24
13	VERITAS® NetBackup/Computer Associates ARCserve 2000	Third-party P/N
14	SANworks Management Appliance	189715-001
15	Secure Path 3.1A	165989-B21

Table 2, Bill of Materials for SAP/Oracle Backup Solution

Table Explanation

Item 1: EVM Software – EVM software is necessary on every client and server on the SAN that needs to communicate or coordinate communication with hosts or storage devices. EVM Server is installed on the SAN Appliance Manager. EVM Agent is installed on the dedicated DL380 backup server and on the Production Database server.

Item 2: Controller Firmware – One instance of ACS is necessary 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 supports cloning and snapshots as well.

Item 3: OS Platform Kit CDROM – “StorageWorks Solutions Software RA8000/MA8000 Fibre Channel Solution Software NT, Windows2000, ProLiant®, x86” includes drivers and installation files for the specific operating system and host involved.

Item 4: Disk Array – For our tests we used EMA12000/4214 and MA8000/4254, each with two HSG80 controllers. On both subsystems, the controllers are configured in MultiBus Failover mode.

Item 5: Disks – The choice of disk type does not affect the functionality of this solution. Two factors to consider for disk type are rpm and the number of spindles. For applications like SAP/Oracle, the access pattern is random. Also, it is desirable to use as many spindles as possible to increase parallel access to the volume as opposed to fewer disks. Presently, the largest number of total sets that EVM can address is 30. In our test we created six mirror sets, each with three members, and then striped them together to form one logical unit.

Item 6: Backup Server – The backup server must be powerful enough to maintain performance levels during a backup/restore operation. For our tests, we used a single ProLiant DL380 Server.

Item 7: Fibre Channel Host Bus Adapter (HBA) – The Windows EVM host agent supports the following Fibre Channel (FC) host bus adapters (HBA): Emulex LP8000 or LP952 with miniport driver version 4.52a8. The 4.52a8 driver update kit can be downloaded from <http://www.compaq.com/products/storageworks/adapters/emulexdownload.html>. Support for the Emulex port driver will be available in a future release of EVM.

Item 8: Fibre Channel Interconnect – SAN interconnect switches allow for significantly better performance, whereas loops (hubs) present significantly limited performance and connectivity options. However, since they are newer, switches currently have limited supported configurations where loops are supported by a variety of components. 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.

Item 9: FC Optical Short Wave GBIC – GigaBitInterConnect (GBIC) – GBICs don't come with the hubs or switches and must be ordered separately. SAN switches require the GBIC, 380561-B21.

Item 10: Fibre Channel Optical Cable – Order the length needed. The basic part number is 234457 and the suffix B2x denotes the length in meters. Suffix B21 is 2, B22 is 5, B23 is 15, B24 is 30, and B25 denotes a 50-meter cable. Each controller comes with one five-meter cable. All other cables will have to be ordered separately.

Item 11: Modular Data Router -- Compaq StorageWorks Modular Data Router (Fibre Channel Tape Controller). SCSI attached tape libraries require a 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.

Item 12: MSL5026 Super DLT 110/220 Library – Compaq SDLT 110/220, Compaq 35/70 GB DLT tape drives

Item 13: Third-Party Backup Software – Computer Associates ARCserve 2000 or VERITAS NetBackup.

Hardware Configuration Overview

Compaq ProLiant Server

- Database server (PL 8500)
- Application server (DL360, DL380)
- Backup server (DL380)

RAM

4096 MB: Database server
2048 MB: Application servers
512 MB: Backup server

Software Configuration

- Windows 2000 Service Pack 2
- Windows 2000 Resource kit
- Fibre Channel Solution Software
- Secure Path 3.1A
- Oracle Database Version 8.1.7.2.3
- SAP R/3 Version 4.6C

- VERITAS NetBackup Version 3.4
- SAP SSQJ Tool
- Standard SAP Benchmarks

Storage Subsystems

Symbios Logic 875XS D, 2280X PCI SCSI Adapter	Compaq 18.2 GB 10,000 RPM Ultra3 SCSI (internal hard drives) Compaq 36.4 GB 10,000 RPM Ultra3 SCSI (internal hard drives)
Compaq KGPSA-BC, Fibre Channel Host Bus Adapter	Compaq StorageWorks SAN Switch 16 x 2
(2) Compaq StorageWorks HSG80 Array Controller ACS v.8.6P	(2) Compaq StorageWorks RAID Array MA8000 Compaq DS-RZ1DD-VW 18.2 GB 10,000 RPM UltraSCSI 16-bit
Compaq StorageWorks Modular Data Router (Fibre Channel Tape Controller)	Fibre Channel to SCSI tape controller with (2) 2-port SCSI modules
Compaq StorageWorks MSL 5026 Tape Library	Compaq Super DLT 110/220 tape drives

MSL5026 Tape Library Notes

Two MSL5026 tape libraries are configured with a pass through mechanism and connected to a modular data router (MDR) using a SCSI connection. The MDR is connected to a SAN-switched via Fibre. Using this configuration we obtained the best backup/restore throughput and reduced the backup time. The use of tape libraries in a SAN environment provides the freedom to locate the libraries and the media in a different location than the storage server or the disk storage components with the application data. It also offers the possibility to move application data directly from the HSG80 storage system to the tape storage system without using the traditional communication networks.

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 1, the SAN Fabric solution simplifies management by expanding the number of storage systems to be managed and by facilitating deployment of storage among servers. It also provides far richer data protection options than traditional storage. FC-SW allows multiple pairs of nodes to communicate with each other simultaneously. As more nodes are added, the aggregate throughput capability can increase incrementally.

Storage Subsystem Configuration Notes

- Fibre Channel Loops
- Uses Compaq Secure Path for Windows 2000
- Two subsystems have dual HSG80 controllers configured in multiple-bus failover mode
- Created 8 LUNs (D1-D8) with 2 active controller host ports.

- Set the units' preferred path. Use the preferred path unit parameter to assign units to specific controllers. Units should be distributed equally across controllers for load balancing. Below is an example for one subsystem. This should be repeated for the other subsystems.
- Set preferred units D1-D4 to controller A
- Set preferred units D5-D8 to controller B

Host Bus Adapter Configuration Performance Considerations

As shown in Figure 1, the configuration 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.

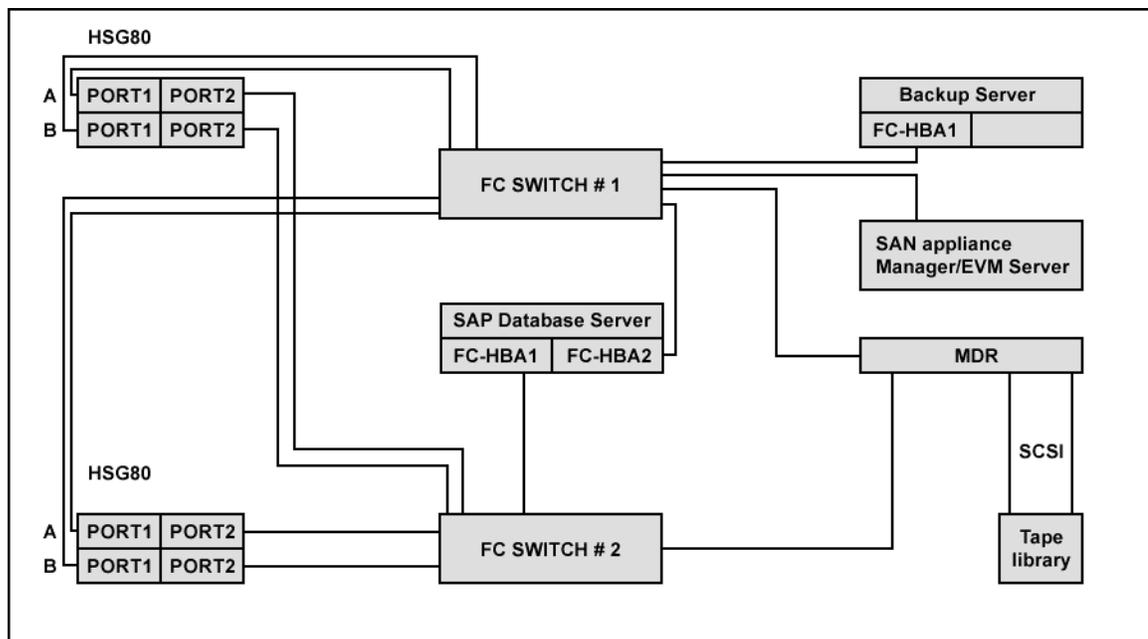


Figure 1, Fibre Channel Configuration

EMA12000/MA8000 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 purpose. 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.

Use a serial terminal (VT100) connected to the first HSG80 serial port to configure the controller like mentioned in the following paragraph.

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 procedures.

MultiBus FailOver Configuration for EMA12000/4214

1. Install cache memory, 512 MB, for best performance.
2. Remove the program card ESD cove, and insert the controller's program card. Replace the ESD cover.
3. Turn on the power to the enclosure by turning on the AC input boxes.
4. Establish a local connection to the controller. The maintenance port provides a convenient way to connect a PC or terminal to the controller so that one can troubleshoot and configure it. This port accepts RJ-232 jacks default values of 9600 baud using 8 data bits, 1 stop bit, and no parity.
5. If the controller reports a node ID of all zeros (0000-0000-0000-0000), set the subsystem worldwide name (node ID) to the worldwide name that came with your subsystem.
6. SET THIS_CONTROLLER SCSI_VERSION=SCSI-3
7. Set the port topology for each port.
SET THIS PORT_1_TOPOLOGY=FABRIC
SET THIS PORT_2_TOPOLOGY=FABRIC
SET OTHER PORT_1_TOPOLOGY=FABRIC
SET OTHER PORT_2_TOPOLOGY=FABRIC
8. SET MULTIBUS_FAILOVER COPY=THIS_CONTROLLER
9. SET THIS_CONTROLLER MIRRORED_CACHE
10. SET THIS_CONTROLLER TIME=DD-MMM-YYYY: HH:MM:SS
11. Set up the battery discharge timer:
 - a) Start Frutil using the following command: Run FRUTIL
Frutil displays the following: Do you intend to replace this controller's cache battery?
Y/N
 - b) Enter Y(es).
Frutil displays a three-step procedure and prompts you to press return.
 - c) Press return.
12. Display details about the controller you configured. Use the following command.

```
ESG-W2K-B-BOT>Show This controller Full
Controller:
    HSG80 ZG95005283 Software V86P-1, Hardware E11
    NODE_ID = 5000-1FE1-0007-A620
    ALLOCATION_CLASS = 0
    SCSI_VERSION = SCSI-3
    Configured for MULTIBUS_FAILOVER with ZG95114434
    In dual-redundant configuration
    Device Port SCSI address 6
    Time: 16-JAN-2002 10:35:16
    Command Console LUN is LUN 0 (NOIDENTIFIER)
Host PORT_1:
    Reported PORT_ID = 5000-1FE1-0007-A621
    PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address = 011800
Host PORT_2:
    Reported PORT_ID = 5000-1FE1-0007-A622
    PORT_2_TOPOLOGY = FABRIC (fabric up)
    Address = 011C00
    NOREMOTE_COPY
Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
    CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
Mirrored Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
Battery:
    UPS = DATACENTER_WIDE
```

Note: The same MultiBus failover configuration procedure to be repeated for MA8000/4254

CLI Procedure

The steps that follow show how to configure RAID 1+0 LUN on EMA12000/4214 and MA8000/4254 each with two HSG80 Controllers. In our test we created the following RAID configuration for SAPDATA1 to SAPDATA6, straight mirror sets, RAID 1+0, for redo log, SAP kernel, SAP archive, and mirror sets RAID 1.

The approach is the same for the other SAPDATA, from SAPDATA 2 to SAPDATA6. Please refer to Appendix F for further details.

Configuring Storage Sets

Following is an example showing how to create striped mirror sets (RAID 1+0)

Create Mirror Sets

```
HSG80>ADD MIRRORSET M1A DISK60000 DISK50000 DISK40000
HSG80>ADD MIRRORSET M1B DISK30000 DISK20000 DISK10000
HSG80>ADD MIRRORSET M1C DISK60100 DISK50100 DISK40100
HSG80>ADD MIRRORSET M1D DISK30100 DISK20100 DISK10100
HSG80>ADD MIRRORSET M1E DISK60200 DISK50200 DISK40200
HSG80>ADD MIRRORSET M1F DISK30200 DISK20200 DISK10200
```

Create Stripe Set

```
HSG80>ADD STRIPESET R1M1 M1A M1B M1C M1D M1E M1F
```

Init Stripe Set

```
HSG80>INITALIZE R1M1 SAVE_CONFIGURATION
```

Create Unit

```
HSG80>ADD UNIT D1 R1M1 WRITEBACK_CACHE
```

Use the above procedure to create additional (RAID 1+0) units for SAPDATA

HBA Connections

Install KGPSA driver on server

Restart server

Show Connections

```
HSG80>Show connections
```

Rename connections to match OS and any unique identification methods (e.g. SAPSERVER1, SAPSERVER2) using "rename !newconXX"

Confirm connection OS type is correct. If necessary change with

```
HSG80>set SAPSERVER1 Operating system=WinNT
```

The same should be repeated for other connections.

Note: Bringing servers online one at a time facilitates the easy identification of connections to MA8000

Add Connections to Appropriate Units

```
HSG80>Set d1 enable=SAPSERVER1, SAPSERVER2, EVMSERVER
```

The same should be repeated for other units also.

EVM – SAP Network Operation

An EVM network is a network that consists of EVM host computers connected to a common LAN and to a Fibre Channel SAN of the EVM-compatible storage subsystem. EVM consists of one server and an unlimited number of clients.

EVM server software (on a Management Appliance) manages the EVM network and EVM jobs, and provides the Web browser interface. The actual production computers that are involved in EVM jobs are called EVM hosts. Storage subsystems are the other major components of an EVM network. Storage subsystems provide all of the storage units (disks) that are used in EVM jobs.

In its most compact configuration, an EVM network can consist of EVM server software (on a Management Appliance) that is connected to a storage subsystem. For most typical production environments, however, EVM networks include one or more EVM host computers.

An EVM network can contain multiple storage subsystems. Although an EVM server can manage an unlimited number of subsystems, there can be EVM performance and responsiveness issues when the number of storage subsystems becomes large. Please refer to the EVM 2.0 Release Notes for more information. EVM is compatible with Transparent and MultiBus Failover configurations.

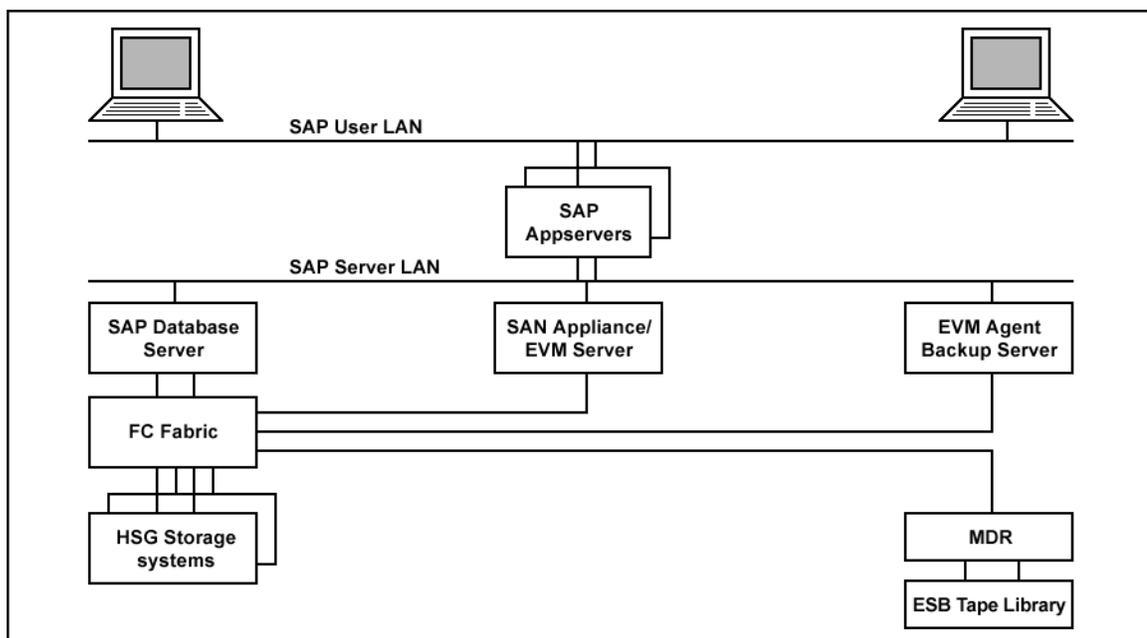


Figure 2, Sample EVM Network

A SANworks Management Appliance is an important component in an EVM network. The EVM server runs on a Management Appliance. If the EVM server stops, there is no way to browse to the EVM network or to run EVM jobs. EVM communicates with storage subsystems via the HSG Element Manager that runs on the Management Appliance. If the HSG Element Manager stops, EVM cannot display correct information and cannot run EVM jobs.

Figure 2 shows the current SAN configuration for an SAP-specific EVM network. The main SAP system is connected to the StorageWorks disk subsystem via an FC switch to the SAN.

Depending on the operating system being used, the StorageWorks unit may be configured with Transparent or MultiBus Failover. Windows NT and Windows 2000 support both failover modes.

The Windows solution discussed in this paper requires CCL to be enabled. EVM Agent needs to be installed on the SAP Database System.

EVM Server will be installed on the SAN Appliance Manager, which enables you to create and modify EVM jobs. EVM Agent needs to be installed on the backup server. The backup server operates and controls the tape library connected to the SAN via Modular Data Router or directly attached via SCSI. By installing EVM Agent, you can execute EVM jobs from the command line on the backup server. EVM Agent submits the jobs to EVM Server for further processing. The third-party backup software needs to be installed on the backup server.

Please refer to the EVM installation guides.

The Compaq SANworks Management Appliance and Open SAN Manager (OSM) provide the starting point for accessing the EVM graphical user interface. A normal Internet browser like Internet Explorer can be used to connect to the Management Appliance.



Figure 3, Management Appliance Splash

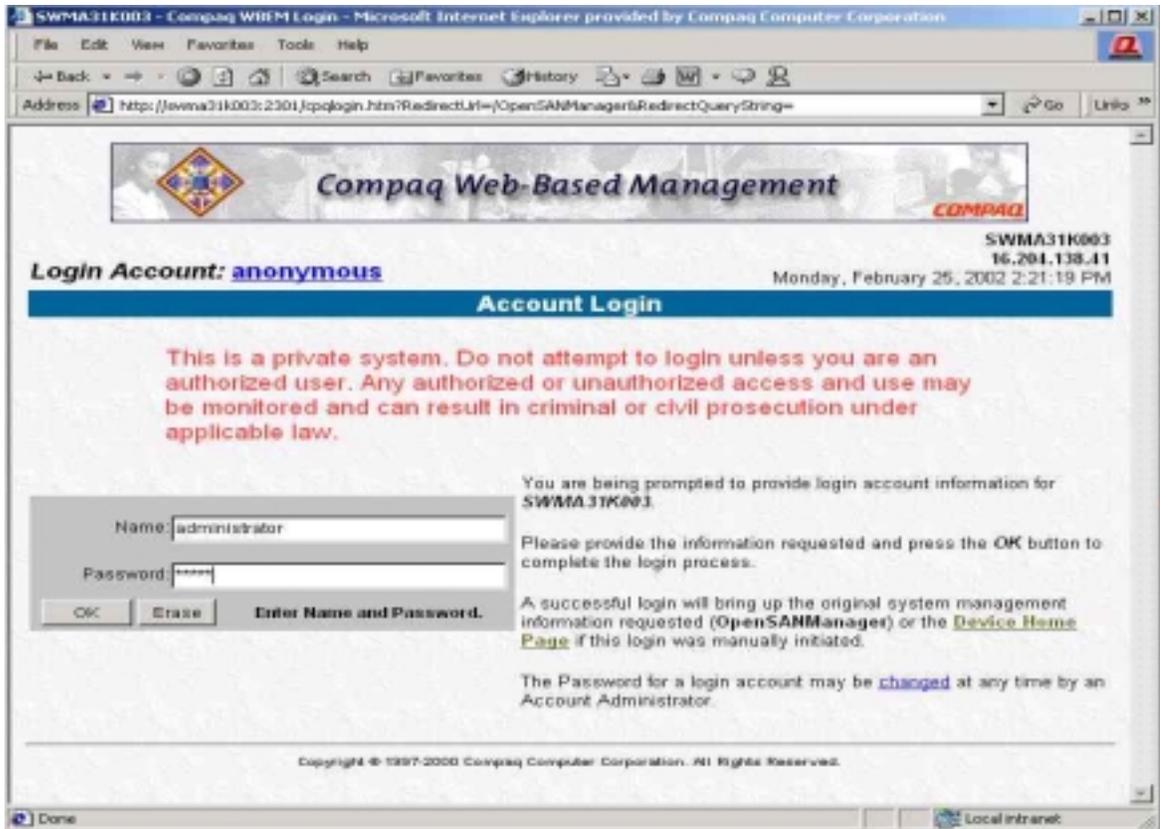


Figure 4 Open SAN Manager Login Screen

The EVM Launch Page summarizes key features and allows you to start the EVM graphical user interface.

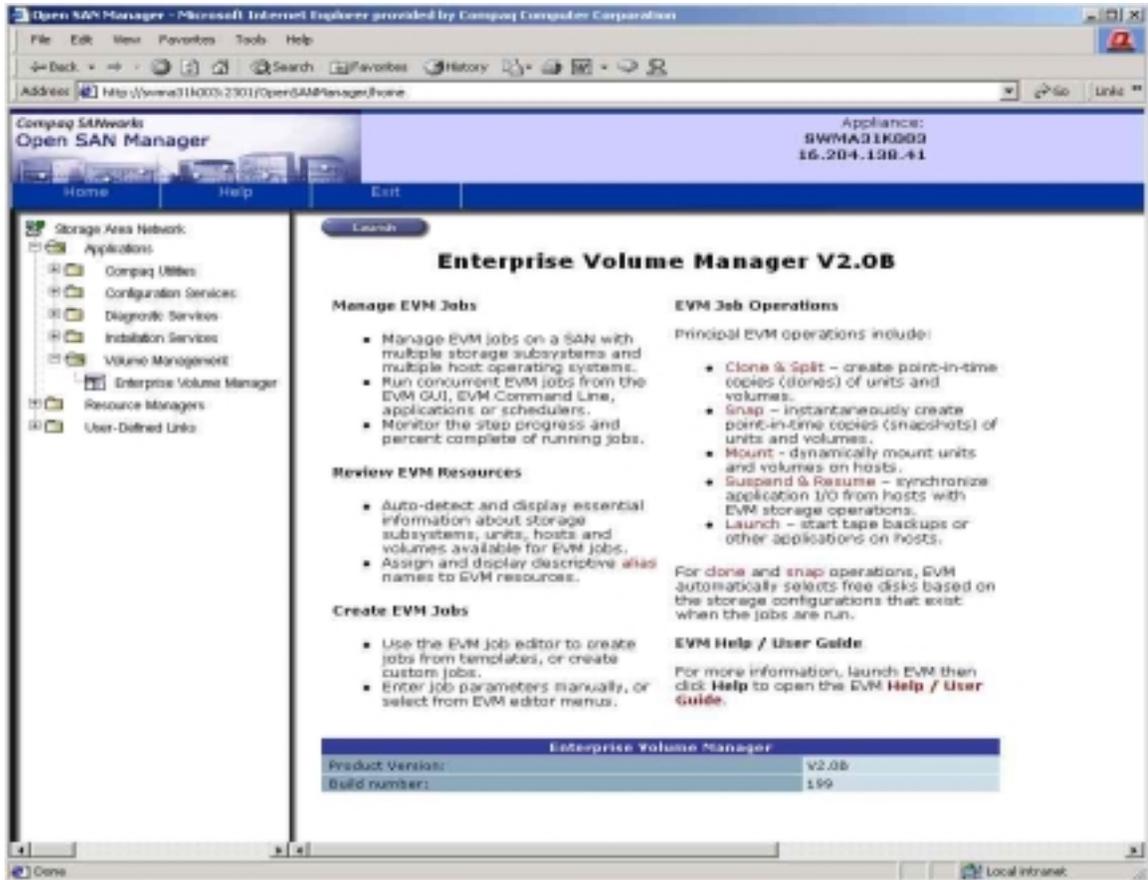


Figure 5 EVM Launch Page

You use the EVM Launch Page to start one or more sessions of the EVM graphical user interface. When Enterprise Volume Manager is launched, the EVM homepage appears.

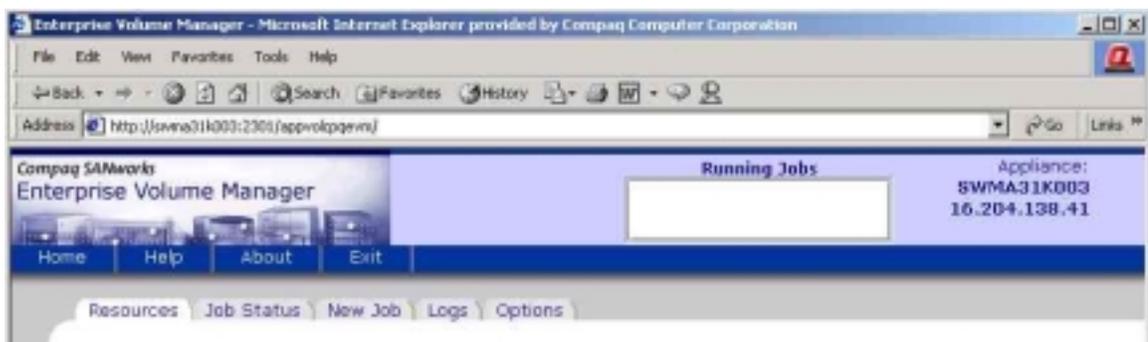


Figure 6 EVM session pane

The upper portion of the display in Figure 6 is called the *session pane*. Features and uses of the session pane include:

- Running Jobs window - view a list of running jobs
- Home button - refresh the EVM display and return to the EVM Resources page

- Options button - display the EVM Configuration page
- Help button - display the EVM Help/User Guide (Multi-OS)
- About button - display EVM version information
- Exit button - end the EVM session

The lower portion of the display is organized as a set of tabbed pages. You can use the pages as follows:

- EVM Resources page - view the active components of the EVM network
- EVM Jobs Status page - view a list of EVM jobs and select job commands, such as *execute* and *validate*
- EVM Job Monitor page - view the status of a running EVM job
- EVM Job Add page - use a specialized EVM editor to create a new EVM job
- EVM Job Edit page - use a specialized EVM editor to edit an EVM job
- EVM Job Details page - view the line-by-line details of an EVM job
- EVM Log Reporting page - view logs of EVM events

Each tabbed page includes page help and message bars

EVM Communication Flow

Figure 7 illustrates the EVM communication flow in a server acting as a EVM Server, typically the SAN Appliance Manager.

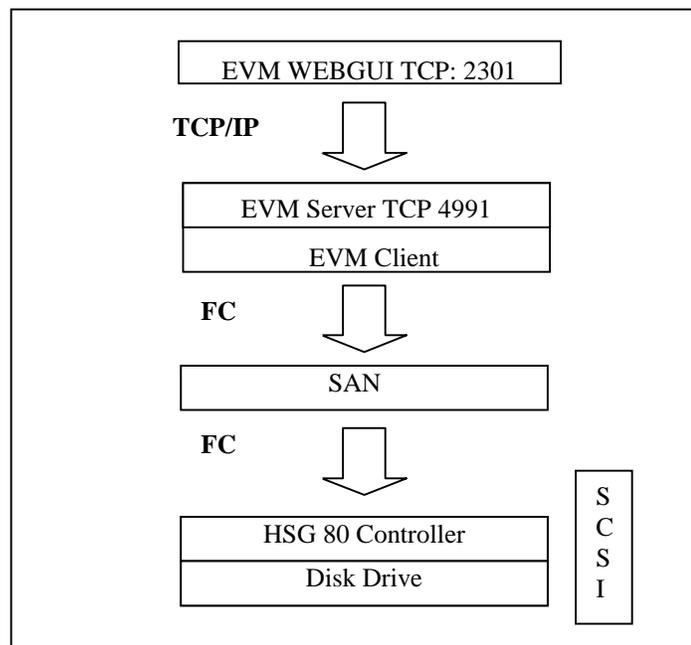


Figure 7, EVM Communication Flow

A user connects to the EVM server via TCP port 2301. The EVM server generates the EVM-related Web pages and presents these pages to the connecting user. While EVM is starting up, the server component of EVM queries the EVM client through TCP port 4991. The EVM client, in turn, communicates with the StorageWorks unit via the HS Series Agent in-band, the fiber channel and the SAN, in order to assemble the complete configuration of the StorageWorks unit.

The EVM Web interface is mainly used for creating, modifying, deleting, executing and visually monitoring EVM/snapshot jobs. Once the job is created, the EVM command line utility can be used for processing the clone or snapshot jobs.

The EVM Command Line (EVMCL) interface allows you to issue EVM job commands outside of the browser-based graphical user interface. With this capability you can create and save jobs using the graphical user interface, then execute them later from the OS command line, batch or script files. When used in conjunction with a scheduling application, the EVM Command Line feature allows maximum flexibility in the execution of EVM jobs.

While designing a batch or script (SAP backup job) file to execute an EVM job, you can include conditional statements that examine the job return codes that are issued by EVM.

When run in the default mode (Immediate switch off), EVM monitors the job progress. When a job runs to successful completion, EVM returns a code 0. If an error event occurs, EVM will issue return codes 3, 5, 6, 7. Codes 3, 5, 6 and 7 indicate that the EVM job did not run successfully. The reason is indicated in the code value.

When run with the immediate switch on, EVM does not monitor the job, but immediately issues job return code 1 to the batch file. No events or errors of any type are reported by EVM.

Note: The /I switch is not supported if an EVM job is used to split or resynchronize Business Continuance Volumes (BCV) within an SAP backup job!

EVMCL provides the basis for the SAP split mirror and resynchronization jobs described in this paper.

EVM Jobs

An EVM/SAP backup job consists of an EVM job that replicates storage units and the SAP BRBACKUP utility that performs a backup of the replicated units on the backup server instead of on the SAP production database server.

In Figure 8, a “clone” job has been created on the EVM server with the EVM WEBGUI. BRBACKUP initiates the EVM job. This job replicates storage units and mounts on the backup server. The SAP BRBACKUP utility continues to perform the backup on the backup server instead of on the SAP production server.

The BRBACKUP process on the backup server is started by a scheduler. BRBACKUP puts all tablespaces in backup mode. Setting the parameters in `init<SID>.sap` from `backup_type` to `online_split` will allow you to execute the batch file where the EVMCL command can be executed to submit an EVM job to EVM Server. After successful execution of the EVM job, EVMCL returns a status code 0 to BRBACKUP, which then continues to back up the database on the BCV. If the EVMCL return code is not 0, the backup operation terminates with an error.

After normalization of all units, issue the command *alter system suspend* to stop I/O activity on the database. After splitting the BCV unit, execute *alter system resume* to resume I/O activity on the database.

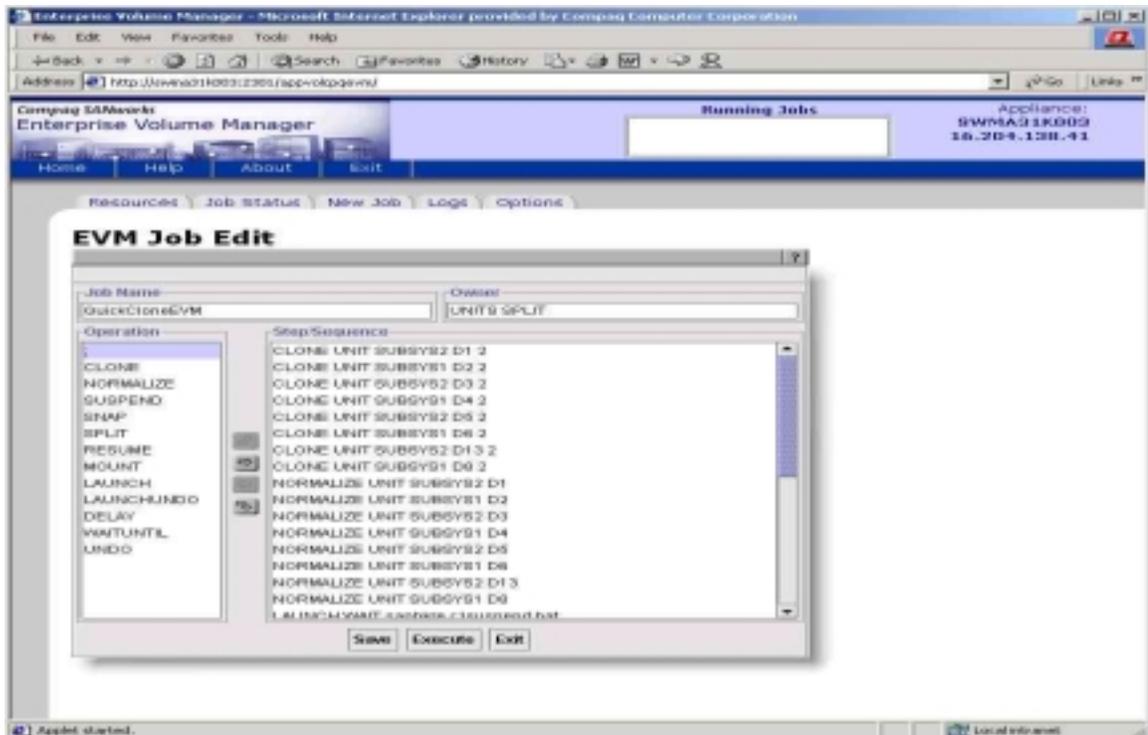


Figure 8, EVM Job Operations

Before the split operation, execute the optional file system flush (sync.exe) command on the SAP production server to flush the data from all possible cached data. Sync is not a standard Windows 2000 command, and can be downloaded from www.sysinternals.com. The sync is freeware, no support is available from Compaq.

Note: Using Oracle8i SUSPEND/RESUME functionality, you can suspend I/O to the database, then split the mirror and make a backup of the split mirror. This feature, which complements the hot backup functionality, allows you to quiesce the database so that no new I/O can be performed. You can then access the suspended database to make a split of the database without I/O interference.

EVM Replication Methods

EVM provides three different options for creating a Business Continuance Volume (BCV) as a simple single-line operation in an EVM job. EVM supports the clone, quickclone and snapshot methods.

Clone

An EVM clone job creates a physical point-in-time copy of a storage unit. The *clone* operation brings in a member to a mirror set unit. Important factors are the capacity and number of disks being normalized.

For each unit to be cloned, EVM automatically checks the RAID level, determines the correct procedures, then searches in the subsystem and selects free disks for the additional (BCV) member. EVM then copies the data and normalizes the disks in the mirror set. When normalization is complete, all members of the mirror set have identical data and EVM splits the additional (BCV) mirror from the mirror set.

Source Unit is a Three-member Mirror Set (RAID 1, RAID 1+0). To achieve near instantaneous cloning with EVM, a source unit should be configured as a three-member mirror set and be fully normalized. In this configuration, EVM instantly splits the third member to create the BCV unit. This approach to instant cloning also applies to mirror sets with four or more members. Please also refer to quickclone.

Source Unit is a Two-member Mirror Set (RAID 1, RAID 1+0). EVM automatically adds a third member to the mirror set and normalizes the members. When the members are normalized, EVM splits the third member to create the new BCV unit.

Source Unit is Not Mirrored (RAID 0 and JBOD). EVM automatically creates a single-member mirror set and then adds a second member. When the members are normalized, EVM splits the second member to create the new BCV unit. Compaq and SAP do not recommend using RAID 0 or JBOD configurations in productive environment due to the risk of losing data in the event of disk failure.

Before execution of the clone operation, ensure that a sufficient number of free disks is available. For example, RAID 1 with two-member mirror sets and a disk capacity of 18.2 GB requires one free 18.2-GB disk or larger.

Example:

EVM JOB includes,

1. CLONE SUBSYS1 D1 2
2. NORMALIZE SUBSYS1 D1
3. SPLIT D1 D21
4. MOUNT SUBSYS1 EVMAGENT 1 D21 F:

As shown in Figure 9, Unit D1 is a two-member mirror set. The clone operation brings in a free disk as a third member to the mirror set. Normalize will wait until the third member synchronizes completely with the mirror set. The EVM Split operation removes the disk from the mirror set and creates BCV unit D21. BCV unit D21 is mounted on the EVMAGENT server.

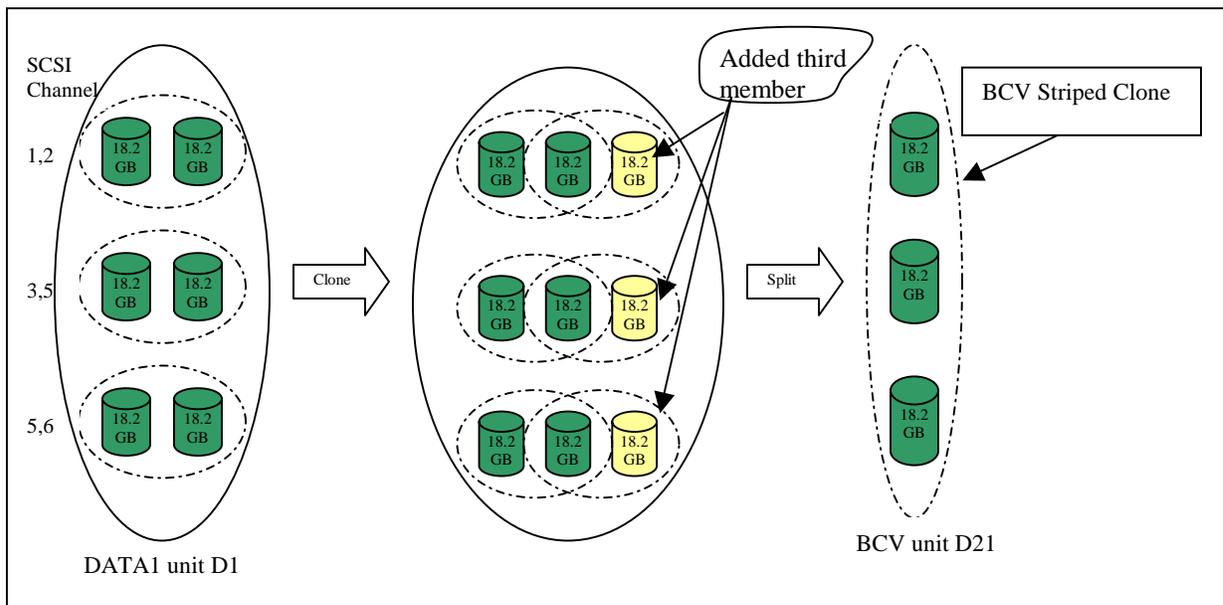


Figure 9, Clone Overview

Note Figure 9 shows a two-member mirror set. The clone operation adds a third member to the mirror set. Assuming a moderate I/O load on the storage controller, and an adequate configuration of disks in the array, a replication rate of 60 GB per hour can be achieved. The time between the beginning and end of the backup operation should be minimal. During job execution all the tablespaces will be in backup mode. Online backup allows the database to remain online and available for user transactions. During this time the database relies on the redo and archive logs.

Quickclone

This refers to the practice of pre-configuring a unit so that an EVM job can make a clone BCV unit nearly instantaneously. The required pre-configuration for a quick clone unit is a three-member mirror set. When a clone operation is executed, EVM first checks the source unit configuration. If more than three members exist in the mirror set, EVM immediately proceeds to the next step in the job. With a quickclone, the normalize operation typically finds the unit in a *normal* state and immediately proceeds to the next step. If that step is a split, then EVM immediately splits off a member to create the BCV unit.

Units that are configured as RAID 3/5 cannot be replicated with the clone and quickclone methods.

Quickclone is ideal if the backup or offline process should not affect the production units. Quickclone can also be used for nearline backup for instant recovery purposes and is recommended for SAP write-intensive systems.

Example:

EVM JOB includes

1. CLONE SUBSYS1 D1 2
2. NORMALIZE SUBSYS1 D1
3. SPLIT D1 D21
4. MOUNT SUBSYS1 EVMAGENT 1 D21 F:

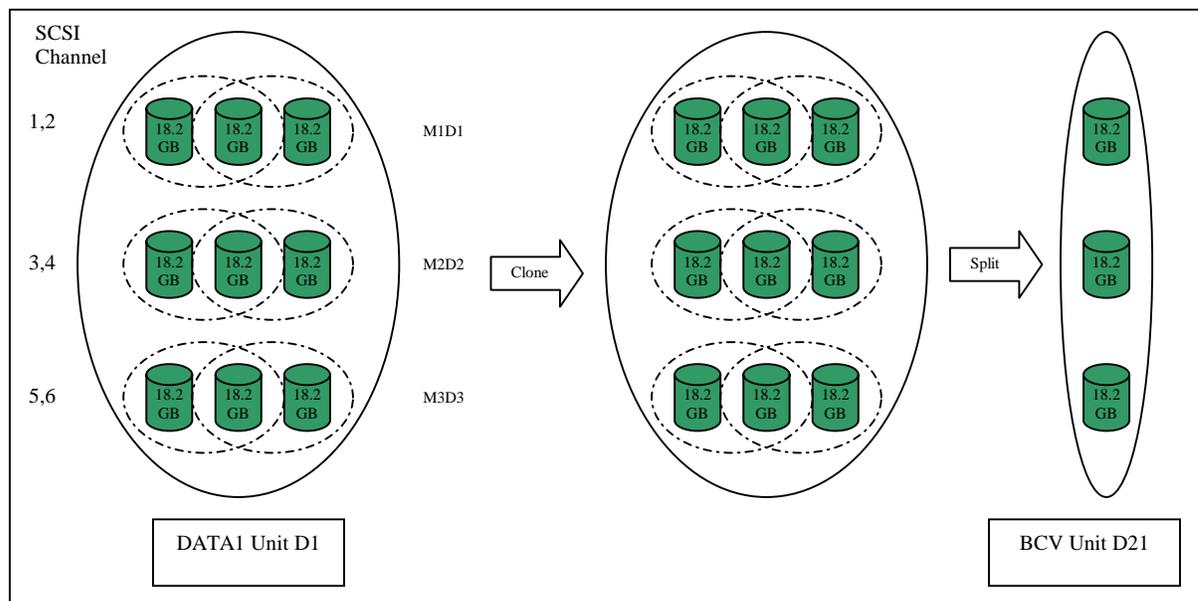


Figure 10, Quickclone Overview

As shown in Figure 10, Unit D1 is a three-member mirror set. The clone operation checks the redundancy of the disks to see if more than two disks are already present in the mirror set. Then it automatically switches to the next operation in the job. Normalize checks the normal state of the mirror set and switches to the next operation in the job. The EVM Split operation removes the disk from the mirror set, and then creates stripe sets of the disks and creates BCV unit D21. BCV unit D21 is mounted on the EVM Agent/backup server.

Note Figure 10 is a three-member mirror set. The quickclone operation checks the redundancy of the mirror set. During job execution all tablespaces are in backup mode. An online backup allows the database to remain online and available for user transactions. The system relies on the redo and archive logs generated during the backup to restore the database to a consistent state. Execution of the EVM job between the tablespace in backup mode and end backup mode will take significantly less time compared to the clone method.

Snapshot

Snap is an EVM job that instantly creates a virtual, point-in-time copy of a storage unit. The resulting unit is called a snapshot. EVM provides snapshot functionality as a simple, single-line operation in an EVM job. All of the underlying StorageWorks CLI commands that are required to make a snapshot are automated and incorporated into EVM. When you include snap operations in a job, EVM automatically makes snapshot BCV units when the job is executed. An EVM job can also dynamically mount BCVs on a computer using the drive letters or mount points that you specify. Or, you can have the EVM job leave BCV units unmounted.

To support the EVM snap operation, the storage subsystem and ACS must be configured as follows:

- Fibre Channel switched fabric topology
- Dual HSG controllers in Transparent or MultiBus Failover mode
- 512 MB of cache per controller
- Mirrored cache

Rules for operation of snapshot:

- You cannot create more than four snapshot units per controller pair (subsystem).
- You cannot create a snapshot of a snapshot.
- You cannot create a snapshot of a unit that already has a snapshot.
- A snapshot unit and its source unit must be online to the same controller.
- A snapshot unit and its source unit must be in the same unit number range of 0 to 99 or 100 to 199 (Transparent failover only).

The logic to work within these rules is integrated into EVM. When EVM indicates that a snapshot is not possible, or if an EVM job that involves a snapshot fails, it is likely that one or more of these rules has not been satisfied.

Before execution of the snapshot operation, ensure that there is sufficient storage capacity.

Example

Job includes: NORMALIZE SUBSYS1 UNIT D1
 SNAP SUBSYS1 D1 D21
 MOUNT SUBSYS1 D21 EVMAGENT I:

As shown in Figure 11, unit D1 (RAID 5) has a total disk capacity of 67.8 GB. NORMALIZE makes the state of the unit normal before executing the next operation in the job. SNAP automatically searches in the storage subsystem for free data capacity, allocates a single or multiple disks (36.4 x 2) and creates the BCV (D21) unit. At the same time the bit map is created in the controller cache. The MOUNT operation makes the BCV (D21) unit available for the host EVM agent with drive letter I.

If an update request comes for production unit D1 after creation of the snapshot BCV unit D21, the data block on D1 will first be copied to BCV unit D21 and the bit map will be updated with the changed block number on the controller cache. After this, the changed data block is written to the unit D1.

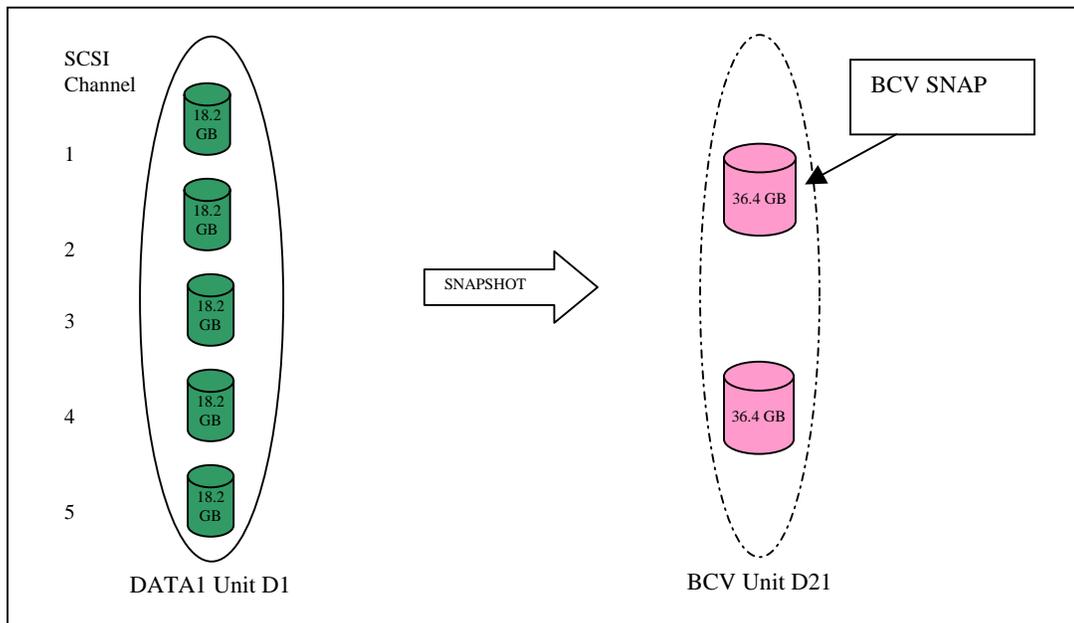


Figure 11, Snapshot Overview

During the backup process, the controller will ensure that the data is read from D1 or D21 based on the controller bit map. If the blocks are not changed on the production unit during the backup process, the data block is read from production unit D1.

Supported Versions for Clone and Snapshot

Following are the supported Compaq StorageWorks subsystem with HSG80 Fibre Channel controllers

ACS Version	Supports
8.5 F	Clone
8.5 P	Clone and Snapshot
8.5 S	Clone and Snapshot
8.6-1 F	Clone
8.6-1 P	Clone and Snapshot
8.6-1 S	Clone and Snapshot

Table 3, Supported ACS Versions

RAID LEVELS	Supports
RAID 0 and JBOD	Clone and Snapshot
RAID 1	Clone and Snapshot
RAID 1 + 0	Clone and Snapshot
RAID 3/5	Snapshot

Table 4, RAID Levels

Solution Planning and Sizing

Before installing the EVM software make sure that all the resources and SAN infrastructure satisfy the general EVM configuration rules in order to successfully integrate EVM into the SAP application.

EVM Network Configuration Rules and Minimum Requirements

One EVM component, called the EVM server, manages the EVM network and EVM jobs and provides the Web browser interface to users. The EVM server is installed on a Management Appliance. Another EVM component, called an EVM agent, is installed on host computers and runs the EVM jobs created for volume management. StorageWorks subsystems provide the storage units used in EVM jobs. When planning an EVM V2.0 network, keep the following configuration rules in mind.

Each EVM network requires the following:

- A Management Appliance with Open SAN Manager (OSM) software version 1.0B or newer and Element Manager for StorageWorks HSG version 1.0B-Patch or later for HSG80 controllers installed.
- A Management Appliance and at least one EVM Agent node must be connected to at least one EVM-compatible StorageWorks subsystem (all containing Compaq StorageWorks HSG80 Array Controllers and using a StorageWorks Solution Software Kit 8.5 C, 8.6 or newer)

- Minimum Web browser Microsoft Internet Explorer 5.01, JRE version 1.3.0_02 or Netscape Communicator 4.77, JRE version 1.3.0_02 should be installed on EVM server and EVM Agents.
- An EVM network can support from 1 to 25 StorageWorks subsystems.
- At least one EVM server and one EVM client should be available in the EVM network.
- Please do not use “C!” or “S!” at the beginning of a container name, and do not use “!C” or “!S” at the end of one. Do not change the name of any container created by the EVM network that begins or ends with these sequences. Changing the name of a container created by EVM may impact the ability of EVM to track that container. If a container name that uses the “C!...! C” or “S!...!S” format is changed, EVM will not recognize that container correctly. Since EVM safety rules prohibit the EVM program from deleting anything that EVM storage did not create, the container will not function properly.
- The required SANworks Element Manager for EVM V2.0 is called Element Manager for StorageWorks HSG version 1.0B-Patch or newer. However, the element manager provided with the Management Appliance is Element Manager for StorageWorks HSG version 1.0A. An Element Manager update is required to upgrade from version 1.0A to 1.0B-Patch.
- One EVM-compatible StorageWorks RAID Array 8000 (RA8000), Modular Array 8000 (MA8000), Enterprise Storage Array 12000 (ESA12000), Enterprise Modular Array 12000 or 16000 (EMA12000/16000) subsystem, should be configured in the SAN.
- Command Console LUN (CCL) enabled, if the EVM Agent is installed on the dedicated EVM backup server. No other EVM client is required for the Rapid Restore and Non-Disruptive Backup solution.
- The EVM client is only needed on the SAP system when EVM jobs are started and scheduled on the SAP production system.
- If SecurePath is installed on the SAP production system, only 7 SCSI LUNs can be used for disk volumes due to the 8 LUN limitation of SecurePath.

Snapshot Requirements and Limitations

The EVM snapshot feature uses the hardware level snapshot functionality provided by the HSG80 controller and enabled by the ACS. To support snapshot, the HSG80 controller and the ACS require the following configuration attributes:

- Minimum ACS Version 8.5 P or S
- Fibre Channel switched fabric topology
- Dual HSG80 controllers in Transparent or MultiBus Failover mode
- 512 MB of cache per controller
- Mirrored cache

The ACS imposes rules on the creation of snapshots. EVM contains logic to work within these rules so that the user does not need to worry about them. If EVM indicates that a snapshot is not possible or if an EVM job involving a snapshot fails, snapshot limitations may be the reason.

These limitations include:

- No more than 4 snapshots are allowed per controller pair.

- A snapshot of a snapshot is not allowed.
- Each unit is allowed only one snapshot.
- A snapshot and the source unit must be online to the same controller.
- A snapshot and the source unit must be in the same unit number range (0 to 99 or 100 to 199).

Planning to Use EVM with Data Replication Manager

When used with Data Replication Manager (DRM), EVM supports clone and snapshot replication of target site units and clone replication of initiator site units. EVM does not support snapshot replication of units at the initiator site.

Additionally, the DRM configuration must use the synchronous transfer mode. This ensures that data is written to the target site disks before an EVM BCV is created. Refer to DRM documentation for information on how to enable synchronous mode.

EVM Configuration Rules for the SAP System

- At least one EVM host agent as backup server to be installed for backing up SAP database.
- Compaq Web-based management agents should have latest Java Plug-in installed on the EVM server and on any computer used to browse to EVM.
- No additional EVM software required for SAP production system.
- SAP Database system and EVM Agent should have the same operating system, service packs, and patches.
- Static IP addresses and DNS have to be used and properly configured on all EVM network nodes.
- EVM backup server should have the same Oracle database version as the SAP system.
- The current solution requires one, and only one, FC-HBA on SAN Appliance Manager connected to SAN switches.
- The backup server must be connected to the SAN switches. The supported backup software has to be installed on the backup server.
- The backup server should have SAP backint certified backup software installed.

SAP Job Rules

- BCV units of all SAPDATA and redo log volumes must be created.
- All BCV units on the EVM backup server should have the same drive letters as in the SAP production system.
- Additional jobs for the archive logs to be created when EVM online split job takes place in EVM backup server.

Example

Production	Backup Server	Units	Drive Letters
Sapdata1	Sapdata1	D1	M
Sapdata2	Sapdata2	D2	N
Saplog	Saplog	D3	P

For example, mount snapshot units D1 D2 D3 as M N P on the EVM backup server.

Start “brbackup –c –a” to back up the data files and to archive redo logs.

EVM Supported New Features

EVM is redesigned using a distributed architecture that enhances the scalability and functionality of EVM networks for SAP with the following configurations.

- The EVM 2.0 server runs on a Compaq SANworks Management Appliance as a component of the suite of Open SAN Manager applications.
- The EVM 2.0 server supports multiple storage subsystems, like MA 8000, RA 8000, EMA 12000, ESA 12000 and EMA 16000.
- The EVM 2.0 server simultaneously supports hosts with different operating systems, including Windows NT, Windows 2000, Sun Solaris.
- The EVM 2.0 server supports an unlimited number of licensed hosts.
- The EVM 2.0 server simultaneously supports a mix of subsystems running ACS 8.5 and 8.6-1.

Example 1: Non-Disruptive Backup EVM SAP Solution

The SAP backup/restore functions are among the most important operations in the production environment. EVM’s clone and snapshot 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 24/7 production environment. Due to the fact that SAP is a write-intensive system, we strongly recommend cloning. RAID 1 or RAID 1+0 is only supported for clone operations and also offers better write performance for SAP applications.

Minimum Configurations

The first example is a standard SAP/Oracle implementation implemented using ACS 8.6P and which supports 84 disks in EMA 12000 or ESA 12000. If the database size is 305 GB, you need 3 striped mirror sets: 3 x 18 disks x 18.2 GB. This example shows a minimum configuration with one HSG80 Array dual-redundant controller with the minimum of 512 MB cache installed. Such a minimum configuration can have up to 305 GB of SAP/Oracle tablespace.

Description	Raid	Disks	Formatted Capacity	SAPDATA
Oracle Tablespace files	R 1+0	(18X18.2GB)X3 Units	305 GB	SAPDATA1 SAPDATA2 SAPDATA3
Online Redolog Volume	R1	(3X18.2GB)X2 Units	16.8 GB	
MirrLog Volume	R1	(3X18.2GB)X2 Units	16.8 GB	
Archive Volume	R1	2X36 GB	33.9 GB	
SAP Volume	R1	2X18.2	16.8 GB	
Oracle Binaries	R1	2X18.2	16.8 GB	

Table 5, SAP/Oracle Minimum Configuration

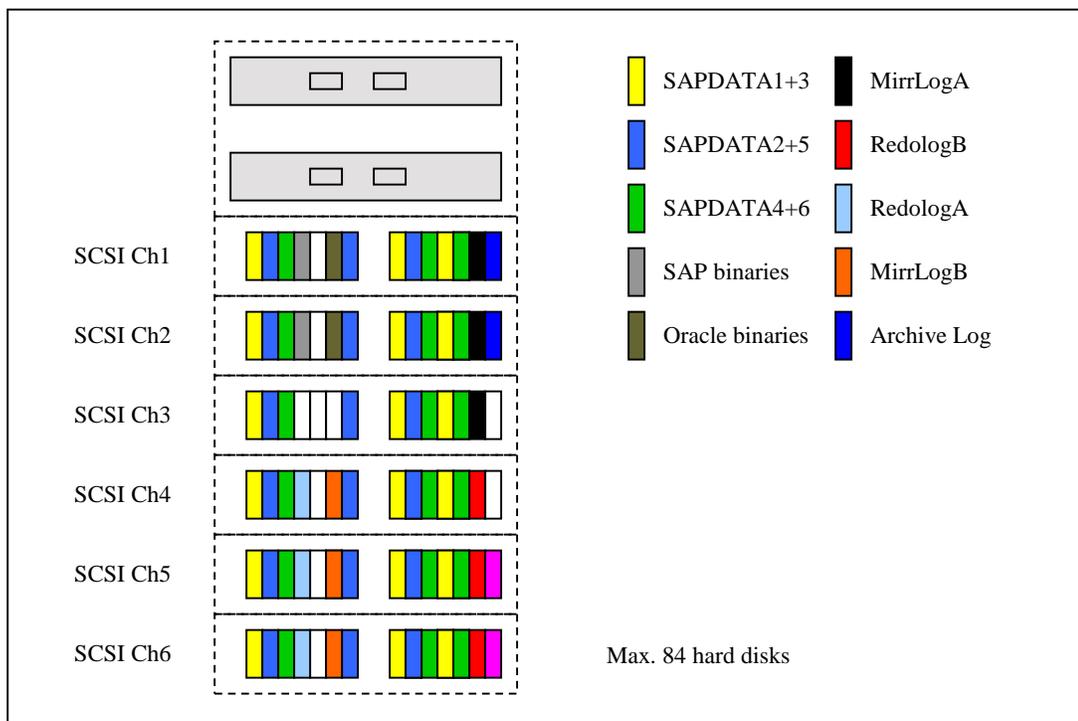


Figure 12, SAP/Oracle Minimum Configuration

Note: Figure 12 shows how to configure the StorageWorks unit for EVM usage. Quickclone is recommended method for write-intensive SAP systems!

Medium Configuration

The second example is a standard SAP/Oracle implementation implemented using EMA 12000 or ESA 12000. If the database size is 606 GB, you need two striped three-member mirror sets; for one striped three-member mirror set you need 18 disks x 18.2 GB. You also need two striped three-member mirror sets and for one striped three-member mirror set you need 9 disks x 36 GB. The example shows a medium configuration with one HSG80 Array Controller with the minimum of 512 MB cache installed. Such a medium configuration can have up to 606 GB of SAP/Oracle tablespace.

Description	Raid	Disks	Formatted Capacity	SAPDATA
Oracle Tablespace files	R 1+0	18X18.2GB =101.7 GB X 3	303 GB	SAPDATA1 SAPDATA3 SAPDATA5
	R 1+0	9X36 GB =101.7 GB X 3	303 GB	SAPDATA2 SAPDATA4 SAPDATA6
Online Redo log Volume	R1	3X18.2 GB X 2	16.8 GB	
Mirror RedoLog	R1	3X18.2GB X 2	16.8 GB	
Archive Volume	R1	2X36 GB	33.9 GB	
SAP Volume	R1	2X18.2 GB	16.8 GB	
Oracle Binaries	R1	2X18.2 GB	16.8 GB	

Table 6, SAP/Oracle Medium Configuration

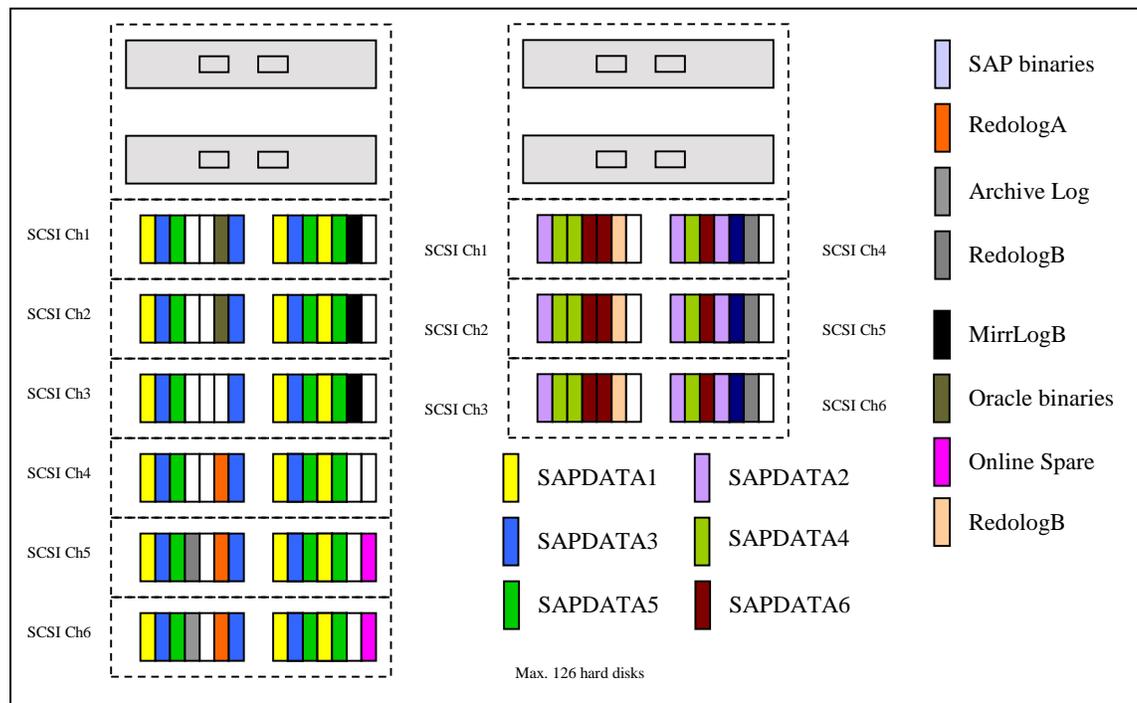


Figure 13, SAP/Oracle Medium Configuration

Note The example in Figure 13 shows how to configure the StorageWorks unit for EVM with six volumes for SAPDATA. Quickclone is recommended for write-intensive SAP systems.

Large Configuration (1.2 TB SAP/Oracle DB)

The third example is a standard SAP/Oracle implementation, implemented using two EMA12000/4214. For 1.2 TB size of database, you need 18 disks per unit to create a three-member striped mirror set so you need 54 disks x 36 GB. Also you need 18 disks per unit to create a three-member striped mirror set so you need 54 disks x 36.4 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 1.2 TB of SAP/Oracle tablespace.

Description	RAID	Disks	Formatted Capacity	SAPDATA
Oracle Tablespace files	R 1+0	(18 x 36.4 GB)X3 Units	609 GB	SAPDATA1
				SAPDATA3
				SAPDATA5
	R 1+0	(18 x 36.4 GB)X3 Units	609 GB	SAPDATA2
				SAPDATA4
				SAPDATA6
Online Redo Log Volume	R1	4 x 18.2 GB	16.8 GB	
MirrLog Volume	R1	4 x 18.2 GB	16.8 GB	
Archive Volume	R1	2 x 36 GB	33.9 GB	
SAP Volume	R1	2 x 18.2 GB	16.8 GB	
Oracle Binaries	R1	4 x 18.2 GB	16.8 GB	

Table 7, SAP/Oracle Large Configuration

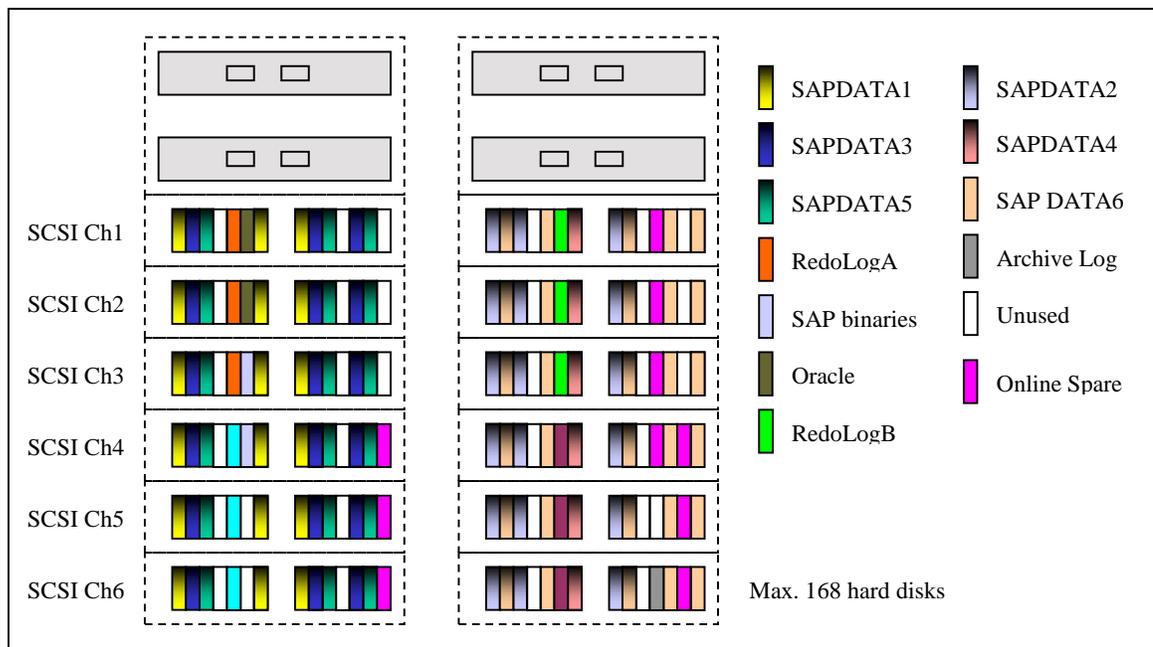


Figure 14, SAP/Oracle Large Configuration

Note The above configuration example shows how to configure the StorageWorks unit using EVM. SAPDATA1, SAPDATA2, SAPDATA3, SAPDATA4, SAPDATA5, SAPDATA6 volumes could be used! Cloning is recommended for write-intensive SAP systems!

Snapshot Minimum Configuration

In Figure 15, SAP/Oracle Minimum Configuration, RAID 3/5 is used as the disk stripe set. Because of the RAID 3/5 level, the capacity is higher, but the BCV creation permits only EVM snapshots. A RAID 3/5 may be selected if the SAP application is not so write-intensive and a higher capacity is needed at a lower price. The modular dual SCSI 4254 disk enclosure shelves are used. The maximum number of disks in this configuration is 42 (1-inch).

The snap volume will be removed after the backup process. This solution is very space efficient, but offers poorer I/O performance than a full RAID 1+0 configuration.

Description	Raid	Disks	Formatted Capacity	SAPDATA
Oracle Tablespace files	R 3/5	(5X36.4GB)X2 Units	270 GB	SAPDATA1+3+5 SAPDATA2+4+6
Redolog Volume	R1	(2X18.2GB)X2 Units	16.8 GB	
Mirror Log Volume	R1	(2X18.2GB)X2 Units	16.8 GB	
Archive Volume	R1	2X36GB	33.9 GB	
SAP Kernel Volume	R1	2X18.2GB	18.2 GB	
SAP Data BCV	R3/5	10 X 36.4 GB	270 GB	
Redo Log BCV	R1	(2X18.2GB)X2 Units	16.8 GB	
Mirror LOG BCV	R1	(2X18.2GB)X2 Units	16.8 GB	
Oracle Binaries	R1	2X18.2GB	16.8 GB	

Table 8, SAP/Oracle Minimum Configuration for Snapshot

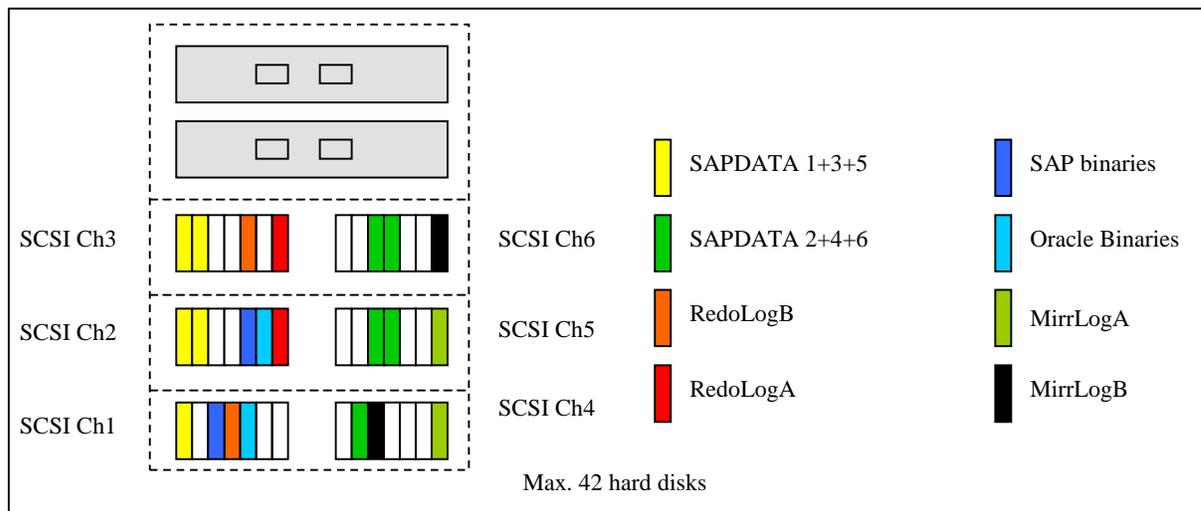


Figure 15, SAP/Oracle Minimum Configuration for Snapshot

Note: This is just an example of how to configure the StorageWorks unit for EVM. Snapshots are not recommended for write-intensive SAP systems!

Snapshot Medium Configuration

In medium configuration for snapshot RAID 3/5 is used as the disk stripeset. Because of the RAID 3/5 level, the capacity is higher, but the BCV creation only permits EVM snapshots. A RAID 3/5 may be selected if the SAP application is not so write-intensive and a higher capacity is needed at a lower price.

The snap volume will be removed after the backup process. This solution is very space-efficient, but offers poorer I/O performance than a full RAID 1+0 configuration!

Description	Raid	Disks	Formatted Capacity	SAPDATA
Oracle Tablespace files	R3/5	(12 x 18.2 GB)X3 Units	559 GB	SAPDATA1+3 SAPDATA2+5 SAPDATA4+6
Online Redolog Volume	R1	(2 X 18.2 GB)X2 Units	16.8 GB	
Mirror Redolog Volume	R1	(2 X 18.2 GB)X2 Units	16.8 GB	
Archive Volume	R1	2 X 36.4 GB	33.4 GB	
SAP Volume	R1	2 X 18.2 GB	16.9 GB	
Oracle Binaries	R1	2 X 18.2 GB	16.9 GB	
Data BCV	R 3/5	21 X36.4 GB	609 GB	
Redo Log BCV	R1	2 X 18.2 GB	16.8 GB	
Mirror Log BCV	R1	2 X 18.2 GB	18.2 GB	
Spare disk		2 X 36.4 GB	72.8 GB	

Table 9, SAP/Oracle Medium Configuration for Snapshot

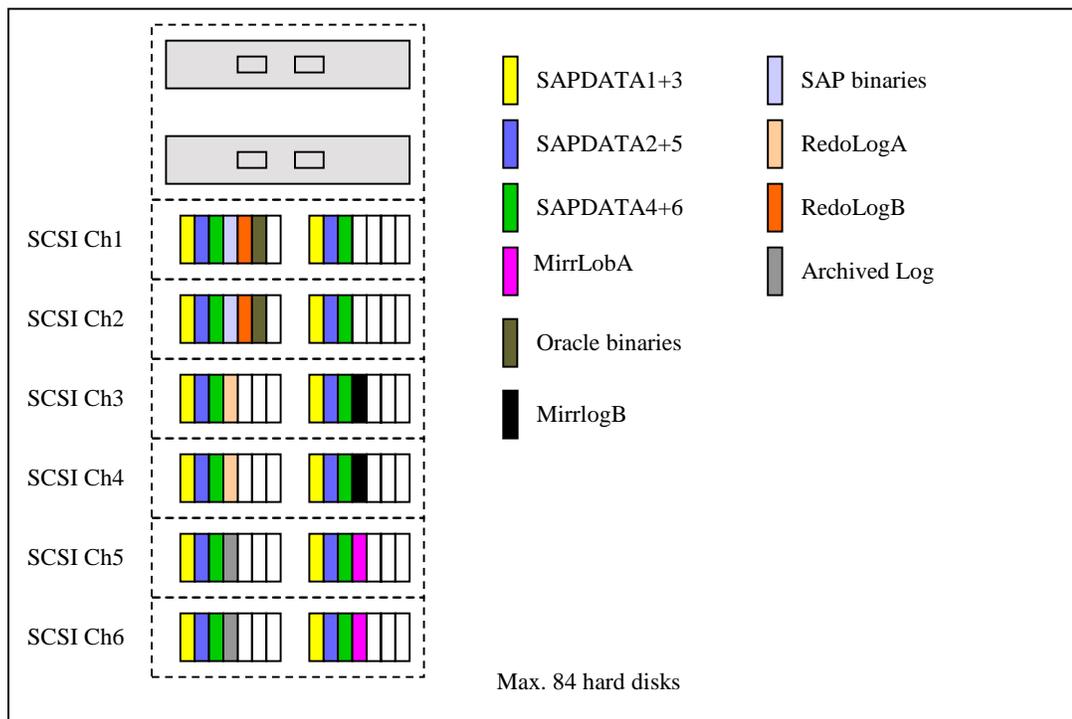


Figure 16, SAP/Oracle Medium Configuration for Snapshot

Example 2: Rapid Recovery Solution for SAP Oracle Database

This solution allows us to rapidly recover data from the BCV unit. This form of recovery is known as “Rolling Clones” and helps us to reduce the recovery time, which would be much longer if we were restoring from tape. The database on the production server can be restored and recovered rapidly as long as the BCV unit remains available.

Figure 17 illustrates how EVM creates the first clone while SAP Server is running.

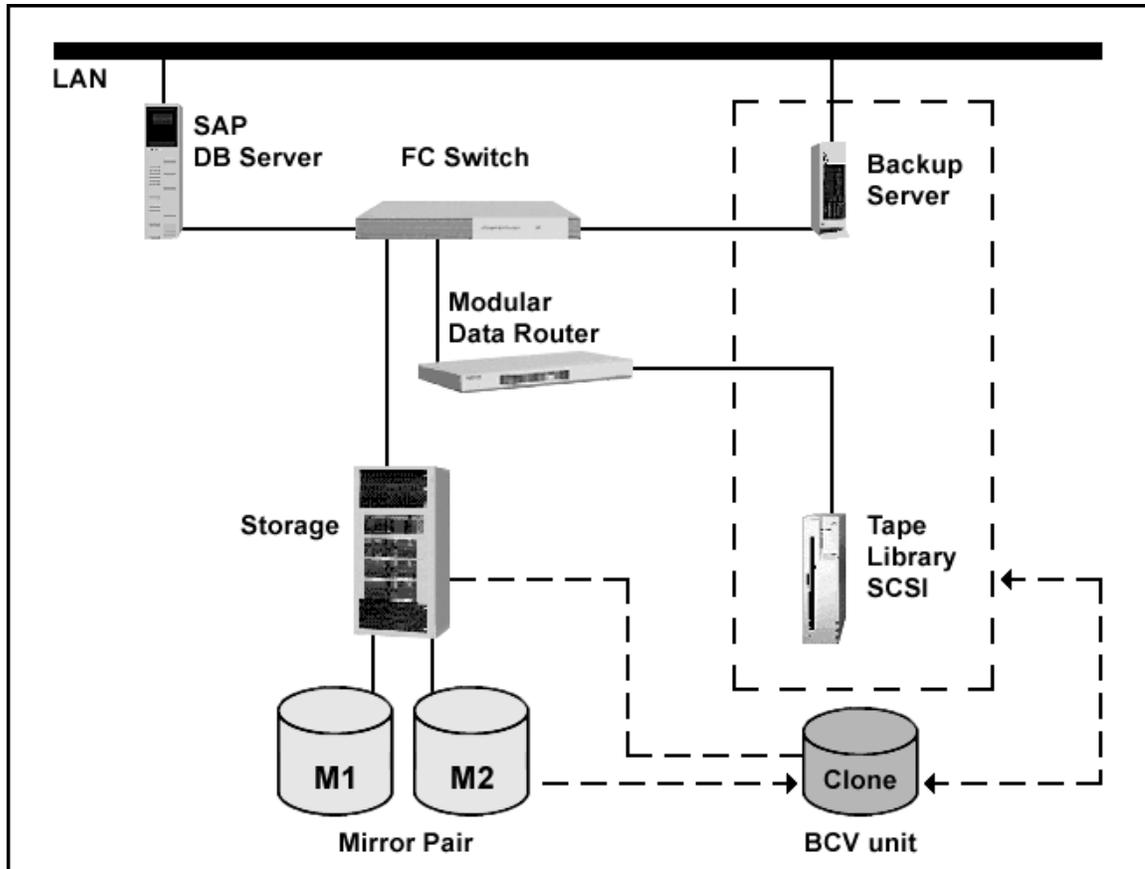


Figure 17, EVM Rapid Recovery Overview

When the BCV is established, you can recover rapidly by redirecting the BCV unit to the production server rather than having to do a lengthy tape restore. It is possible to recover the full SAP database directories, individual tablespace or individual files from the BCV unit.

While the decision to implement SAP recovery requires human intervention, the recovery process itself can be scripted and tested using the SAP BRRESTORE utility.

SAP/EVM Backup Process Flow

An SAP offline split backup occurs when the database backup is in offline mode, in other words, the database is shut down to perform the BCV split. Once the split has finished, BRBACKUP restarts the production database immediately and then the physical backup is performed on the backup host. This is the simplest backup strategy and results in a consistent point-in-time database backup; however, the database must be shut down, which results in loss of productive user work.

An online split backup is more complex than offline split backup as database is running. Specifically, the BRBACKUP tool in the SAPDBA tool suite alters the tablespaces into begin backup mode. This mode prepares the data files for backup by forcing all committed transactions to be recorded in the archive logs therefore reducing I/O to the tablespaces. Once the BCVs are split from their standard devices, BRBACKUP alters the tablespaces to end backup mode. Oracle then automatically applies all committed transactions from the archive logs to the tablespaces, all the while allowing users to continue work without interruption. 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 types of split mirror backups. In general, BRBACKUP performs the following steps:

1. Creates a copy of the database control file and places it into %SAPBACKUP%.
2. If performing an offline split, shuts down the database. If performing an online split, checks that archive logging is enabled, then alters the tablespaces into begin backup mode.
3. Splits the BCVs with the split command file designated in the initSID.SAP configuration file. The split command file is tailored for each EVM installation and contains the Enterprise Volume Manager Component for Windows 2000 EVM jobs, and Windows NT commands required to:
 - Alter System suspend
 - Flush the file systems on the Standard host
 - Split the BCVs
 - Alter System Resume
 - Mount the BCVs on the backup host (EVM agent).
 - Optionally run chkdsk on the drives to ensure the disks are clean.
4. If performing an offline split, restarts the database. If performing an online split, uses Oracle Server Manager to alter the tablespaces into end backup mode.
5. Begins backing up the data files, redo logs, and control files residing on the mounted devices on the backup host.
6. Resynchronizes the BCVs with the original mirrorsets. The command for doing this is specified in the initSID.SAP configuration file.

The resynchronization command file is tailored for each EVM installation and contains the EVM command, required to:

- Unmount the BCVs on the backup host.
- Re-join the BCVs with their production unit.

Online and Offline Backup Process

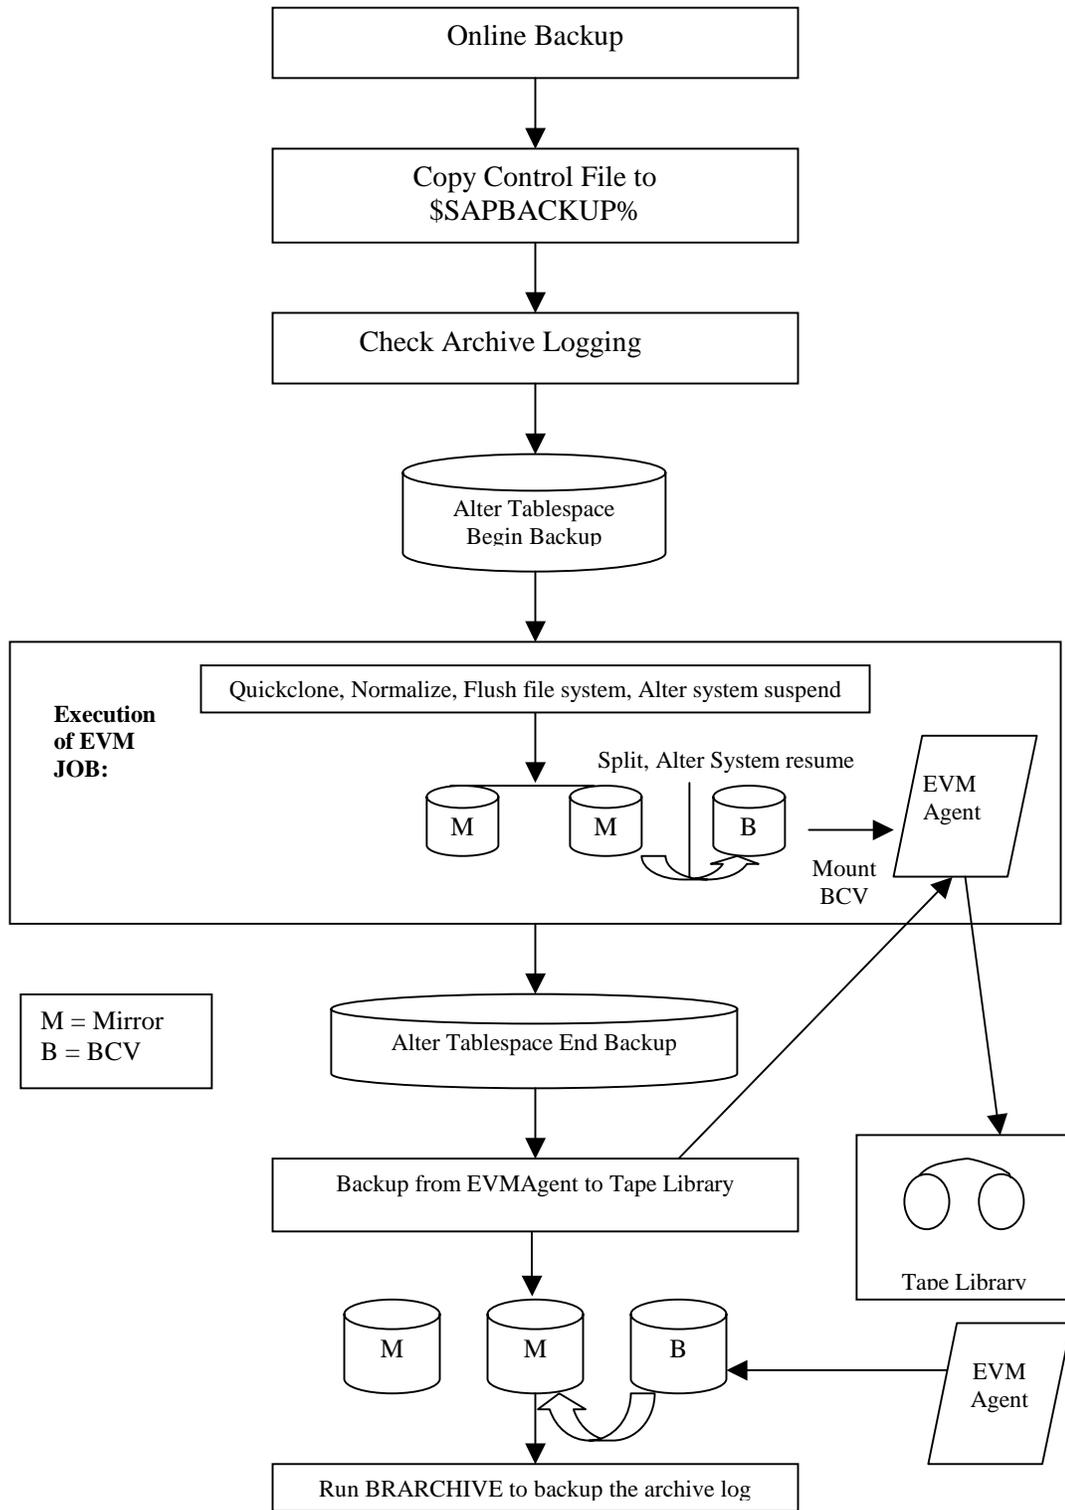


Figure 18, SAP/ORACLE Split Mirror Online Backup

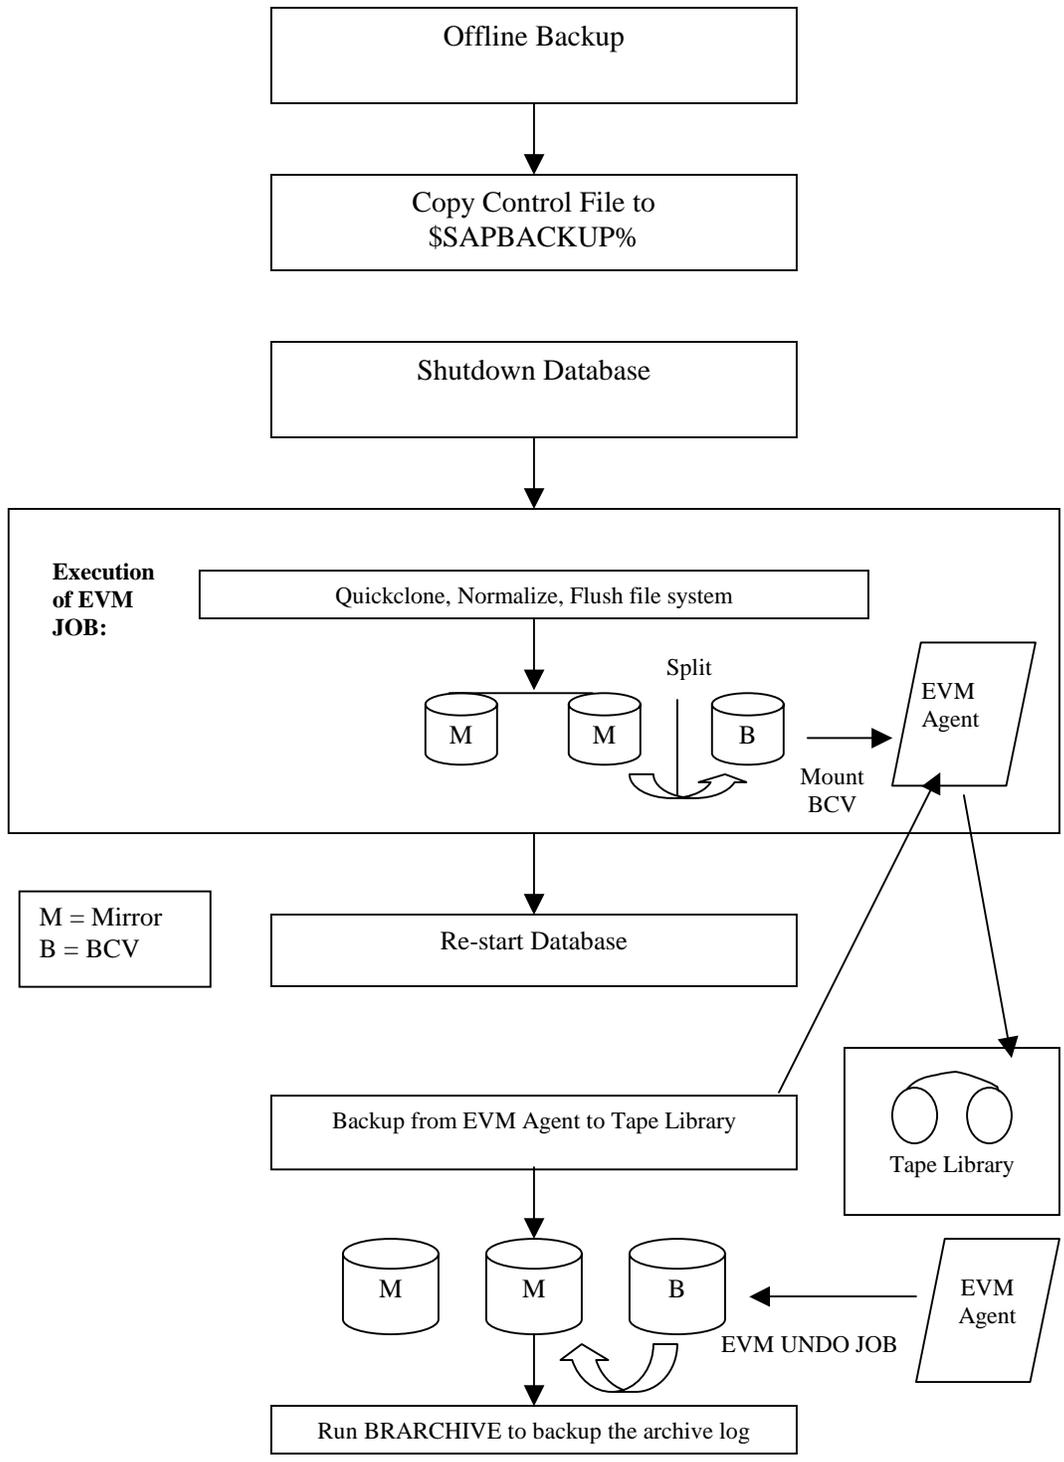


Figure 19, SAP/ORACLE Split Mirror Offline Backup

Installation of EVM

Server

The EVM server is installed on a Management Appliance and controls all activities within the EVM Network. Management Appliance acts as an interface between EVM, storage and host. The EVM host agents perform the EVM jobs, as directed by the EVM server. Refer to the EVM V2.0 Installation Guide for installing EVM Server.

Host Agent

EVM agent software needs to be installed on the backup host and on the SAP system. The implementation will not affect the production environment. Please refer to <http://www.compaq.com/products/storageworks/adapters/emulexdownload.html> for HBA drivers.

Storage

If the existing subsystem does not have free resource to create EVM BCVs, then please refer to the sizing section to change or upgrade your existing storage.

Please follow the EVM installation guide for Windows 2000 when you install EVM Server on SAN Appliance and on EVM agents.

Note: An EVM network can support from 1 to 25 StorageWorks subsystems.

Refer to Compaq's EVM required installation guide
<http://www.compaq.com/products/sanworks/evm/documentation.html>

Environment Variables and File Systems

Table 10 lists examples of the environment variables and file systems required on both the standard and backup hosts. Most of the environment variables are set up automatically on the standard host during R/3 installation. However, the remaining variables need to be added manually to the SIDadm account on the backup host.

Environment Variable	File System	Function
ORACLE_HOME	E:\Oracle\<SID>\817	Oracle8 installed on both hosts
ORACLE_SID	SID	Oracle8 installed on both hosts
TNS_ADMIN	E:\Oracle\<SID>\Network\ADMIN	Oracle8 installed on both hosts
Path	E:\Oracle\<SID>\817\bin C:\WINNT\system32 C:\Program Files \Compaq\SanWorks\Enterprise Volume Manager\bin K:\usr\SAP\SID\sys\exe\run	
DBS_ORA_TNSNAME	SID	R/3
SAPDATA1	I:\Oracle\SID\sapdata1	Set on both hosts
SAPDATA2	F:\Oracle\SID\sapdata2	Set on both hosts
SAPDATA3	J:\Oracle\SID\sapdata3	Set on both hosts
SAPDATA4	G:\Oracle\SID\sapdata4	Set on both hosts
SAPDATA5	O:\Oracle\SID\sapdata5	Set on both hosts
SAPDATA6	P:\Oracle\SID\sapdata6	Set on both hosts
SAPARCH	N:\Oracle\SID\SAParch	R/3, installed on standard host and shared mount to backup host on N:
SAPBACKUP	N:\Oracle\SID\SAPbackup	R/3, installed on standard host and shared mount to backup host on N:
SAPCHECK	N:\Oracle\SID\SAPcheck	R/3, installed on standard host and shared mount to backup host on N:
SAPDATA_HOME	V:\Oracle\SID	R/3
SAPEXE	K:\usr\SAP\SID\sys\exe\run	R/3 locally on the backup host. Only the BR tools need to be in this directory
SAPLOCALHOST	Sapbkre or evmagent	R/3 Set to local hostname
SAPREORG	N:\Oracle\SID\SAPreorg	R/3, installed on standard host and shared mount to backup host on N:
SAPTRACE	N:\Oracle\SID\SAPtrace	R/3, installed on standard host and shared mount to backup host on N:

Table 10, Environment Variables

The %SAPBACKUP% directory resides on the N: drive of the standard host (sapbkre) and is read-write share mounted on the N: drive of the backup host (EVM agent). Since the archive log destination (LOG_ARCHIVE_DEST) is the same on both hosts, initSID.ora requires no edits. The R/3 installation was SAP-compliant.

Note: The latest versions of any SAP R/3 programs are available on SAPserv3/4 at the SAP Net support site and can be downloaded with anonymous ftp

User Accounts and Groups

User named SIDadm in the Windows 2000 domain will be created during SAP R/3 installation.

Perform the following:

1. To ease system administration, assign Administrator group privileges to the SIDadm user.
2. On both hosts, use the Windows 2000 User Manager to create two local user groups:
ORA_SID_DBA
ORA_DBA
3. Add SIDadm to these local groups on both hosts.

Installing Oracle

Standard/Database Host

A split mirror backup configuration requires a properly maintained Oracle8 password file and additional initialization parameters on the standard host. Perform the following on the standard host:

1. Create the Oracle8 password file at %ORACLE_HOME%\database:
C:\>orapwd file=%ORACLE_HOME%\database \pwdSID.ora password=<the password for the internal user>
2. Ensure that the following lines exist and are uncommented in the Oracle initialization file at %ORACLE_HOME%\database \initSID.ora:
remote_login_passwordfile =exclusive
remote_os_authent =true
If you make changes to the initSID.ora file, you must restart the database for the changes to take effect.
3. Set the authorizations of the Oracle user SYSTEM:
SVRMGR30>grant SYSOPER to SYSTEM;
SVRMGR30>grant SYSDBA to SYSTEM;
4. The R/3 installation creates Oracle user ops\$SIDadm. If you use Oracle 8.0.5, alter this user with the following commands:
SVRMGR30>alter user ops\$SIDadm identified externally default
tablespace PSAPUSER1D temporary tablespace PSAPTEMP quota
unlimited on PSAPUSER1D quota unlimited on PSAPTEMP;
SVRMGR30>grant connect,resource,dba to ops\$SIDadm;
If you use Oracle 8.1.5 or higher, you also need to assign the sapdba role to the domain OPS\$user, for example:
SVRMGR30>grant sapdba to "SAPDOMAIN \OPS\$SIDADM ";

See R/3 Note Number 0095719, "SAPDBA/BRBACKUP: ORA-1031 after Oracle8 Migration" for more information.

Backup/EVM Agent Host

A split mirror backup requires the Oracle Client software as well as the Oracle Server Manager on the backup host to connect to the production database through SQL*Net. Perform the following on the backup host:

1. Log into the Windows 2000 Administrator account.
2. Refer to the R/3 installation documentation to perform a standard Oracle installation on the backup host into the appropriate Oracle directory (e.g. E:\Oracle\>\817).
3. Log out of the Windows 2000 Administrator account and log into the SIDadm account on the backup host.
4. Set up all the environment variables listed in Table 10, Environment Variables using Settings >Control Panel >System >Environment so that both hosts have the same environment variables for Oracle and SAP (except for SAPLOCALHOST, which defaults to the current hostname).

NET8 Configuration

Since the database resides on the Standard host, the TNS listener service (OracleTNSListener80) needs to be running. R/3 starts the listener automatically, and you can stop and restart OracleTNSListener80 in Settings >Control Panel >Services. The R/3 installation creates the listener.ora and sqlnet.ora files in %ORACLE_HOME%\network\admin (%TNS_ADMIN% as defined in the system environment variables).

1. On the standard host, ensure that the contents of sqlnet.ora include:
AUTOMATIC_IPC =ON
SQLNET.AUTHENTICATION_SERVICES =(NTS)
SQLNET.EXPIRE_TIME =0
2. Both the database and backup hosts require a tnsnames.ora file. The R/3 installation creates this file automatically. Copy this file from the directory %ORACLE_HOME%\network\admin (%TNS_ADMIN% as defined in the system environment variables) on the database host to the backup host. The listener.ora file is not needed on the backup host unless starting a second database instance after the BCVs are split (a task that is not described in this paper).

Configuring BRBACKUP

The steps below describe how to configure the SAP backup tool.

1. Copy all the SAPDBA BR utility files located in %SAPEXE%(in this case, K:\usr\SAP\SID\SYS\exe \run) on the Standard host into the same directory on the backup host. Note that all the other \usr\SAP sub-directories are empty. This allows for the SAP-related environment variables to be identical on both machines.
2. Create the Oracle instance on the backup host:
C:>oradim80 -new -sid <SID>
-pfile=%ORACLE_HOME%\database \initSID.ora -intpwd=<same password as the internal user on the Standard host>
3. In Settings >Control Panel >Services, the service OracleServiceSID is set to start manually by default. Press the Startup button and change this service to start automatically at system boot time.

4. Create the %ORACLE_HOME%\scripts and %ORACLE_HOME%\scripts\logs directories. Create the do_split.cmd and do_resync.cmd files in the Appendix and place them in the scripts directory.
5. Edit the SAP backup configuration file initSID.SAP in the directory %ORACLE_HOME%\database.
The following are the relevant lines to modify:
#Set the following value to be the desired backup type.
#Note that N: drive is shared from the Standard host
backup_type =offline_split /online_split
backup_mode =all
backup_root_dir =N:\Oracle\SID\SAPbackup
compress_dir =N:\Oracle\SID\SAPreorg
archive_copy_dir =N:\Oracle\SID\SAPbackup
split_cmd = "N:\Oracle\scripts\do_split.cmd "
resync_cmd = "N:\Oracle\scripts\do_resync.cmd "
primary_db =SID
6. Copy the Oracle parameter file initSID.ora to the backup host in %ORACLE_HOME%\database. It is good practice to keep the originals of both the initSID.ora and initSID.SAP files on the Standard host and distribute current copies of these files to the backup host as needed.

Testing BRBACKUP

After performing this setup, you can start BRBACKUP for a test run on the backup host.

1. Log in as Windows 2000 with user SIDadm and issue the command C:\>brbackup

Note: Unlike a Standard backup scenario, it is not possible to start BRBACKUP with OPS\$authentication. There is no way to avoid sending the system password across the network. For example, Oracle does not accept the following over the network:

```
CONNECT /@SID as SYSOPER
```

Therefore Oracle will not accept a connect request such as:

```
C:\>brbackup -u /
```

2. Once you test that BRBACKUP runs, you can use the batch files do_split.cmd and do_resync.cmd to run the EVM for Windows 2000 commands for flushing, splitting, mounting and unmounting during BRBACKUP.

See the output of BRBACKUP for the contents of the do_resync.cmd and do_split.cmd files. Instructions for configuring EVM BCVs and a sample of the BRBACKUP output for online split mirror backups can be found in the Appendix.

Restoring Lost Files

Restoring lost files requires manual intervention. The volume to which to restore data has to be presented to the Backup Server. The Backup Server then restore the file or files from tape to the volume. After successfully restore, the volume than 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 over the network.

In some cases it may be necessary or beneficial to use EVM to restore lost files. In this case, the suggested procedure is to:

1. Create a new volume of sufficient size through CLI and make it exclusively visible to the backup server by setting the appropriate access IDS, offsets or both.
2. Reboot the backup server, then mount and format the new volume.
3. Restore the data from tape to the new volume.
4. Make the new volume exclusively visible to the Oracle DB server.
5. Reboot the Oracle DB server and mount the new volume.
6. Manually copy the data from the new volume over to the original source volume.

Alternatively, it is possible to restore data from tape directly to the original unit by presenting the unit to the backup server. Once the data is restored on the unit, the unit can be presented back to the production server.

Applying the Redo Logs

This step in the recovery process will depend on what is being restored and whether the database needs to be open during recovery. Oracle provides methods for performing data file recovery, tablespace recovery and database recovery. For more information please refer Oracle backup/recovery documentation.

Appendix A: EVM-SAP Backup Scripts

Script: do_split.cmd

```
@echo off
ECHO #INFO INFO INFO
ECHO #INFO          Compaq Enterprise Volume Manager
ECHO #INFO Will start doing Quick clone, File system Flush, splitting and mounting
ECHO #INFO Please Wait !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!.....
ECHO #INFO INFO INFO
evmcl swma31k003 execute QuickCloneEVM null > quickcloneevm.log
if %errorlevel% GTR 0 goto error
:ok
Echo #INFO QuickcloneEVM job is executed successful
goto end
:error
c:\backupevm\altresum.cmd > c:\backupevm\altresum.log
Echo #INFO QuickcloneEVM job failed look at the Log for detail information
:end
```

Script: do_undo.cmd

```
@echo off
ECHO #INFO INFO INFO
ECHO #INFO          Compaq Enterprise Volume Manager
ECHO #INFO Will start undoing Quick clone JOB
ECHO #INFO Please Wait !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!.....
ECHO #INFO INFO INFO
evmcl swma31k003 undo QuickCloneEVM null >> quickcloneevm.log
if %errorlevel% GTR 0 goto error
:ok
Echo #INFO QuickcloneEVM job is executed undo successful
goto end
:error
c:\backupevm\altresum.cmd > c:\backupevm\altresum.log
Echo #INFO QuickcloneEVM job failed look at the Log for detail information
:end
```

Output Log from BRBACKUP

```
BR051I BRBACKUP 4.6D (16)
BR055I Start of database backup: bdhclrgo.anf 2002-01-14 11.18.30
BR319I Control file copied to R:\oracle\EVM\sapbackup\cntrl.dbf
BR101I Parameters
oracle_sid      EVM
oracle_home     V:\oracle\EVM\817
oracle_profile  V:\oracle\EVM\817\DATABASE\initEVM.ora
sapdata_home   V:\ORACLE\EVM
sap_profile     V:\oracle\EVM\817\DATABASE\initEVM.sap
backup_mode    ALL
backup_type     online_split
```

```

backup_dev_type   util_file
util_par_file     V:\oracle\EVM\817\DATABASE\initEVM.utl
primary_db        EVM
split_cmd         C:\backupevm\do_split.cmd
resync_cmd        C:\backupevm\do_undo.bat
system_info       evadm EVMAGENT Windows NT 5.0 Build 2195 Service Pack 2 Intel
make_info         NTintel OCI_816_SHARE Oct 11 2001
BR116I ARCHIVE LOG LIST before backup for database instance EVM/PRIMARY
Database log mode      Archive Mode
Automatic archival    Enabled
Archive destination   N:\oracle\EVM\saparch\EVMarch
Archive format        ARC%S.%T
Oldest online log sequence  737
Next log sequence to archive  740
Current log sequence    740      SCN: 3041260
Database block size    8192      Thread: 1
Oracle version         8.1.7.0.0
BR280I Time stamp 2002-01-14 11.18.35
BR315I 'Alter tablespace PSAPBTABD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.35
BR315I 'Alter tablespace PSAPBTABI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.35
BR315I 'Alter tablespace PSAPCLUD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.35
BR315I 'Alter tablespace PSAPCLUI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.35
BR315I 'Alter tablespace PSAPDDICD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPDDICI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPDOCUD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPDOCUI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPEL46CD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPEL46CI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPES46CD begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPES46CI begin backup' successful
BR280I Time stamp 2002-01-14 11.18.36
BR315I 'Alter tablespace PSAPLOADD begin backup' successful
.
.
.
.
BR280I Time stamp 2002-01-14 11.18.38
BR315I 'Alter tablespace PSAPTEMP begin backup' successful
BR280I Time stamp 2002-01-14 11.18.38
BR315I 'Alter tablespace PSAPUSER1D begin backup' successful

```

BR280I Time stamp 2002-01-14 11.18.39
BR315I 'Alter tablespace PSAPUSER1I begin backup' successful
BR280I Time stamp 2002-01-14 11.18.39
BR315I 'Alter tablespace SYSTEM begin backup' successful
BR280I Time stamp 2002-01-14 11.18.39
BR295I Splitting mirror disks...
BR278I Command output of 'C:\backupevm\do_split.cmd':
#INFO INFO
#INFO Compaq Enterprise Volume Manager
#INFO Will start doing Quick clone, File system Flush, splitting and mounting
#INFO Please Wait !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!.....
#INFO INFO
#INFO QuickcloneEVM job is executed successfully
BR280I Time stamp 2002-01-14 14.01.56
BR296I Split of mirror disks successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPBTABD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPBTABI end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPCLUD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPCLUI end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPDDICD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPDDICI end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPDOCUD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPDOCUI end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPEL46CD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPEL46CI end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPES46CD end backup' successful
BR280I Time stamp 2002-01-14 14.01.56
BR317I 'Alter tablespace PSAPES46CI end backup' successful
BR280I Time stamp 2002-01-14 14.01.57
BR317I 'Alter tablespace PSAPLOADD end backup' successful
.
.
.
.
BR280I Time stamp 2002-01-14 14.01.57
BR317I 'Alter tablespace PSAPTEMP end backup' successful
BR280I Time stamp 2002-01-14 14.01.57
BR317I 'Alter tablespace PSAPUSER1D end backup' successful
BR280I Time stamp 2002-01-14 14.01.57
BR317I 'Alter tablespace PSAPUSER1I end backup' successful

```

BR280I Time stamp 2002-01-14 14.01.57
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       N:\oracle\EVM\saparch\EVMarch
Archive format            ARC%S.%T
Oldest online log sequence 738
Next log sequence to archive 740
Current log sequence      741      SCN: 3048750
Database block size      8192      Thread: 1
Oracle version           8.1.7.0.0
BR118I Tablespaces and data files
PSAPBTABD ONLINE* ONLINE I:\ORACLE\EVM\SAPDATA1\BTABD_1\BTABD.DATA1
          2097160192 7 8 NOLINK FILE
PSAPBTABI ONLINE ONLINE F:\ORACLE\EVM\SAPDATA2\BTABI_1\BTABI.DATA1
          2097160192 6 5 NOLINK FILE
PSAPCLUD ONLINE* ONLINE F:\ORACLE\EVM\SAPDATA2\CLUD_1\CLUD.DATA1
          2097160192 11 5 NOLINK FILE
PSAPCLUI ONLINE ONLINE J:\ORACLE\EVM\SAPDATA3\CLUI_1\CLUI.DATA1
          2097160192 10 9 NOLINK FILE
PSAPDDICD ONLINE* ONLINE I:\ORACLE\EVM\SAPDATA1\DDICD_1\DDICD.DATA1
          2097160192 13 8 NOLINK FILE
PSAPDDICI ONLINE ONLINE O:\ORACLE\EVM\SAPDATA5\DDICI_1\DDICI.DATA1
          2097160192 12 14 NOLINK FILE
PSAPDOCUD ONLINE* ONLINE O:\ORACLE\EVM\SAPDATA5\DOCUD_1\DOCUD.DATA1
          2097160192 15 14 NOLINK FILE
PSAPDOCUI ONLINE ONLINE J:\ORACLE\EVM\SAPDATA3\DOCUI_1\DOCUI.DATA1
          2097160192 14 9 NOLINK FILE
PSAPEL46CD ONLINE* ONLINE
.
.
.
.
F:\ORACLE\EVM\SAPDATA2\SOURCED_1\SOURCED.DATA1 2097160192 27 5 NOLINK FILE
PSAPSOURCEI ONLINE ONLINE
          G:\ORACLE\EVM\SAPDATA4\SOURCEI_1\SOURCEI.DATA1 2097160192 26 6 NOLINK
          FILE
P:\ORACLE\EVM\SAPDATA6\STABD_1\STABD.DATA1 2097160192 5 15 NOLINK FILE
PSAPSTABI ONLINE ONLINE O:\ORACLE\EVM\SAPDATA5\STABI_1\STABI.DATA1
          2097160192 4 14 NOLINK FILE
PSAPTEMP ONLINE* ONLINE P:\ORACLE\EVM\SAPDATA6\TEMP_1\TEMP.DATA1
          2097160192 3 15 NOLINK FILE
PSAPTEMP ONLINE* ONLINE P:\ORACLE\EVM\SAPDATA6\TEMP_2\TEMP.DATA2
          2097152000 30 15 NOLINK FILE
PSAPTEMP ONLINE* ONLINE P:\ORACLE\EVM\SAPDATA6\TEMP_3\TEMP.DATA3
          2097152000 31 15 NOLINK FILE
PSAPUSER1D ONLINE* ONLINE G:\ORACLE\EVM\SAPDATA4\USER1D_1\USER1D.DATA1
          2097160192 28 6 NOLINK FILE

```

PSAPUSER1I ONLINE ONLINE F:\ORACLE\EVM\SAPDATA2\USER1I_1\USER1I.DATA1
2097160192 29 5 NOLINK FILE
SYSTEM ONLINE* SYSTEM I:\ORACLE\EVM\SAPDATA1\SYSTEM_1\SYSTEM.DATA1
2097160192 1 8 NOLINK FILE
BR119I Redo log files
L:\ORACLE\EVM\ORIGLOGA\LOG_G11M1.DBF 20979712 11 11 INUSE NOLINK FILE
M:\ORACLE\EVM\MIRRLOGA\LOG_G11M2.DBF 20979712 11 12 INUSE NOLINK FILE
L:\ORACLE\EVM\ORIGLOGB\LOG_G12M1.DBF 20979712 12 11 INUSE NOLINK FILE
M:\ORACLE\EVM\MIRRLOGB\LOG_G12M2.DBF 20979712 12 12 INUSE NOLINK FILE
L:\ORACLE\EVM\ORIGLOGA\LOG_G13M1.DBF 20979712 13 11 INUSE NOLINK FILE
M:\ORACLE\EVM\MIRRLOGA\LOG_G13M2.DBF 20979712 13 12 INUSE NOLINK FILE
L:\ORACLE\EVM\ORIGLOGB\LOG_G14M1.DBF 20979712 14 11 INUSE NOLINK FILE
M:\ORACLE\EVM\MIRRLOGB\LOG_G14M2.DBF 20979712 14 12 INUSE NOLINK FILE
BR120I Control files
I:\oracle\EVM\sapdata1\cntrl\cntrlEVM.dbf 8036352 0 8 NOLINK FILE
F:\oracle\EVM\sapdata2\cntrl\cntrlEVM.dbf 8036352 0 5 NOLINK FILE
J:\oracle\EVM\sapdata3\cntrl\cntrlEVM.dbf 8036352 0 9 NOLINK FILE
BR280I Time stamp 2002-01-14 14.01.58
BR057I Backup of database: EVM
BR058I BRBACKUP action ID: bdhclrgo
BR059I BRBACKUP function ID: anf
BR110I Backup mode: ALL
BR077I Database file for backup: R:\oracle\EVM\sapbackup\cntrl.dbf
BR061I 40 files found for backup, total size 78007.891 MB
BR143I Backup type: online split
BR130I Backup device type: util_file
BR109I Files will be saved by backup utility
BR126I Unattended mode active - no operator confirmation required
BR280I Time stamp 2002-01-14 14.01.58
BR229I Calling backup utility...
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... P:\ORACLE\EVM\SAPDATA6\ES46CD_2\ES46CD.DATA2
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... P:\ORACLE\EVM\SAPDATA6\TEMP_2\TEMP.DATA2
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... F:\ORACLE\EVM\SAPDATA2\SOURCED_1\SOURCED.DATA1
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... P:\ORACLE\EVM\SAPDATA6\STABD_1\STABD.DATA1
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... I:\ORACLE\EVM\SAPDATA1\SYSTEM_1\SYSTEM.DATA1
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... G:\ORACLE\EVM\SAPDATA4\USER1D_1\USER1D.DATA1
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... J:\ORACLE\EVM\SAPDATA3\LOADI_1\LOADI.DATA1
#SAVED.... NS1011045722

```

BR280I Time stamp 2002-01-14 15.01.13
#FILE..... G:\ORACLE\EVM\SAPDATA4\POOLI_1\POOLI.DATA1
#SAVED.... NS1011045722
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... F:\ORACLE\EVM\SAPDATA2\BTABI_1\BTABI.DATA1
#SAVED.... NS1011045722
.
.
.
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... P:\ORACLE\EVM\SAPDATA6\ES46CD_1\ES46CD.DATA1
#SAVED.... NS1011047569
BR280I Time stamp 2002-01-14 15.01.13
#FILE..... G:\ORACLE\EVM\SAPDATA4\ES46CI_1\ES46CI.DATA1
#SAVED.... NS1011047569
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... V:\oracle\EVM\817\DATABASE\initEVM.ora
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... V:\oracle\EVM\817\DATABASE\initEVM.dbf
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... V:\oracle\EVM\817\DATABASE\initEVM.sap
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... V:\oracle\EVM\817\DATABASE\initEVM.utl
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... N:\oracle\EVM\sapreorg\reorgEVM.log
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... N:\oracle\EVM\sapreorg\structEVM.log
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... R:\oracle\EVM\sapbackup\bdhclrgo.anf
#SAVED.... NS1011049279
BR280I Time stamp 2002-01-14 15.10.07
#INLOG.... R:\oracle\EVM\sapbackup\backEVM.log
#SAVED.... NS1011049279
*****
BR280I Time stamp 2002-01-14 15.10.07
BR232I 8 of 8 files saved by backup utility
BR230I Backup utility called successfully
BR280I Time stamp 2002-01-14 15.10.08
BR297I Resynchronizing mirror disks...
BR280I Time stamp 2002-01-14 15.10.08
BR298I Resynchronization of mirror disks successful
BR056I End of database backup: bdhclrgo.anf 2002-01-14 15.10.07
BR280I Time stamp 2002-01-14 15.10.08
BR052I BRBACKUP terminated successfully

```

Appendix B: Oracle Architecture on Windows 2000 for SAP

An Oracle instance on Windows 2000 is a single process composed of multiple threads. Thus on Windows 2000, DBWR, LGWR, ARCH and so on are all internal threads of the Oracle process.

- An area of memory available to all Oracle sessions is known as the System Global Area (SGA). This area of memory includes recently accessed data blocks (the buffer cache), SQL, SAP objects, and transaction information (the redo log buffer). The SGA may also contain session information.
- Several tasks that perform dedicated database activities, including the database writer (DBWR), redo log writer (LGWR), system monitor (SMON), process monitor (PMON), and log archiver (ARCH). These tasks are known as background tasks (although they are also often referred to as background processes).
- SAP, Oracle data files that contain the tables, indexes, and other segments that form the Oracle instance.
- Redo logs, which record critical transaction information required for roll-forward in the event of instance failure.
- A SQL*Net listener task that establishes connections from external systems.

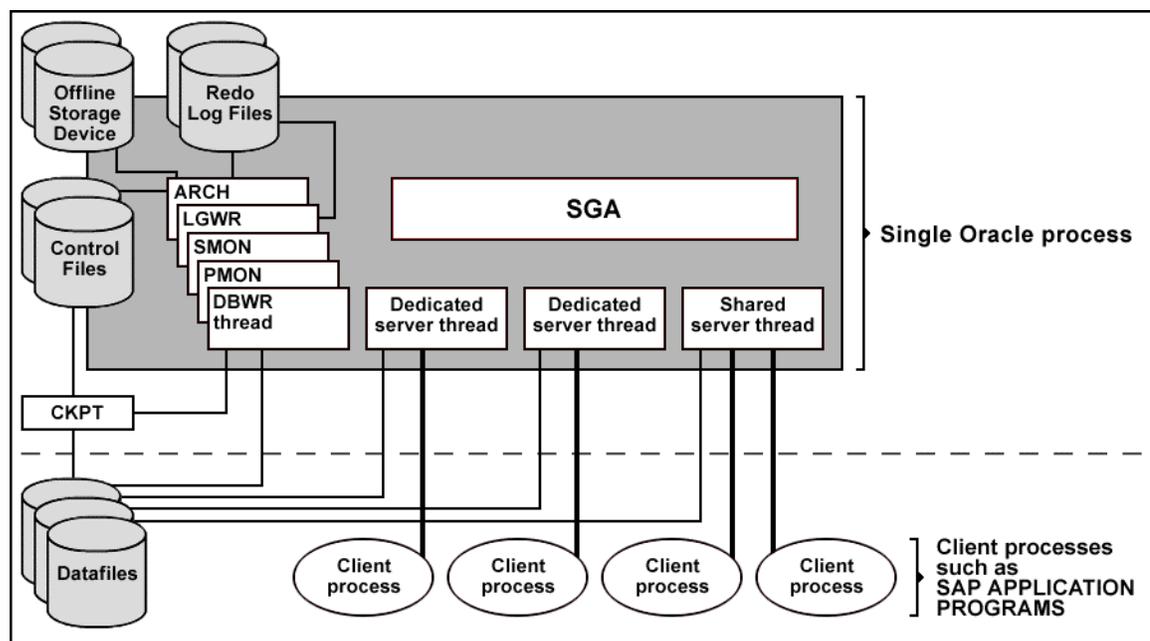


Figure 20, Oracle Overview for SAP R/3

Legend:

ARCH: Archiver	LGWR: Log Writer
SMON: System Monitor	PMON: Process Monitor
DBWR: Database Writer	CKPT: Checkpoint

The thread architecture for Oracle has been proven to be very stable, maintainable and efficient. In addition, there are other benefits that arise from the thread architecture. These include faster operating system context switches amount threads (as opposed to process), a much simpler SGA allocation routine which does not require the use of shared memory, faster spawning of new connections since threads are created more quickly than processes, decreased memory usage since threads share more data structures than processes do, and finally a perception that a thread-based model is somehow more “Windows 2000-like” than a process-based one.

Database and redo log files are generally implemented using the operating system’s native file system or raw disk partitions; port-specific differences at the file level are relatively minor. However, the memory and process structure of an Oracle instance will vary significantly depending on how the operating system implements process and memory management.

In an SAP production database environment the database system must always run in ARCHIVELOG mode. An archiver thread archives a completed online redo log file into the offline redo log file in the archive directory - \ORACLE\\saparch

The Oracle database is stored in 8-KB blocks in data files on disk. In order to accelerate read/write access to the data, these data blocks are cached in the database buffer pool in the production server main memory. The Oracle database management system holds the executable SQL statements in the shared SQL area, which is part of the shared pool.

SAP Kernel Processes

Any computer is capable of running one or more application instances. The main purpose of a dialog instance is to intercept requests for work from the SAP clients and to execute ABAP programs to service the requests for data. In addition, a dialog instance contains a dispatcher task and a set of work processes (WPs).

Each SAP R/3 Dialog, Batch, Update, and Spool work process in an SAP instance uses TCP/IP and TNS listener to:

- Connect to the database as one SAP R/3 user
- Handle database requests for the different R/3 system users
- Communicate between the application servers and the database
- Communicate between the client and the application servers

The dispatcher on an SAP application server receives requests from the SAP users as in Figure 21. In cases where a computer is running more than one dialog instance, there is one dispatcher for each dialog instance.

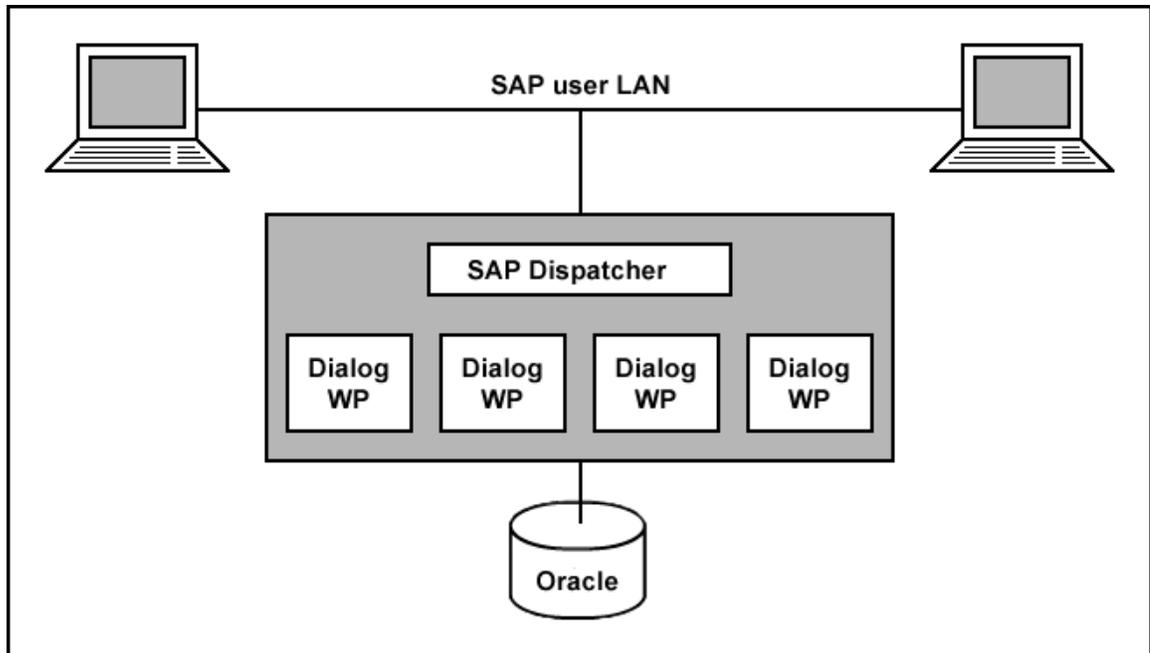


Figure 21, SAP Overview

The following table provides more details about the function and distribution of each of the processes running within a system. An instance is a set of processes and shared memory, which is a part of the system. In all cases, the network interconnection between the mySAP.com components is made using the TCP/IP protocol.

Process	Description	Distribution in System
Dispatcher	<ul style="list-style-type: none"> Queues user requests for processing by dialog processes. The user transaction is handled by the next available dialog process Users are not allocated to a specific dialog process 	1 per instance
Dialog	<ul style="list-style-type: none"> Processing of ABAP/4 programs Interaction with users. 	Multiple per instance
Batch	<ul style="list-style-type: none"> These processes are used to put long running queries into the background and can be scheduled. Batch processes are basically the same entity as a dialog process, but do not accept user input. 	Multiple per instance
Update	<ul style="list-style-type: none"> Modifications to the database changes are managed by update processes and can be done asynchronously to the dialog processes. This improves the user's perceived response time. 	Multiple per instance depending on number of dialog / batch processes and application
Enqueue	<ul style="list-style-type: none"> SAP implements its own data lock manager, known as the enqueue process. 	1 per system
Spool	<ul style="list-style-type: none"> The SAP system contains a comprehensive print spooling system. 	1 per instance
Message Server	<ul style="list-style-type: none"> Manages communication between the work processes, within an instance and between separate instances. 	1 per system

Table 11, SAP Process Overview

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.

EVM 2 Clone/Split supports RAID 1 or RAID 1+0. With HSG80 robust RAID 1+0 technology and its use of striping, hot spots can be eliminated 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” 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. 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: Data Flow from Oracle to SAP

One of the confounding aspects of the Oracle/SAP architecture is that many layers exist in the SAP environment. On its way to the SAPGUI client, Oracle data must be read from disk into the local disk cache, transferred onto the database server file system cache, and then transferred into the Oracle buffer cache. Once the data reaches Oracle, it is shipped to the application server, where it is cached again before making the final trip to the SAPGUI client.

The memory buffers are used to save data from a prior disk read so that they do not have to be reread the next time the data is required. In an Oracle/SAP configuration, there are many layers of data buffers, and each buffer caches much of the same data as shown in Figure 22.

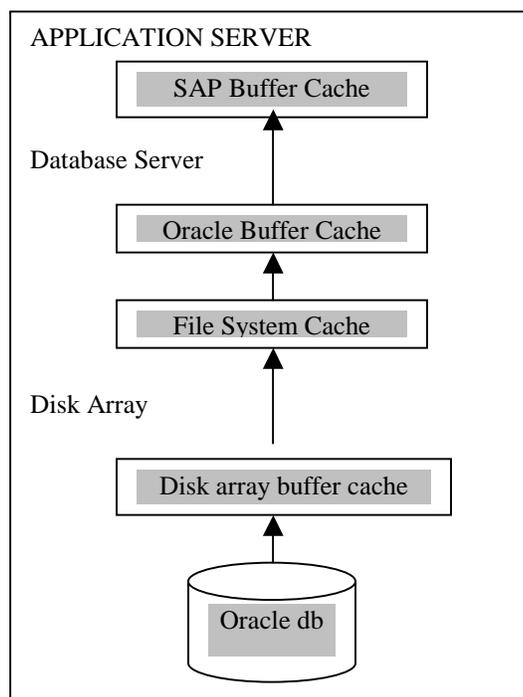


Figure 22, SAP Buffer Overview

However, this redundant caching does not mean that all of the extra buffers are wasted. As data makes the trip from the disk to the SAP client, the data cached in each buffer is successively refined. For example, at the disk level, a physical I/O may result in an entire track of data being stored on the disk array cache. As the data reaches Oracle, only a single database block is cached. Once the data reaches the SAP application server, only specific row information will be stored in the SAP buffer.

Appendix E: 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 Windows 2000 integrates the Oracle backup/recovery capabilities of the following SAP Tools with the backup and recovery management.

- The program BRBACKUP brings database servers on-line or off-line, checks the status of SAP files, and places database tablespaces into BACKUP mode to guarantee their data consistency. BRBACKUP provides on-line or off-line backup of control files, data files, and on-line 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 Windows 2000 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 Windows 2000 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 23 illustrates the components and architecture for the SAP/Oracle backup solution.

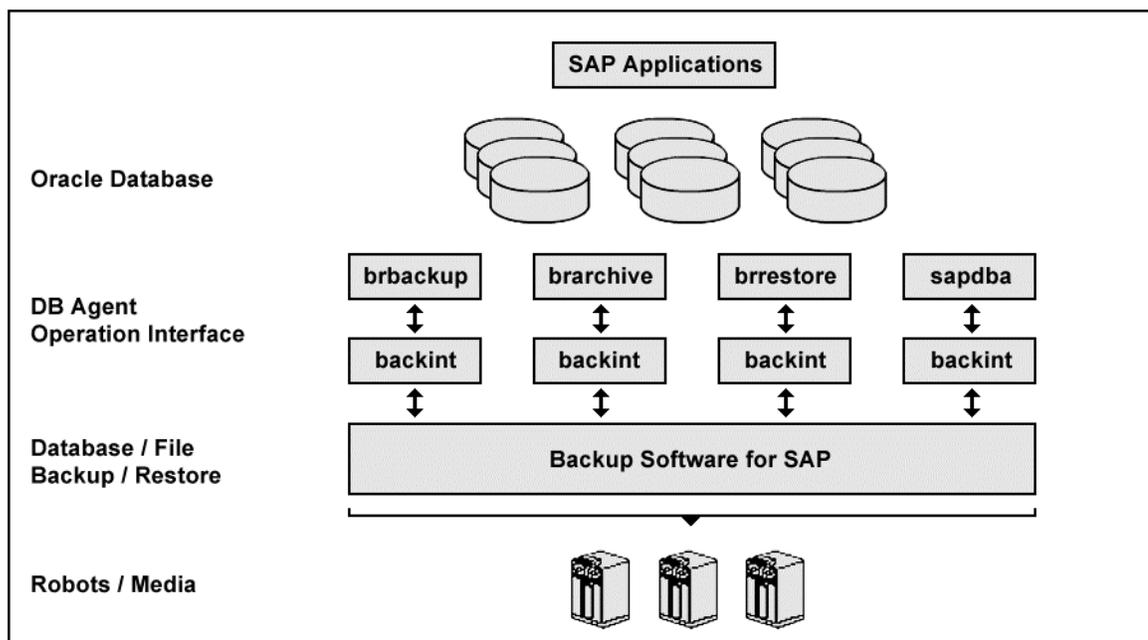


Figure 23, SAP/Oracle Backup Overview

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 Windows 2000 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.

Sequence of Operation

Figure 24 is an example of VERITAS NetBackup sequence operation. NetBackup users or schedules start database backups or restores by selecting an SAP script. A NetBackup process called bphdb starts the SAP script on the client. The SAP application then starts the requested operation on the databases.

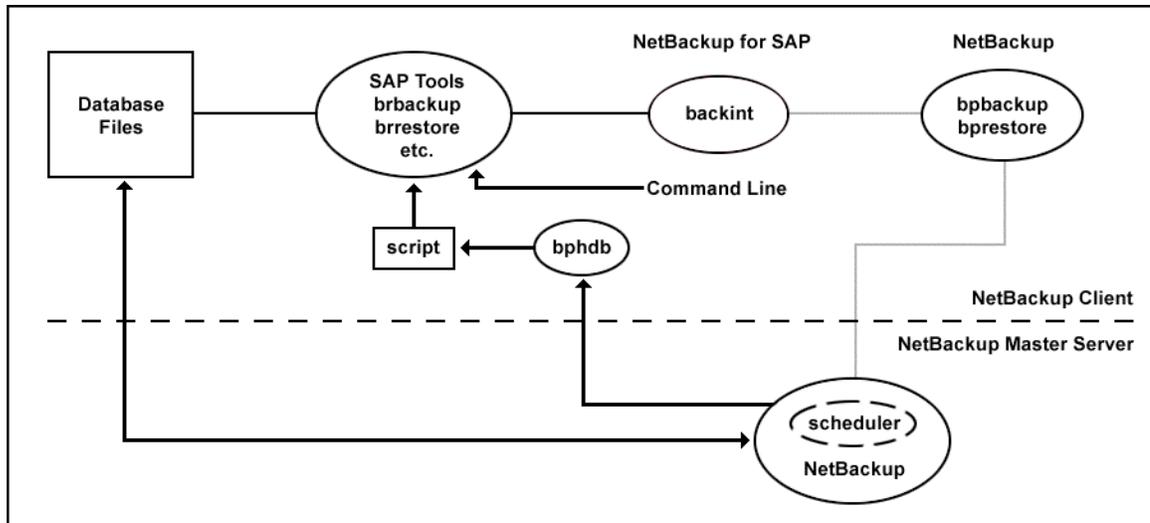


Figure 24, Example of VERITAS NetBackup

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

A restore works in essentially the same manner except that the NetBackup for SAP on Windows NT backint interface issues a bprestore command, causing NetBackup to retrieve the data from secondary storage and send it to the client.

Appendix F: HSG80 Volume Layout for SAP R/3-ORACLE

Depending on the number of disks and the size of the SAP/Oracle database one has to make decision how to configure the disk subsystem. For performance reason we recommend RAID 1+0. In our test we have ensured the availability of the disk subsystem with the following options.

- Online spare disks
- Redundant power supplies
- Dual array controllers
- Mirrored cache modules
- Redundant external cache batteries

We have installed redundant HSG80 Array Controllers. This ensures that if one controller fails, the other controller will take over the subsystem. We have configured two cache modules to mirror each other to enable snapshot capabilities. The subsystem is configured in MultiBus Failover mode in order to have a redundant access path to the storage, and the additional hardware gives you the option to balance the load over the two access paths from server to storage. The following subsystem HSG80 configuration example with the SAP server has two host bus adapters installed, two SAN 16 switches and dual HSG80 Array Controllers. This effectively results in no single point of failure from the server down to the storage.

Configuration Rules

Before you configure your controller, review these configuration rules and ensure your planned configuration meets the following requirements and conditions:

- Maximum 128 assignable, host-accessible units
- Maximum 512-GB unit capacity
- Maximum 20 RAID 5 storage sets
- Maximum 30 RAID 5 and RAID 1 storage sets
- Maximum 45 RAID 5, RAID 1, and RAID 0 storage sets
- Maximum 8 partitions per storage set or individual disk
- Maximum 6 members per mirror set
- Maximum 14 members per RAID 5 storage set
- Maximum 24 members per stripe set
- Maximum 48 physical devices per striped mirror set

The HSG80 controller has six SCSI device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (The StorageWorks Command Console GUI calls the device ports “channels.”).

StorageWorks Enterprise Modular Array 12000 / 4214 Disk Layout

In our test we have used the following StorageWorks disk subsystem design/configuration. It will help you to design the StorageWorks and unit creation.

		Target Numbers													
		0	1	2	3	4	5	8	9	10	11	12	13	14	15
Device port numbers	1	M1B 100	M1D 101	M1F 102	M3B 103	M3D 104	M3F 105	M5A 108	M5C 109	M5D 110	M5E 111	M5F 112	SP 113	SP 114	SP 115
	2	M1B 200	M1D 201	M1F 202	M3B 203	M3D 204	M3F 205	M5A 208	M5C 209	M5D 210	M5E 211	M5F 212	SP 213	SP 214	SP 215
	3	M1B 300	M1D 301	M1F 302	M3B 303	M3D 304	M3F 305	M5A 308	M5C 309	M5D 310	M5E 311	M5F 312	SP 313	SP 314	SP 315
	4	M1A 400	M1C 401	M1E 402	M3A 403	M3C 404	M3E 405	SP 408	M5B 409	SP 410	SP 411	SP 412	SP 413	SP 414	SP 415
	5	M1A 500	M1C 501	M1E 502	M3A 503	M3C 504	M3E 505	OBIN 508	M5B 509	SAP 510	MRA 511	MRB 512	RDA 513	SP 514	SP 515
	6	M1A 600	M1C 601	M1E 602	M3A 603	M3C 604	M3E 605	OBIN 608	M5B 609	SAP 610	MRA 611	MRB 612	RDA 613	SP 614	SP 615

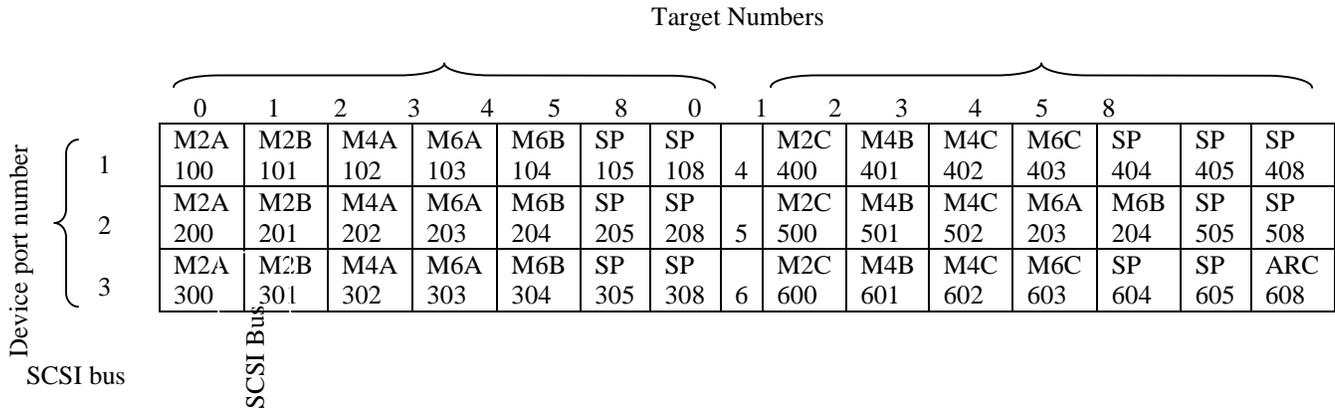
SCSI bus

SCALE:

100>DISK10000,101>DISK10100,200>DISK20000,208>DISK20800
303>DISK30300,404>DISK40400,505>DISK50500,606>DISK60600 SP>Spare disks

RAID	Member	Unit	Description	Disks	Tot. cap.
R 1+0	M1*	D1	SAPDATA1	18 x 18.2 GB	101.7 GB
R 1+0	M3*	D3	SAPDATA3	18 x 18.2 GB	101.7 GB
R 1+0	M5*	D5	SAPDATA5	18 x 18.2 GB	101.7 GB
R1	RDA	D7	REDO LOGA	2 x 18.2 GB	18.2 GB
R1	OBIN	D9	ORACLE BIN	2 x 18.2 GB	18.2 GB
R1	SAP	D10	SAP KERNEL	2 x 18.2 GB	18.2 GB
R1	MR A	D12	MIRROR A	2 x 18.2 GB	18.2 GB
R1	MR B	D13	MIRROR B	2 x 18.2 GB	18.2 GB

StorageWorks Modular Array 8000 / 4254 Disk Layout



SCALE:

1001DISK10100,404>DISK40400,202>DISK20200,505>DISK50500,303>DISK30300,606>DISK60600,SP>SPARE

RAID Set	Member	Unit	Description	Disk Cap.	Tot. Cap.
R 1+0	M2*	D2	SAPDATA2	9 x 36.4 GB	109.2 GB
R 1+0	M4*	D4	SAPDATA4	9 x 36.4 GB	109.2 GB
R 1+0	M6	D6	SAPDATA6	9 x 36.4 GB	109.2 GB
R1	RDB	D8	REDOLOGB	2 x 18.2 GB	18.2 GB
R1	ARC	D11	ARCHIEVE	2 x 36.4 GB	36.4 GB

References

1. Oracle 24 x 7 Tips and Techniques by Venkat S. Devraj, Oracle Press
2. Oracle SAP Administration by Donald K. Burleson O'Reilly press
3. VERITAS NetBackup for SAP System Administrator Guide
4. Oracle 8 DBA handbook by Kevin Loney Oracle Press