

hp storage solutions

january 2003



technical
blueprint

nondisruptive backup for Oracle 9i on Unix

table of contents

executive summary	2
solution overview	2
test environment	3
notes on installation	4
performance results	5
creating and running EVM jobs	8
sample EVM scripts used in testing	12
guidelines for cloning an Oracle 9i storage unit with EVM	14
EVM job interactions with Oracle 9i applications	14
supported configurations	15
supported Web browsers	15
EVM agent requirements	16
for more information	17
hp StorageWorks	17
Enterprise Volume Manager	17
Enterprise Backup Solution	17
Backup software	17
hp services	17

executive summary

This guide describes a quick and complete solution for backing up an Oracle 9i database with minimal disruption to end users and production activities. Using the capabilities of HP StorageWorks Enterprise Volume Manager (EVM) V2 and HP StorageWorks RAID arrays, database administrators can back up databases that are multiple terabytes in size while avoiding both the usual application downtime and the substantial performance degradation that other solutions entail. Testing results show the effectiveness of this solution in providing a reliable backup infrastructure for end-to-end protection of Oracle 9i data.

This solution is based on the cloning capabilities of EVM. Because of its performance and availability characteristics, cloning is the recommended method for data replication of large databases that are read/write intensive. By using integrated technology components in a Storage Area Network (SAN), EVM helps protect Oracle 9i data while allowing applications to remain online with unimpeded performance. A clone is a physical point-in-time copy of an Oracle 9i storage unit and is available as a resource within your overall backup strategy.

This guide provides the following information:

- Tests conducted on an Oracle 9i database under the Sun Solaris 8 operating system in a SAN using EVM
- Performance results
- Best practices for deploying and managing EVM backup jobs

The intended audience is experienced Oracle administrators, who understand how to work with clones, using array controllers in heterogeneous computing environments.

solution overview

Typically, Oracle databases and applications need to be available 24 x 7, yet prudent IT management principles require backup copies for disaster recovery or when testing new applications. In the past, databases were traditionally backed up using offline methods. However, because offline procedures seriously disrupt service to end-users, these procedures are unacceptable in environments that require 24 x 7 availability. For large databases, the impact on uptime can be many hours or even days. By leveraging the capabilities of EVM in a Fibre Channel SAN environment, very large Oracle databases that span multiple storage subsystems—even those separated geographically—can be backed up quickly and with minimal disruption.

Database vendors have developed tools that can perform online backups during normal operations. These tools typically reduce downtime. For example, Oracle 9i uses Recovery Manager (RMAN), which also handles other tasks related to backup management. With RMAN, administrators can keep a database up and running during backup, although performance still will be affected. (Detailed information about RMAN is available in the *Oracle 9i Recovery Manager User's Guide*.) In contrast, by using EVM, performance efficiency is diminished only while the clone is actually being split off.

By testing with a simulated banking application load, HP has demonstrated that significant application uptime benefits result when using EVM along with HP StorageWorks Enterprise Backup Solution (EBS) to replicate and back up an Oracle 9i database under Sun Solaris 8.

This solution includes the features of EVM cloning technology and customizable scripts to automate the backup operations of Oracle 9i databases. With EVM, you can create, run, and manage automated storage replication jobs, as well as link them with external

jobs. EVM eliminates the need for a backup window by backing up from the clone independently from application processing. In addition EBS provides a consolidated backup infrastructure, composed of multiple servers connected over a high-speed, independent Fibre Channel SAN to centrally attached SCSI tape libraries. With this approach, EBS achieves flexible data protection that shrinks backup windows from hours to minutes.

The testing for this solution was conducted using the HP StorageWorks Modular Array (MA) 8000 and HSG80 controllers. However, the solution will also achieve the same gains using other HSG80-based storage arrays: Enterprise Modular Array (EMA) 12000, RAID Array (RA) 8000, and Enterprise Storage Array (ESA) 12000.

This solution provides the following benefits to database administrators:

- Complete backup infrastructure for Oracle 9i that has been tested to produce known results
- Data protection and data integrity for multiple terabyte-size databases
- Investment protection through the maximization of existing hardware and software
- Elimination of need for additional products and services (except those required for EVM in a SAN-based infrastructure)
- Maximized application availability by offloading the backup responsibility of the application server
- Increased profit potential by eliminating planned and unplanned downtime
- Simplified backup management by providing the tools and information needed to implement a proven, pre-tested, pre-configured solution, saving time and money
- Sample scripts that you can customize for your environment for rapid implementation
- Integration with third-party tape backup applications—specifically tested with VERITAS NetBackup DataCenter

test environment

This section describes the configurations used for the testing. The main purpose of these tests was to show the functionality of EVM and demonstrate how it can be used in an Oracle environment. Therefore, no effort was made to achieve maximum absolute backup performance. Instead, using the following configurations, a number of different scenarios were tested to show the relative performance results of backing up Oracle databases in a variety of ways.

- Database server: The database server running Oracle was a Sun 4500 system running under the Sun Solaris 8 operating system. To generate a load on this system, a widely used benchmark—TPC-B—was chosen, which provided a banking application environment with simulated write-intensive transactions. A 23-GB Oracle database was created for this application, and database files were placed on a RAID 1+0 storage set consisting of 8 x 18-GB 10000-rpm disks. Redo logs were placed on another RAID 1+0 storage set consisting of 4 x 9-GB 10000-rpm disks.
- Load generator: A Windows 2000-based HP ProLiant system was used for load generation, using a widely accepted load simulator. This allows the number of users, wait times, and other parameters to be easily set for the creation of a suitable load. The software also collects transactions-per-second (TPS) and timing data during simulations and stores these in a format that is easy to import into other applications.
- Backup server: The database server, running VERITAS NetBackup DataCenter in an EBS configuration with a small DLT-based tape library connected to the SAN, was used to back up the Oracle database. Because a clone can be mounted onto another server, the use of a separate backup server to run the backup software would have

improved the results slightly for the clone-based backup. However, because backing up without a clone must be done from the primary Oracle database server, no such improvement opportunity exists when using RMAN for backup.

This solution was tested using VERITAS NetBackup DataCenter V3.4 software for the actual backup operations. Other backup software packages, such as HP OpenView Storage Data Protector and Legato NetWorker, offer similar functionality and would expect to exhibit similar performance when run under the same conditions as this test.

- Enterprise Volume Manager: EVM is browser-based storage management software that facilitates controller-based clone operations. You can create, run, and manage automated storage replication jobs. EVM runs on the management appliance. Designed to connect directly to the SAN fabric, the management appliance performs management functions without involving host computers.

notes on installation

- Oracle: Tests were performed using Oracle 9i release 2 for Sun Solaris.
- VERITAS NetBackup: The Oracle module for VERITAS NetBackup was employed to enable NetBackup to operate smoothly with RMAN, when conducting online backup tests without EVM.
- HSG80 controllers and HSG80 firmware V8.6S-4 was installed.
- Enterprise Volume Manager: When installing EVM, be sure you have a thorough understanding of EVM operations. Follow the guidelines provided in the EVM Version 2.0D (Multi-OS) Installation Guide (Part Number AA-RTB5A-TE). Also be sure to review the accompanying release notes.

A typical EVM network is illustrated in Figure 1. The EVM server is installed, runs on the management appliance (#1 in Figure 1), and manages the EVM network and jobs. EVM agents are installed on host computers (#2 and # 3 in Figure 1) and run the EVM jobs created for volume management. In this network, server #3 is configured as the backup server, so the database server is relieved of that workload (although the testing was done with a single server performing as both database server and backup server).

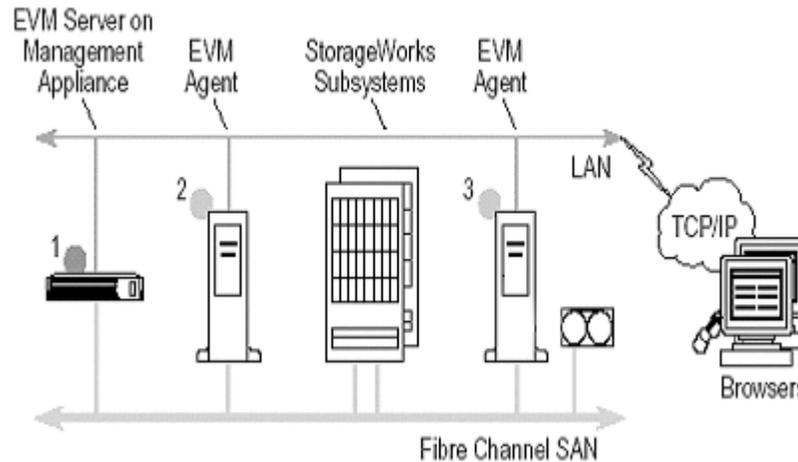


figure 1. typical EVM network

performance results

Oracle 9i databases are typically large, and backup operations can be potentially disruptive to applications and users. The tests imposed a heavy simulated workload on an Oracle database, 23 GB in size, configured with RAID 1+0 storage arrays. Figures 2 and 3 illustrate application availability results during backup of these databases while the application continued to run.

These figures show measured response times from simulated transactions running against the Oracle database. Figure 2 shows response times with a medium transaction workload running on the database, while Figure 3 shows response times under a heavy transaction workload. The curve labeled TP shows the response times from the transaction-processing load without any other EVM activities running. The RMAN curve shows the response times when an RMAN-controlled backup job is started just after the simulated load starts running.

The CLONE curve shows the response times when the EVM job (shown in Figures 5a and 5b) is started just after the simulated load starts running. During the first 5 to 10 minutes, EVM creates the clones, mounts them, and starts the backup. This can be seen from the spikes in response times at the beginning of the job. Splitting the clones and mounting them on the system slows down response times for a few minutes. The simulation runs for about 70 minutes; both RMAN and CLONE finished the backup in about 1 hour. In the CLONE tests, the mirrorsets were initially configured with three members, which enabled EVM to split off the third member almost instantaneously. Using two-member mirrorsets also would have been possible. In such a case, EVM would have to add the third member and normalize it before splitting it off. These activities incur limited performance penalty because they are accomplished entirely by the storage controllers.

Another test with a high transaction load was performed to demonstrate the impact of the normalization phase. The normalization of the third member took about 40 minutes. During this time, the average response time was 86 ms (Figure 4). Note that starting with a two-member mirrorset implies a normalization during the CLONE/NORMALIZE commands. With a three-member mirrorset, normalization occurs during the EVM UNDO job. This means that it is possible to choose the time when it is most convenient to perform the normalization, based on application availability and performance requirements.

Average ms response - medium load

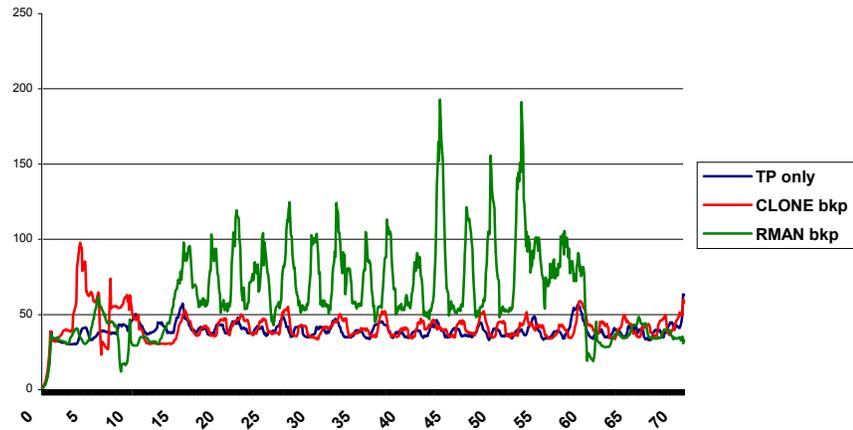


figure 2

Average ms response - high load

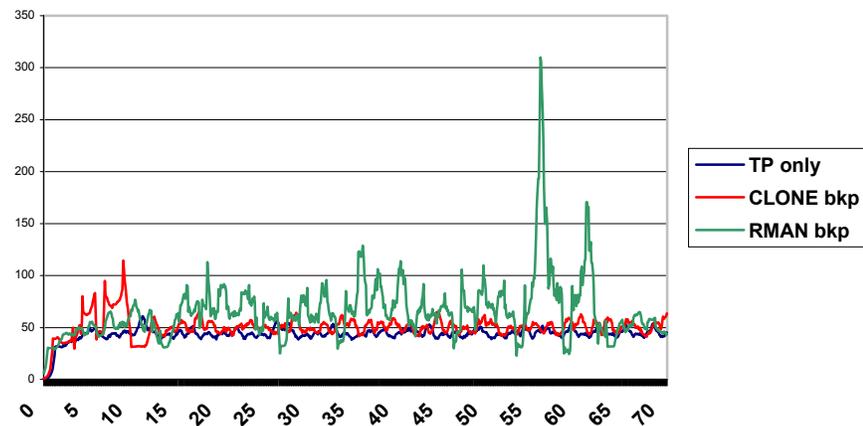


figure 3

The average response times shown in these figures are moving geometric averages over 1 minute for each case. The response times during a CLONE backup are only slightly higher than response times without any backup running. When using RMAN for backup, the database must be in online backup mode for the full duration of the backup operation, which creates a greater impact on performance. Table 1 summarizes the performance impacts of backing up from a clone and backing up using RMAN, compared to the response times of running the application transactions alone, with no backup running.

table 1. average response times during backup

	medium transaction load	heavy transaction load
TP only, no backup	41 ms	46 ms
backup from clone	43 ms	52 ms
backup with RMAN	127 ms	137 ms

Normalization impact

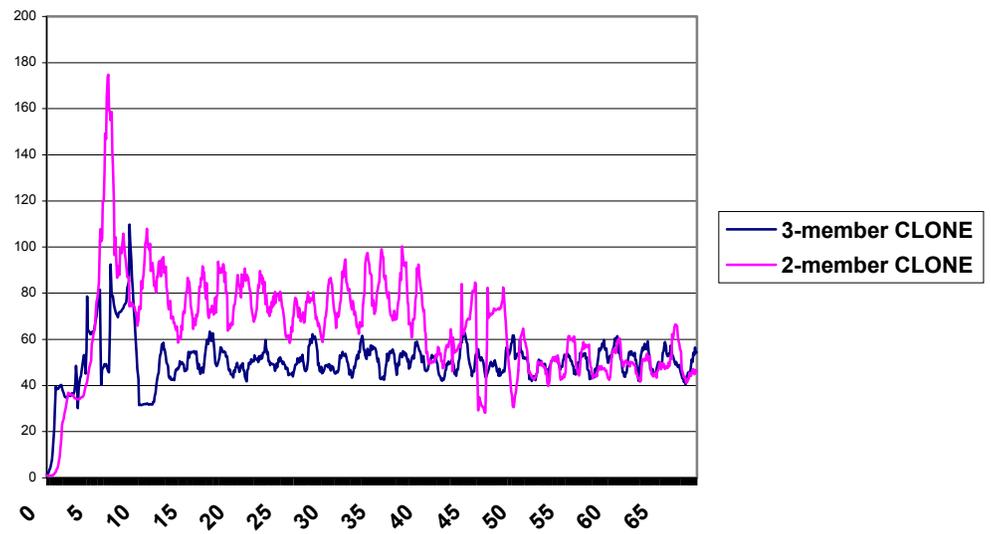


figure 4

This solution provides the key advantage of allowing EVM to perform LAN-free backups. Backup data is moved across the SAN and kept isolated from the general-purpose LAN. As a result, network performance degradation is eliminated during backup.

creating and running EVM jobs

EVM manages storage unit replication jobs, automatically executing the correct commands to create clones, as well as mounting them on a server. The clones are then available for a variety of purposes, including backup or integrity checks. Clones also can remain unmounted for use as near-line backups similar to snapshots.

For performance purposes, a clone can be built using fast or normal modes. Fast mode builds the clone at 20 MB/s (72 GB/hr), while normal mode builds it at 7 MB/s (25 GB/hr). Normal mode runs at a lower priority than the host I/O and is the default and recommended setting.

When replicating an Oracle 9i storage unit, the most important factors are the capacity and number of disks being normalized. For a moderate I/O load on an HSG80 controller and an adequate configuration of disks in the array, a replication rate of 60 GB/hr can be achieved.

By making clones part of your backup strategy, you can back up directly to tape or across the SAN. Because clones are offline, users are protected from performance degradation.

Table 2 provides an overview of the features and applications of the cloning replication method.

table 2. overview of features and applications

feature	EVM clone
storage unit replication	Creates physical copy of storage unit
production impact	During creation and normalization, original unit remains online
units supported	JBOD, RAID 0, RAID 1, or RAID 1+0
applications	Backups, data warehousing, testing, and system migrations

EVM V2 makes it easy to create and manage storage unit replication jobs on an HP StorageWorks subsystem. EVM incorporates a multi-threaded job engine that allows multiple jobs to be created and/or executed simultaneously. To start EVM, browse to the management appliance and log on to Open SAN Manager. Select Enterprise Volume Manager under Volume Management. Most jobs can be created with the Create Job wizard. Many jobs can be completed in just four simple steps, while more complex jobs may require six steps.

EVM incorporates and automates all the underlying HP StorageWorks command line interface (CLI) commands to create clones. Figures 5a and 5b illustrate a typical command set.

The CLONE command creates a mirrored disk from which one member can be split off. If the disk already has three members, the operation is almost instantaneous; if it has two, then EVM adds and synchronizes a third member. The NORMALIZE command waits for the mirrorset to synchronize. The SUSPEND command starts a command script to run SQL and puts the database in backup mode. The SPLIT command splits off a member from the mirrorsets and creates a new unit. The RESUME command runs a script to take the database out of backup mode. The MOUNT command mounts the new unit. Finally, the LAUNCH command runs the actual backup script.

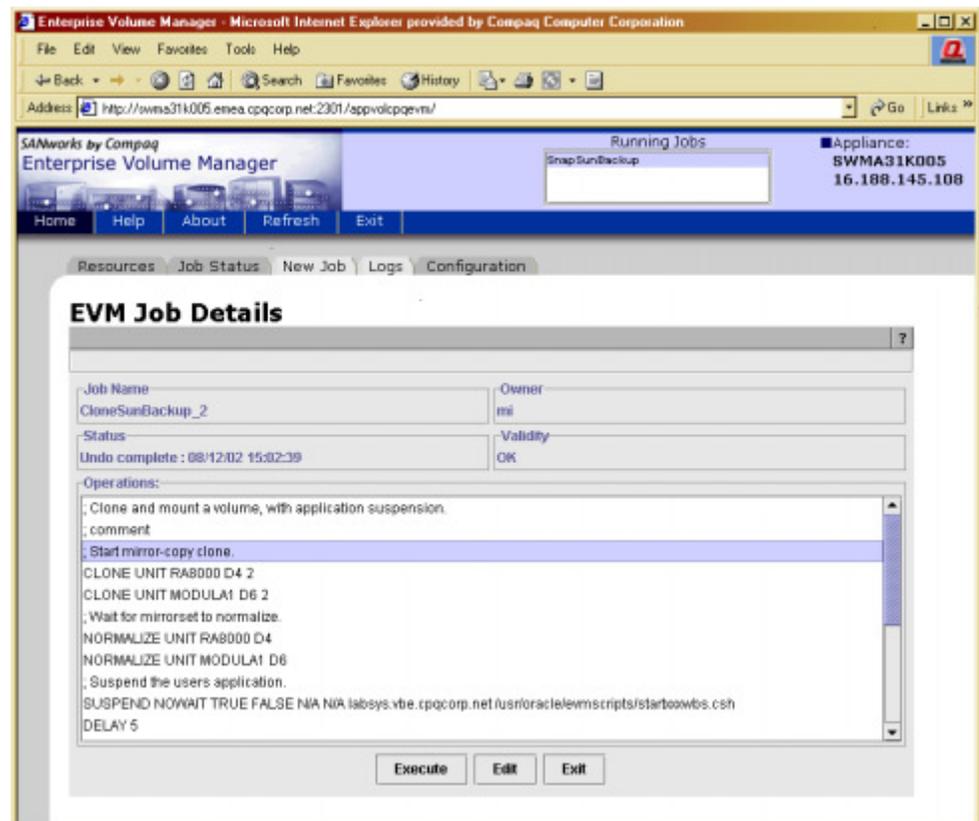


figure 5a

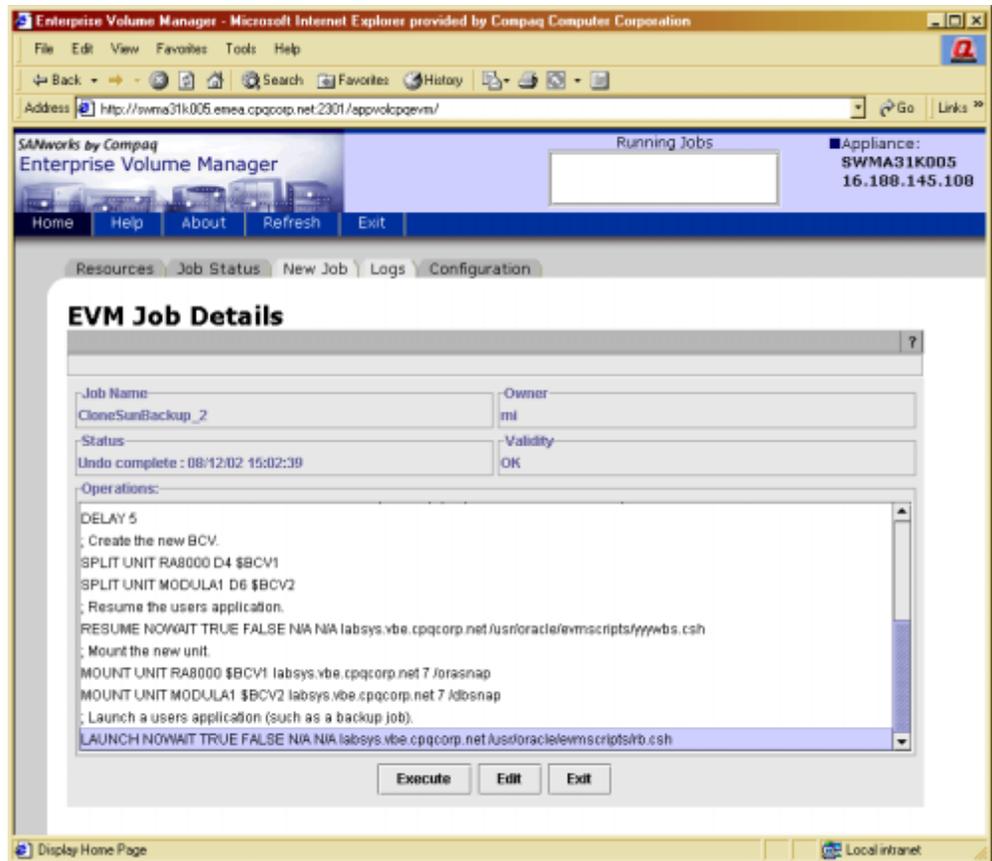


figure 5b

In the example shown in Figures 5a and 5b, the last argument of the SUSPEND and RESUME directives is a standard path\filename specification, which could be a batch file, a script, or an executable. In this example, the directives contain SQL scripts that are used to execute Oracle commands, but other ways exist to achieve the desired results. The SUSPEND script will checkpoint the database and place table spaces in online backup mode, and the RESUME script will reverse these actions. Using this method ensures that the backup contains a valid backup of the database.

EVM also makes it easy to check the status of a job, as shown in Figure 6. The job can be executed, inspected, edited, validated, or deleted.

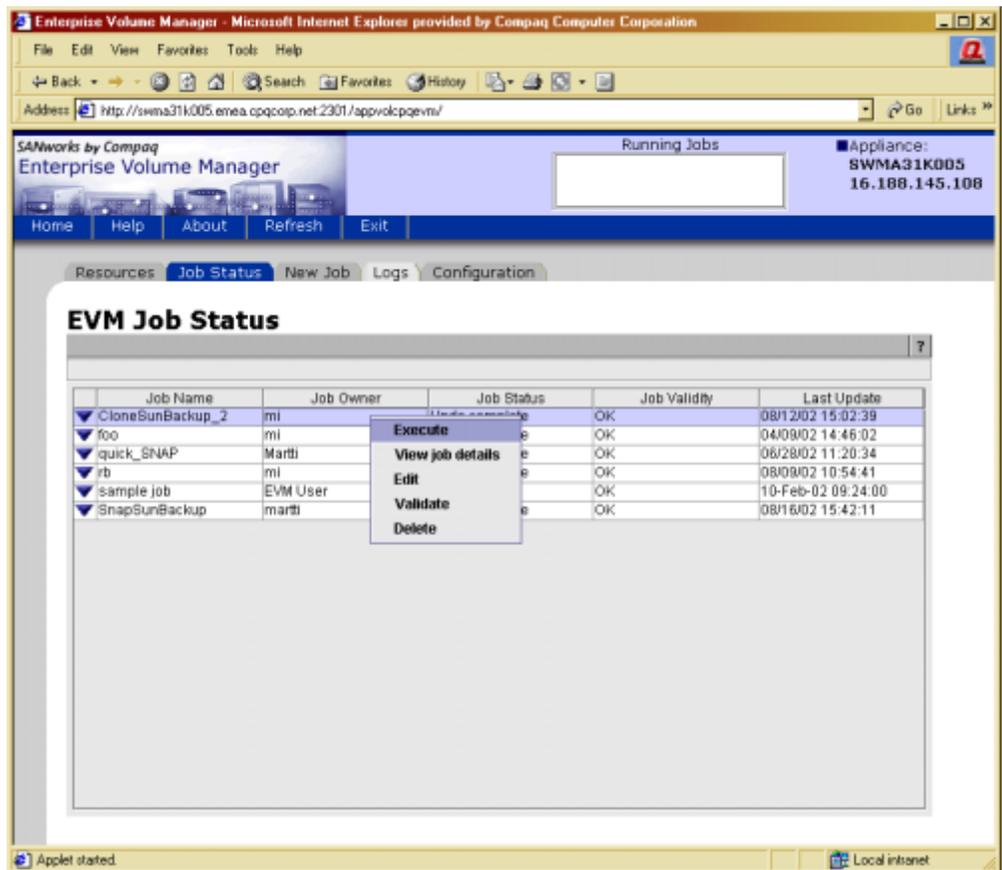


figure 6

sample EVM scripts used in testing

This section is a listing of all the script files that were used during the tests. They are provided as examples only; the same results can be obtained in many different ways. The EVM job command SUSPEND calls the startxxxwbs.csh script, which starts another script, xxxwbs.csh, to run in the background:

startxxxwbs.csh

```
#!/usr/bin/csh
/usr/oracle/evmscripts/xxxwbs.csh &
```

xxxwbs.csh

```
#!/usr/bin/csh
su - oracle -c "./evmscripts/suspendresume.csh"
```

xxxwbs.csh runs another script, called suspendresume.csh. xxxwbs also sets the user ID to "oracle" to have the correct user environment.

suspendresume.csh

```
#!/bin/csh
setenv ORACLE_SID TCR
sqlplus system/manager >/usr/oracle/evmscripts/archbefore.log << EOF
connect / as sysoper;
archive log list;
exit;
EOF
sqlplus system/manager >/usr/oracle/evmscripts/suspend.log << EOF
alter tablespace system begin backup;
alter tablespace undotbs begin backup;
alter tablespace users begin backup;
alter tablespace tools begin backup;
alter tablespace indx begin backup;
host /usr/oracle/evmscripts/waitforsplit.csh
alter tablespace system end backup;
alter tablespace undotbs end backup;
alter tablespace users end backup;
alter tablespace tools end backup;
alter tablespace indx end backup;
alter system archive log current;
exit;
EOF
sqlplus system/manager >/usr/oracle/evmscripts/archafter.log << EOF
connect / as sysoper;
archive log list;
exit;
EOF
```

Suspendresume Csh starts by making a log of the existing archived logs. Then it puts all active tablespaces in online backup mode and executes the waitforsplit.csh script. The suspendresume.csh script then waits until waitforsplit.csh returns, indicating that the clones have been split off.

waitforsplit.csh

```
#!/usr/bin/csh
touch /usr/oracle/evmscripts/dummy_file
while -f /usr/oracle/evmscripts/dummy_file
  sleep 5
end
```

After the SPLIT command occurs, the RESUME command calls the yyywbs.csh script, which signals a successful split (Figure 5b). (SPLIT and RESUME are EVM commands.)

yyywbs.csh and rmfile.csh

```
#!/usr/bin/csh
su - oracle -c "./evmscripts/rmfile.csh"

# !/usr/bin/csh
rm /usr/oracle/evmscripts/dummy_file
```

The yyywbs.csh script causes waitforsplit.csh to continue. Then the suspendresume.csh script takes all tablespaces out of backup mode and logs the existing archive logs. A system manager then can determine what logs were archived during the time the database was in backup mode.

Finally the EVM job mounts the cloned disks and starts the backup by calling the rb.csh script:

rb.csh

```
#!/usr/bin/csh
/usr/opensv/netbackup/bin/bpbackup -c Oracle -w -f
/usr/oracle/evmscripts/TCR_files_list.txt
```

TCR_files_list.txt

```
/orasnap/admin/TCR/pfile
/orasnap/oradata
/dbsnap/TCR
/data2
/usr/oracle/evmscripts/arch*.log
```

The cloned disk */orasnap* contains all Oracle-system-related data; the cloned disk */dbsnap* contains all database files; the */data2* disk contains archived redo-logs; and the *arch*.log* files show what was logged during the backup.

TCR_files_list.txt is a list of the directories and files to backup.

guidelines for cloning an Oracle 9i storage unit with EVM

For each Oracle 9i storage unit to be cloned, EVM automatically checks the RAID level, determines the correct procedures, and searches the subsystems to select free disks for the clone, which acts as an additional mirror. EVM then directs the controllers to copy the Oracle data and normalize the disks in the mirrorset, including the clone. When normalization is complete, all members of the mirrorset have identical data, and EVM splits the clone from the mirrorset.

Depending on the type of source unit, EVM applies slightly different replication procedures. It is important to consider the differences in the following procedures when planning EVM jobs and storage configurations.

- If source unit is a three-member mirrorset (RAID 1, RAID 1+0):
To achieve near-instantaneous cloning with EVM, a source unit should be configured as a three-member mirrorset and be fully normalized. In this configuration, EVM instantly splits off the third member to create the clone. This approach to instant cloning also applies to mirrorsets with four or more members. Because the mirrorset has more than two members, it also enables maximum hardware reliability and best performance.
- If source unit is a two-member mirrorset (RAID 1, RAID 1+0):
EVM automatically adds a third member to the mirrorset and normalizes all the members. When the members are normalized, EVM splits off the third member to create the new clone.
- If source unit is not mirrored (RAID 0 and JBOD):
EVM automatically creates a single-member mirrorset and then adds a second member. When the members are normalized, EVM splits the second member to create the new clone.
- If source unit is RAID 3/5:
Cloning RAID 3/5 units is not supported.
- If source unit is a clone:
The source unit or the volume in the clone and split operation can be a clone.

Caution: When planning and coordinating EVM jobs, remember that the intended source clone in one job could be deleted by an UNDO of another job.

EVM job interactions with Oracle 9i applications

EVM provides a simple yet powerful set of features that allows storage unit replication jobs to be seamlessly combined with Oracle 9i applications. For example, it is easy to create the following jobs:

- EVM-Oracle jobs that generate near-line backup units of Oracle 9i databases on the subsystem
- EVM-Oracle backup jobs that copy online Oracle 9i databases and mount them on backup servers for automated tape backup
- EVM-Oracle jobs that copy Oracle 9i databases and mount them on another server for data mining or testing

By using batch files and scheduling applications, it is possible to create job sequences that run 24 hours a day, seven days a week. EVM also lets administrators perform backup operations on a clone that has been mounted on a host other than the production server, provided that host resides on the same SAN. Through this non-disruptive method of backup, EVM eliminates device contention between the production workload and the backup operation, and offloads significant CPU overhead associated with the backup process from the production server to a dedicated backup host.

supported configurations

EVM can be configured in a network of EVM-enabled computers, which can participate in automated storage unit replication jobs across as many as 25 subsystems equipped with HSG80 controllers, including EMA or ESA 12000 systems, as well as Modular Array or RAID Array 8000 Fibre Channel systems. EVM also can be used in the same way with the HP StorageWorks Enterprise Virtual Array and HSV110 controllers. Each EVM network and its associated subsystems can be remotely managed using a Web browser included with EVM 2.0.

The following list contains the minimum configuration requirements. Each EVM 2.0D network requires the following:

- Management Appliance with Open SAN Manager V1.0C and the HSG80 Element Manager V1.0D
- EVM V2.0D server software installed on the management appliance
- The management appliance and at least one EVM Agent node must be connected to at least one EVM-compatible HP StorageWorks subsystem (all containing HSG80 array controllers and using HP StorageWorks Solution Software Kit 8.6B or later)
- Supported ACS versions for the HSG80 controller (8.6-4 and 8.7)
- Access to a qualified Web browser (see ["supported Web browsers"](#))

supported Web browsers

Table 3 identifies the supported browsers by the minimum browser and Java runtime environment (JRE) plug-in version numbers for each host operating system platform. Supported browser versions match the management appliance browser requirements.

NOTE: Refer to the EVM V2.0D Release Notes for the latest information on browsing to an EVM network by other supported operating system computers. Each Web browser requires the JRE plug-in to properly display EVM data. See Enterprise Volume Manager Version 2.0D (Multi-OS) Network Administration Guide (Part Number AA-RR0VB-TE).

table 3. supported web browsers

host operating system	minimum browser version	required JRE version
Windows 2000 Windows NT	Microsoft Internet Explorer 5.01 Netscape Communicator 4.77	1.3.1_01
Sun Solaris	Microsoft Internet Explorer 5.0 for UNIX Netscape Communicator 4.77	1.3.0_04

EVM agent requirements

Agent requirements include required and optional attributes. Table 4 describes the requirements for each host operating system platform.

table 4. EVM agent requirements

host operating system	agent attributes
Sun Solaris 2.6, 7(32-bit or 64-bit), 8(32-bit or 64-bit)	<p><u>Required:</u></p> <ul style="list-style-type: none"> • Connection to a LAN (common with EVM Server and other EVM Agents) • Connection to the HP StorageWorks subsystem with the Jaycor (JNI) Fibre Channel adapter FC-1063 or FC64-1063 with a driver version provided by the HP StorageWorks Solution Software Kit 8.6B • HP StorageWorks Secure Path Version 3.0A for Sun Solaris or higher

for more information

hp StorageWorks

Overview

www.hp.com/go/storage

Enterprise Volume Manager

<http://h18000.www1.hp.com/products/sanworks/evm/index.html>

Using links from the EVM overview page you can find all the following information and more:

- EVM product overview
- EVM Version 2.0D (Multi-OS) Installation Guide (Part Number AA-RTB5A-TE)
- Enterprise Volume Manager Version 2.0D (Multi-OS) Network Administration Guide (Part Number AA-RR0VB-TE)
- EVM V2.0D Release Notes (Part Number AA-RTB6A-TE)

Enterprise Backup Solution

<http://h18000.www1.hp.com/products/storageworks/ebs/>

Using links from the EBS overview page, you can find all the following information and more:

- EBS product overview
- EBS compatibility matrix
- EBS sizer tool
- EBS Installation and Startup Service

Backup software

VERITAS:

www.veritas.com

hp services

Installation and Startup Service for Enterprise Backup Solution (EBS):

http://h18005.www1.hp.com/services/storage/st_enterprise.html

Let us know what you think about the technical information in this document. Your feedback is valuable and will help us structure future communications. Send your comments to: Oracle.Storage.Solutions@hp.com

All brand names are trademarks of their respective owners.

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

© 2003 Hewlett-Packard Company

01/2003