



hp StorageWorks rapid restore for Microsoft SQL Server 2000 solution

Using hp StorageWorks enterprise volume manager V2.0D and hp StorageWorks enterprise virtual array V2

table of contents

executive summary	2
about Vsnap and snapclones	3
features and benefits	4
solution overview	4
solution components	
best practices	
test environment	6
test configuration	
hardware and software components	
performance results	8
solution setup and installation instructions	10
working with EVM	12
preparing to back up and restore	
creating snapclones with EVM	
cleaning up EVM jobs	
automating EVM and the backup process	
backing up to tape	
sample backup plan	
restoring SQL server using snapclones and Vsnap	
related solutions	22
for more information	22

Help us improve our technical communication. Let us know what you think about the information in this document. Your feedback is valuable and will help us structure future communications. Please send your comments to:

Microsoft_Storage_Solutions@hp.com

executive summary

This guide gives a quick and complete method for recovering clustered Microsoft SQL Server 2000 databases with minimal disruption. Using HP StorageWorks Enterprise Volume Manager (EVM) V2.0D and HP StorageWorks Enterprise Virtual Array (EVA) V2, customers can create Vsnap and snapclone of their SQL Server databases and resume full operation of their clustered SQL Server environments in minutes. A Vsnap is a virtually capacity-free snapshot that is a space-efficient point-in-time copy of the data. A snapclone is a complete physical point-in-time copy of a data volume. Both Vsnap and snapclones can be used to perform an **extremely rapid restoration** of application environments.

The intended audience is made up of experienced SQL database administrators who understand how to work with snapshots and clones using array controllers. Performance testing provided numerous insights into best ways to use Vsnap and snapclones with EVA for rapid recovery of SQL Server data. The information in this guide will enable SQL database administrators to do the following:

- Determine whether Vsnap or snapclones are more suitable for their environment
- Create Vsnap or snapclones of the SQL database
- Restore data from Vsnap or snapclones
- Create scripts to implement automatic backups from snapclones
- Back up from Vsnap or snapclones to tape

Administrators face lengthy delays when restoring a SQL environment from tape. This solution, which is based on creating Vsnap and snapclones with EVM in a SAN configuration, reduces the restore time from hours to minutes and causes minimal disruption to end users. You can achieve a highly available SQL environment with a complete set of hardware, software, and supporting utilities. In addition, the time savings can be dramatic when restoring your SQL environment, as illustrated in the following diagram:

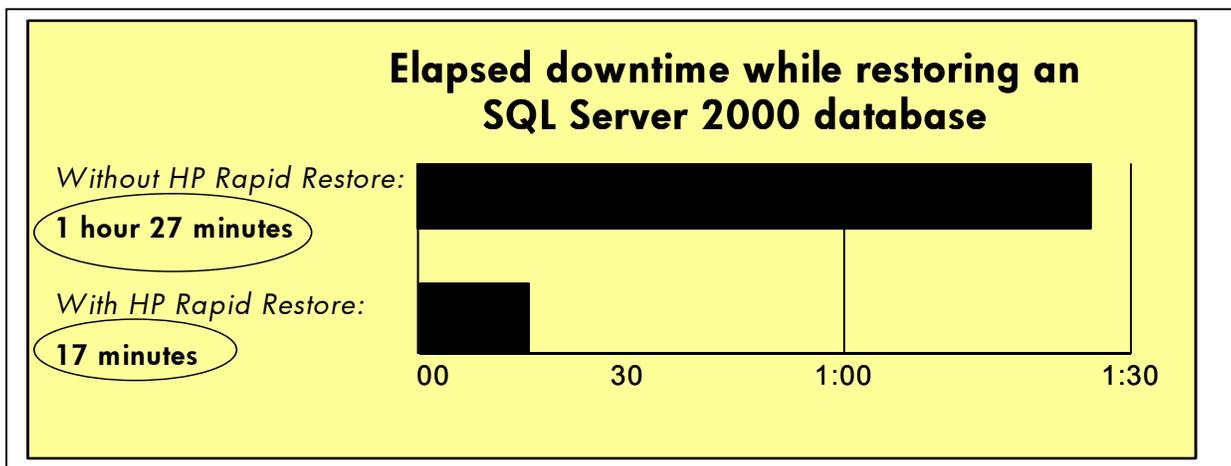


figure 1 – Comparison of disk-to-disk restoration times (snapclone) using SQL Server 2000, 400GB database, and 200 users

By using this solution to augment your tape-based recovery plan, you can recover your entire application environment rapidly, decreasing the amount of application downtime. When you create a Vsnap or snapclone of your database, you have an exact point-in-time copy of your production data. You can offload operations from the production server. These operations include tape backup, which is discussed in this guide, as well as testing, data mining, and system maintenance. The solution also utilizes a Microsoft-supported interface and is engineered within Microsoft support guidelines.

“Microsoft SQL Server 2000 customers put their trust in the performance and reliability of the database. HP StorageWorks solutions for SQL Server 2000 help provide even more robust storage solutions for our customers. Rapid Restore for SQL Server and Virtualized Storage Management for SQL Server 2000 are pretested and automated solutions that take advantage of our highly scalable and versatile storage management software. Together, HP StorageWorks and SQL Server 2000 offer enhanced application availability and organizational agility in fast-paced environments.”

Stan Sorensen, Microsoft Corporation
Director, SQL Server Product Marketing

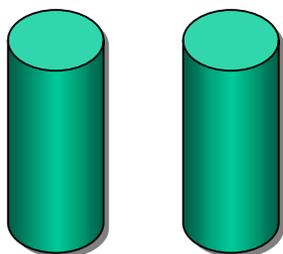
about Vsnaps and snapclones

With EVA, you can create traditional snapshots, demand-allocated snapshots (**Vsnaps**), and snapshots that normalize into **snapclones**. The traditional snapshot requires that you reserve and set aside space equal to the size of the original active virtual disk (Vdisk). Data is not written into this reserved space until necessary. As data changes in the original active Vdisk, the original data is written to the snapshot. Virtually capacity-free Vsnaps and virtually instantaneous snapclones are the powerful recovery technologies used in the Rapid Restore for SQL Server solution.

Vsnaps allow on-the-fly space allocation. You can replicate data instantly by taking a “picture” of the data within seconds, without reserving an equal amount of capacity. Similar to traditional snapshots, Vsnaps are point-in-time copies of data volumes although space is only used as the original virtual disk data changes. Once the snapshot is taken, the Vsnap takes up very little capacity. As users write to the database, the Vsnap grows as the Vdisk changes. With Vsnaps, the amount of disk capacity used by the copy only grows as data in the production volume changes over time, resulting in the most cost-efficient use of storage space.

Snapclones are complete physical point-in-time copies of data volumes and are ideal for reducing I/O loads on the production volumes. These copies can be used for backups to tape, disk-based backup, or application testing and data mining that will not impact production data. Snapclones are snapshots that normalize into clones. A snapclone has the same properties as a snapshot, so it is instantly available. The normalization process of copying the data from the original LUN happens in the background. Similar to a snapshot, a snapclone can be instantly mounted or backed up to tape. However, with snapclones, once the normalization process is complete, you have an exact copy of the database or LUN you wanted to clone. This is especially useful for fast restores if a catastrophic failure happens. Snapclones can be taken in **any redundancy level (Vraid 0,1,or 5)** and they do not require **extensive advanced preparation**. Long “re-sync” times are also eliminated with snapclones since a current clone copy can be available for use in moments rather than hours.

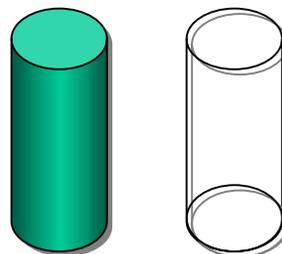
Traditional Array Snapshot



Production
Volume
50 GB

Copy
Volume
(50 GB
Reserved)

Enterprise Vsnap Snapshot



Production
Volume
50 GB

Copy
Volume
(~0 GB
Reserved)

figure 2 – Comparison of traditional snapshot to Vsnap snapshot

features and benefits

This solution provides the following benefits to administrators:

- Dramatically improved restoration times for SQL Server 2000 databases
- Best practices for maximizing SQL Server 2000 availability during database recovery scenarios
- Simplified implementation and management, including script examples
- Data integrity and engineering within Microsoft guidelines
- Investment protection by leveraging existing HP hardware and software, supporting multiple configurations, and providing interoperability with future products
- Integration with third-party tape backup applications—specifically tested with VERITAS NetBackup and HP OpenView Storage Data Protector

solution overview

The Rapid Restore Solution for SQL Server 2000 is a validated and fully integrated configuration that provides a SAN-based backup-and-restore infrastructure for end-to-end data protection. This solution leverages the capabilities of EVA V2, EVM V2.0D, and the HP StorageWorks Enterprise Backup Solution (EBS), as illustrated in Figure 3 below.

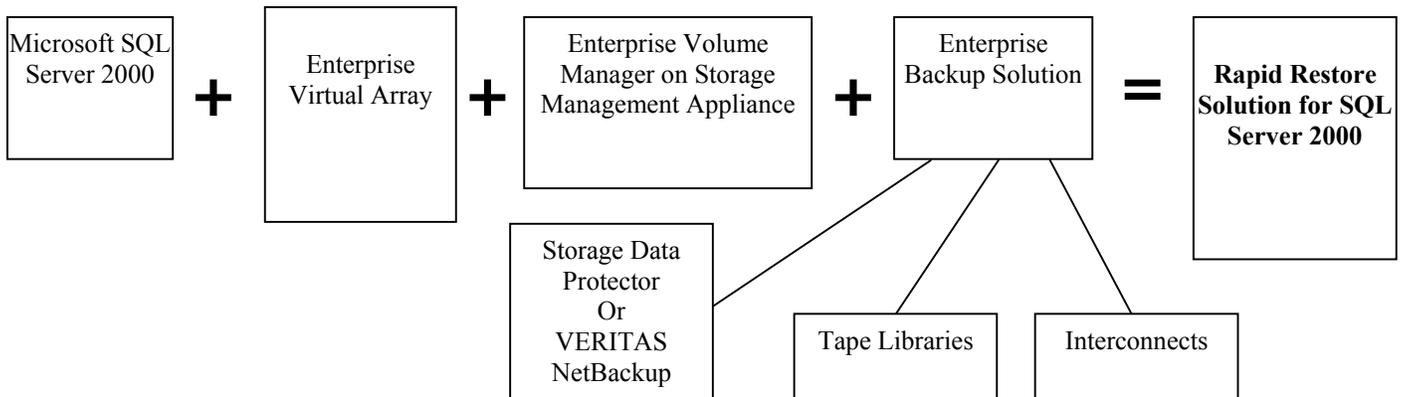


figure 3 – Solution building blocks

solution components

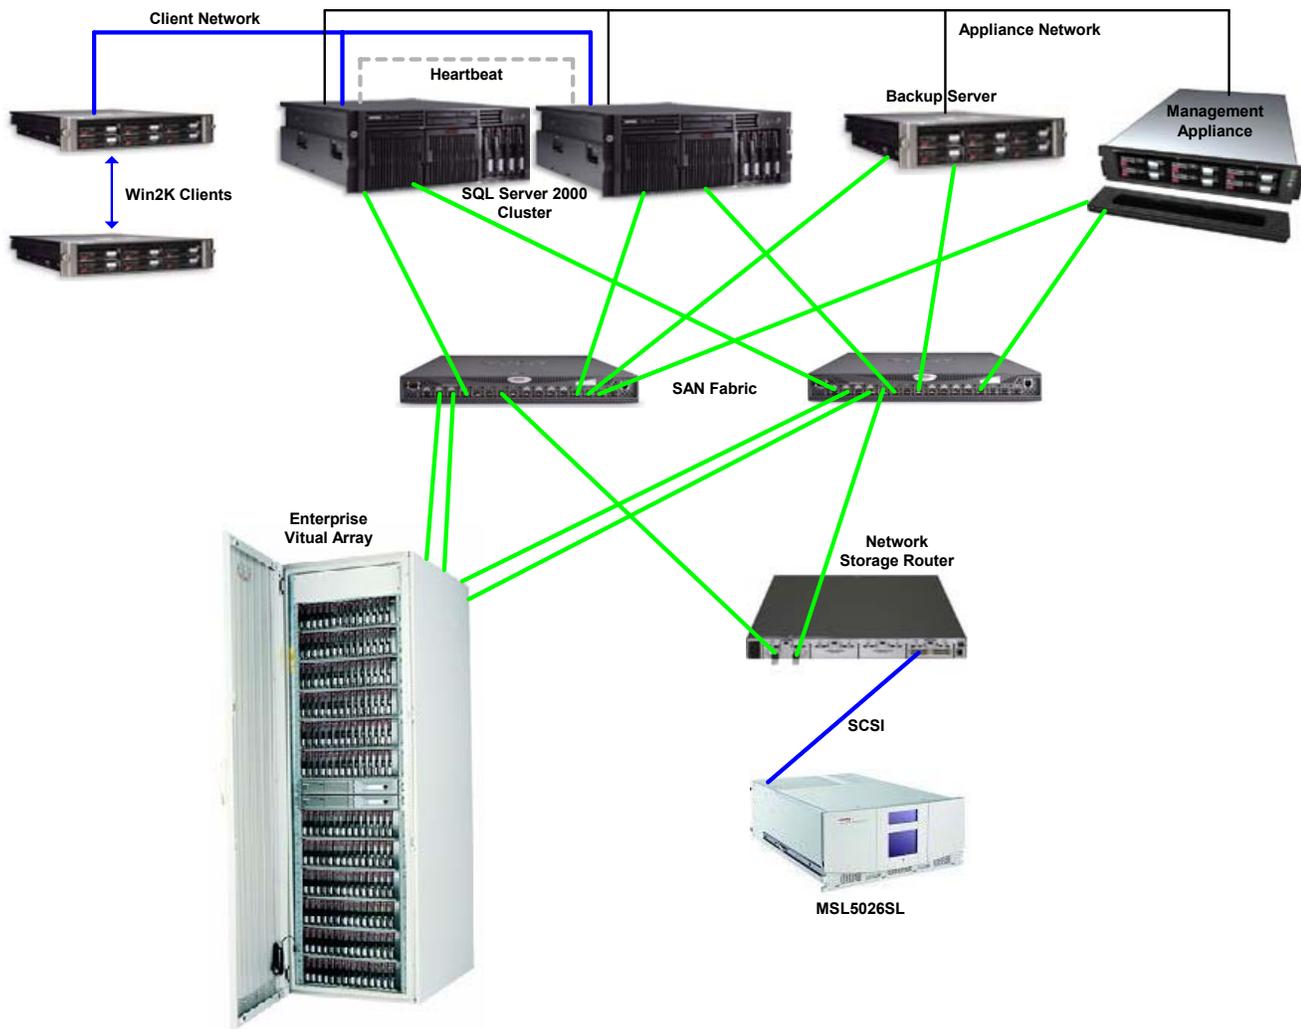


figure 4 – Diagram of components

The Rapid Restore for SQL Server 2000 Solution consists of these basic components:

- Microsoft SQL Server 2000
You can configure the SQL Server 2000 application and database on any Windows 2000 server and use an MSCS configuration for improved application availability.
- Enterprise Virtual Arrays
The SQL Server database volumes must be located on the Enterprise Virtual Arrays. Vsnapshots and snapclones are created on the virtual array using virtual controller software (VCS), which operates the storage subsystem. EVA is the newest generation of disk arrays that is enabled by virtualization technology. Based on the HSV controller, EVA is a high-performance, high-capacity and high-availability “virtual” RAID storage solution that eliminates the time, space, and cost boundaries of traditional storage. EVA can save significant disk space and improve disk utilization efficiency because it does not reserve the same amount of disk capacity as the production volume being copied.
- Enterprise Volume Manager V2.0D
Enterprise Volume Manager is browser-based storage management software that facilitates controller-based clone operations to make a block-to-block copy of a storage volume. With EVM, you can create, run, and manage automated storage replication jobs, as well as link them with external jobs.

You must have an EVM host agent on each database server to utilize controller-based database cloning. The HP

OpenView Storage Management Appliance is the operating platform for EVM, which offloads the processing from the user host systems.

- Enterprise Backup Solution (EBS; optional)
EBS provides a consolidated backup infrastructure comprising multiple servers connected over a high-speed Fibre Channel SAN to centrally attached SCSI tape libraries. EBS significantly shrinks backup windows with a flexible approach to data protection.

Enterprise Backup Solution with HP OpenView Storage Data Protector V5.0 or VERITAS NetBackup V3.4 provides SAN-based backup-and-restore operations for snapclone based tape backups. You must use a SAN attached backup configuration when using this solution for clone-based offline tape backups. An approved third party backup application, with associated SQL Server agents and SAN agents, provides backup management. A separate backup host with the backup application loaded and viewable to the storage subsystem is required to act as the dedicated backup server. For granular backups and restores, a copy of the SQL Server 2000 application must be loaded on the backup server.

best practices

- Traditional configuration rules apply when you start planning your storage for your SQL Server database. Use RAID 0 or 1+0 for databases with high read and write activity and RAID 5 for databases with low write activity.
- When configuring data and log files, it is an HP best practice and Microsoft recommendation to use different disk groups for log files and data files. Using this strategy, the log files will always be located on different volumes than the data files.
- Keep log files on LUNs that have been configured using RAID 1+0 because of heavy write activity.
- HP recommends that you set the protection level to a minimum of single redundancy in all disk groups. Plan on using as much space as the parent Vdisk can hold, and if possible add two extra spindles for single redundancy.
- If you are expecting very heavy user activity, use preferred paths whenever possible to direct I/O flow across both controllers.
- EVM 2.0D Host Agent Release Notes recommend logging out of the backup host when running EVM scripts that mount volumes.

test environment

This section describes the baseline configuration and provides listings of the hardware and software components used in the solution.

test configuration

Two ProLiant DL580 servers, each having two FCA2101 (2Gb HBA) adapters, powered the SQL Server (active/passive) cluster. A ProLiant DL380 was used as a backup server to host the backup applications and manage data protection for the tape library. The SAN was configured with two redundant 2Gb 16-port SAN switches and managed by the HP OpenView Storage Management Appliance. The storage consisted of an Enterprise Virtual Array with 3.1TB of disk storage and was backed up with an MSL5026 SDLT tape library.

The SQL database of 370GB was configured with two data volumes of 100GB and 200GB, and one log volume of 100GB. One disk group for the data files and one disk group for the log files were used to distribute the I/O across the EVA Fibre Channel disks. The I/O load was generated with TPC-C transaction processing.

The Ethernet was set up and configured for cluster operations using Cluster Administrator. Three different submasks were used for the Ethernet configuration. A submask was set up for the heartbeat between the two cluster nodes. The heartbeat was configured as a private network for cluster communication only. A second submask was set up for client access to the application servers. It was configured as a public network. The third submask was set up for the Storage Management

Appliance and its applications. This network was set up as a mixed network so the applications running on the Storage Management Appliance could operate corresponding agents on the hosts connected to the SAN.

hardware and software components

Here is a list of the hardware and software components used in the test configuration. Keep in mind that new versions of these components frequently become available and the listed parts may not be orderable.

Table 1 – Hardware Components

ProLiant Servers Hardware	Part Number
ProLiant DL580 servers –	(2) 155618-003
Quad Pentium III Xeon, 900MHz Processors, 512K level 2 cache (256K per processor)	(4) 222627-B21
4GB RAM Memory	(2) 189083-B21
18GB Ultra 3, 15K rpm drives in each server	(2) 188122-B22
FCA 2101 (LP952) HBA	(4) 245299-B21
Backup Server	
ProLiant DL380 server –	193706-001
Dual Pentium III 1GHz Processors, 512K level 2 cache (256K per processor)	187602-B21
2GB RAM Memory	128280-B21
Thermal Upgrade Kit	210818-B21
FCA 2101 (LP952) HBA	(2) 245299-B21
Storage Management Appliance	
HP OpenView Storage Management Appliance	189715-002
Storage Hardware	
HP StorageWorks Enterprise Virtual Array – 2C6D-A	283194-B21
42U M3220 Controller assembly with dual HSV110 controllers and 6 M5214 dual fibre loop 14-bay drive enclosures.	
36GB 10K rpm dual-port 2Gb/sec FC-AL 1-inch drive	(84) 238590-B21
Enterprise v1.0 to v2.0 Upgrade Kit (Vixel Switches)	283266-B21
2Gb 16-Port SAN Switch	(2) 240602-B21
2Gb SFF-SW Transceiver kit	(32) 221470-B21
Modular Data Router	163083-B21
MSL5026SL RM SDLT Minilibrary (2 drives)	231892-B22
30m LC-SC Multi-Mode Fibre Channel Cable	221691-B26
15m LC-SC Multi-Mode Fibre Channel Cable	221691-B23
5m LC-LC Multi-Mode Fibre Channel Cable	221692-B22
15m LC-LC Multi-Mode Fibre Channel Cable	221692-B23

Table 2 – Software Components

Software	Part Number
Windows 2000 Advanced Server with SP2	Third Party
Microsoft Cluster Server – MSCS	Included
Microsoft SQL Server 2000 Enterprise Edition with SP2	3 rd Party
HP OpenView Storage Management Appliance software V2.0	QB-6M4AA-SH
HP StorageWorks Secure Path 4.0 (5 Licenses)	231292-B22
HP StorageWorks EVM 2.0d (Starter kit, 5 host licenses)	263670-B22
WNT/W2K KIT V2.0 ENT VIR ARY (for Windows NT 4.0/2000)	250195-B22
Snapshot License for VCS V2.0 Dual HSV Controllers up to 3.1 TB	253256-B22
VCS PKGV2.0 Dual HSV Controller (base controller software)	250203-B23
HP OpenView Storage Data Protector 5.0 with SQL Server 2000 Data Agents	
Starter pack	B6961AA
SAN drives	B6953AA (one per drive)
Online extension for SQL agents	B6965BA
TPC-C Benchmark Kit for SQL (load generator)	Third Party
VERITAS NetBackup V 3.4 Server Software and SQL Data Agents	Third Party

performance results

A baseline configuration was developed to determine the impact of snapclones and Vsnap on the SQL Server 2000 database. All performance tests were compared against this baseline configuration, which included a clustered SQL Server 2000 without EVM installed, 200 users, and a 50% CPU load against the database servers. TPC-C benchmarking was used to create the OLTP load on the database and the performance numbers were collected using Windows Performance Monitor. The numerical results are displayed in Table 5, "System and Application Performance Results." The following configurations were tested:

- **Baseline with EVM – test 1**

EVM V2.0D was installed into the SAN environment and a baseline test was performed to determine the impact of EVM on the SQL Server and SAN environment.

Results: Introducing EVM to the SQL Server SAN environment did not have an impact on the performance of the environment.

- **Creating snapclones – test 2**

Creating snapclones with Enterprise Virtual Array involves utilizing available disk space in the disk group of the parent virtual disk to create the virtual disk copy. Using snapclones in this solution will consume 1.2TB of data, including the mirrored database of 400GB.

Results: The performance counters showed a slight degradation in the performance of the database by about 14% while the snapclones were being built. The snapclone creation took about one minute and 50 seconds. During this time, the SQL Server database remained online, database reads were unaffected, and database writes were temporarily queued until the creation was complete. During the snapclone building process, all transactions occurred normally. Creating and building snapclones has a minimal effect on the SQL Server and EVA performance.

Table 3 shows the snapclone building times for different size database volumes.

Table 3 – Snapclone Building Times

LUN Size/Type	Normalization Time wo/Load	Normalization Time w/Load
100GB Log	74 min	132 min
100GB Data	80 min	97 min
200GB Data	137 min	135 min
Total 400GB	281 min	339 min

- **Creating Vsnap – test 3**

Four Vsnap were created for each virtual disk. The performance impact was measured while changes were being made to the SQL Server database. Similar to the previous test, this test ran until the data consumed about 1.2TB, including the database.

Results: The performance counters showed a slight degradation in the performance of the database by about 12% while the Vsnap were growing. The Vsnap creation took about two minutes and ten seconds. During this time, the SQL Server database remained online, database reads were unaffected, and database writes were temporarily queued until the creation was complete. During the Vsnap growing process, all transactions occurred normally. The growth of the Vsnap is unpredictable and determined solely by the number of database changes for a given load. Administrators in the SQL Server environment should characterize this number and plan the Vsnap cycle accordingly. Creating and building Vsnap has minimal effect on the SQL Server and EVA performance.

- **Backing up to tape – test 4**

For comparison purposes, a tape backup was made of the SQL Server database while it was online. The time it took to create the tape backup was compared to the time it took to back up the snapclones to tape. Both tests used the tape backup applications (Storage Data Protector and Veritas NetBackup) and SQL Server data agents when necessary.

Results: The time to perform a full online tape backup of the database was six hours and ten minutes, or 278 times longer than the time needed to create the snapclone.

Although online tape backups have a minimal effect on SQL Server performance, they are an inefficient method for frequent backups because of the time required for completion and the potential contention with user activity. Snapclone-based backups also have minimal impact on the application performance but require much less time for completion.

- **Restore Operations – test 5**

Using snapclones or Vsnap to restore a SQL Server 2000 database can be done with the Windows drag and copy functions. HP recommends that you take the time to copy the Vsnap files onto the original volumes.

Results: The use of snapclones with EVA and EVM proves to be the fastest and most reliable means of restoring SQL Server 2000 databases. See Table 4.

Table 4 – Time Needed for Different Restore Methods

SQL Server Restore Method	Time (hrs)
SQL Server Snapclone Restore	17min

SQL Server native Restore	1:27
Windows File Drag and Copy	17min – 2hrs (depending on DB file size)
SQL Server Tape Restore	22 hours

Table 5 - System and Application Performance Results

	Baseline without EVM	Baseline with EVM	Creating snapclones	Creating Vsnap
User Connections	200	200	200	200
Benchmark Profile	OLTP	OLTP	OLTP	OLTP
Length of Test	4 hours	4 hours	6 hours	6 hours
# of Log Disk	1	1	1	1
# of DB Disk	2	2	2	2
Database pause time	n/a	n/a	1min 50sec	2min 10sec
Processor Utilization				
Processor Time (%)	67.6	51.9	62	64.6
Process % CPU Time – sqlservr	na	246.7	241	248
Memory Utilization				
Available Bytes (MB)	1631	na	160	1585
Pages/Sec	0.076	na	0.074	1.007
Logical Drive Utilization				
Total Avg. Disk Bytes/Transfer	8412	8390	8415	8412
Total Disk Write Bytes/Sec	12413923	8291141	6534211	7292705
Total Disk Read Bytes/Sec	25293760	18117173	9821049	14454407
Total Current Disk Queue Length	179	179	188	159
Total % Disk Time	225	154	90	114
Total Disk Transfers/sec	4482	3143	2305	2585
SQL Server				
Transactions/Sec	177	117	101	103
Active Transactions	156	150	160	159
Log Flushes/Sec	36.6	23	18	22
Buffer Cache Hit Ratio	89	89	89	89
Page Reads/Sec	3028	2173	1577	1730
Page Writes/Sec	1466	980	735	860
Cache Hit Ratio	99	99	99	99

solution setup and installation instructions

Caution: Many of the commands and procedures described require database administrator privileges and could disrupt database operation if used incorrectly. If you are not familiar with these commands and functions, consult your application and documentation before proceeding. As with any changes to your database configuration or backup plan, be sure to test these changes thoroughly before implementing them in a production environment.

The following components must be installed and configured prior to implementing this solution (shown in proper order of installation):

- Be sure the Storage Area Network is set up with the following components:
 - HP OpenView Storage Management Appliance software V2.0
 - HSV Element Manager V2
 - Snapshot licenses for HSV V2 controllers
 - EVM V2.0D Server
 - HP StorageWorks Secure Path Manager V4.0
 - Enterprise Virtual Array V2

Configure disk groups and virtual disks. Use separate disk groups for log files and data files. In large database configurations, separating the log and data files helps to spread I/O out across spindles.

- Be sure the servers are set up appropriately:
 - HBAs and NICs for storage and network connectivity
 - The Ethernet was set up and configured for cluster operations using three different submasks. A submask was set up for the heartbeat between the two cluster nodes. The heartbeat was configured as a private network for cluster communication only. A second submask was set up for client access to the application servers. It was configured for client access only or public network. The third submask was set up for the Storage Management Appliance and its applications. This network was set up as a mixed network so the applications running on the Storage Management Appliance could operate corresponding agents on the hosts connected to the SAN.:
 - FC2101 HBAs
 - Win2k Kit for EVA V2
 - EVM 2.0D Host Agent for Windows 2000
 - Secure Path V4.0 server on hosts
- Connect SAN components:
 - Connect hosts to the SAN switch.
 - Add host connections to hosts folder on the HSV Element Manager.
 - Present virtual disks to hosts for cluster quorum disk and SQL server files and databases.
 - Be sure to perform full formats on the newly presented volumes. For best results, present volumes to only one host and use that host to format volumes. Once the volumes are formatted, they can be presented to the remaining host.
 - If LUNs were previously created and formatted, this process is not necessary. Just present them to the hosts. To add LUNs as cluster resources, each host must see them as the same drive letter. Be sure to reboot if changes were made to drive letters.
- Set up the cluster and SQL server.

You can follow the Microsoft Guidelines for installing cluster services:

Step-by-Step Guide to Installing Cluster Services

<http://www.microsoft.com/windows2000/techinfo/planning/server/clustersteps.asp>

- a. Install MSCS. Choose **Start > Control Panel > Add/Remove Programs > Add/Remove Windows Components** on the first node.

- b. Install SQL virtual server.

SQL Server 2000 Failover Clustering

<http://www.microsoft.com/sql/techinfo/administration/2000/failovercluster.asp>

- HP recommends installing SQL files on a separate partition from the production database.
- Once the virtual server is installed and running, add the physical disk resources for your production database. Remember to label them as they appear in Disk Manager.

- Configure Secure Path and EVM

*For naming purposes, add an entry in your DNS table for the Storage Management Appliance.

- a. Configure the Secure Path logon and client table on each host. Choose **Start > Programs > Secure Path > SecurePathCfg**.
 - Set a common password that Secure Path Manager will use to poll each client machine.
 - If the client is a cluster member, in the Clients Configuration Table (in no particular order), add the name of the Storage Management Appliance, and then enter the cluster node name for each Secure Path client followed by the cluster name as follows: node_name:cluster_name.
 - If the client is not a member of the cluster, then just add the host name and the Storage Management Appliance names in the Client Config Table.
 - From the Storage Management Appliance, start Secure Path from the Network Utilities folder and create a profile. Add the name, password, and each host as they were listed in the Client Config Table and save the profile.
- b. Configure EVM

Important: Read EVM 2.0D Server and Host Agent Release Notes before configuring EVM. There are configuration limitations that must be acknowledged in order for the EVM configuration to work.

- Install EVM host agents on all client machines and choose the cluster option for all cluster nodes.
- From the Storage Management Appliance, use Windows Explorer and choose **Start > Program Files > HP > SANworks > Enterprise Volume Manager > bin > hosts**.
- Make sure you see all your host and cluster entries in the hosts folder. There should be a .HST and .VOL file for each host in your EVM network.
- Initially after the EVM Switchboard service is restarted (on the client machines), the EVM Server will acquire the information from the host machines and there will be a .DATA file in the hosts folder. Once all the data from the hosts has been collected, the .DATA files will go away.
- Start EVM from the Volume Management folder.
- Under the Resources tab, your hosts and cluster should show up in the top view pane and the HSV subsystem should show up in the bottom view pane. If the Resources window hangs, EVM might not be able to log on to the HSV subsystem.
- To verify that EVM is using the correct logon, open up the evm.cfg file in the EVM bin folder and verify the correct StorageId and StoragePassword. If either is incorrect, make the appropriate changes and restart the Switchboard service on the Storage Management Appliance. The ID and password should match that of the Storage Management Appliance logon.

working with EVM

This section describes how to perform the following tasks:

- Preparing to back up and restore using DButil
- Creating snapclones with EVM
- Cleaning up EVM jobs
- Automating EVM and the backup process
- Backing up to tape
- Restoring SQL Server using snapclones and Vsnap

preparing to back up and restore

SQL Server 2000 uses DButil to integrate with EVM. DButil is an application that is used to pause the SQL Server database, to resume the database after the snapshot is taken, and as a restore method if needed. DButil runs from a command line and must be initiated before you start the backup process. DButil maintains consistency throughout the backup and restore process. It places database files in a consistent state before creating a clone. Using this utility, the SQL Server database is disk write-queued, suspending write activity to the disks at the time of clone creation.

During the brief period when the SQL Server database is suspended, SQL Server read access is completely unaffected, and read requests are processed as they would be normally. Writes may continue to be submitted, but they are queued. The user may see a brief delay on write requests while waiting for the database to resume. Suspending the database does not affect the normal transaction isolation level. Users are not logged out of the database, and active transactions can continue. All held transactions are automatically rolled forward (that is, executed in order) once the database is resumed. Clone creation is therefore transparent to the user.

Example scripts are provided as batch files that you can use to automate these processes. You can download the Script Kit from the following website, which also contains the DButil application:

<http://www.compaq.com/products/storageworks/solutions/request/rrsql2kit.html>

DButil uses `-pre`, `-post`, and `-restoresns` commands to pause, resume, and restore SQL databases. The syntax is as follows:

DButil `-pre <sqlserver/instance_name> <database_name> <username> <password>`

DButil `-post <sqlserver/instance_name> <database_name> <username> <password>`

DButil `-restoresns <sqlserver/instance_name> <database_name> <username> <password>`

When DButil runs, a metadata file is created that is required when using the `-restoresns` command, which puts the SQL Server database in a state where transaction logs can be applied for recovery. DButil places this metadata file in the location from where it is run. To maintain the file, run the DButil command against the disk that is going to be cloned. This will create the metadata file on the clone. The DButil commands are run from batch files and are called `pre.bat` and `post.bat`. The batch files are very simple to create. For example, if your database data files are on your F: drive and the logs are on your G: drive, place DButil and the pre and post batch files in a directory on your F: drive. The pre and post batch files will look similar to these below.

Pre.bat

```
@ echo off
cd \dbfreeze \
dbutil -pre SQLEVM MyDatabase sa password
exit
```

Post.bat

```
@ echo off
cd \dbfreeze\
dbutil -post SQLEVM MyDatabase sa password
exit
```

For these files, the default SQL Server instance name is SQLEVM without an instance name and the database name is MyDatabase. The sa login was used along with the password "password."

HP recommends that you run DButil on the original LUN that is being cloned. In this way, the metadata file is created on the original LUN. Then the metadata file exists on the clone when it is split off. When DButil is run, information is passed to the SQL Server from the metadata file. The SQL Server uses this information to create a Virtual Device (placing the database in a loading state). The clone database disks can then be mounted under the assigned drive letters. After the restore command has run and the clone has been mounted to the correct drive letters, the option to restore sequential transaction logs is available. For each transaction log restore, select the No Recovery option for all but the very last transaction log.

creating snapclones with EVM

Snapshots and snapclones can be created with EVM 2.0D and EVA to guarantee safe data storage. After opening EVM, create a new job name and owner in the New Job window. Be sure to include comments in your scripts to let others know the purpose of the commands. The EVM commands are listed in the left side of the New Job window. The most common commands are SUSPEND, SNAP, RESUME, MOUNT, and UNDO. In this solution, the LAUNCH and DELAY commands are also used.

Here's the EVM Command Sequence used in this solution:

- SUSPEND – The first command in the job should be the SUSPEND command. The SUSPEND command calls the pre.bat file that pauses the SQL database and places it in a consistent state for backup.
- SNAP – The SNAP command is the command available to replicate volumes on an EVA. Within the snap command are options for the type of snapshot for the job. For this solution, the SNAPCLONE and DEMAND_ALLOCATED snap options were used.
- RESUME – Once the SNAP process has completed, the RESUME command must be run to call the post.bat file to resume activity within the SQL database.
- MOUNT – Once the volumes have been replicated, the snapclone can be mounted on a backup host for tape backup or an offline server for test purposes.
- DELAY – The DELAY command provides time for the Windows host to recognize the new volumes.
- LAUNCH – The LAUNCH command is used to start a backup to tape using the backup application's command line utility.

Here is a sample EVM script.

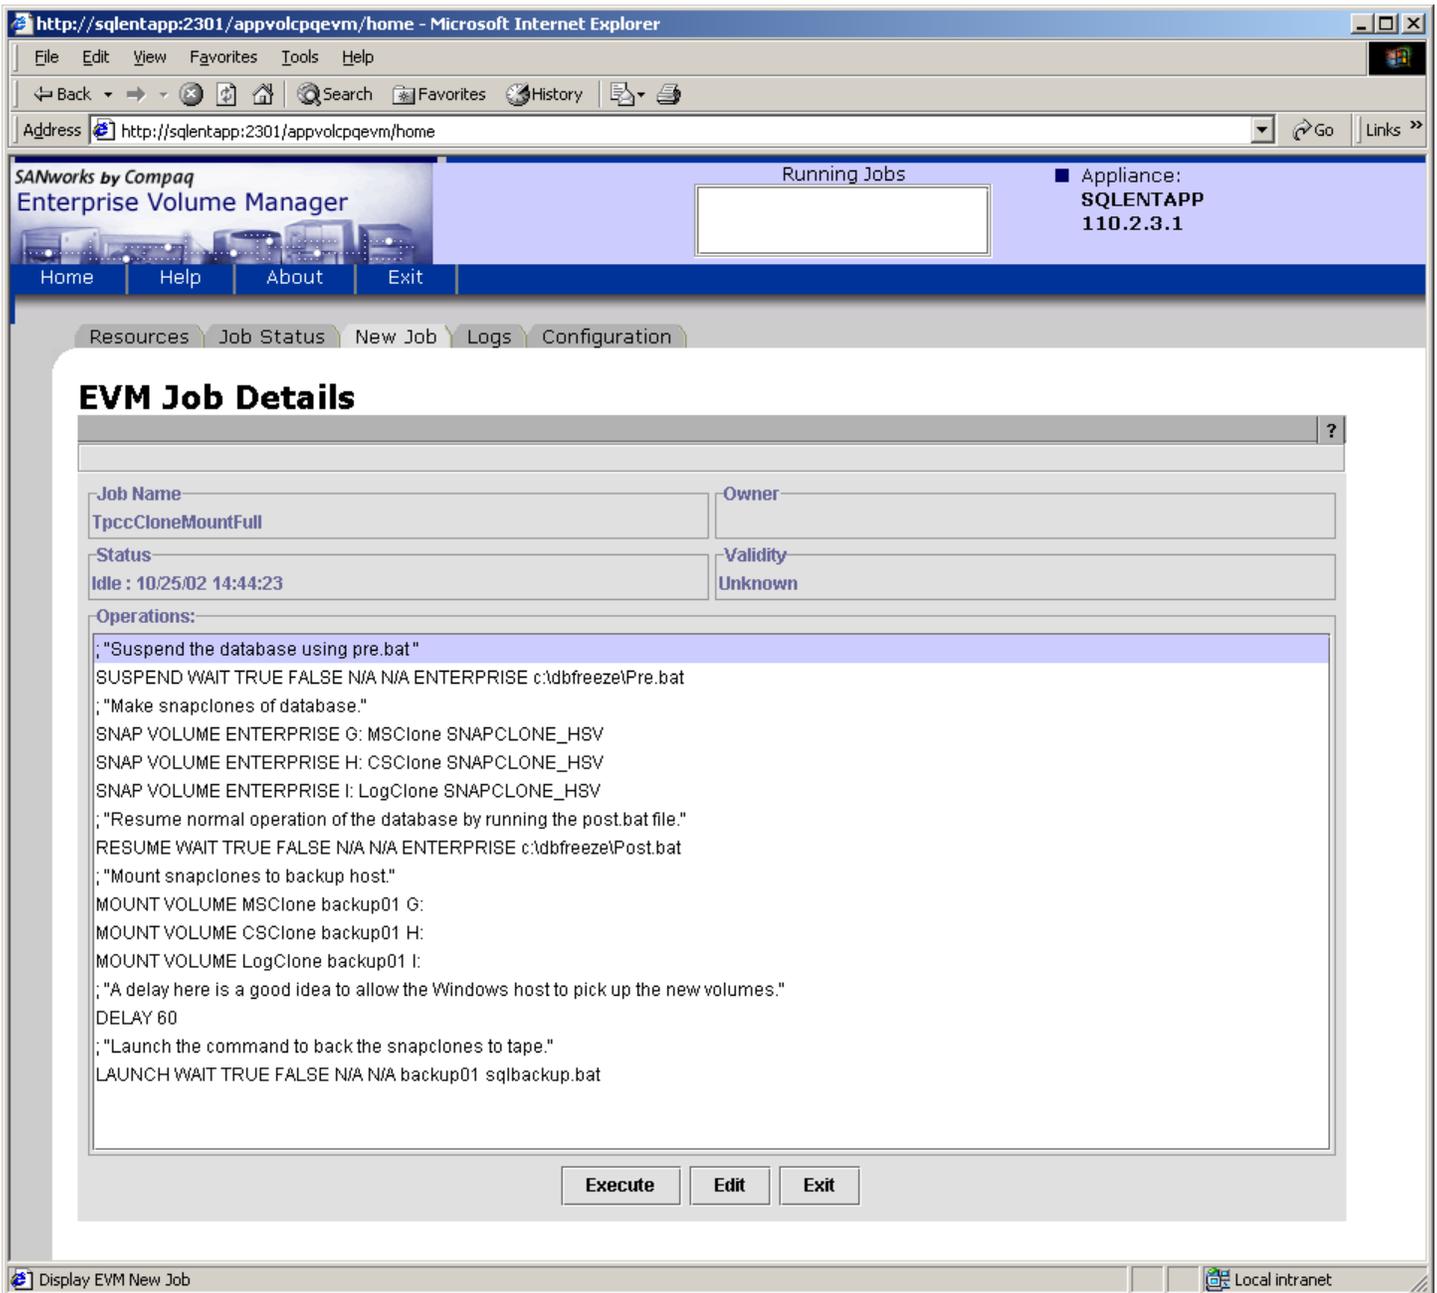


figure 5 – Sample EVM script

This EVM job suspends the SQL database, takes a snapclone of three database volumes, resumes the SQL database, mounts the snapclones to a backup host, and launches a backup job to tape. The same job could be used to create snapshots by changing the type in the SNAP commands to DEMAND_ALLOCATED_HSV.

Note: If there is an EVA hardware failure during the period when the snapclone is building or normalizing, the clone data can become corrupt or lost.

Best Practice: EVM 2.0D Host Agent Release Notes recommend logging out of the backup host when running EVM scripts that mount volumes.

cleaning up EVM jobs

It's important to recycle EVM jobs to keep the job list as clean as possible. In this way, it will be easier to manage job schedules and perform restore operations in the fastest amount of time. To minimize the number of EVM jobs in the joblist, it is a good idea to create, save, and schedule only the number of jobs needed to have two copies of the database available at all times. Once you have created the jobs, you can schedule them to run so that the first job is deleted before it is scheduled to run again. See the "sample backup plan" section later in this document.

To UNDO an EVM job, do the following:

1. Click the Job Status tab.
2. Right-click the job name and click UNDO.

Once the UNDO completes, you can run the job again to create another copy.

Caution: If you use the snapclones for a restore operation and the snapclones are running as the recovered database, do not click UNDO. If you run UNDO while the snapclones are running, you will lose the database. See the "restoring SQL server using snapclones and Vsnap" section for more information.

automating EVM and the backup process

The EVM scripts can be automated to a set schedule with little user intervention. HP OpenView Storage Management Appliance Software V2.0 runs on the Storage Management Appliance and can call an EVM script. Once set up, the same EVM scripts can be scheduled to run once, daily, weekly, biweekly and so on. (Storage Management Appliance Software V2.0 will be available late in 2002.)

Windows Task Scheduler can also be used to automate the backup process. Choose **Start > Accessories > System Tools** to locate the Task Scheduler. This solution used Task Scheduler to start the batch files used to automate EVM jobs. EVM includes a command-line utility called evmcl.exe that is located in the EVM/bin directory. Evmcl.exe can be run using batch files in the following manner:

```
Evmcl.exe <san_appliance_ip_address> execute or undo <job_name>
```

A sample batch file to start a previously created EVM job looks like this:

```
@ echo off

cd \Program Files\HP\SANworks\Enterprise Volume Manager\bin

evmcl.exe 110.2.3.1 execute EVMSnap1

exit
```

backing up to tape

Tape backup remains an integral part of the backup process. HP OpenView Storage Data Protector V5.0 and VERITAS NetBackup V3.4 were used as tape backup applications for this investigation. Each application was configured with SQL Server 2000 Data Agents running on the SQL Server host servers and the backup host running on the application manager. Be sure to check for the latest Storage Data Protector patches. You can find software patches at the following website:

HP OpenView software patches

http://support.openview.hp.com/patches/patch_index.jsp

- Backing up with Storage Data Protector

When configuring devices, it is important to disable any HP Tape drivers that were installed previously. Select **Start > Administrative Tools > Computer Management > Devices** and disable the tape and medium devices.

Once you disable the devices, you can configure them with HP OpenView Storage Data Protector Device Manager.

Storage Data Protector was configured only as a tape backup application. The HP OpenView Storage Data Protector Cell Manager Starter Pack was installed on the backup host, and the SQL Server 2000 Data Agents were installed on the SQL Server cluster nodes.

EVM uses its Launch command to integrate with Storage Data Protector. The LAUNCH command initiates a batch file created in the EVM/bin folder. A Storage Data Protector backup specification was created in order to call the backup job name using the Storage Data Protector command line. Once the replicated volumes were mounted, the batch file started the backup specification and ran the backup job. The virtualization technology built into EVA allows for instant backup of a Vsnap that is building or a snapclone that is normalizing.

The batch file below is initiated by the EVM LAUNCH command and starts a predefined Storage Data Protector backup specification.

```
@ echo off
cd \Program Files\OmniBack\bin
omnib -datalist "datalist_name"
exit
```

- Backing up with VERITAS NetBackup V3.4

For information on EVM and VERITAS NetBackup V3.4 integration, you can review the following document:

Enterprise Backup Solution with VERITAS NetBackup Datacenter v3.4
[13ZJ-1200A-WWEN.pdf on ftp.compaq.com](http://ftp.compaq.com/13ZJ-1200A-WWEN.pdf)

Similar to Storage Data Protector, NetBackup requires that the backup job be configured before EVM initiates it. When including a batch file in the EVM job creation form, the following command was used to initiate a backup job:

```
c:\veritas\netbackup\bin\bpbackup -i -c <class>
```

The above example assumes that NetBackup has been installed in the default directory. The two switches (-i, -c) tell NetBackup to start the job immediately and to back up by class. The field <class> refers to the NetBackup class in which the EVM client resides.

sample backup plan

A solid backup plan is the key ingredient to successful recovery in the event of a database disaster. The backup plan depends on the database administrator and the operating environment. HP recommends keeping a minimum of two clones available at all times and cycling every 48 hours. With this approach, disks recycle automatically and provide rapid restoration for up to 48 hours. Here is a sample backup plan:

Monday

Undo Clonejob1
Run Clonejob1
Tape Backup Optional
Log backups throughout day

Tuesday

Undo Clonejob2
Run Clonejob2
Tape Backup Optional
Log backups throughout day

Wednesday

Undo CloneJob1
Run Clonejob1
Tape Backup Optional
Log backups throughout day

Thursday

Undo Clonejob2
Run Clonejob2

Friday

Undo Clonejob1
Run Clonejob1

Saturday

Undo CloneJob2
Run Clonejob2

Tape Backup Optional
Log backups throughout day

Tape Backup Optional
Log backups throughout day

Tape Backup Optional
Log backups throughout day

Sunday

Full tape backup of online SQL Server database and log files.

Snapclones can also be used in the following ways:

- As a backup for the SQL Server
- To run multiple operations, such as nondisruptive verification and offline testing
- To enable granular restores by running the SQL stored procedure `sp_attach_db` or by using the SQL Enterprise Manager GUI

restoring SQL server using snapclones and Vsnap

This section describes restore methods using snapclones and Vsnap. Each restore method depends on the type of backup that was done, what type of failure occurred, the condition of the database and logs at the time of failure, and the time available for recovery.

Database recovery options:

- Restore database from snapclone (no further transaction log backups can be applied). This option is best when you wish to restore the database to the point in time when the snapclone was created.
- Restore database from snapclone and apply transaction logs. This option is best when you want to restore incremental transaction log backups.
- Restore database from Vsnap using the Windows drag-and-copy method. This option is best when you are using small database volumes (less than 10GB) and can move the data over quickly.

Be sure to choose the appropriate restore method. Choosing the incorrect restore procedure may prevent transaction logs from being applied for recovery and could result in being unable to restore to the latest log backup. Keep in mind that you can also restore the database from tape. This option is best when time is not an issue and you want a full database restore from tape.

Scenario 1: Database disks have crashed and transaction logs are intact or the cloned transaction logs are current. You decide to restore the database from a snapclone.

1. Back up transaction logs.
2. Detach the failed database from the SQL Server. You can write a script using `sp_detach_db` and Query Analyzer or through the Enterprise Manager GUI.
3. Stop the SQL Server cluster service. You can just stop the service or take the SQL Server service offline using Cluster Administrator.
4. Remove database volumes as SQL Server cluster dependencies in Cluster Administrator, and delete database cluster disks.
5. Swap the corrupt database disks with the snapclone volumes by un-presenting the volumes using the HSV Element Manager on the Storage Management Appliance. Open the corrupt active virtual disk and click the Unpresent tab. Unpresent the corrupt active disks from the cluster nodes.
6. To present the snapclones to the cluster hosts, open the snapclone active virtual disk and click the Present tab. Select the cluster nodes one at a time.
7. Change the drive letters on the new volumes to reflect the drive letters of the original database volumes using Computer Manager – Disk Management.

8. If the transaction logs are intact, then retain the original log disk and mount the snapclones of the data (.mdf) files only. If the transaction logs are damaged, go back to step 5 to swap the snapclone of the log files.
9. Add the snapclone volumes as cluster resources and SQL Server cluster dependencies.
10. Start the SQL Server cluster service.
11. Attach the database with the new volumes and the existing/original logs using Enterprise Manager or sp_attach_db through Query Analyzer.
12. Go back to EVM. On the Job Status screen, delete the job that originally created the snapclones. It is important that you delete this job so that it can never be undone.

Caution: Clicking UNDO will delete the volumes and you will lose your database.

In cluster configurations, the fastest way to restore the SQL Server resource disks is to use the same drive letters as the original disks. Figure 6 shows the SQL Server cluster disk resources.

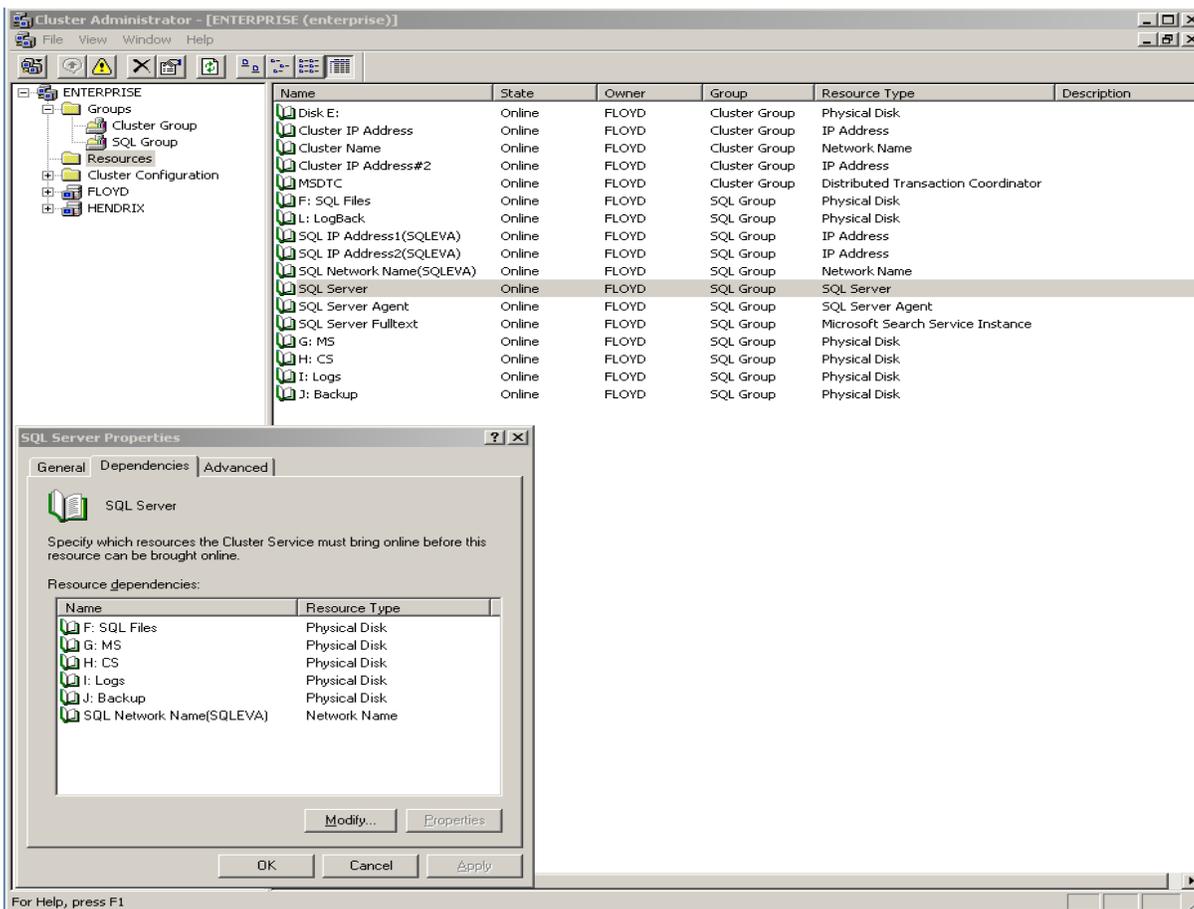


figure 6 – Cluster Administrator

Scenario 2: Data and logs need to be restored. Multiple transaction log backups need to be restored to roll forward the database with it is attached. You decide to restore the database from a snapclone and apply transaction logs.

1. Back up transaction logs.
2. Detach the bad SQL database using Enterprise Manager or Query Analyzer.
3. Stop the SQL Server cluster service.

4. Remove database volumes as SQL Server cluster dependencies in Cluster Administrator, and delete database cluster disks.
5. Swap the corrupt database disks with the snapclone volumes by un-presenting the volumes using the HSV Element Manager on the Storage Management Appliance. Open the corrupt active virtual disk and click the Unpresent tab. Unpresent the corrupt active disks from the cluster nodes.
6. To present the snapclones to the cluster hosts, open the snapclone active virtual disk and click the Present tab. Select the cluster nodes one at a time.
7. Change the drive letters on the new volumes to reflect the drive letters of the original database volumes using Computer Manager – Disk Management.
8. Add the snapclone volumes as cluster resources and SQL Server cluster dependencies.
9. Start the SQL Server cluster service.
10. Run the DButil –restoresns command using the script below to put the database into a loading state. In the loading state, the database is ready to accept a restore of the transaction logs. The DButil –restoresns command uses the meta file to put the database in a loading state. If you are restoring a SQL Server database to a new SQL server with a different SQL server name, you need to edit the name of the meta file to match the new name of the SQL server.

DButil –restoresns <sqlserver/instance_name> <database_name> <username> <password>
11. Use Enterprise Manager to verify that the database is in a loading state and apply the sequential transaction logs. You can use Enterprise Manager, Query Analyzer, or a third party backup application. When using Enterprise Manager or Query Analyzer, be sure to use the No Recovery option for each log backup but the last, and choose the With Recovery option for the last log backup to make the database operational. See Figure 7.
12. Go back to EVM. On the Job Status screen, delete the job that originally created the snapclones. It is important that you delete this job so that it can never be undone.

Caution: Clicking UNDO will delete the volumes and you will lose your database.

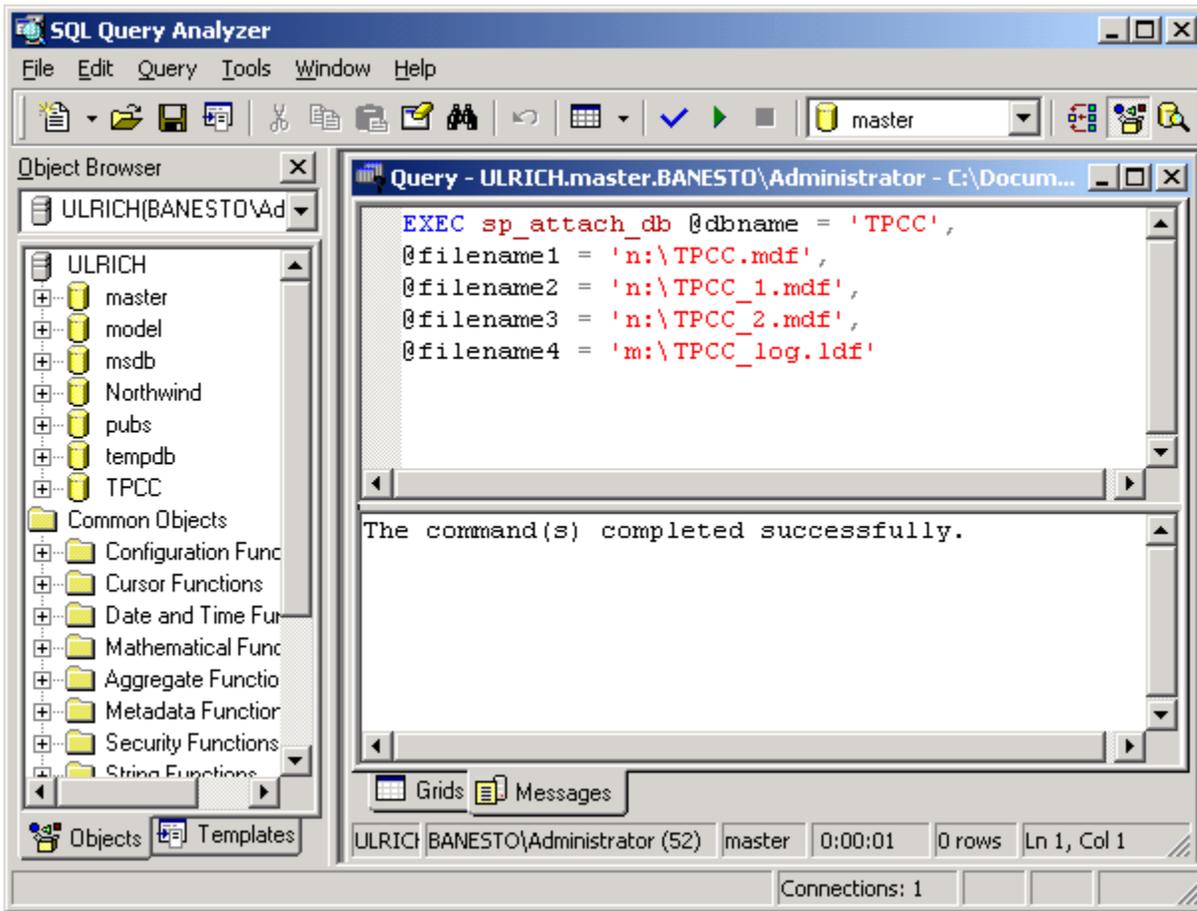


figure 7 – Query Analyzer

Scenario 3: Data and logs need to be restored. There is a valid Vsnap as a recovery device and Windows drag-and-copy is the best means to transfer the files due to the minimal size of the Vsnap. You decide to restore the database from the Vsnap using Windows drag-and-copy.

1. Back up transaction logs.
2. Detach the bad SQL database using Enterprise Manager or Query Analyzer.
3. Stop the SQL Server cluster service.
4. Remove database volumes as SQL Server cluster dependencies in Cluster Administrator, and delete database cluster disks.
5. From the active cluster node, assign alternate drive letters to the original database volumes.
6. Present the Vsnaps to the active cluster node.
7. Copy files from the Vsnaps to the original database volumes.
8. Unpresent the Vsnaps from the active cluster node.
9. Reassign the original drive letters to the original database volumes, and add as SQL Server cluster resources.
10. Start the SQL Server cluster resource.

Note: If transaction logs are going to be applied, then go to Step 9 in restore scenario #2. If you are not applying transaction logs, then proceed with this next and last step.

11. Attach the database with the newly copied volumes.

related solutions

Rapid Restore for SQL Server 2000 on HSG80 Storage Solution

www.compaq.com/products/storageworks/solutions/rapidrestoresql/index.html

Virtualized Storage Management for SQL Server 2000

<http://www.compaq.com/products/storageworks/solutions/vsm-sql2000/index.html>

for more information

For additional information on the HP StorageWorks Rapid Restore for SQL Server 2000 Solution

www.compaq.com/products/storageworks/solutions/rapidrestoresql/index.html

For the script kit and DBUtil utility for HP StorageWorks Rapid Restore for SQL Server 2000

<http://www.compaq.com/products/storageworks/solutions/request/rrsql2kit.html>

For Microsoft SQL Server information:

<http://www.microsoft.com/sql/default.asp>

For HP StorageWorks Enterprise Virtual Array information and documentation:

<http://www.compaq.com/products/storageworks/enterprise/index.html>

For HP OpenView Storage Management Appliance information and documentation:

<http://www.compaq.com/products/sanworks/managementappliance/index.html>

For HP StorageWorks Enterprise Volume Manager documentation:

<http://www.compaq.com/products/sanworks/evm/index.html>

For EVM integration with Enterprise Backup Solutions, including VERITAS NetBackup:

<http://www.compaq.com/products/storageworks/ebs/documentation.html>

For HP StorageWorks Secure Path information and documentation:

<http://www.compaq.com/products/sanworks/secure-path/documentation.html>

For HP OpenView Storage Data Protector information and documentation:

<http://www.openview.hp.com/products/dataprotector/index.asp>

For VERITAS NetBackup information and documentation:

<http://www.veritas.com >> Products>

To provide feedback to HP about this document:
Microsoft_Storage_Solutions@hp.com

trademark and date information

All brand names are trademarks of their respective owners.
Technical information in this document is subject to change without notice.
© Copyright Hewlett-Packard Company 2002
11/02