hp storage solutions

april 2003

implementation blueprint

hp StorageWorks rapid recovery for Exchange 2000

Using hp StorageWorks business copy EVA V2.1 and hp StorageWorks enterprise virtual array V2.0A

table of contents

executive summary

The Rapid Recovery for Exchange 2000 solution is a quick and complete method of recovering Microsoft Exchange 2000 databases with minimal disruption to Exchange services. The solution uses **snapclones** as a source for recovery and automated offline backups. A snapclone is a complete physical point-in-time copy of a data volume. Snapclones can be used to perform an **extremely rapid restoration** of Exchange Information Stores.

Using HP StorageWorks Business Copy EVA V2.1 and HP StorageWorks Enterprise Virtual Array (EVA) V2.0A, customers can create snapclones of their Exchange Server databases and resume full operation of their Exchange Server environments in minutes.

By using the information in this solution, customers will be able to do the following:

- Create snapclones of the Exchange 2000 databases
- Restore data from snapclones
- Create scripts to implement automatic backups from snapclones
- Backup from snapclone to tape
- Use either clustered or standalone Exchange 2000 servers

Administrators face lengthy delays when restoring an Exchange environment from tape. This solution, which is based on creating snapclones with StorageWorks Business Copy EVA (BC EVA) in a SAN-based configuration, reduces the restore time from hours to minutes and reduces disruption to Exchange users. These time-savings can be dramatic, as illustrated in the following diagram:

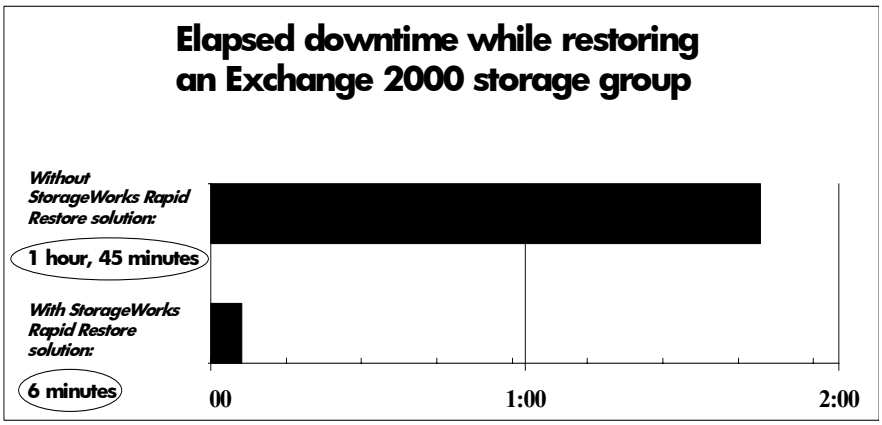

figure 1. comparative restoration times – Rapid Recovery for Exchange 2000 solution as compared to tape (single Exchange 2000 storage group with 2 databases, 3000-user load, logs not applied)

By using this solution to augment your tape-based recovery plan, you can recover your Exchange Information Stores rapidly, decreasing the length of application downtime. Since a snapclone provides an exact copy of your production data, you can offload operations, such as tape backups, testing, and maintenance, from the production server to a backup server. And the entire solution is engineered within Microsoft support guidelines.

Here's what Microsoft says…

"By partnering with companies like Hewlett-Packard, the worldwide prime integrator for Exchange 2000, we are able to deliver highly beneficial storage solutions to our customers. The **Rapid Recovery for Exchange 2000** and Virtualized Storage Management for Exchange 2000 solutions take advantage of highly scalable and easy-to-manage storage systems and management software, as well as the company's expertise in designing, deploying, and managing Exchange 2000 infrastructures. These HP StorageWorks solutions provide a great way for customers to maximize the availability of their Exchange 2000 environment."

Kevin McCuistion, Microsoft Corporation
Group Product Manager, Exchange

features and benefits

This solution provides the following benefits to the Exchange administrator:

- Dramatically improved restoration times for Exchange 2000 databases
- Best practices for maximizing Exchange 2000 availability during database recovery
- Simplified implementation and management, including automation examples
- Software engineered within Microsoft guidelines
- Investment protection by leveraging existing HP hardware and software, supporting multiple configurations, and providing interoperability with future products
- Integration with the leading tape backup applications for Exchange —specifically tested with HP OpenView Storage Data Protector V5.0 and VERITAS NetBackup V3.4/ V4.5

about clones, snapshots, vsnaps, and snapclones

With EVA, you can create traditional snapshots, demand-allocated snapshots (vsnaps), and snapshots that normalize into snapclones. Virtually instantaneous **snapclones** are the powerful recovery technology used in this Rapid Recovery for Exchange 2000 solution.

Snapclones

Snapclones are complete physical point-in-time copies of data volumes and are ideal for reducing I/O loads on the production volumes. Snapclones are snapshots that unshare (a function similar to normalization) into clones. A snapclone has properties that are similar to 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 unsharing 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 (Virtual RAID [VRAID] 0,1,or 5), and they do not require extensive advanced preparation. Resynchronization times are also eliminated with snapclones since a current snapclone copy can be available for use in moments rather than hours, with the unsharing process occurring in the background.

Snapshots

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. Snapshots can be created quickly and are immediately available for use. Since snapshot volumes share data with the original production volume (those blocks that have not changed), they can only be used to recover from data loss or corruption as a result of writes to the original database.

Virtual Snapshots

Vsnaps are similar to snapshots, but allow on-the-fly space allocation. A vsnap is a virtually capacity-free snapshot that is a space-efficient point-in-time copy of the data. You can replicate data instantly by taking a "picture" of the data within seconds and without reserving storage capacity equal to the production volume. With vsnaps, 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.

Clones

A clone is a complete physical copy of the original data. Clones require the allocation of storage capacity that is equal to the size of the production volume. And clone creation is not immediate, as the original data must be copied to the clone volume. You can create a triple mirror set (RAID set with an additional mirror copy) up front and eliminate the time required to create the original clone volume. However, updating the RAID set with the clone volume requires time for resynchronization with the production volume.

See table 1 for a comparison of clones, snapshots, vsnaps, and snapclones.

table 1. comparison of clones, snapshots, vsnaps, and snapclones

	Clones	**Snapshots**	**Vsnaps**	**Snapclones**
Array (Controller) Support	EMA (HSG80)	EMA (HSG80) EVA (HSV110)	EVA (HSV110)	EVA (HSV110)
Storage Capacity Allocation	Equal to the production volume	Equal to the production volume	On-the-fly storage allocation based on writes to the production volume	Equal to the production volume
RAID Level	RAID 0, 1, 5	VRAID 1, 5	VRAID 0, 1, 5	VRAID 0, 1, 5
Usage for data recovery	Physical copy of the production volume, available for point in time recovery. Requires re-synchronization.	Virtual copy of the production volume, available immediately for point-in-time recovery. Not available to restore in cases where production volume is lost.	Similar to snapshot. Only uses storage capacity required to save changes, more capacity efficient.	Physical copy of the production volume, available immediately for point-in-time recovery. Does not require re-synchronization.

solution overview

The Rapid Recovery Solution for Exchange 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 the EVA V2.0A, BC EVA V2.1, and the HP StorageWorks Enterprise Backup Solution (EBS) as illustrated in figure 2.

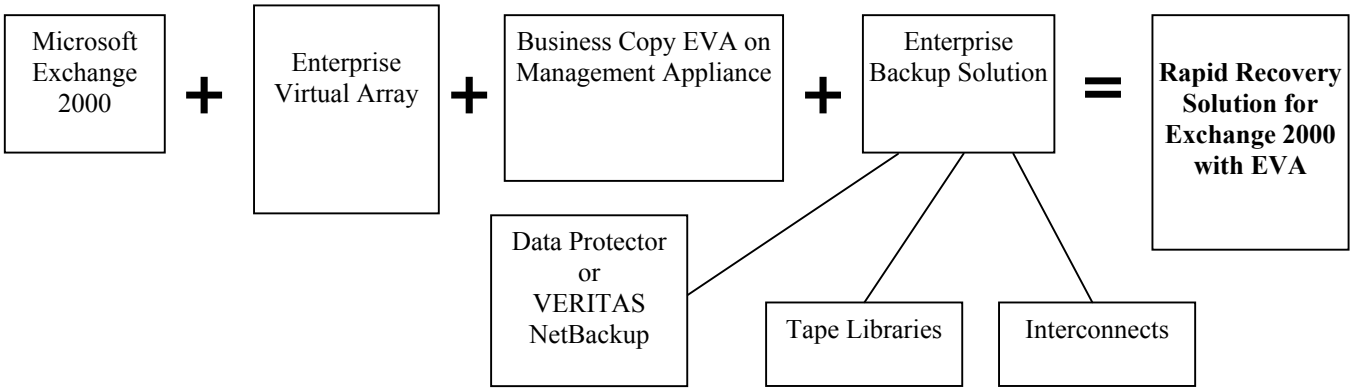

figure 2. Rapid Recovery for Exchange 2000 solution building blocks

solution components

The Rapid Recovery for Exchange 2000 Solution consists of these core components:

- Microsoft Exchange 2000
  You can configure the Exchange 2000 application on any Windows 2000 server and use a Microsoft Cluster Server (MSCS) configuration for improved application availability.
- Enterprise Virtual Array (EVA)
  The Exchange 2000 database and log volumes must be located on the Enterprise Virtual Array. Snapclones are created on the virtual array using virtual controller software (VCS), which operates the storage subsystem. Based on the HSV110 controller, the 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.

  Note that, should the storage capacity requirements be met, the eva3000 array can also be used in conjunction with this solution configuration.
- Business Copy (BC EVA)
  Business Copy is browser-based storage management software that facilitates controller-based clone operations to make a block-to-block copy of a storage volume. With BC EVA, you can create, run, and manage automated storage replication jobs, as well as link them with external jobs.

  You must have a Business Copy EVA host agent on each client to utilize controller-based database cloning. The HP OpenView Storage Management Appliance is the operating platform for BC EVA manager, which offloads the processing from the host systems.
- SAN-Based Backup Infrastructure
  Enterprise Backup Solution with HP OpenView Storage Data Protector V5.0 or VERITAS NetBackup V3.4 / V4.5 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 snapclone-based offline tape backups. 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 Exchange 2000 application must be loaded on the backup server.

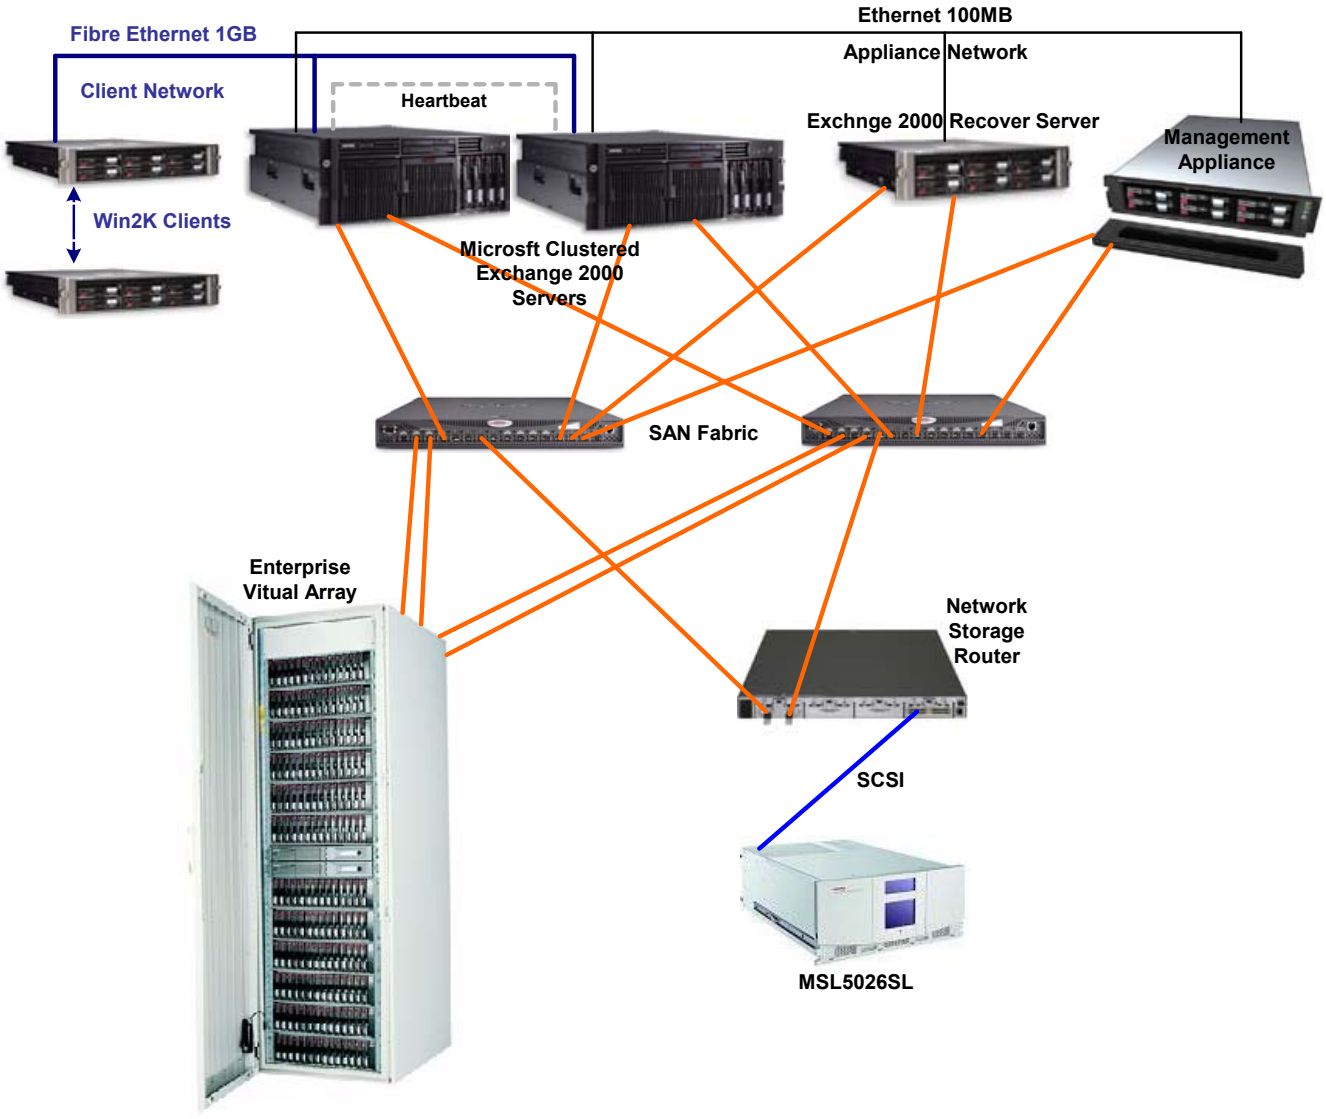

figure 3. Rapid Recovery for Exchange 2000 solution configuration (as tested)

best practices

- Traditional Exchange 2000 Server configuration rules apply when you start planning your storage. Use VRAID 1+0 for EDB, STM, and log files. The EVA uses VRAID technology.
- 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 from the data files.
- Keep log files on volumes that have been configured using VRAID 1 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 virtual disk 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.
- Always configure disk groups with an even number of drives.
- Always attempt to configure disk groups with drives of identical or similar capacity and performance.
- Whenever possible, strive for disk groups that are multiples of the number of shelves, or if not that, multiples of six or eight drives.
- BC EVA 2.1 Host Agent Release Notes recommend logging out of the backup host when running BC EVA scripts (jobs) that mount volumes to it.

test environment

This section describes the environment used to test the Rapid Recovery for Exchange 2000 solution. It describes the fully redundant no-single-point-of-failure configuration and lists the hardware and software components used in the solution.

network infrastructure

The Ethernet was set up and configured for cluster operations using three different network segments. A network segment 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 Fibre network segment was set up for client access to the application servers. It was set up for client access only and configured as a public network. The third network segment was set up for the Management Appliance and its applications. This network was set up as a mixed network, so that the applications running on the Management Appliance could operate corresponding agents on the hosts connected to the SAN. Three additional IP addresses were created for the virtual cluster environment.

The SAN was configured with two redundant 2-GB 16-port SAN switches and managed by the Management Appliance. Two Fibre Channel paths for each server offered a fully redundant configuration.

Exchange recovery server

A dedicated ProLiant DL380 G2, loaded with Microsoft Windows 2000 Advanced Server, was used for SAN-based tape backups. Exchange 2000 Server was installed on this server. Using this method allows for individual Exchange Server mailbox restores, offline integrity checks, and testing new versions of software. By mounting the snapclones on the backup server, they could be backed up to a tape library without impacting the production Exchange Server. VERITAS NetBackup V3.4 and HP OpenView Data Protector V5.0 were installed and used as the backup applications. A HP MSL5026 tape library using SDLT 110/220 tape drives was used to back up the databases and logs.

enterprise virtual array

The Enterprise Virtual Array (EVA) consists of three main components:

- HSV Element Manager 2.0A (Build 255)—the user interface that communicates with the HSV controllers that control and monitor the EVA.
- VCS 2.02—the firmware that allows the EVA to communicate with the HSV Element Manager and controls the HSV 110 controller.
- Hardware—the physical pieces that constitute the EVA, such as drive enclosures, switches, and racks.

The EVA in the test setup had (84) 36.4-GB disks (3.1 TB). Note that an EVA can be expanded up to 240 disks, with a maximum capacity of 17.5 TB, using 72-GB drives on a single controller pair. Refer to the BC EVA user guide in order to understand the available EVA configuration options.

virtual RAID technology

A virtual disk is a simulated disk drive created by the HSV 110 controllers as storage for one or more hosts. Virtual disk characteristics provide a specific combination of capacity, availability, performance, and accessibility. The controller pair simulates these characteristics by deploying the disk group specified for the virtual disk. The host computer sees the virtual disk as "real," exactly as it would see a physical disk with the same characteristics.

The maximum number of virtual disks allowable is 512. A single virtual disk can be presented to multiple hosts. The maximum size of a virtual disk is 2047 GB. The maximum number of total presentation is 8192 GB.

There are three types of virtual disks:

- Active member of a virtual disk family—a virtual disk (LUN) that is accessed by one or more hosts for storage. An active member of a virtual disk family is automatically created whenever a new virtual disk family is created. An active virtual disk member and its snapshot, if they exist, constitute a virtual disk family.
- Snapshot virtual disk—a virtual disk that reflects the contents of another virtual disk at a particular point in time. A snapshot operation can only be done to an active member of a virtual disk family. A snapshot is intended to be temporary. A maximum of seven snapshots per virtual disk can exist at one time.
- Virtual disk snapclone, or virtually instantaneous snapclone—a virtual disk that is an exact copy of another virtual disk at a particular point in time. A snapclone operation can only be done to an active member of a virtual disk family. The snapclone, like a snapshot, reflects the contents of the source virtual disk at a particular point in time. Unlike the snapshot, the snapclone is an actual clone of the source virtual disk and becomes an independent active virtual disk member of its own disk family when all the data is replicated.

  Within the Virtual Disk Folder in the Navigation pane are virtual disk families. You can click on any one of the virtual disk families within the Virtual Disk Folder to see the properties page for a specific virtual disk family (see Figure 4).

The EVA does not use traditional RAID sets by disk capacity size and type of RAID protection. Rather, a virtual disk is defined, drawing its capacity from a designated pool of storage called disk groups. Virtualization allows data to be spread across more spindles (disks), dramatically improving performance. Because of storage pooling virtualization, the EVA is able to support multiple virtual disks of varying capacity and virtual RAID (VRAID) types within a single storage group. In addition, all virtual disks in a disk group spread their capacity across all the physical disks that contribute to that pool.

For example, a traditional RAID 1 (mirrored) configuration for 36 GB of data space requires two 36-GB disks to be assigned as RAID 1 (see figure 4). So when the data, which might be in a high-performance application, is read or written, there are only two disks involved, thus limiting performance. With virtualization, this same 36-GB RAID 1 requirement is spread across eight or more disks in a disk group. Volume performance is greatly improved, because more spindles are used.

Virtualization also allows data to be redistributed across physical disks within a virtual disk if an activity occurs that causes a change to the disk group. As a result, the EVA is able to utilize an on-the-fly leveling algorithm to balance performance without interrupting ongoing workloads. This process redistributes each virtual disk's blocks evenly across as many spindles as the virtual disk's redundancy type allows. This leveling process is activated whenever the EVA detects an opportunity to improve utilization, such as a change in the number of disks in the pool.

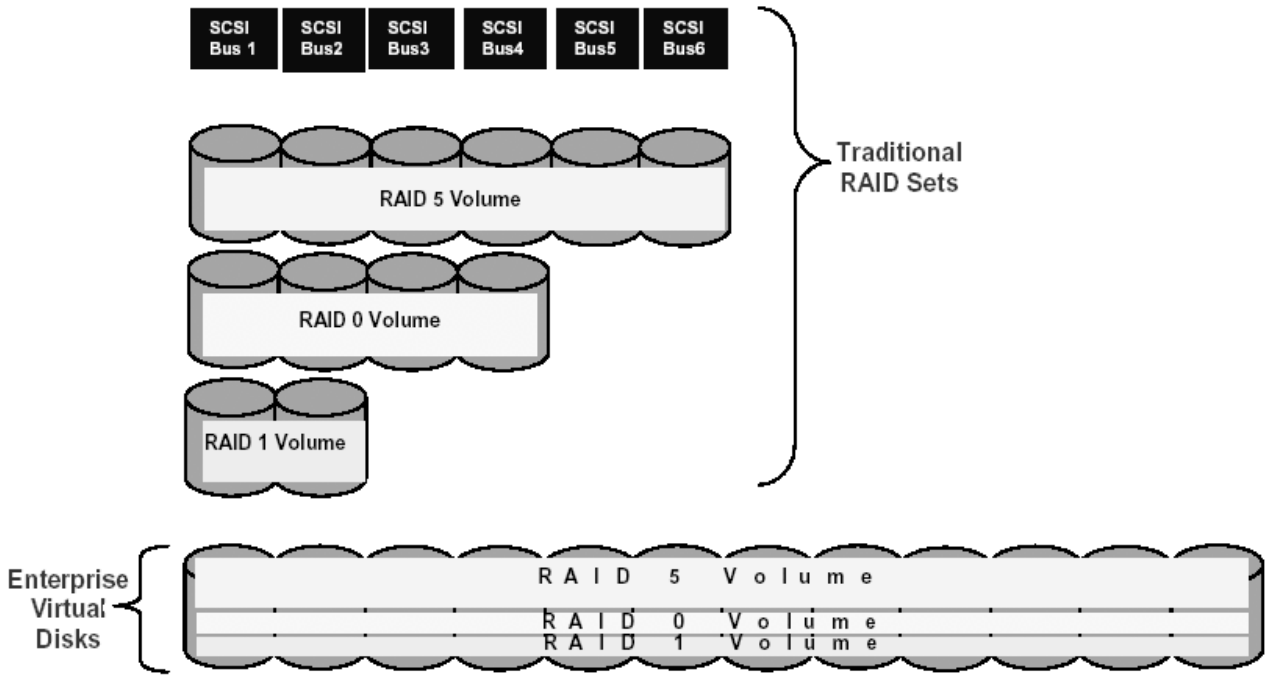

figure 4. virtual RAID

For a full description and installation documents for HSV Element Manager and EVA, visit these URLs:

http://www.hp.com/products/sanworks/managementappliance/index.html
http://www.hp.com/products/storageworks/enterprise/index.html

Exchange servers

The Exchange servers used Microsoft Windows 2000 Advanced Server SP3 and Microsoft Exchange Server 2000 Enterprise Edition SP3.

Two ProLiant DL580 servers, each having two FCA2101 (2Gb HBA) adapters, powered the Exchange Server (active/passive) cluster.

Each server used HP Secure Path 4.0A to manage a high-availability multiple-path Fibre Channel connection to the SAN.

Microsoft Load Simulator

The tests were conducted using Microsoft Load Simulator (LoadSim) load generator with a MMB2 (MAPI Messaging Benchmark) profile. LoadSim is a tool for simulating a medium corporate e-mail user using Microsoft Outlook 2000. Its purpose is to enable a Windows 2000 machine, called a LoadSim client, to simulate multiple Exchange Server users. MMB2 measures throughput in terms of a specific profile of user actions, executed over an 8-hour working day. The mail tasks measured in these tests included send and receive, browse, read, and forward, as well as scheduling tasks, logon/logoff, and distribution list usage.

Exchange virtual servers

Two virtual Exchange cluster servers were created (VS1-EXCH, VS2-EXCH). Each virtual server had a single storage group with two databases. The storage groups were labeled as SG1 and SG2. The first storage group contained two databases (SG1DB1, SG1DB2) and a public folder store (Public (VS1-EXCH)). The second storage group contained two databases (SG2DB1, SG2DB2). Using Microsoft LoadSim, 3000 MMB2 users were generated and

were equally distributed in each database. Each storage group contains 1500 mailboxes (750 mailboxes in each database).

Circular logging was disabled in both storage groups so that the transaction logs could be rolled forward during a restoration.

The /3GB switch was added to the boot.ini file. If your Exchange server has more than 1 GB of RAM, you need to add the /3GB switch to the startup line of the Windows boot.ini file. This change is required so that virtual memory for Exchange 2000 can take advantage of the additional RAM.

See Microsoft Knowledge Base article Q266096 (Exchange Server Requires /3GB Switch with More Than 1 Gigabyte of Physical RAM) for detailed instructions.

http://support.microsoft.com/default.aspx?scid=kb;en-us;291988

Note: Some software applications will fail during installation because of the /3GB switch. The error message displays **out of memory**. If the administrator receives this message during installation of an application the /3GB switch should be turn off until the software can be installed. Since the switch is configured in the Boot.ini file Windows 2000 will need to be rebooted before the software can be install.

Exchange virtual disk configuration

This solution followed Microsoft's recommendation of separating the EDB and STM files from the transaction logs by placing them on separate virtual disks on the HP Enterprise Virtual Array (EVA). The virtual disks were created from a single EVA disk group of 24 disks. A disk group is first created through the HSV Element Manager. You must decide what size the disk group should be, based on the current space needed plus additional space for sparing and replication. Below is the Exchange storage group virtual disk structure. Each database had an EDB and STM file associated with it.

G:\SG1\exchsrvr\mdbdata\SG1DB1\	Data size 21.6GB, Disk size 75GB
G:\SG1\exchsrvr\mdbdata\SG1DB2\	Data size 21.7GB, Disk size 75GB
G:\SG1\exchsrvr\mdbdata\Public\	Data size 52MB, Disk size 75GB
H:\ SG2\exchsrvr\mdbdata\SG2DB1\	Data size 21.5GB, Disk size 75GB
H:\ SG2\exchsrvr\mdbdata\SG2DB2\	Data size 21.6GB, Disk size 75GB
I:\SG1log\	Data size 25MB, Disk size 25GB
J:\SG2log\	Data size 25MB, Disk size 25GB

A better performance model would be to move the logs to their own disk group, but there is a cost. The performance increase would come from the separation of sequential and random I/Os. The increased cost of this model would be wasted disk group capacity. When a disk group is created, it starts with an eight-disk minimum. In this solution, placing two 25-GB logs would be a waste of disk space. Using 36-GB drives, each log would need two disks for VRAID 1, plus one disk for overhead and one for spare.

Note: When you create a very large virtual disk (over 100 GB), HP recommends that you *do not* choose the **Present to Host** option on the **Create a Virtual Disk Family** page in the HSV Element Manager. Instead, wait until after the virtual disk is created, then go to the **Virtual Disk Active Properties** page and choose **Present**.

clients

There were seven clients connected to the Fibre gigabit Ethernet network. They all had identical configurations. Microsoft LoadSim was used to generate a medium MAPI load. The load was balanced across six clients with 450 users each, and a seventh client acted as a control unit with 300 users.

LoadSim generates the user directory database and then initializes the test by populating the Exchange storage group database. A test typically lasts for eight hours. During a test, the simulated users log on to the Exchange

server and call MAPI functions to create, send, delete, and otherwise process messages and attachments. The performance data is written to a log file, and the user actions are logged to the file for later analysis. After a test, LSLOG.exe, a command-line utility packaged with LoadSim, was used to parse the loadsim.log file and calculate the 95th percentile response time. A successful test must satisfy the Microsoft Exchange Server UPS Policy Guidelines V2.0A4. See table 4 for client scores.

hardware and software components

Table 2 lists the hardware and table 3 the 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. Refer to the Rapid Recovery for Exchange 2000 Technical Blueprint for the latest suggested configuration details.

table 2. hardware components

Server Hardware	**Part Number**
HP ProLiant DL580 servers – (2) Intel Xeon 700 MHz, 512MB ECC - SDRAM, 2MB Level 2 Cache, Integrated 10/100 NIC	(2) 155617-001
2GB RAM Memory	189083-B21
HP NC6136 Gbit Server NIC	(2) 203539-B21
36GB Ultra 3, 15K rpm drives in each server	(4) 232916-B22
FCA 2101 (LP952) HBA	(4) 245299-B21
Backup Server	
HP ProLiant DL380 G2 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	245299-B21
Management Appliance	
Storage Management Appliance II v2.0 SP1)	189715-002
Storage Hardware	
Enterprise Virtual Array 5000 – 2C6D-B, 42U M3220 Controller assembly with dual HSV110 controllers and 6 M5214 dual Fibre loop 14-bay drive enclosures	283198-B21
36GB 10K rpm dual-port 2Gb/sec FC-AL 1in drive	(84) 238590-B21
Enterprise v1.0 to V2.0a Upgrade Kit (Vixel Switches)	283266-B21
2Gb 16-Port SAN Switch	(2) 240602-B21
2Gb SFF-SW Transceiver kit	(32) 221470-B21
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
Backup Hardware	
Modular Data Router	163083-B21
MSL5026SL RM SDLT Minilibrary (2 drives)	231892-B22

table 3. software components

Software	Part Number
Windows 2000 Advanced Server with SP3	Microsoft Reseller
Microsoft Cluster Server – MSCS	Included
Microsoft Exchange Server 2000 Enterprise Edition with SP3	Microsoft Reseller
HP OpenView Storage Management Appliance Software V2.0 SP1 http://www.compaq.com/products/sanworks/managementappliance/index.html	Web download
HP StorageWorks Secure Path 4.0A (per License) HP StorageWorks Secure Path 4.0A (5 Licenses)	165989-B23 231292-B23
HP StorageWorks Business Copy EVA v2.1 HP StorageWorks Business Copy EVA V2.1a HP StorageWorks Business Copy Upgrade UI EVA v2.1 HP StorageWorks Business Copy Upgrade UI EVA V2.1a	326719-B21 326719-B22 331246-B21 331246-B22
WNT/W2K KIT V2.0 ENT VIR ARY (for Windows NT 4.0/2000)	250195-B22
HP StorageWorks Business Copy 4TB LTU EVA 5000 (License only)	326724-B21
VCS PKGV2.0A Dual HSV Controller (base controller software)	250203-B24
HP OpenView Data Protector 5.0 with Exchange 2000 Data Agents Starter pack SAN drives Online extension for Exchange agents	 B6961AA B6953AA (one per drive) B6965BA
VERITAS NetBackup V 3.4 Server Software and Exchange Data Agents	Third Party

performance results

Performance results from this solution are a way to measure the integration of the components listed above. The various charts, tables, and graphs below show the results from the various tests.

The overall performance for the HP Rapid Recovery for Microsoft Exchange 2000 solution performance passes the MMB2 test as specified by Microsoft. There was no significant hardware or software degradation during any of the tests. Microsoft MMB2 results are shown in table 11, at the end of this section.

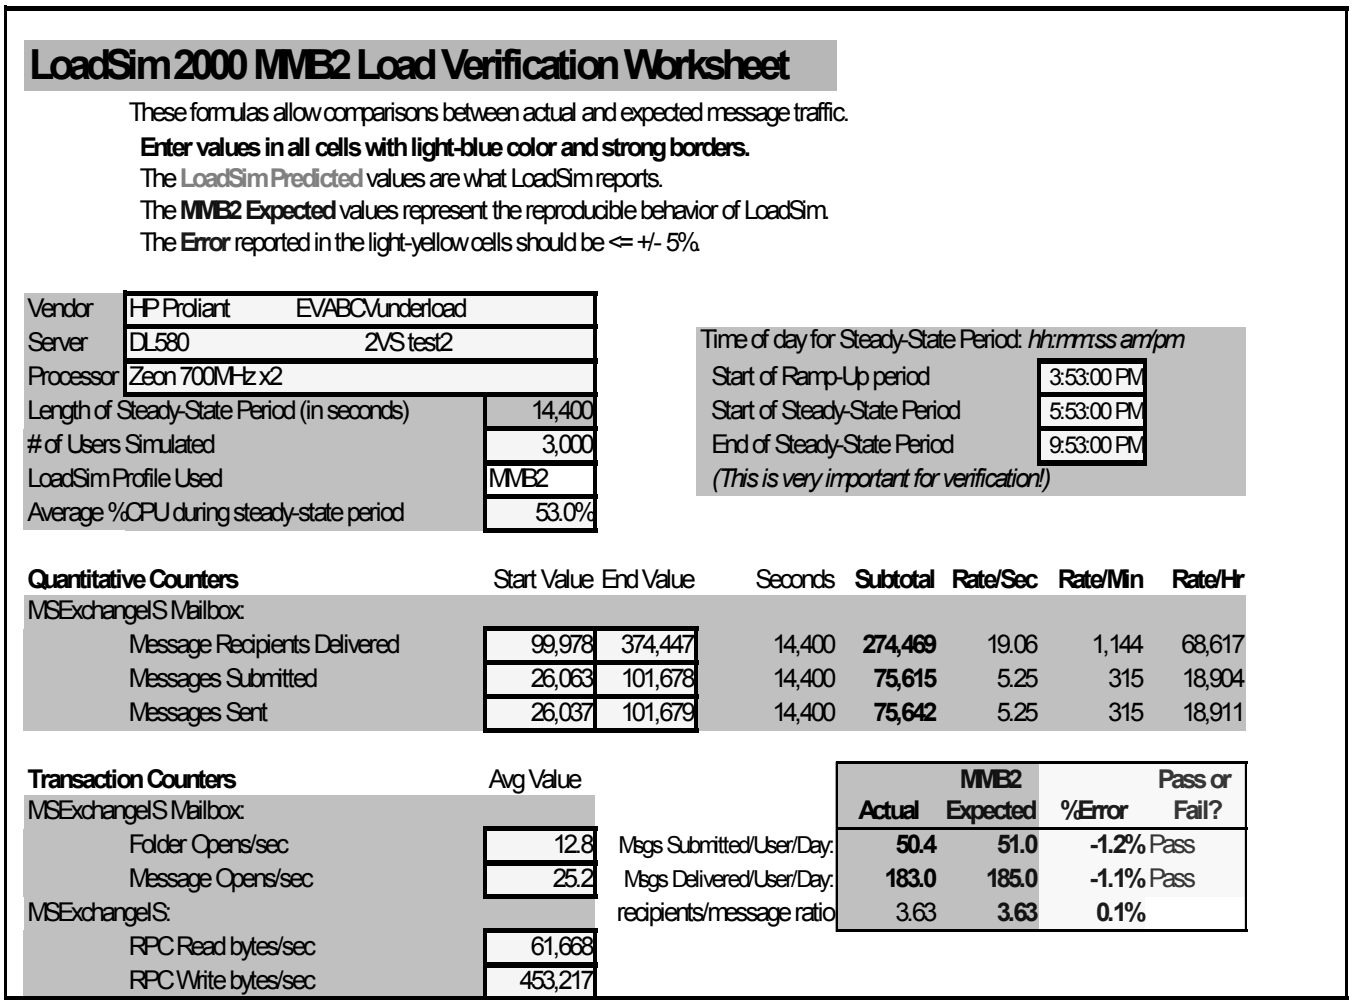

LoadSim 2000 MMB2 Load Verification Worksheet

These formulas allow comparisons between actual and expected message traffic.
Enter values in all cells with light-blue color and strong borders.
The **LoadSim Predicted** values are what LoadSim reports.
The **MMB2 Expected** values represent the reproducible behavior of LoadSim.
The **Error** reported in the light-yellow cells should be <= +/- 5%

Vendor	HP Proliant EVABCVunderload
Server	DL580 2VS test2
Processor	Zeon 700MHz x2
Length of Steady-State Period (in seconds)	14,400
# of Users Simulated	3,000
LoadSim Profile Used	MMB2
Average %CPU during steady-state period	53.0%

Time of day for Steady-State Period: *hh:mm:ss am/pm*

Start of Ramp-Up period	3:53:00 PM
Start of Steady-State Period	5:53:00 PM
End of Steady-State Period	9:53:00 PM

(This is very important for verification!)

Quantitative Counters	Start Value	End Value	Seconds	**Subtotal**	**Rate/Sec**	**Rate/Min**	**Rate/Hr**
MSExchangeIS Mailbox:							
Message Recipients Delivered	99,978	374,447	14,400	**274,469**	19.06	1,144	68,617
Messages Submitted	26,063	101,678	14,400	**75,615**	5.25	315	18,904
Messages Sent	26,037	101,679	14,400	**75,642**	5.25	315	18,911

Transaction Counters	Avg Value
MSExchangeIS Mailbox:	
Folder Opens/sec	12.8
Message Opens/sec	25.2
MSExchangeIS:	
RPC Read bytes/sec	61,668
RPC Write bytes/sec	453,217

	Actual	**MMB2 Expected**	**%Error**	**Pass or Fail?**
Msgs Submitted/User/Day:	**50.4**	**51.0**	**-1.2%**	Pass
Msgs Delivered/User/Day:	**183.0**	**185.0**	**-1.1%**	Pass
recipients/message ratio	3.63	**3.63**	**0.1%**	

figure 5. LoadSim worksheet

The LoadSim worksheet in figure 5 is an extraction of a spreadsheet that Microsoft has created for the purpose of testing load verification. Each Exchange Server test that is reported in this solution had its Windows performance numbers entered into a similar spreadsheet to determine if any hardware or software components have caused a disconcerting load. Each test in this solution has passed within a 5% tolerance that Microsoft has established. All LoadSim worksheet data was re-entered into table 11.

Figure 5 reflects the second snapclone test while under load using two virtual Exchange servers. All the numbers for this test were entered into table 11. For example, rate per hour for message recipients delivered can be found in the far right column of table 11. In table 11 the hourly transaction load from each test should not vary more or less than 5% between tests. The significance of these numbers reflects the operation of LoadSim and the effects it has on the Exchange servers.

Additionally, the response time score as shown in table 4 represents a 95th-percentile score of the clients' measured test run. A response time score of 1000 ms or less is considered acceptable for e-mail users using Exchange Server's MAPI protocol.

table 4. client scores

	(1) MSCS Baseline with 1 Virtual Server	(1b) MSCS Baseline with 1 Virtual Server	(2) MSCS SnapClone with 1 Virtual Server		(1c) MSCS Baseline with 2 Virtual Servers	(2a) MSCS SnapClone with 2 Virtual Servers	(2b) MSCS SnapClone with 2 Virtual Servers
Category	95th percentile	95th percentile	95th percentile		95th percentile	95th percentile	95th percentile
SEND	125	125	157		156	172	188
READ	63	62	63		78	78	94
REPLY	63	62	63		63	78	94
REPLY ALL	78	63	78		78	93	109
FORWARD	78	63	78		78	93	109
MOVE	109	109	125		125	141	172
DELETE	47	47	47		62	63	63
S+ CHANGE	172	156	172		157	203	235
DELIVER	812	765	921		254	311	295
LOAD IMSG	31	16	47		31	31	31
RESOLVE NAME	16	16	16		16	16	16
SUBMIT	78	63	78		78	94	110
LOAD ATTACH	235	234	265		250	282	281
EMPTY FOLDER	1719	1656	1687		1844	1875	2062
OPEN MSG STORE	32	31	32		32	31	32
LOGON	5328	5297	5360		5344	5422	5438
BROWSE CALENDAR	157	141	187		172	218	234
MAKE APPOINTMENT	360	344	406		391	484	531
REQUEST MEETING	531	500	593		562	671	719
HANDLE MEETING REQUEST	313	297	343		329	406	438
HANDLE MEETING RESPONSE	94	94	109		109	125	141
JOURNAL APPLICATIONS	234	188	219		219	219	250
CREATE CONTACT	140	125	156		141	156	203
BROWSE CONTACTS	312	297	344		328	422	453
**Weighted Avg Score-----> **	**108**	**103**	**117**		**123**	**139**	**156**

test 1—baseline with BC EVA

BC EVA V2.1 was installed in the SAN environment and a baseline test performed to determine its impact on the Exchange Server and SAN environment. The test ran for 8 hours, during 4 hours of which the test was sampled to determine if any part of the configuration was stressing under the load.

The baseline had two components:

- Baseline 1a had a single-cluster virtual server but did not have BC EVA installed.
- Baseline 1b had a two-cluster virtual server installed. BC EVA agents were running.

results: Introducing BC EVA to the Exchange Server SAN environment did not have an impact on the performance of the environment. The baseline performance numbers were acceptable under Microsoft MMB2 guidelines. See performance numbers in table 11. Figure 6 shows the Windows performance monitor and the 4-hour steady state that was sampled during each test.

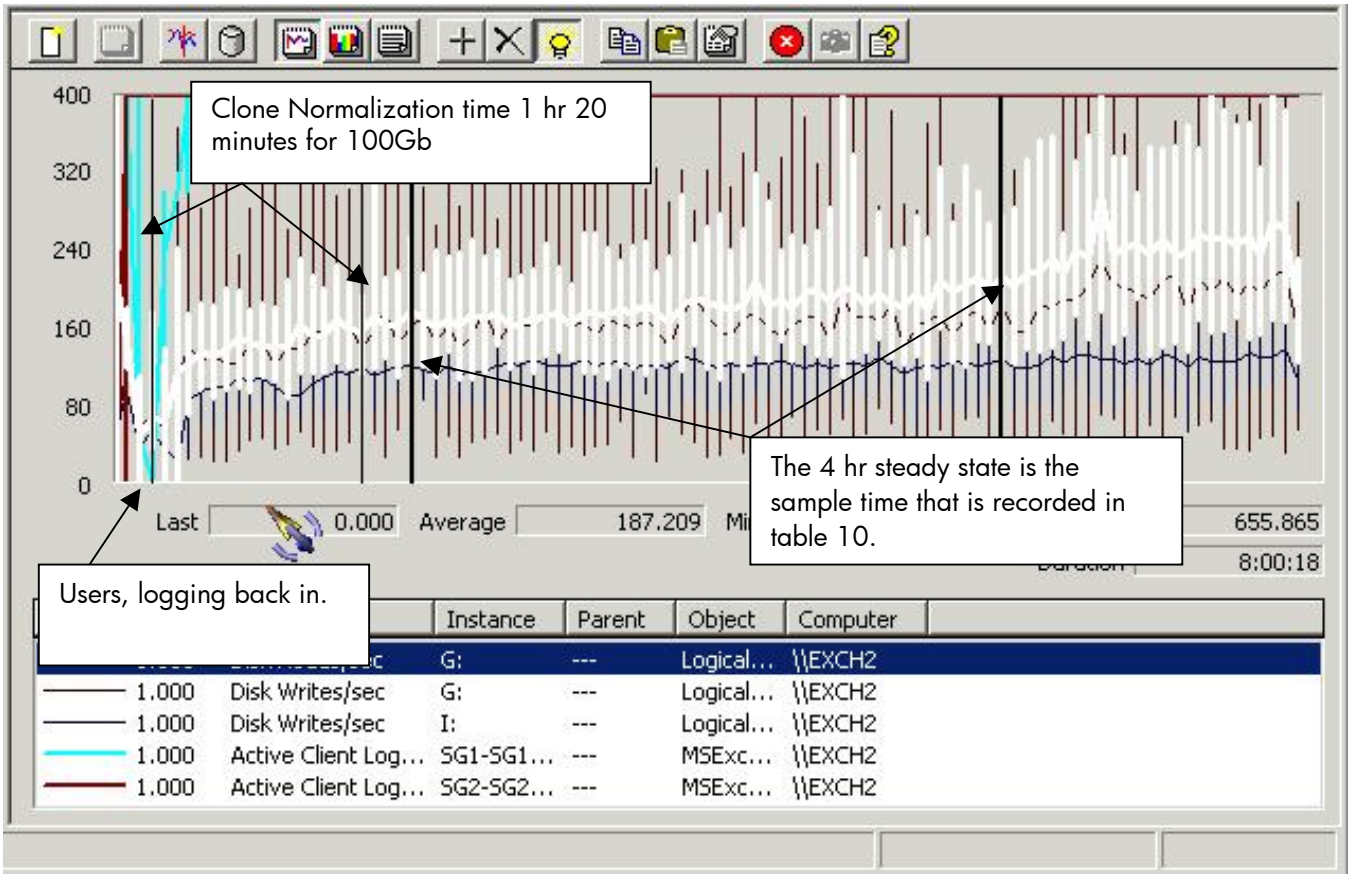

figure 6. MS Windows performance monitor

test 2—creating snapclones with a 3000 user load

Creating snapclones with HP StorageWorks Business Copy EVA involves utilizing available disk space that comes from the disk group of the parent virtual disk.

The snapclone process made a copy of two databases and logs from the SG1 storage group on virtual server EXCH1. Because the databases and logs were on separate LUNs, two snapclones were created in a single BC EVA job.

Before a snapclone can be activated, the data for the storage group has to be in a stable state. The data in an Exchange server is never in a stable state of transition, because data transactions from cache to logs to the message categorizer to storage are always in a state of flux. The databases in SG1 were stopped, which places the data in a stable state by dismounting all databases in SG1. During that process, the other Exchange storage group on virtual server EXCH2 stayed on line. Half the users (1500) in this Exchange configuration stayed on line during the test.

A snapclone on an EVA uses a two-step method of creating an image of a virtual disk (volume). The first step is the creation of a virtually instantaneous snapshot, with immediate access to this image. The second step is an automatic morphing process (BC EVA calls it unsharing) that creates a clone (BCV) or separate replicated image of the host volume. The snapshot element is then deleted once the unsharing is completed. In order to take advantage of this technology, BC EVA was used to automate the process. For more details on this process, see solution setup and installation instructions.

During this process, BC EVA also presented the snapclones to an Exchange recovery server, and it launched custom scripts to dismount and mount the databases, enable consistency tests on the snapclones using Exchange Server utility commands, and run the backup application.

Performance results for this test were only interested in the snapclone creation process.

results: The performance numbers for Test 2 are recorded in table 11 and the process time in table 5. The snapclone portion of the BC EVA job took 4.6 minutes. The highlighted section in table 5 breaks down each process of the BC EVA job for the snapclone. This time can vary, depending on the number of logged-on users and the current load of Exchange. This test was repeated several times to accumulate a weighted average.

Table 11 shows the impact of the Exchange Server environment during this test under load. The logical drive utilization section in table 11 shows the disk performance of a 4-hour sampling time. All the numbers are acceptable, as required to pass Microsoft MMB2 testing. There was an expected performance impact during the normalizing process. (See table 5) Disk queue lengths were increased by 15%. During normalization, Exchange operations continued: users were logging back in from being down because of the dismounting of one of the storage groups, two clones were being created, and LoadSim was generating mail messages. This was similar to a real production environment; therefore, the performance impact should be expected.

Normalizing will take some disk performance away in order to copy the disk. However, the average disk queue length to expect should be less than the number of spindles in the physical device. Since each virtual disk shares the load across all spindles in a disk group, the number of disks used for the disk group in this solution was 24. Table 11 shows that the average disk queue length was 3 during the snapclone test.

table 5. snapclone process time

Process	Time/Sec.
BC command process, cumulative	93
Snapclone (snapshot)	7
SG1 Dismount/Mount	60
BCV mount to Host	120
BC Delay command	60
DB integrity check	330
Launch Backup job	80
BC Job completion time in minutes	**12.5**
Normalizing BCVs average	18.3MB/sec

The BC EVA job in this solution averaged 12.5 minutes to complete. There are varying circumstances that will change job times that BC EVA has no control over. BC EVA has to wait for the HSV controller to complete commands that were issued. If the controller is busy, commands could be queued momentarily. BC EVA has limited control over scripts that launch and fail. Each script should be thoroughly tested before it is launched from BC EVA. Adding delays to the BC EVA job will prolong its completion. Issuing scripts that have to wait for an integer or to finish will prolong the BC EVA job. For example, there is a script included in this job that runs ESEFILE for database integrity and ESEUTIL for database consistency. There is no way to know how long this test will last. If there are problems with the database, the time could extend longer than that exhibited in this test. Another possible delay is network communication. When BC EVA mounts disks to a host, it has to wait for the host to complete the task.

test 3—backing up to tape

This test backed up the snapclone while the BCV was unsharing. It compared online to offline tape backup processes. This test was not designed to get the best performance but to show that BC EVA could back up the BCV from command line scripts. Both backup applications are sophisticated EBS solutions. Each backup application can be tweaked to get better performance results. A MSL5026 library was used during testing. Only one of the SDLT tape drives was used.

Custom backup scripts were designed to be launched by BC EVA and executed on the Exchange recovery server. The backup scripts were designed to back up the BCVs created by the snapclone. In order for scripts to work from BC EVA, the backup application must accept commands from the Windows command interpreter. The scripts can also be used as a stand-alone batch job to start a backup. The backup scripts in this job used the launch-and-forget BC EVA method. BC EVA launched the backup script and finished the job without waiting for the backup to complete. Both tests used HP OpenView Data Protector and VERITAS NetBackup applications*.

***Note:** All VERITAS NetBackup testing was performed with Enterprise Volume Manager 2.0D, the previous version of Business Copy 2.1.

table 6. backup times

Tape Operation	Store Size	Backup Rate		Backup Time
HP OpenView Data Protector - online	43.6 GB	15 MB/sec	908 MB/min	48 min
VERITAS NetBackup V 3.4 - online	78 GB	21.5 MB/sec	1.3 GB/min	63 min
Launch DP scripts using BC EVA offline	46 GB	14 MB/sec	836 MB/min	55 min
W2K Drag/Drop - offline	41.6 GB	20.4 MB/sec	1.2 GB/min	34 min

results: The integration between BC EVA and the backup script worked as expected. W2K drag/drop, which was a copy to disk, was placed on the chart to give a comparison of a tape backup rate to a copy to EVA virtual disk (VRAID 1). The BC EVA job launches the **omniback.cmd** batch file, which can be found in the Rapid Recovery Script Kit. The batch job issued commands to the Data Protector from the Windows command interpreter on the Exchange recovery server. As can be seen in scripts, the batch job does a full backup to the BCVs on devices O and P. The batch job took 55 minutes to back up both volumes.

test 4—restore operations

This test shows the time differentiator between traditional tape and BCV restores. A traditional tape restore requires a tape device that restores a previous tape backup. The BCV is a copy of the original LUN created by the BC EVA snapclone. The Exchange store that was recovered came from the previous backup test. A single storage group, consisting of two databases and logs, was restored, while the remaining storage group stayed on line.

All tapes and BCVs were restored on the Exchange recovery server before they were presented to the Exchange host that was the target. Each restore from tape copied data to a Fibre Channel disk that was presented to the Exchange recovery server. Normally a restore of an Exchange store would be copied to the target Exchange host. The administrator would have to keep 50% of the storage group LUN free in case a restore was needed. This is a wasteful method of managing storage. With the HSV110 controller, storage capacity can be effectively managed to reduce wasted disk space. Additionally, working on a recovery server can reduce production server resources.

Many elements may affect restore time, such as the following:

- Service Level Agreement (SLA) requirements
- The administrator's knowledge of restore procedures
- The number of users logged in at the time of the restore
- The number of LUNs that are presented to the host Exchange Server that will be used for recovery
- The amount of data that is to be recovered
- Data integrity and consistency checks of the restored Exchange stores
- Hardware devices and backup applications
- Data storage capacity

Generally, a restore policy that has an SLA governs the required restore time. As archived data storage becomes larger, the task becomes difficult to achieve with the traditional tape restore method. This test did not have an SLA requirement.

Traditional tape restore methods use backup/restore applications that try to capitalize on getting maximum throughput from SCSI bandwidth. There are various methods of decreasing the backup/restore time, such as data compression or use of parallel tape devices. In this test, each backup application used the default compression ratio and also used a single tape device.

The tape restore used two backup applications, VERITAS NetBackup and HP OpenView Data Protector. Table 7 shows the results using each backup application.

table 7. tape restore time

Tape Restore to Exchange Recovery Server	Store Size	Restore Rate		Restore Time
VERITAS NetBackup V 3.4	78 GB	15 MB/sec	917 MB/min	85 min
HP Data Protector 5.0	43.6 GB	11 MB/sec	660 MB/min	66 min

Results: The results from table 7 are consistent with traditional tape restore times. The restore time reflects the average of multiple test runs of each backup application. Each backup application used standard out-of-the-box parameters and did not try to get the best performance from the tape operation. Additional time should be included when using the Exchange utilities ESEFILE and ESEUTIL, which should be used to check database integrity after restoring from tape. This step would add more time to the overall restore from tape.

Table 8 shows the point-in-time restore results of the BCV method. Two databases and logs were restored. The BCV method used selective presentation to move the BCV to the Exchange host target. The original storage group was dismounted by stopping the virtual Exchange server cluster attendant, and the volumes were then un-presented. The BCVs were presented to the Exchange server and then mounted. ESEFILE and ESEUTIL were done during the BC EVA backup job, so they did not need to be repeated. In addition, the logs were not rolled forward. Comparing both charts using the column for restore rate per minute shows the BCV method outperforming the tape restore by better than 10 to 1.

table 8. BCV restore time

Present BCVs to host Exchange Server	Store Size	Restore Rate		Restore Time
Solution BCV to host server restore	43.6 GB	181 MB/sec	11 GB/min	4 min

Note: Because a BCV represents the entire volume, not just the data, the restore time of a BCV stays the same even if the size of the stores is increased, whereas the tape restore time increases as the size of the store increases.

Note: Before the original Exchange stores are un-presented, the disk signature and disk ID must be recorded. The disk signature and disk ID on the incoming BCVs must be modified in order for the cluster disk resource to accept these new disks. A disk signature utility such as HP Writesignature or Microsoft Dumpcfg can be used to make the change. HP Writesignature utility V1.0 can be found in the Rapid Recovery script kit. Microsoft Dumpcfg comes in the W2K Resource kit.

test 5—hp OpenView automation manager

In this test, Automation Manager V2.0 (AM) was used to launch BC EVA jobs. The task at hand was to identify any issues with the integration between AM and BC EVA once a policy was created. AM policies were tested using a single occurrence and a daily scheduled occurrence. The same job that was used in the snapclone test was used for this test. For further information on AM go to the HP OpenView Automation Manager Solution Configuration section and visit: ftp://ftp.compaq.com/pub/products/sanworks/techdoc/managementappliance/AA-RQ62C-TE.PDF

results: There were no integration issues between AM and BC EVA. The AM policy worked seamlessly. There was no significant time added to the BC EVA job completion time (see table 10). Table 10 reflects the completion of AM and BC EVA, since BC EVA does not wait for normalizing to complete. That process was not included in table 10.

table 10. hp OpenView automation manager process time

BC EVA Job	Time (min)
Execute	12.6
Undo	4.1
Total time	**16.7**

table11. system and application performance results

	(1a) MCCS Baseline 1	(1b) MCCS Baseline 2	(2a) MSCE snapclone with 2 Virtual Servers	(2b) MSCS snapclone with 2 Virtual Servers
Users	3000	3000	3000	3000
Benchmark Profile	Medium	Medium	Medium	Medium
Protocol	Outlook MAPI	Outlook MAPI	Outlook MAPI	Outlook MAPI
Length of Test	8 hours	8 hours	8 hours	8 hours
Steady state that was measured	4 hours	4 hours	4 hours	4 hours
# of Log Disk	(24) VRaid 1	(24) VRaid 1	(24) VRaid 1	(24) VRaid 1
# of DB Disk	(24) VRaid 1	(24) VRaid 1	(24) VRaid 1	(24) VRaid 1
Response Time (milliseconds)	103	123	139	156
Average Local Delivery Time	0	0	0	0
Transaction Load (hourly)				
Messages Submitted	19152	19504	19219	18904
Message Recipients Delivered	70420	70855	69773	68617
Messages Sent	19254	19503	19220	18911
Transaction Load (per second)				
Message Opens/Sec	29	30	29	25.2
Folder Opens/Sec	12.5	12.7	12.5	12.8
RPC Read Bytes/Sec	67544	70308	68652	61668

RPC Write Bytes/Sec	543739	553632	539034	453217
Transaction Queues				
IS Send Queue Size Average Total	3.5	2.7	2.7	2.5
Processor Utilization				
System Processor Utilization (%)	44	55	54	53
System Processor Queue Length	2	3.9	3.5	3.2
System Context Switches/Sec	2811	3989	3887	3841
Process % CPU Time – Store	77	86	84	82
Memory Utilization				
Available Bytes	2.6GB	2.5GB	2.5GB	2.5GB
Pages/Sec	0.137	0.053	0.070	0.071
Process Working Set Bytes – Store	1GB	1.1GB	1.1GB	1.3GB
Process Virtual Bytes – Store	1.9GB	1.9GB	1.9GB	4.1GB
Logical Drive Utilization				
IS Database Disk Reads/Sec G:	188	197	186	184
IS Database Disk Writes/Sec G:	123	127	167	165
IS Database Average Disk Queue Length G:	2.6	2.2	3	3.1
IS Log Disk Writes/Sec I:	127	123	124	122
IS Log Average Disk Queue Length I:	0.16	0.159	0.173	0.195
IO Reads Bytes/Sec – Store	2.9MB	3.3MB	3.1MB	3MB
IO Writes Bytes/Sec – Store	3.9MB	4.4MB	4.2MB	4MB
Sampling taken during the normalizing of BCVs	**(1a) MCCS Baseline 1**	**(1b) MCCS Baseline 2**	**(2a) MSCE snapclone with 2 Virtual Servers**	**(2b) MSCS snapclone with 2 Virtual Servers**
IS Database Disk Reads/Sec G:	N/A	N/A	150	150
IS Database Writes/Sec G:	N/A	N/A	145	143
IS Database Avg. Disk Queue Length G:	N/A	N/A	3.7	3.3
IS Log Disk Writes/Sec I:	N/A	N/A	103	105
IS Log Average Disk Queue Length I:	N/A	N/A	0.163	0.175
IO Reads Bytes/Sec – Store	N/A	N/A	2.5MB	2.4MB
IO Writes Bytes/Sec – Store	N/A	N/A	4MB	3.7MB

solution setup and installation instructions

Caution: Many of the commands and procedures described in these instructions require Exchange administrator privileges. If used incorrectly, they could cause unwanted disruptions to Exchange operation. This document assumes you are familiar with Windows 2000 server clustering and domain network concepts. Although this document mentions Microsoft Active Directory and Microsoft Domain Controller, discussion of these products is limited. All Microsoft and HP products have online documentation, to which this document refers as appropriate.

As with any changes to your Exchange 2000 configuration or backup plan, be sure to test these changes thoroughly before implementing them in a production environment.

Follow these steps to set up and install this solution:

I. Set up the storage area network (SAN)

1. Be sure the SAN is set up with the following components. Install them in the order listed.
    a. HP StorageWorks Enterprise Virtual Array 5000 and HP StorageWorks VCS 2.02 or higher with BC EVA (Snapshot) licenses on HSV110 controllers
    b. HP ProLiant servers
    c. HP Storage Management Appliance II
    d. Network infrastructure
    e. HP OpenView Storage Management Appliance software V2.0A or higher on HP Storage Management Appliance II

       Download this upgrade from the following site:

       http://h18000.www1.hp.com/products/sanworks/softwaredrivers/managementappliance/index.html
    f. HP HSV Element Manager V2.0A on SAN HP Storage Management Appliance II
    g. HP StorageWorks Business Copy V2.1 Server
    h. HP StorageWorks Secure Path Manager V4.0A (Installed on Appliance)
    i. HP StorageWorks Secure Path V4.0A agent
2. Configure Quorum, EVA disk groups, and EVA virtual disks. Use separate virtual disks for logs and database files.

II. Set up the Exchange 2000 servers

1. Install Microsoft Windows 2000 Advanced Server with Service Pack 3 (SP3).
2. Install FC2101 HBAs.
3. Install the Windows 2000 platform kit for HP StorageWorks Enterprise Virtual Array. The CD-ROM includes Fibre Channel drivers, Storage System Scripting Utility (SSSU), and documentation.
4. Set up three Ethernet networks, as follows:
    a. One network for the heartbeat between the two cluster nodes. Configure the heartbeat as a private network for cluster communication only.
    b. One network for client access to the application servers. Configure it as a public network.
    c. One network for the HP OpenView Management Appliance and its applications. Set it up as a mixed network. Its purpose is to separate from the users' network any traffic generated by the Management Appliance and any host that has a BC EVA agent and needs to be connected in the same network.

5. Install cluster service on both servers, ensuring that both servers are identical in every aspect of the installation.
6. Install HP StorageWorks Business Copy 2.1 Host Agent for Windows 2000.

III. Configure Secure Path

1. Configure the Secure Path logon and client table on each host. (Choose **Start > Programs > Secure Path > SecurePathCfg**.)
2. Set a common password that the Secure Path manager will use to poll each client machine.
3. In the Clients Configuration Table, enter data as follows:
    a. If the client is a member of the cluster, enter the name of the HP OpenView Storage Management Appliance; enter the cluster node for each Secure Path client followed by the cluster name, as follows: **node_name:cluster_name**.
    b. If the client is not a member of the cluster, enter the names of the host and the HP OpenView Storage Management Appliance.
4. From the HP OpenView 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 Configuration Table, and save the profile.

IV. Connect SAN components

1. Connect the server hosts and DNS server to the SAN switch. For each server host, use a dual-path standard Fibre Channel configuration.
2. Add host server connections to each host's folder on the HP OpenView HSV Element Manager.
3. Using HP OpenView Element Manager, present the virtual disks (LUNs) to both hosts for the cluster quorum disk and Exchange 2000 stores.
4. Be sure to perform full formats on the newly presented LUNs.

    Note: For best results, initially present the virtual disks to only one host and use that host to format the LUNs. Once the LUNs 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.

V. Set up the Exchange 2000 cluster server

Note: This is only a *summary* of how to set up Exchange cluster service. For the full installation process, use the Microsoft installation guides, available from the following site:

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

1. Verify that Network News Transfer Protocol (NNTP) is installed before attempting to install Exchange 2000 Server. By default, Internet Information Services (IIS) is installed with Windows 2000 Server; however, NNTP is not. NNTP is required for Exchange 2000 Server installation, but is not supported in the cluster.
2. Verify that Simple Mail Transfer Protocol (SMTP) is installed.
3. Make sure that you have a pool of static IP addresses available to you when you create the Exchange virtual servers.
4. Make sure that the cluster service account is a member of the Built-in/Administrators group on each node and that it is granted "Exchange Full Administrator" privileges in the Exchange organization.

5. If you are installing Exchange 2000 Server for the first time in your organization, run ForestPrep and DomainPrep.
6. Make sure that all shared disks, including the quorum disk, are physically attached to a shared bus. Verify that disks attached to the shared bus can be seen from all nodes. Do this at the host adapter setup level. Refer to the HP FC2101 documentation for adapter-specific instructions.
7. Be sure that all shared disks are configured as basic (not dynamic) disks, and that all partitions on the disks are formatted as NTFS with a single partition.
8. Install Microsoft Cluster Server (MSCS). You must install the same version of Exchange 2000 Enterprise Server on all nodes in the cluster. Choose **Start > Control Panel > Add/Remove Programs > Add/Remove Windows Components** on the first node. Repeat the process on the second node but not at the same time. For consistency, install Exchange 2000 Server on the same drive and in the same folder on each computer and shared storage device in the cluster. The default installation folder (for program files, which are not shared) is the local system drive. Follow Microsoft's guidelines in *A Step-by-Step Guide to Installing Cluster Services*:

   http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/exchange/exchange2000/deploy/depopt/sp2clustgd/1-intro.asp

Note: Make sure the latest service packs for Exchange 2000 Server are installed. Consult the Microsoft documentation for the latest service pack releases before installing them on production computers. In our tests, SP3 for Windows 2000 Advanced Server and SP3 for Exchange 2000 were the highest levels installed.

VI. Install Exchange Virtual Server

To create an Exchange 2000 Server cluster, create a Windows 2000 cluster group and then add specific Exchange resources to it. Exchange 2000 Server clusters are referred to as Exchange Virtual Servers. Unlike the physical computer running Exchange 2000 Server, an Exchange Virtual Server is a cluster group that can be failed over if the server itself fails. When the physical computer fails, the remaining nodes in the cluster take over the failed Exchange Virtual Server, and clients can access this server using the same server name.

If you're not familiar with clustering, use the cluster wizard to help install the first Exchange cluster group. (See the Microsoft web site's installation procedures.) The following steps are only a *summary* of how to create an Exchange cluster group manually:

1. Create an Exchange virtual server.
2. Create an IP address resource.
3. Create a network name resource.
4. Add two disk resources to the Exchange cluster group. Label disk resources the same as the volume label and drive letter that are presented in Windows Disk Manager. Be aware that the system attendant has a dependency on the disk resources. If the disk resources go down, the system attendant will stop, which will shut down the exchange resources. The cluster service will try to fault over to the other node. The following section, working with BC EVA, further describes the relationship of BC EVA and the disk resource.
5. Create an Exchange 2000 system attendant resource. When the system attendant is brought on line for the first time, it creates the rest of the Exchange services shown in figure 7. All resources have to come on line before the virtual server is operational. The Exchange resources come on line in a staggered process. Before a restore can take place, the system attendant must be brought down
6. Repeat steps 1-5 to create a second Exchange virtual server.

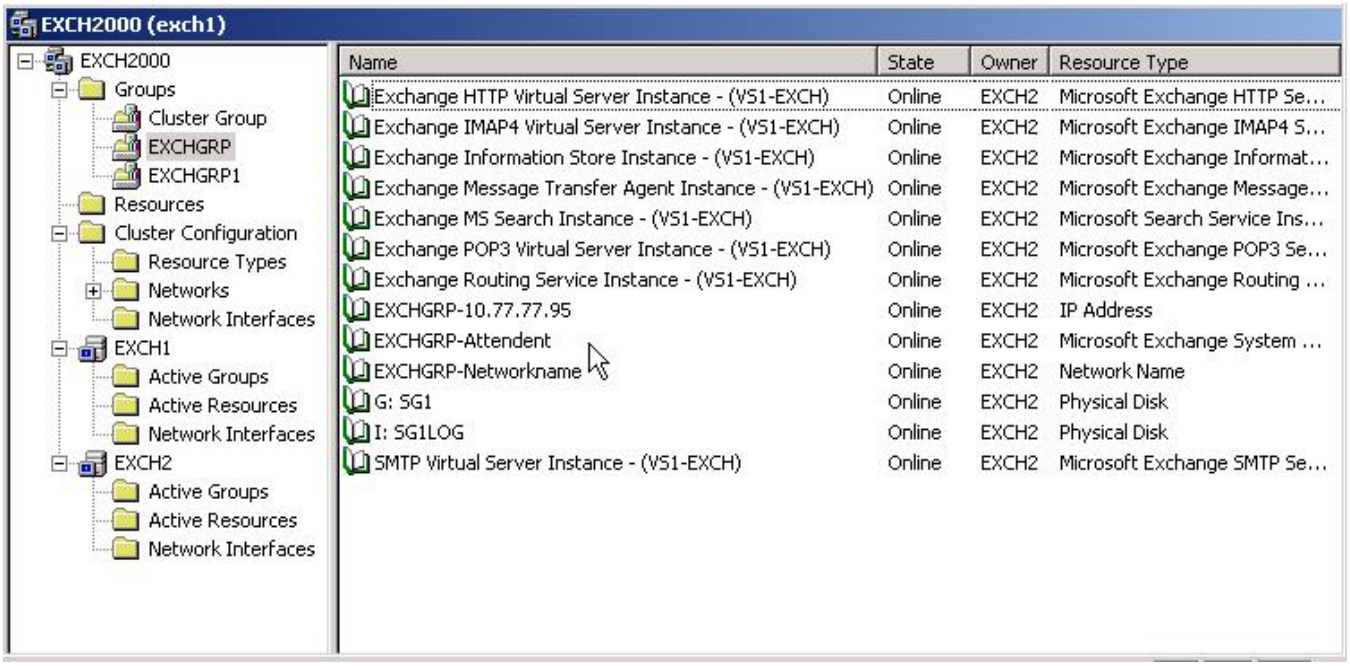

figure 7. cluster administrator

There are many ways to configure Exchange virtual servers. In order to take advantage of BC EVA technology and create as little disruption as possible to the Exchange production environment, special consideration was given to the design of the Exchange cluster groups. Two virtual servers were created to separate the two Exchange disk groups. Therefore, only one virtual server is down at any time during a BCV creation.

This design also takes into account the way the Exchange system attendant reacts to cluster disk resources. If any disk resources in a cluster group are taken away (fail or are forced offline), the system attendant shuts down. In addition, the cluster service tries to cut over to the other node. The reason for this process is the dependency the system attendant has on the disk resources in the cluster group (virtual Exchange server). There is no way to bring down a disk resource without diminishing the cluster services of that disk group and affecting Exchange services for the virtual server.

VII. Configure BC EVA

Important: This document does not cover every aspect of BC EVA. There are many online BC EVA manuals that have extensive custom configuration information. Before configuring BC EVA, consult the installation and configuration documentation for BC EVA 2.1 manager and all release notes, available from the site below. Use that information in conjunction with the information in this document.

http://h18004.www1.hp.com/products/sanworks/BC EVA/documentation.html

VIII. Install BC EVA

1. Install the BC EVA 2.1 manager on the HP OpenView Storage Management Appliance 2.0 or higher. If using clustering choose the cluster option for all cluster nodes. The BC EVA 2.1 host agent for Windows 2000/NT is not fully cluster-aware but does support MSCS (Microsoft Cluster Service) on Windows 2000 Advanced Server and Datacenter. Install BC EVA 2.1 host agents on the cluster servers and the backup server.

2. Before the HP OpenView Storage Management Appliance Software can be launched the Java (JRE) executable that comes with the BC EVA CD should be installed. The Storage Management Appliance software can be launched from Microsoft Internet Explorer version 5.5 and 6.0.26 browsers. This can be done remotely or at the Storage Management Appliance. The current JRE version used in this document is version 1.3.0.1.02. Install the JRE executable on each system that will access the Storage Management Appliance software.
3. After logging in to the Storage Management Appliance select **> Tools > Business Copy > Resources**. The Resource screen will show EVA subsystems and configured disks. The Resource screen will show any BC EVA clients such as cluster severs, virtual server and backup server. The subsystem and any configured disk will also appear. See figure 8.

Troubleshooting Note: During the initial launching of the BCV EVA the Resource window may hang, BC EVA may not be able to log on to the HSV subsystem. To verify that BC EVA is using the correct logon, open the **EVM.cfg** file in the BC EVA Bin folder and verify the correct StorageID and StoragePassword. If either is incorrect, make the appropriate changes and restart the Switchboard service on the Management Appliance. The ID and password should match those of the Management Appliance logon.

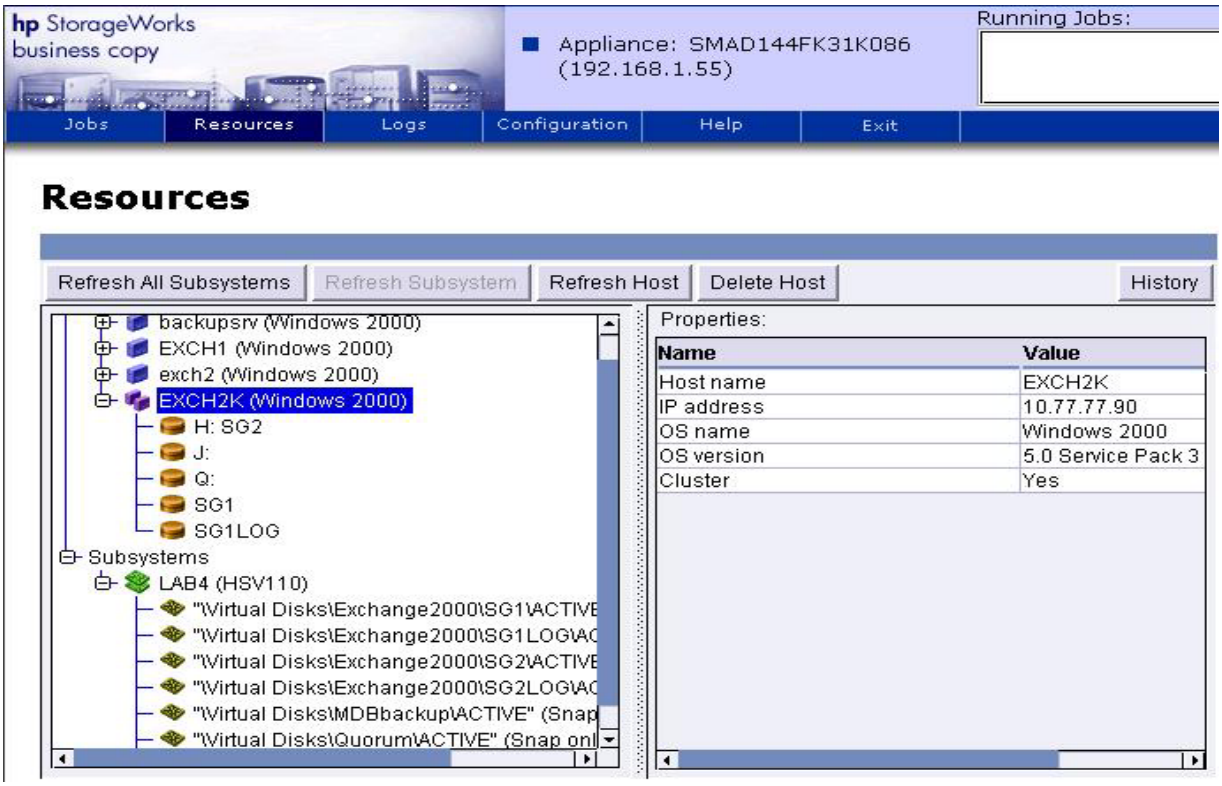

figure 8. Business Copy Resources

Note: The EXCH2k cluster virtual server in figure 8 has various drive letter naming conventions. This was purposely done in order to show how the drives would show up in the BC EVA Resource view. If the administrator takes the cluster disk resource defaults (such as Disk J: or Disk Q:) the drives will appear with J: or Q: example in figure 8. In order to get the H: SG2 example in figure 8, disk resources would need to be modified by adding the "SG2" after the drive letter. The example SG1 in figure 8 would need to be changed

to "G: SG1" as the disk resource name. The quotes are not allowed.

Troubleshooting Note: Here are few troubleshooting techniques that can help if any of the BC clients depicted in figure 8 where grayed out in the resource view.

First try using the **refresh host** hot button on any host that is grayed out. Allow for the switchboard service to refresh. This refresh process can take minutes depending on the other activities that the Storage Management Appliance (SMA) is doing. The next step would be to stop/start the BC EVA switchboard services on the SMA. If that doesn't work go check the ***.hst** and ***.vol** file on the SMA. Initially, after the BC EVA Switchboard service is restarted (on the BC host client machines), the BC EVA manager acquires the information from the BC EVA host client machines and places a temporary ***.data** file in the host directory that resides on the SMA. Once all the data from the hosts has been collected, the **.data** files go away and the HST and VOL files get updated.

Check the modified date on HST and VOL files using explorer to determine the last time the files were touched. This would give an indication if the BC manager is communicating to the BC host clients. The files can be found by selecting explorer **Program Files > Compaq > SANworks > Enterprise Volume Manager Copy > bin > hosts** on the Storage Management Appliance. If the HST and VOL files are not in the host directory there may be a connectivity issue. The next step would be to ping the clients from the Storage Management Appliance. If pinging gets a good return then reinstall the BC EVA host client that is having the problem. For more troubleshooting ideas refer to the BC EVA administrator guide.

working with BC EVA

This section describes how to set up and use BC EVA. It covers the following topics:

- pre-configuring BC EVA
- job performance
- required disk space
- job design rules
- scripts and batch files
- BC EVA switchboard
- creating snapclones with BC EVA
- cleaning up BC EVA jobs
- automating BC EVA

pre-configuring BC EVA

BC EVA uses a job process to create replicas of volumes (i.e., snapclones) in Windows 2000 using the VCS firmware within the HSV controllers. Before starting the job, you need to know some pre-configuration rules of the snapclone method. With HSV storage systems, the following rules apply:

- The VCS snapshot software license equals or exceeds the capacity of the storage system.
- A snapclone is created in the same disk group as the source, but within a different disk family.
- A snapclone of a virtual disk that is a snapshot and/or has a snapshot is **not allowed**.
- No snapclone of a virtual disk is created during the unsharing process with another snapclone. Unsharing is an EVA background process in which data from the initial snapshot phase of the snapclone is copied completely to a virtual disk.

- No snapclone of a snapclone can be created during the unsharing process.
- No snapclone can be created of a virtual disk that is in the process of being deleted.

job performance

BC EVA simultaneously supports up to 25 HSG-based and 16 HSV-based storage systems. However, as the number increases, BC EVA job performance and user interface responsiveness can degrade. Actual performance depends on several factors, including the number of simultaneous jobs, the number of steps and complexity of the jobs, and I/O activity on the storage systems and hosts. Scripts and batch files launched from the BC EVA job add complexity and increase job completion time. See test 3—backing up to tape.

required disk space

A snapclone requires, at a minimum, the same amount of disk capacity as the source disk. BC EVA checks each storage system referenced by the job to determine if there is sufficient disk space.

The HSV110 controller handles all disk assignments and informs BC EVA of whether or not the required space is available.

job design rules

- Do not include more than eight storage systems in a single job
- BC EVA supports up to eight simultaneous jobs
- Do not mix HSG and HSV virtual disks within the same job
- The amount of storage space being replicated in a job cannot exceed the free storage space on the associated storage system
- When creating a job, choose either unit or volume syntax. If both unit and volume syntax are used within the same job, the job will fail.

scripts and batch files

Scripts and batch files can be used to launch various commands to BC EVA host clients. For information about designing scripts, see custom scripts.

BC EVA switchboard

The BC EVA switchboard service periodically monitors and updates BC EVA manager and host client resources, including any changes to host volumes or network connections. The switchboard monitors the following:

- Switchboard startup and configuration
- Device information for connected BC EVA host clients and manager
- Communication between BC EVA host clients and manager

The Switchboard service resides on the BC EVA server and clients. Since the switchboard is a service, it can be stopped or started without rebooting the host server or client.

The switchboard also generates logs of BC EVA activity. The logs are an important tool in troubleshooting BC EVA issues. Note that you will need these logs when working with BC EVA HP Engineering support.

Logs can be viewed from the BC EVA GUI or can be found in the Program Files directory, in **Program Files > Compaq > SANworks > Enterprise Volume Manager Copy > bin > Logs**.

creating snapclones with BC EVA

The following procedure explains how to create snapclones. We recommend that you do not deviate from the sequence of commands as given in steps 4–9. The command files mentioned in this procedure are discussed in scripts.

1. Launch BC EVA from **Tools** in HP OpenView Manager.
2. Select **Create BC EVA Job** in order to make a script.
3. Create a new job name, owner and category in the **New Job** window. Be sure to include comments in your scripts to let others know the purpose of the commands. The BC EVA commands are listed on the left side of the New Job window. The most commonly used commands are SUSPEND, SNAP, RESUME, MOUNT, and UNDO. Rapid Recovery also uses the LAUNCH and DELAY commands. (See figure 9, which lists the BC EVA job that was used for each test.)
4. **SUSPEND**—This should be the first command in the job. **SUSPEND** calls the dismountexch.cmd file that dismounts the Exchange stores and places the virtual Exchange server in a consistent state.
5. **SNAP**—This is the command used to replicate volumes on an EVA. **SNAP** includes options for the type of snapshot. Use the **Snapclone** option.

   This process has two steps: (1) Creation of a virtually instantaneous snapshot, to which you have immediate access, followed by (2) Unsharing, an automatic morphing process that creates a clone (replicated image) of the stores from the snapshot step. The snapshot is deleted once the clone is available.
6. **RESUME**—Once the **SNAP** process has completed, use the **RESUME** command, which calls the mountexch.cmd file to mount (i.e., bring online) the relevant virtual Exchange server.
7. **MOUNT**—This command is optional. It invokes a process that presents the newly created snapclone to a Windows 2000 host. If you use **MOUNT**, present the snapclone to the Exchange recovery server. You must select volume letters, which should be the same as the original volume letters if possible. Once the snapclones are settled on the backup server, back them up to tape or copy them to another hard drive.
8. **DELAY**—This command creates a delay between command processes. Rapid Recovery uses a 60-second delay to allow the Exchange recovery server to settle down the snapclones before proceeding to the next process. This is because BC EVA communicates to the recovery server during the mount process. During that time Windows has to accept or reject the new volumes from BC EVA and assign drive letters twice before completing. The second 60-second delay allows the ESEFILE and ESEUTIL commands to execute before backing up the snapshots.
9. **LAUNCH**—This command is used to run command-line scripts on a Windows host. The scripts can be any command files, such as scripts or batch files. Rapid Recovery uses the **LAUNCH** command to start a backup application to tape. You can use any backup application, as long as it can be launched by the Windows command-line interpreter. Our tests used HP OpenView Data Protector and VERITAS NetBackup. See Scripts 5 and 6 in custom scripts for sample batch files.

Rapid Restore also uses the **LAUNCH** command to execute Exchange ESEFILE and ESEUTIL database utility tools. The Exchange utilities can be launched right after the Exchange recovery server acknowledges the snapclones. The results can be redirected to a file. See Scripts 1, 2, and 3 in custom scripts for sample batch files.

Note: The BC EVA job can be monitored during it's run with the use of the monitor hot button. The progress screen will update each BC EVA command during the run phase allowing the user view the status of the job.

breakdown of the BC EVA job

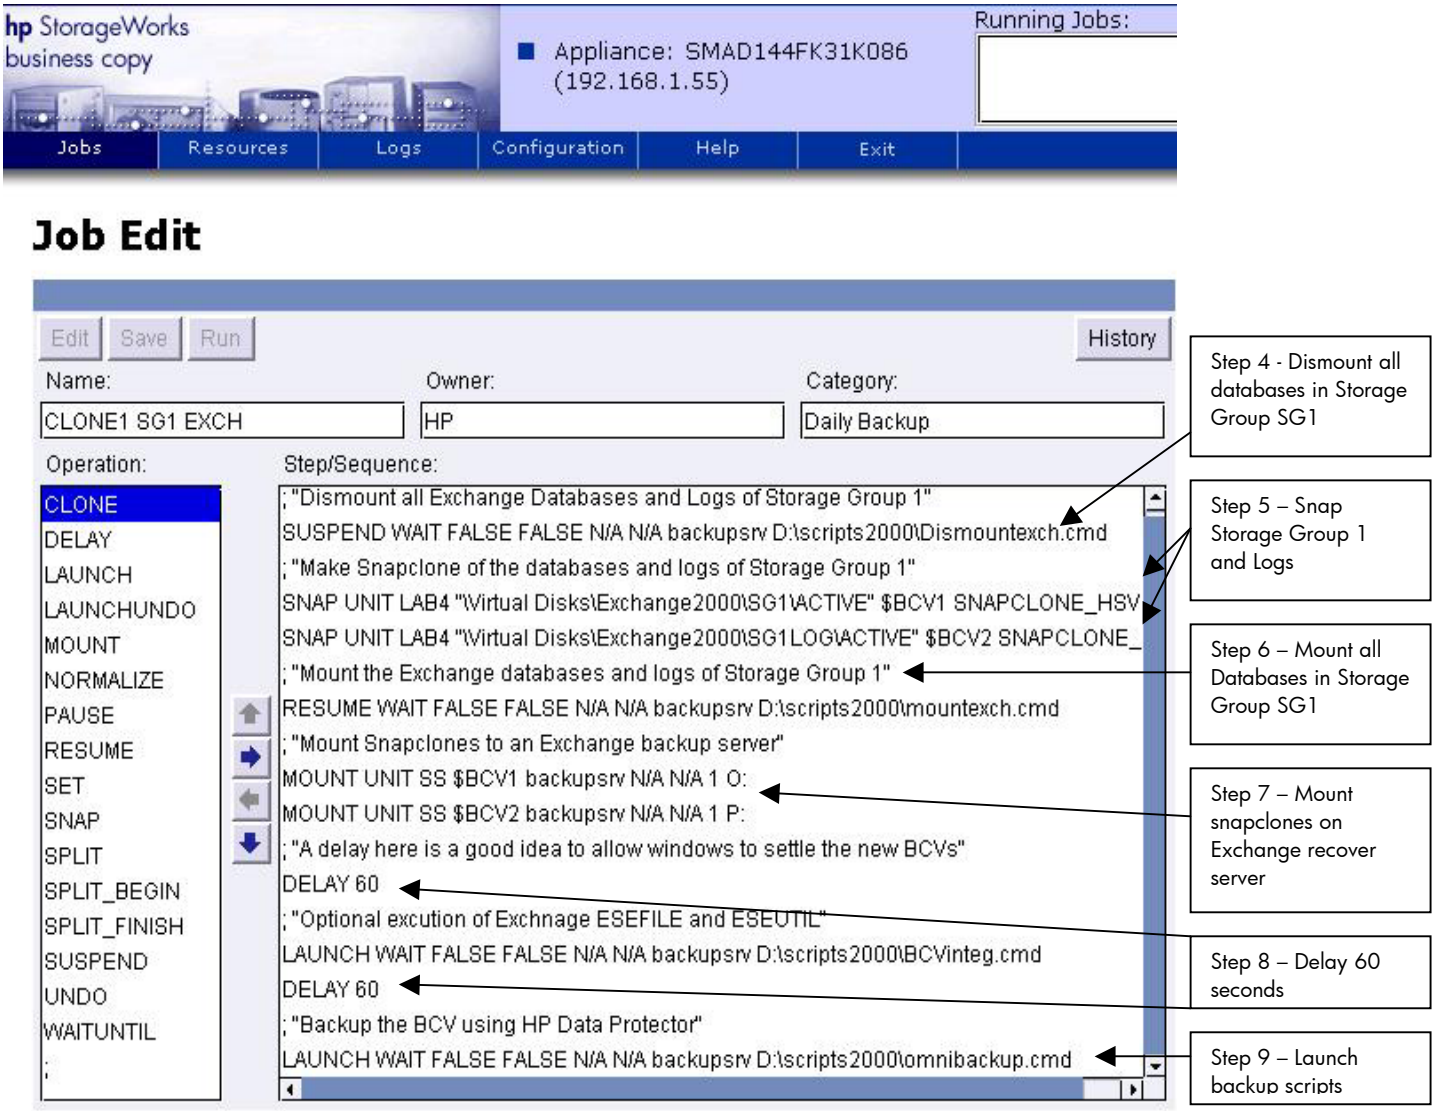

figure 9. sample BC EVA script

- BC EVA runs numerous background command processes during a job. The cumulative background time for the BC EVA job in our tests was 93 seconds (see table 5). This time included waiting for VCS commands to complete and launching scripts and BC EVA processes. Keep in mind that this was an average time. The time will vary depending on the loads of the SAN Appliance, network, clients, HSV110 controller, and BC EVA, and whether Automation Manager is used.
- The snapshot process in our tests took from 5 to 9 seconds to complete, after launching the dismount script for the Exchange store. There were two snapclones, one each for the databases and the logs. The snapclone for the databases, which consisted of EDB, STM, and public files, was 43.8 GB, and the one for the logs was 14 GB, for an 8.3-GB-per-second copy rate.

  Note: The data capacity has very little effect on the snapshot process time. The total HSV110 controller load at the time of creation of the snapclone has greater impact on the snapshot portion of the snapclone.
- Custom scripts launched by BC EVA are used to dismount and mount Exchange databases. The overall time for the dismount/mount of the databases to complete is 40 seconds, as can be seen in the Event Viewer log

shown in figure 10. The time the store is down depends on numerous conditions, such as the number of users logged in to the store and the current load on the Exchange server. In order to shut down a store in Exchange Server, all the databases in a storage group must be brought down. The ESE instance for that storage group shuts down automatically once all databases are dismounted. In our test, 1500 users were logged out of the SG1 storage group.

- The BC EVA mount process that assigns the snapclone to a host (Exchange recovery server) takes 2 minutes to complete (see table 3), regardless of the size of the snapclone. Once presented, the snapclone looks like any other volume (LUN) in a Windows 2000 server. At this stage, the snapclone is still a snapshot and can be backed up to tape or accessed for other uses. The normalization phase can continue while the backup is taking place.
- Unsharing (normalization) speed is about 18.3 MB per second, thus requiring 1 hour and 20 minutes for a 100 GB volume to complete (see table 3). These times will vary based on the size of the virtual disk, VRAID design, current workload on the disk, and disk rebuild policies. When two or more snapclones are being created, the smaller snapclone will complete faster. Note that the clone process (block-to-block copy) copies the entire disk capacity whether the disk is full or not. The BC EVA job will complete before normalization is done.

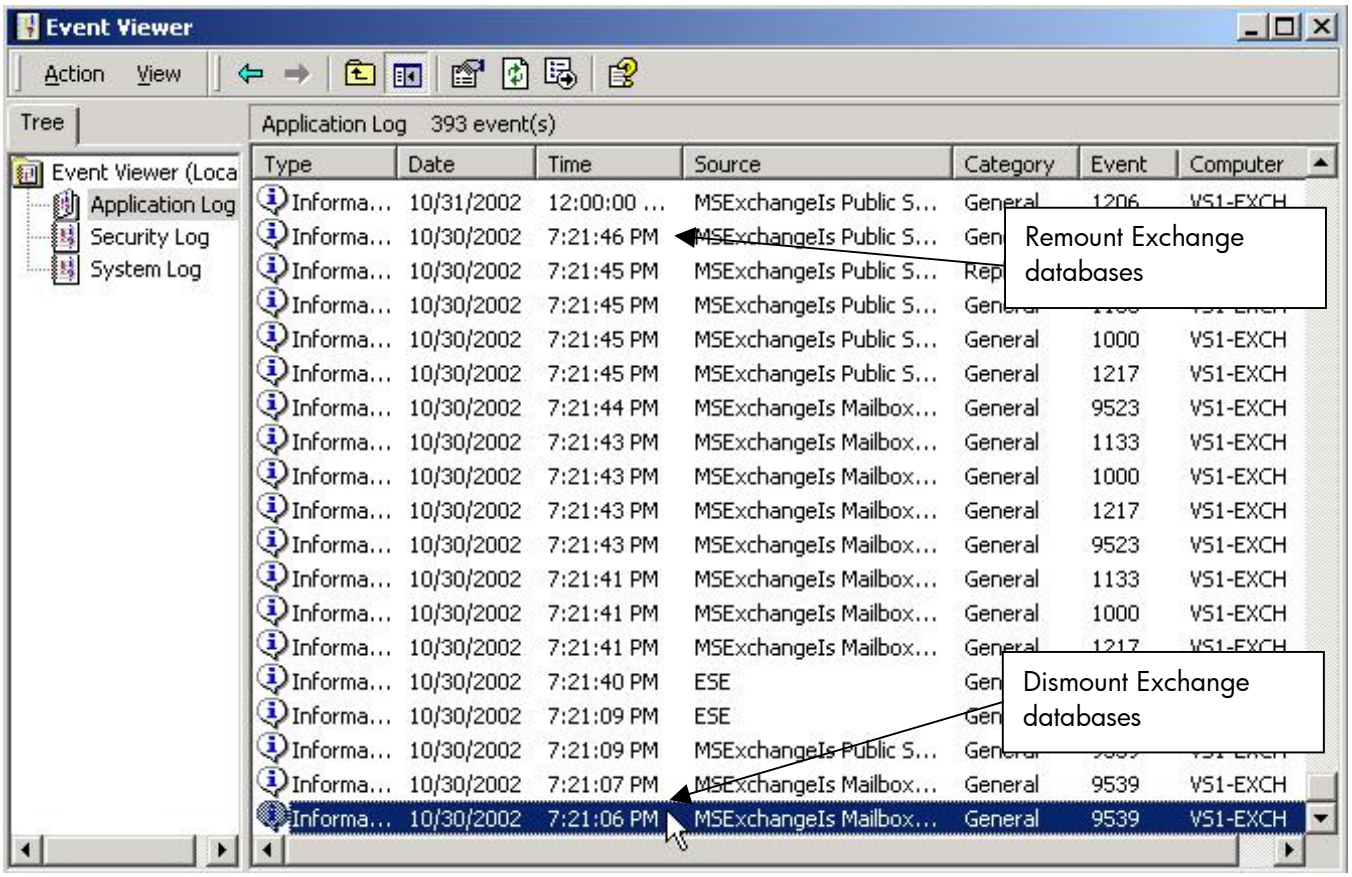

figure 10. Event Viewer

Note: EVA hardware failure during the period when the snapclone is in the snapshot or unsharing mode can cause corruption and loss of the snapclone.

During the snapshot phase, the metadata from the host disk has been replicated and is used to maintain the snapshot. If the host disk experiences a hardware failure, the snapclone could be destroyed.

During the unsharing phase, if a hardware failure occurs on the host volume, the snapclone can become corrupt or lost.

Keep in mind that a disk failure protection level (single or doubled) at the disk group HSV110 controller level should have been selected. With disk failure protection and VRAID1 or VRAID5 data on, an HSV110 controller has a greater chance of surviving disk hardware failures.

Best Practice: BC EVA 2.1 Host Agent Release Notes recommend logging out of the backup host when running BC EVA scripts that automatically mount volumes to the backup server. This will avoid acknowledgement of the pop-up window hardware notification on the backup server. During the presentation of snapclones to the backup server, Windows 2000 has to negotiate with the BC EVA mount drive letter request. Before settling, Windows twice assigns the snapclone a drive letter. If a user is logged on and the user acknowledges the hardware notification, Windows 2000 requests a reboot. If the system is rebooted, the drive letters could change. This change would affect BC EVA jobs and scripts launched by BC EVA. To avoid this inconsistency, you should be logged off of the backup host. This way, Windows ignores the hardware change, so a reboot is not required. If this step cannot be avoided, change the drive letters on the backup server back to their original letters, and then rerun the Undo command to avoid failure of the BC EVA Undo command.

cleaning up BC EVA jobs

It's important to recycle BC EVA jobs to keep the job list as clean as possible. This will make it easier to manage job schedules and perform restore operations in the fastest amount of time. To minimize the number of BC EVA jobs in the job list, 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.

To undo a BC EVA job, do the following:

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

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**. This will cause you to lose the database. For more information, see restoring Exchange stores using snapclones.

automating BC EVA

The BC EVA scripts can be automated to a set schedule with little user intervention using these tools:

- HP OpenView Storage Management Appliance Software V2.0 SP1
- HP OpenView Automation Manager (included in Appliance Software)
- Windows Task Scheduler

hp OpenView Automation Manager solution configuration

HP OpenView Automation Manager (AM) lets a storage administrator automate the management of a SAN. AM runs, controls, and manages predefined polices that you can configure for your environment. AM runs on the Management Appliance II and can call BC EVA jobs (scripts). Once you create a policy, you can schedule the same BC EVA scripts to run as needed: once, daily, weekly, biweekly, and so on. AM has predefined event and action

Perl scripts that are integrated with BC EVA. You can also create and import Perl scripts to create policies specifically for the HP SAN environment. AM can also generate reports.

The following procedure describes how to configure Automation Manager to integrate with Business Copy 2.1 (BC EVA) and schedule it to run a BC EVA job.

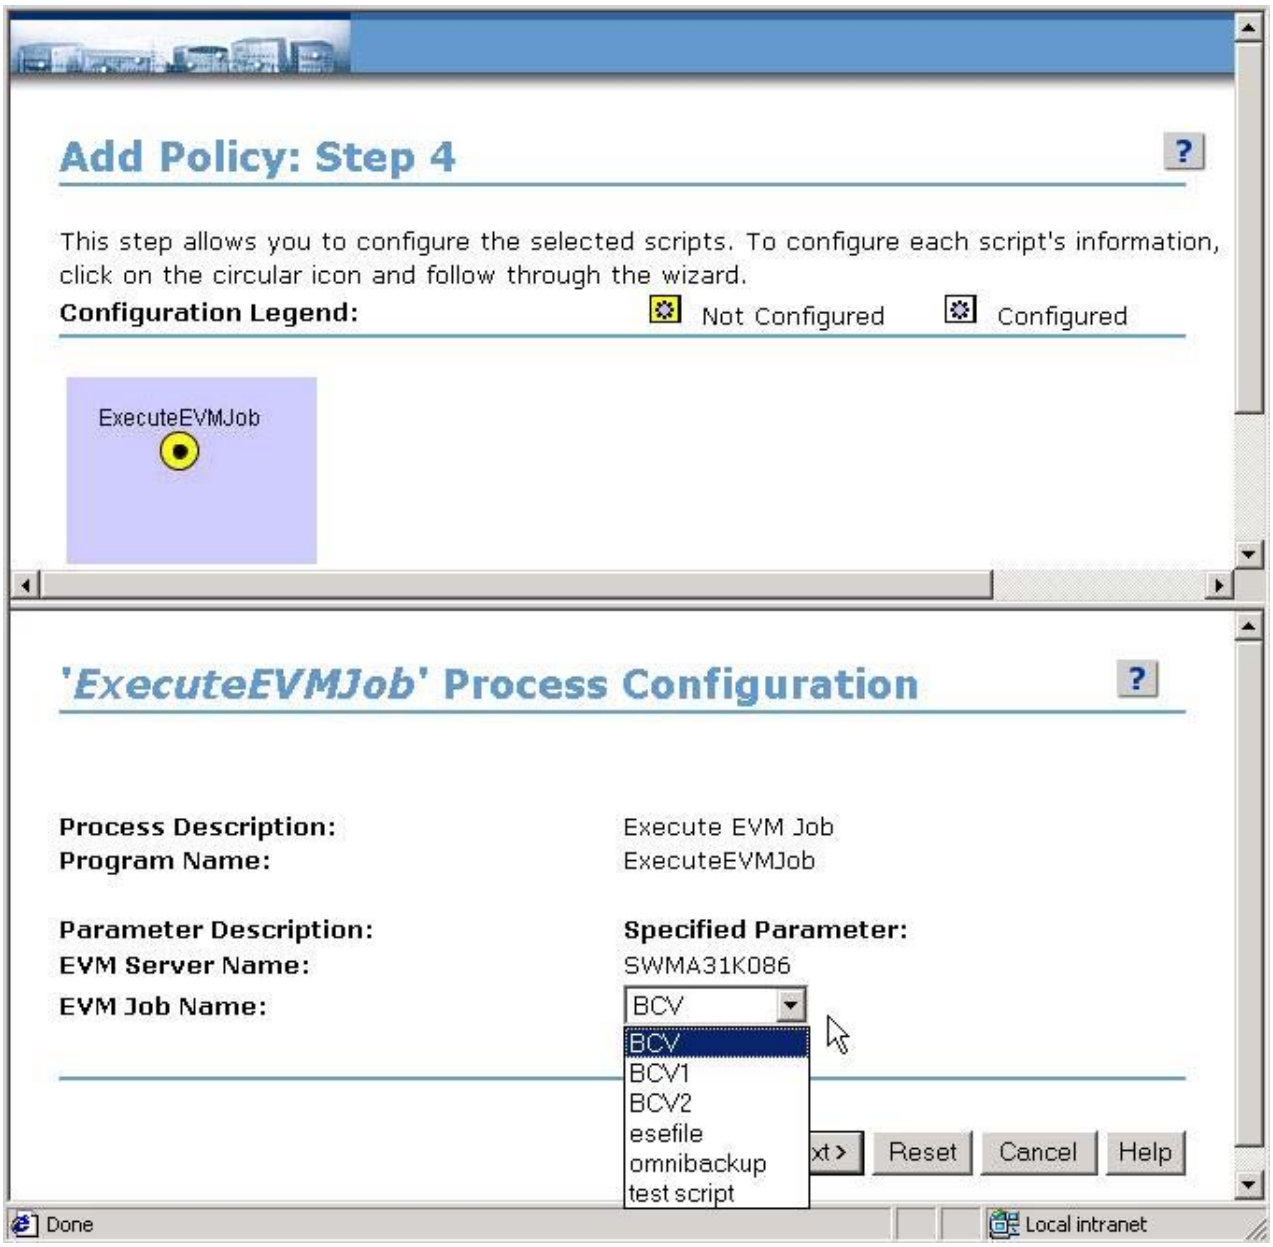

figure 11. hp OpenView Automation Manager

AM has two built-in EVM scripts that can check status and launch BC EVA jobs. Follow these steps to set up a policy in AM:

1. Create a policy in AM by selecting Add Policy. The policy wizard lets you quickly create a scheduled task.
2. To launch a BC EVA job, select and add the predefined Action Script **ExecuteEVMjob**.

3. In the wizard's Step 3, you can organize multiple scripts to have dependency links to each other. However, Rapid Recovery does not use this option.
4. In the wizard's Step 4 (see figure 11), double-click Execute EVM Job. This launches the BC EVA job screen. A pull-down menu shows all BC EVA jobs that you can include in the AM policy.
5. Choose a scheduling option: single run, recurrence, or customized pre-defined calendar.
6. Save and launch the policy.

Each time the **ExecuteBC EVAjob** Perl script launches, it checks to see if the BC EVA job has been executed. If so, it issues an Undo command. Once the Undo command is completed, AM launches a BC EVA Execute command. The time needed to execute varies, depending on the number and size of BCVs that must be undone, but generally it is only a few minutes.

The combination of BC EVA and AM's **ExecuteEVMjob** Perl script have a huge competitive advantage over other scheduling applications in regard to rolling clone (BCV) processes. Rolling clones is an automated technique that is used to keep recycling older clones as new ones are created. In other words, the process keeps a current backup of host data on disk for fast recovery. As an example, you could set up an AM policy that makes two or more clones. Each clone would be active for two days and they would be refreshed every other day. For more information, see job performance.

Windows Task Scheduler

Microsoft Windows 2000 Task Scheduler can also be used to automate the backup process. Choose **Start > Accessories > System Tools** to locate the Task Scheduler. Rapid Restore uses Task Scheduler to start the batch files used to automate BC EVA jobs. BC EVA 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:

BC EVAcl.exe <san_appliance_ip_address> execute or undo <job_name>

Note: The command line syntax for the IP address can be replaced with the SAN Appliance name.

A sample batch file to start a previously created BC EVA job could look like this:

```
REM ** Run a BC EVA job ***
@ echo off
cd \Program Files\Compaq\SANworks\Enterprise Volume Manager\bin
evmcl 10.77.77.70 execute <BC EVA job name>
exit
```

BC EVA backup process

backing up to tape

Tape backup remains an integral part of the data protection process. HP OpenView Storage Data Protector V5.0 and VERITAS NetBackup V3.4 or V4.5 are used as tape backup applications in Rapid Restore. Configure each application with Exchange Server 2000 Agents running on the Exchange virtual server and the backup host running on AM.

backing up with Data Protector

Note: When operating Data Protector, be sure to check for the latest patches for proper operation. You can find software patches at the following website:

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

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 Data Protector Device Manager.

Configure Data Protector only as a tape backup application. Install Data Protector Cell Manager on the backup host, and install the Exchange Server agents on the Exchange virtual servers.

BC EVA job uses its Launch command to integrate with Data Protector. The Launch command initiates a batch file from a directory on the backup server. Using the Data Protector command line, create a Data Protector backup specification to call the backup job name. Once the replicated volumes are mounted, the batch file starts the backup specification and runs 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 BC EVA job Launch command and starts a predefined Data Protector backup specification. For additional script information, see custom scripts.

```
REM *** This file will issue commands to HP OpenView Storage Data Protector and run a ***
REM *** full backup of the below Volumes SG1, SG1LOG ***
****You need to make the scripts generic and call out the optional parameters with <…>

omnib –winfs backupsrv:O: <storage group> –device Drive1 –mode full –protect none –pool <???>

omnib –winfs backupsrv:P: SG1LOG –device Drive1 –mode full –protect none –pool LAB4
```

Note: You can add a debugger command to help initially troubleshoot any problems associated with this batch job. The debugger redirects output to the log directory in the Omniback parent directory. Be aware that the debugger switch increases the backup time when it is turned on.

For additional information on how to use HP Data Protector 5.0 to back up Exchange servers, see the following internal HP Knowledge Brief:

http://dbospseu01.emea.cpqcorp.net/Technology/Documents/Knowledge%20Briefs/Q1FY03/KB0535%20Using%20HP%20Data%20Protector%205.0%20to%20backup%20Exchange%202000%20servers.doc

This is part of a series that includes the following, available from the same site:

- KB0534: HP Data Protector 5.0 Quick-Start Guide for Windows Backup
- KB0535: HP Data Protector 5.0 and Exchange 2000 Backup Basics
- HP Data Protector 5.0 and Windows 2000 Server Recovery

The third Knowledge Brief in the list above explains how to perform a Windows 2000 server recovery using replacement hardware with less downtime than otherwise experienced using traditional system state restore methods. A more comprehensive white paper on Exchange Server disaster recovery will be published on MSTech and HP Active Answers.

backing up with VERITAS NetBackup

For information on BC EVA and VERITAS NetBackup V3.4 integration, see the following document:

Enterprise Backup Solution with VERITAS NetBackup Datacenter v3.4
13ZJ-1200A-WWEN.pdf on ftp.hp.com

Like Data Protector, NetBackup requires that the backup job be configured before BC EVA initiates it. When including a batch file in the BC EVA job creation form, use the following command to initiate a backup job:

c:\veritas\netbackup\bin\bpbackup –i –c <class>

This command 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 BC EVA client resides.

Note: There is a change in the command-line syntax for **bpbackup** from NetBackup 3.4 to 4.5. *Backup policy* has replaced *backup class* in NetBackup terminology. They both mean the same thing; in fact, the old 3.4 syntax still works in 4.5—it's just not documented. In 3.4, the syntax is: **bpbackup –i –c {backup class}**, where **–i** means to start the backup immediately, and **–c** is the backup class to be referenced for backup details. In 4.5, the syntax is: **bpbackup –i –p {backup policy}**.

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. HP Automation Manager was used to test this rolling clone technique. Here is a sample backup plan:

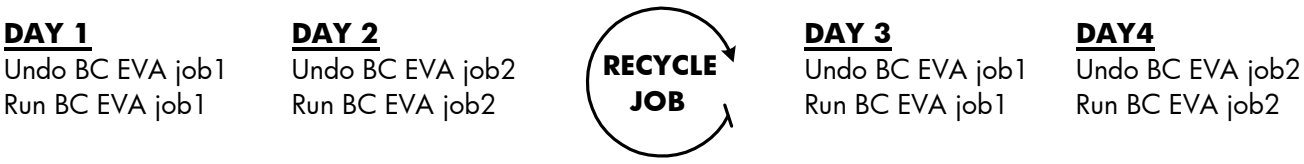

Note: Before running a schedule like this, make sure there is adequate disk capacity in the EVA disk group for multiple snapclones.

Best Practice: Always use Microsoft online backup procedures to back up Exchange EDB, STM, and log files. Use BC EVA for rapid recovery of those files.

Snapclones can also be used for the following purposes:

- Backup for Exchange Server
- Offline testing before upgrades. A BC EVA copy from a live store could be used.
- Defragging a production database on an Exchange recovery server
- Mailbox restores

restoring Exchange stores using snapclones

Restoring Exchange stores from snapclones using Rapid Recovery follows the best practice methods of Microsoft and HP offline restore. (For links to relevant documents, see the websites at the of this paper.) The restore procedure in this section shows how to restore a storage group with two databases and logs in a cluster-active/passive two-node configuration.

Various documents give insight into these restore techniques. To access them, see the following site:

http://support.microsoft.com/default.aspx?scid=KB;en-us;296788&

The first step in restoring is to establish a restore policy that allows for various types of restores, such as a full Exchange rebuild or a single mailbox restore. The customer, hardware, and software govern the restore policy. The backup policy could include a Service Level Agreement (SLA), which is approved by the customer. If there is an SLA, it generally is the determining factor that controls how fast certain data has to be restored. In fact, the 1-hour SLA is the prime reason that the Rapid Recovery for Microsoft Exchange 2000 solution was developed. As the test data for this solution shows, a restore of a storage group took only 4 minutes for 34.6GB of data.

We know that a 100% successful restore requires a good backup. This solution's backup uses BC EVA, Automation Manager, and custom scripts. The next step is to move the backed-up data to the host server (Exchange Server). There are two possible methods: a tape restore, which is traditional, or a BC EVA (snapclone) restore. You must also decide if the data will be moved to a backup server first or directly to the Exchange server.

comparing restore methods

In this solution, both tape and snapclone restore methods were used, for comparison purposes. This section outlines the significant differences between tape and snapclone restore methods.

The task of a restore is to get the Exchange server or the Exchange stores online as fast as possible. Tape restores require time to access the data, while a snapclone already has access.

table 10. restore comparisons

Process	Tape	Snapclone
Average restore time (from table 6)	75min	0
Stop Exchange System Attendant	44sec	44sec
Un-present defective disk	22sec	22sec
Present BC EVA	34sec	34sec
Change disk signature	30sec	30sec
Start Exchange System Attendant	1min	1min
Mount store	1min	1min
Total	**61min**	**4min**

There are several areas where snapclones save time over tape:

- **Data copy:** Table 10 shows the time it takes the data to be copied to disk and the time to move the snapclone to the target exchange server. Using tape, the administrator has to wait 75 minutes before actually starting a recovery.
  - **In comparison, the snapclone already has access to the data and has presented it to the Exchange recovery server. The administrator can immediately begin the restore process.**
- **Data integrity check:** Another step in the restore process is the data integrity check. Once the tape method has copied data to a disk, an integrity check should be done. Even if an integrity test was done during the backup operation, there is no way to tell if the data is still good when it comes back from tape to disk unless it is tested. Many factors can contribute to data corruption from a tape restore, such as dirty communication path, faulty tape, and faulty backup application. Test the store using Exchange ESEFILE and

ESEUTIL utilities. This process will take up additional restore time (not included in table 10) depending on the condition and size of the store.

- **In comparison, the snapclone was tested using ESEFILE and ESEUTIL during the BC EVA backup job process. Since the BC EVA was not modified after the test, there was no need to repeat the integrity check prior to a restore.**

Note: Using ESEFILE and ESEUTIL will impact system performance during their use. Best practice would be to run ESEFILE and ESEUTIL on the Exchange recovery server before moving the recovered store to the production server. During this testing, the performance impact to server resources was as high as 20% when using these utilities. Make sure you fully understand the use of these utilities. They can damage the restored stores if you use them improperly.

- **Backup location:** Another restore consideration should be where the backed-up data should be copied. The stores are generally moved to the original Exchange server or to an Exchange recovery server. If the data is moved to the original host, there needs to be enough disk space on that host to perform the restore task. The rule of thumb is that 50% of the total disk space on the host Exchange server must be reserved for restores. This rule is necessary in this scheme, but a wasteful use of company resources. When using an EVA the ROI is greater, because the extra disks needed for recovery purposes are put in a pool that can be shared by many servers.

  - **In comparison, the snapclone method can be presented to any host in less then a minute. Moreover, disk space does not need to be reserved for a recovery. The 50% rule does not apply. If an Exchange store does need additional space, the virtual disk can be expanded. The snapclone method uses disk resources from a disk group on the EVA. Once the snapclone is deleted, the disk resource is returned to the pool (disk group) to be reused.**

snapclone restore process

In order to control the disk on the EVA, we use the HP SANworks HSV Element Manager, found on the HP OpenView Storage Management Appliance (see figure 12). The HSV Element Manager configures the disk groups, virtual disk, and selective presentation.

figure 12. storage management appliance

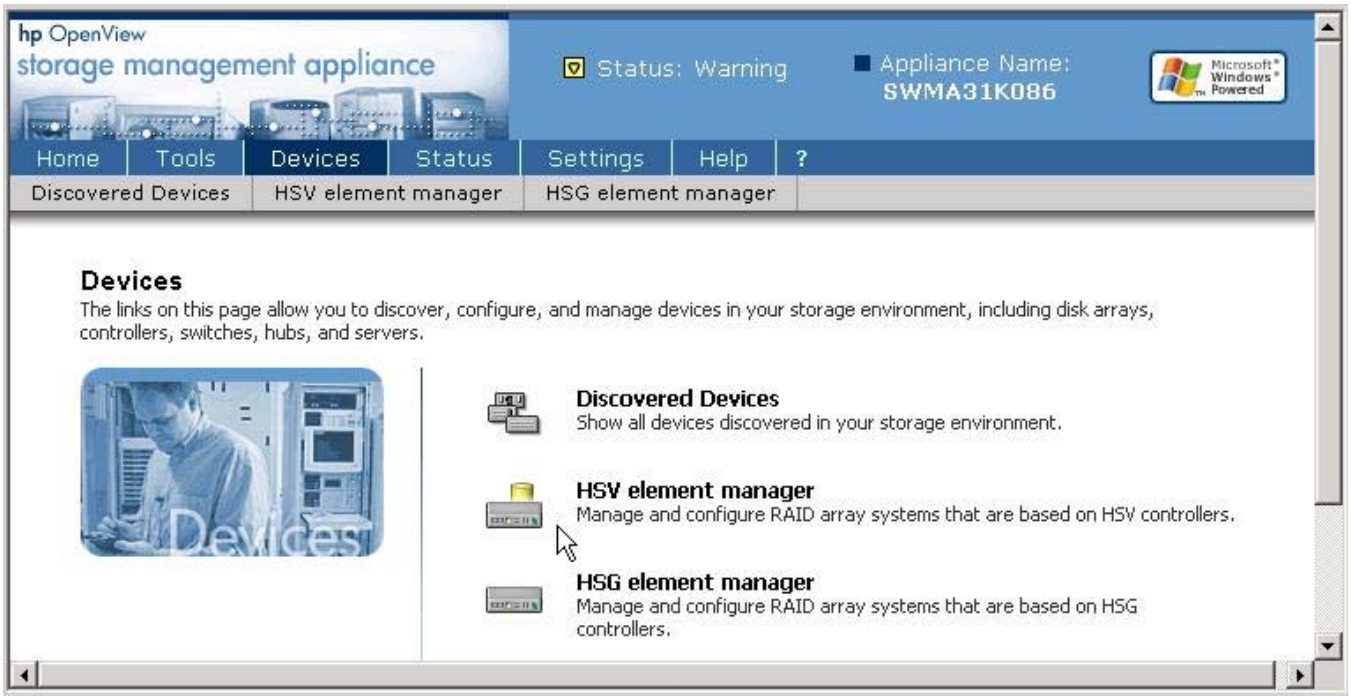

1. Record disk signatures, disk IDs, and disk drive letters of the volumes that will be replaced.

   Use the HP SANworks BC EVA **writesignature.exe** command to display all drives on the active node. The writesignature.exe command comes with the solution script kit:

   http://h18000.www1.hp.com/products/storageworks/solutions/rrex2k/scriptkit.html

   Use the **writesignature –ds** switch to display all active drives, as shown in figure 13.

   Note: Use of the writesignature.exe command is only supported when all cluster resources are on one node. Use the **writesignature.exe** command on the active node in the cluster. Use this command with caution: the –**e**s and –**w**s can change the drive ID and signature.

   Best practice is to record all disk signatures before problems occur on the production Exchange server.

   Note: Another tool that can do the same task is the **dumpcfg.exe**, which is found in the Microsoft Exchange 2000 Resource Kit.

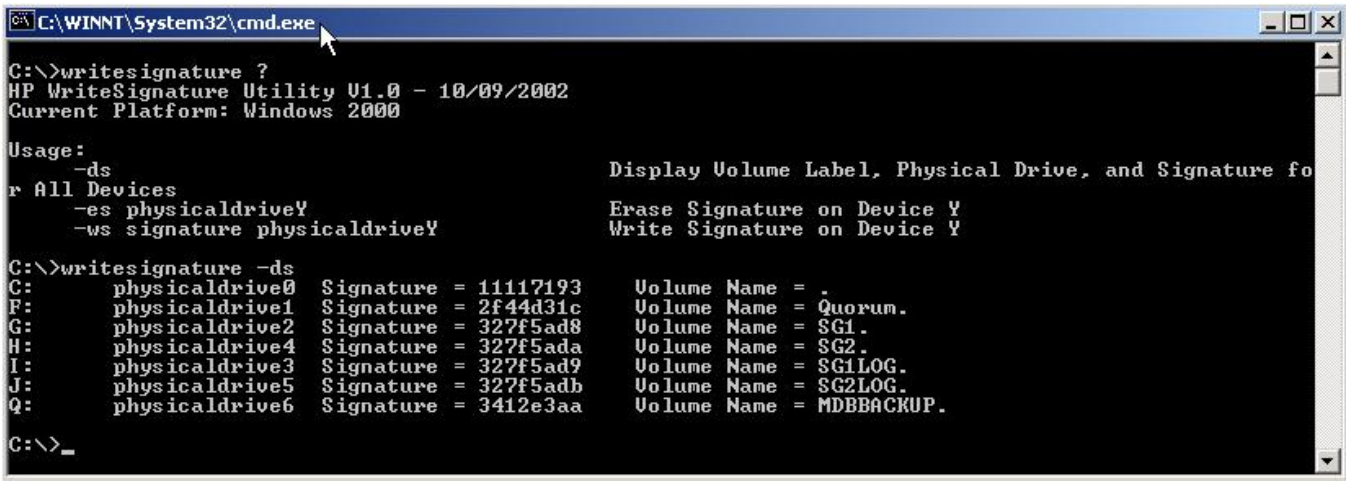

figure 13. hp SANworks BC EVA writesignature program

2. When the BCV (snapclone) is completed, use Cluster Administrator to take down the defective Exchange resource group by taking it offline. This action will dismount the virtual server (VS1-EXCH). Right-click the cluster resource group (EXCHGRP in figure 7) to display a pull-down menu that allows you to take the cluster group resource down. All resources will be down including the disk.

   Note: During this process, do not go into the Exchange Administrator to view the virtual disk that is offline. The Exchange Administrator will hang and eventually time out with an error message. The Exchange Administrator is not aware of the activity from the Cluster Manager.

3. Use HSV Element Manager to remove the defective LUNs from the Exchange server by using selective presentation:

   a. Open the virtual disk folder (Exchange 2000 in figure 14). Find the virtual disks that will be removed and select the **active** virtual disk property.

   b. Click **Unpresent** (see figure 14). Remove all hosts in figure 14. In a cluster scheme, all nodes that are using the virtual disk need to be unpresented.

      Note: Both hosts in figure 14 can be removed at the same time by holding down the Shift key while selecting each.

   c. Click **Finish** to complete the unpresent process.

   d. Repeat Step 3 for each LUN that needs removal.

Note: For this demonstration, SG1 and SG1LOG were unpresented in order to show that any combination, including removing all databases and logs in a storage group, can be restored at the same time. In many cases the logs may be the only element that needs restoring to recover Exchange 2000 stores.

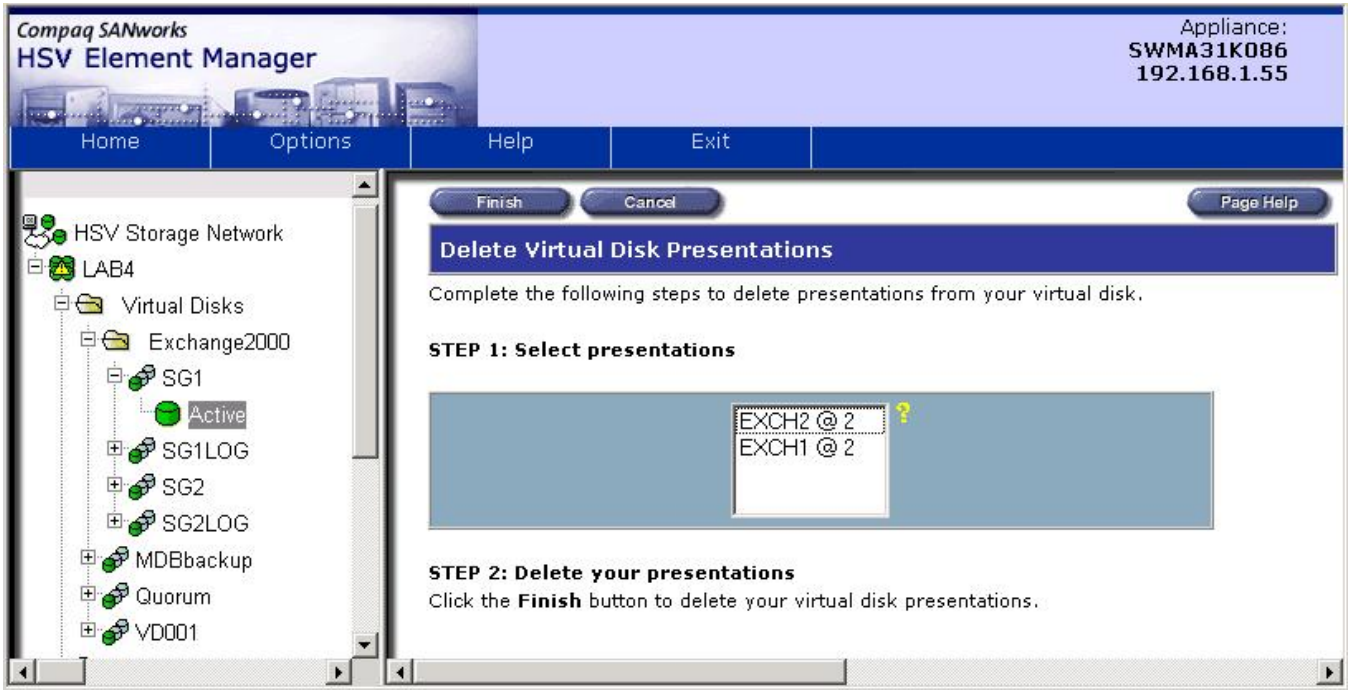

figure 14. unpresented virtual disk

4. Move BCVs to the Exchange server using selective presentation in the HSV Element Manager:
    a. Select the **Active** icon of the BCV to be presented (see figure 15). The virtual disk properties are displayed.
    b. Click **Present**.
    c. In the next screen (figure 16), select the hosts that the New LUNs (BCV) will reside on. Use the same two hosts that were used for the original disks.
    d. Click **Adv(anced) Options**, which lets you select the disk ID (figure 17). The disk ID defaults to the next available disk ID on the selected host. It *must* match the original disk ID that was removed. That disk information should have been retrieved in Step 1. If the disk ID does not match the ID that the cluster disk resource expects when the system attendant starts the exchange, the virtual server will fail.
    e. Click **Finish** to complete presenting the LUNs to the host server.
    f. Repeat Step 4 to select the second cluster server, which needs access to the same LUNs as the first cluster server.

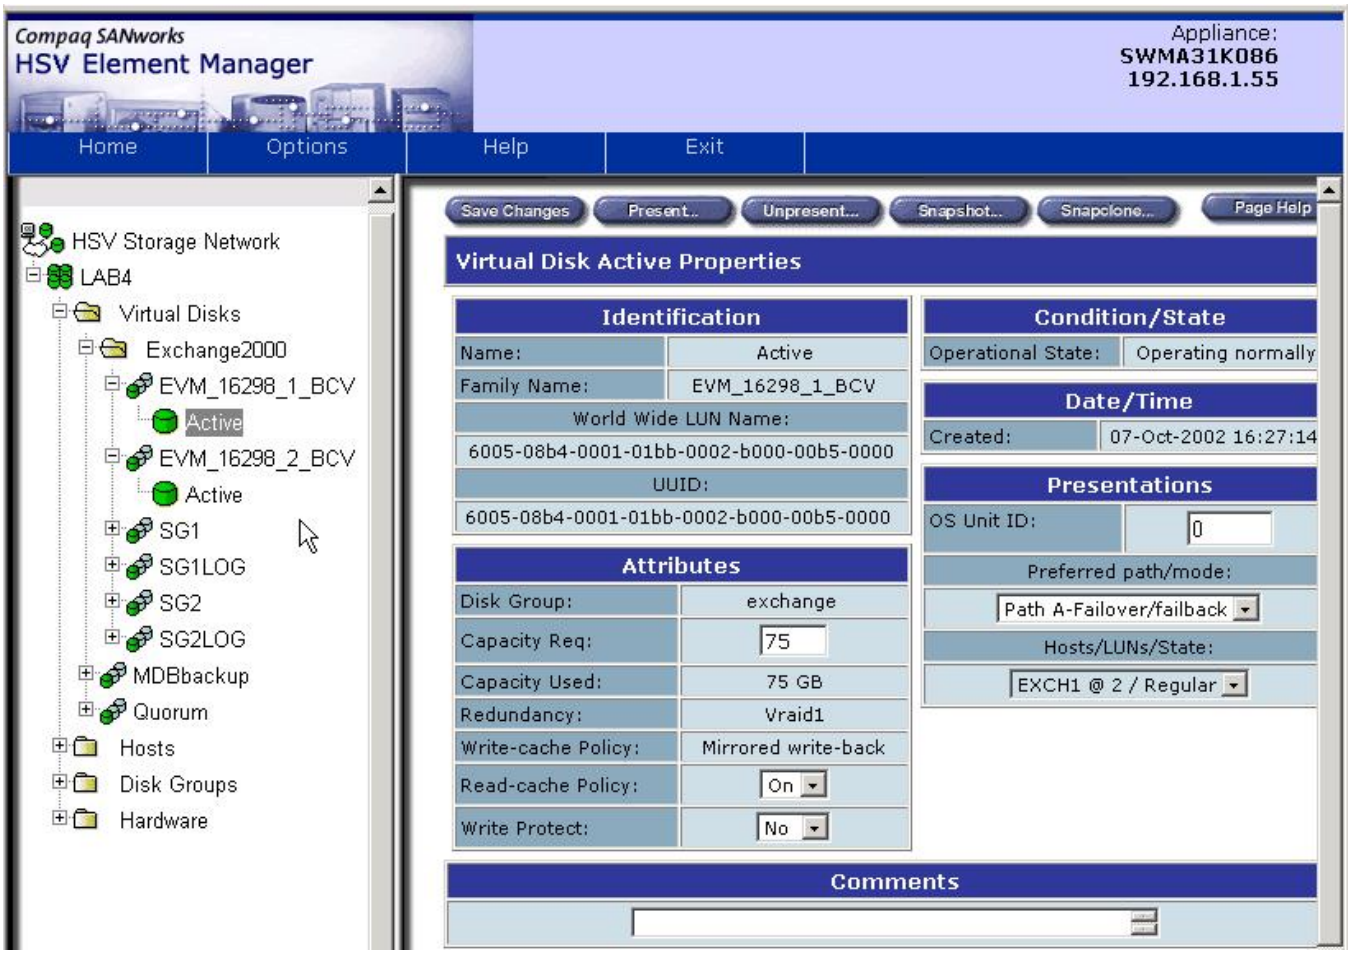

figure 15. HSV element manager

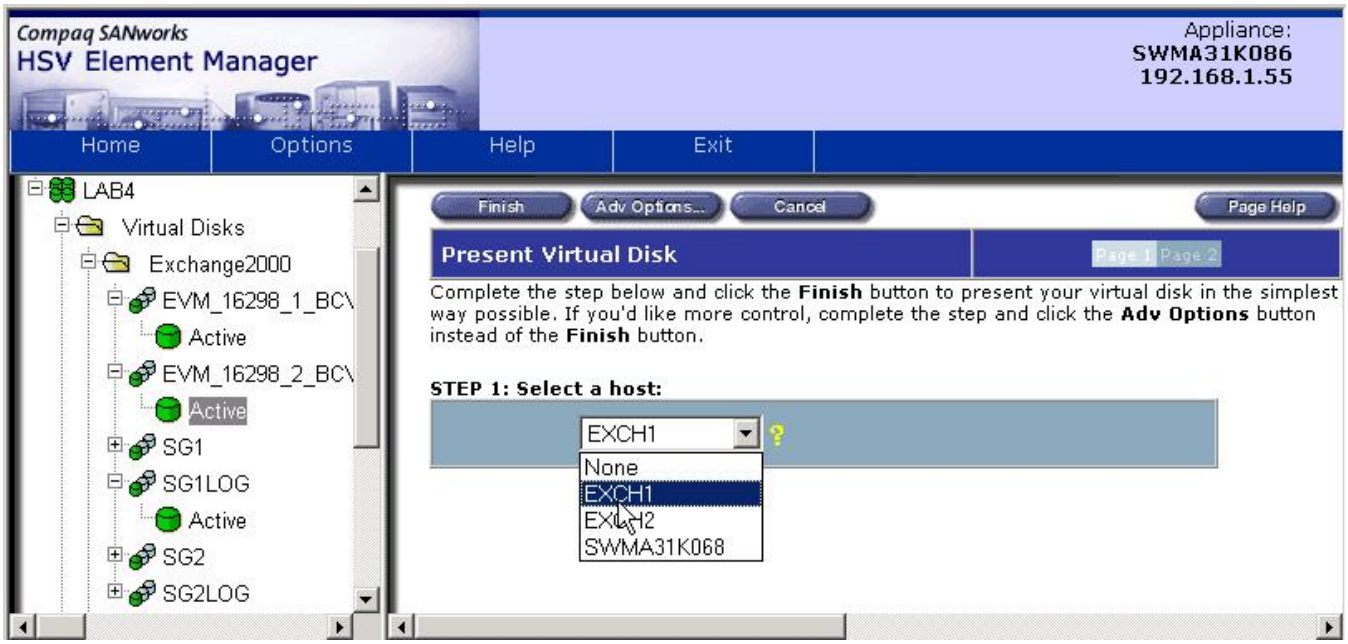

figure 16. presented virtual disk

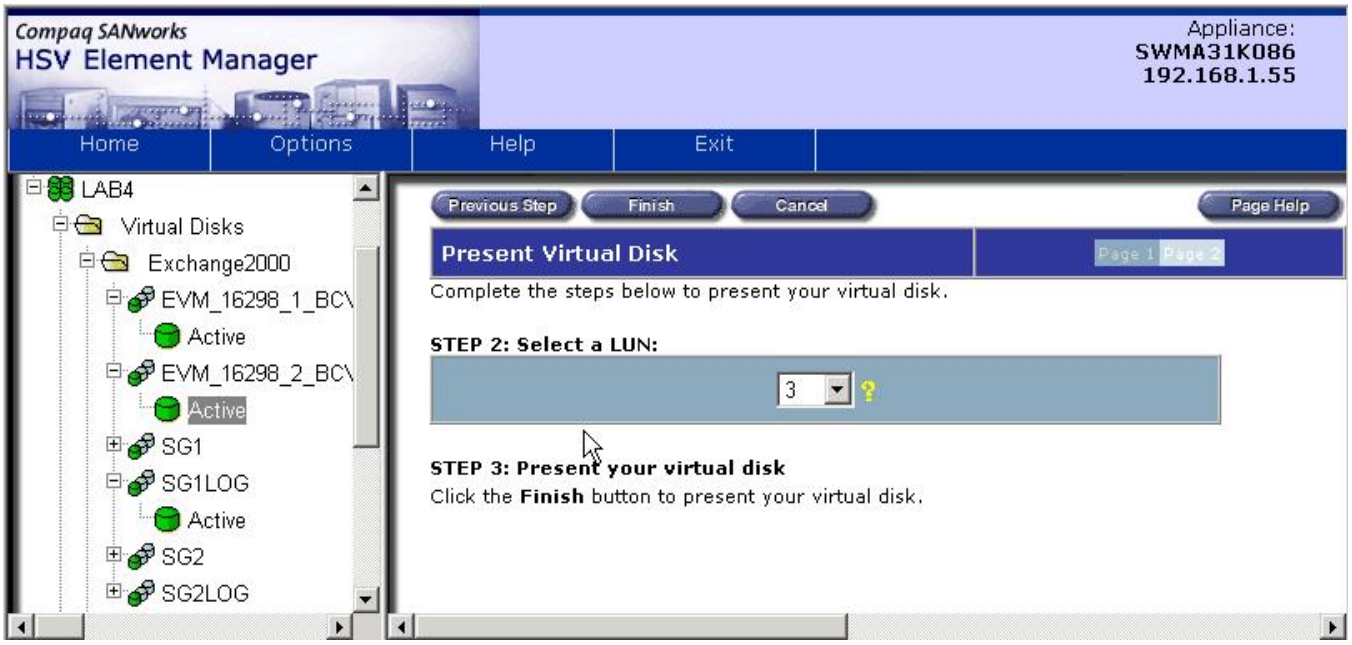

figure 17. LUN ID

5. Using the Cluster Administrator, go to the disk resource in the cluster group that represents the Exchange virtual server that is off line (VS1-EXCH).

   Bring on line *only* the disk resources SG1 and SG1LOG (see figure 7). The disk resources have to be brought on line separately from the rest of the virtual server resources to change the disk signature.

   Note: There is a known problem when replacing a cluster disk with a new disk (normally used in a restore procedure). It occurs because the server that is running Microsoft Cluster Server (MSCS) relies on disk signatures to identify and mount volumes. If a hard disk is replaced or the bus is re-enumerated, MSCS may not find the disk signatures that it is expecting and consequently may fail to mount the disk. Since the Business Continuance Volume (BCV) introduced by BC EVA has a different signature from the disk it is replacing, the cluster resources for the Exchange virtual server will fail when it is brought on line. To resolve this issue, use the HP BC EVA **writesignature** tool to edit the disk signature.

6. Change the drive signature. In this example we continue to use G:SG1 and I:SG1LOG.

   The disks below were the disks that were presented in Step 5. HSV Element Manager has assigned the disk ID and BC EVA has assigned the signatures. Use **wrtesignature –ds** to display the drives.

```
C:\>writesignature -ds
G:      physicaldrive2  Signature = BC EVA500011    Volume Name = SG1.
I:      physicaldrive3  Signature = BC EVA500012    Volume Name = SG1LOG.
```

   Use this command syntax to write to a disk: **writesignature –ws <drive letter>X: <disksignature>**

   The following disk signatures were retrieved from Step 1:

```
C:\>writesignature –ws g: 327f5ad8

Write Signature on g: Successful

C:\>writesignature –ws I: 327f5ad9

Write Signature on i: Successful
```

7. Use the Cluster Administrator to bring the cluster group on line. All resources should come on line.
8. Use Exchange System Manager to verify the stores are available.

scripts

custom scripts

The definition of a BC EVA job is the creation of a process that automatically runs Fibre Channel array (i.e., HSV, HSG) controller commands, which will replicate the host disk on the controller. The BC EVA job itself becomes a sophisticated script, designed using predefined BC EVA processes. You can also use the BC EVA customizable processes to interact with applications such as Exchange, SQL Server, Oracle, and backup products.

The primary BC EVA job commands used in this solution are SUSPEND, RESUME, and LAUNCH. SUSPEND and RESUME are pre- and post-script commands; that is, they allow scripts to interact with an application before and after the BCVs are created. LAUNCH also is used to run a script to interact with an application. LAUNCH has optional WAIT and NOWAIT functions. WAIT lets the process complete before continuing. NOWAIT uses a launch-and-forget method.

There are two tasks in this solution that require custom scripts.

- Stopping and starting the Exchange Server Stores for offline BCV creation
- Backing up the snapclone to tape

Note: All the scripts mentioned here can be found in the Rapid Restore Script Kit, located at:

http://h18000.www1.hp.com/products/storageworks/solutions/rrex2k/scriptkit.html

The scripts are meant as samples; they may require modification to work in your specific environment.

stopping and starting the Exchange Server stores

Custom scripts launched by BC EVA suspend the Exchange store by dismounting all databases in a storage group. These scripts depend on a VBS scripter program. The scripts are placed on the Exchange recovery server in the EVM bin directory. The following scripts are used:

- **Dismountexch.cmd**
- **Mountexch.cmd**
- **BCVinteg.cmd**

example scripts

The **Dismountexch.cmd** and **Mountexch.cmd** scripts rely on the VBS scripter to assimilate the Exchange variables and excute the commands in each script. There are four variables required for each script: the commands **Exchange Server**, **Storage Group**, **Database**, and **Dismount/Mount**.

For example, the **Dismountexch.cmd** file (see Script 1) is broken down as follows.

cscript launches the script to excute the **mounts.vbs** file with the following variables:

- **VS1-EXCH (the Exchange Virtual Server)**
- **SG1 (the storage group)**
- **SG1DB1, SG1DB2 and Public (VS1-EXCH) (the databases)**
- **Dismount**

Note: The variables must be entered as they appear in the Exchange System Manager MMC (case sensitive). If there are any spaces in a variable, surround them with double quotes. See Scripts 1 and 2 for examples.

dismount script

The **Dismountexch.cmd** file (see Script 1) dismounts all the databases in Storage Group 1, including public stores. It includes the following commands:

- Optional **date** command to redirect output to the backup.log file
- Optional **redirect** command placed at the end of each cscript line to post results in a file

script 1—Dismountexch.cmd

```
REM *** Dismount Stores in an Exchange Server ***

date /t time /t >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 SG1DB1 Dismount >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 SG1DB2 Dismount >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 "Public (VS1-EXCH)" Dismount >> c:\temp\backup.log
```

mount script

The **Mountexch.cmd** file (see Script 2) mounts all the databases in Storage Group 1, including public store. It includes the following commands:

- Optional **date** command to redirect output to the backup.log file
- Optional **redirect** command placed at the end of each cscript line to post results in a file

script 2—Mountexch.cmd

```
REM *** Dismount Stores in an Exchange Server ***

date /t time /t >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 SG1DB1 Mount >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 SG1DB2 Mount >> c:\temp\backup.log

cscript mount.vbs VS1-EXCH SG1 "Public (VS1-EXCH)" Mount >> c:\temp\backup.log
```

Exchange integrity check script

The first launch command used in the BC EVA job (see figure 8) runs **BCVinteg.cmd** to test the integrity of the BCVs. BCVinteg.cmd runs two Exchange utilities to test the consistency and integerity of the databases. These utilities are run before the BCVs are backed up to tape.

- ESEUTIL consistency test
- ESEFILE integrity check

Best Practice: When an Exchange store is replicated to other media for the purpose of recovering a host Exchange store, test the integrity of the copy using Microsoft Exchange utilities ESEUTIL and ESEFILE.

script 3—BCVinteg.cmd

```
REM ********* Integrity and consistency check for BCVs  ********
REM ** See readme file in RR script kit for explanation of this batch file or the **
REM ** Exchange Rapid Recovery Implementation Blueprint Document ****

eseutil /mh o:\exchsrvr\mdbdata\SG1DB1\1priv1.edb | find /i "consistent" >> c:\temp\snap.logrem
eseutil /mh o:\exchsrvr\mdbdata\SG1DB1\1priv1.stm | find /i "consistent" >> c:\temp\snap.log
eseutil /mh o:\exchsrvr\mdbdata\SG1DB2\1priv2.edb | find /i "consistent" >> c:\temp\snap.log
eseutil /mh o:\exchsrvr\mdbdata\SG1DB2\1priv2.stm | find /i "consistent" >> c:\temp\snap.log
eseutil /mh o:\exchsrvr\mdbdata\public\pub1.edb | find /i "consistent" >> c:\temp\snap.log
eseutil /mh o:\exchsrvr\mdbdata\public\pub1.stm | find /i "consistent" >> c:\temp\snap.log

esefile /s o:\exchsrvr\mdbdata\SG1DB1\1priv1.edb >> c:\temp\snap.log
esefile /s o:\exchsrvr\mdbdata\SG1DB2\1priv2.edb >> c:\temp\snap.log
esefile /s o:\exchsrvr\mdbdata\public\pub1.edb >> c:\temp\snap.log
```

Visual Basic scripter: Mount.vbs

The **Mount.vbs** file is a Visual Basic scripter that receives input from the scripts created for Exchange. Place **Mount.vbs** in the Exchange recovery server in the bin directory of BC EVA.

script 4—Mount.vbs

```
Set oServer = CreateObject("CDOEXM.ExchangeServer")
Set oFirstStorageGroup = CreateObject("CDOEXM.StorageGroup")
Set oMailboxStoreDB = CreateObject("CDOEXM.MailBoxStoreDB")
Set oPublicStoreDB = CreateObject("CDOEXM.PublicStoreDB")

'Set up Params
        if WScript.Arguments.Count = 4 then
        strServer = WScript.Arguments(0)
        strStorageGroup = WScript.Arguments(1)
        strMDB = WScript.Arguments(2)
```

```
if WScript.Arguments(3) = "Mount" then
        bMount = true
        elseif WScript.Arguments(3) = "Dismount" then
                bMount = false
        else
                WScript.echo "4th parameter must be either Mount or Dismount"
        end if
else
        WScript.echo "Usage Error, must enter 4 parameters"
        Call Usage
        WScript.quit
end if

'Open up the server
oServer.DataSource.Open (strServer)

'Get an array of StorageGroup objects
oStorageGroups = oServer.StorageGroups

'Get the count of SGs
cSGs = UBound(oStorageGroups) - LBound(oStorageGroups)

'Go through the SGs and look for the Storage group the user wants
bFound = false
for iSGs = 0 to cSGs
        oFirstStorageGroup.DataSource.Open (oStorageGroups(iSGs))
        if oFirstStorageGroup.Name = strStorageGroup then
                bFound = true
                exit for
        end if
next
if not bFound then
        WScript.echo "Storage group not found"
        WScript.quit
end if

'Get an array of MBX DBs
oDBs = oFirstStorageGroup.MailboxStoreDBs
```

```
'Get count of DBs
cDBs = UBound(oDBs) - LBound(oDBs)

'Go through the MBX DBs and look for the one the user wants to dis/mount
bFound = false
for iDBs = 0 to cDBs
        oMailboxStoreDB.DataSource.Open (oDBs(iDBs))
        if oMailboxStoreDB.Name = strMDB then
                bFound = true
                exit for
        end if
next

'If we found the DB the user wants to Mount/Dismount
'If not found try the PF DBs
if bFound then
        if bMount then
                oMailboxStoreDB.Mount
                strActionDone = "Mounted"
        else
                oMailboxStoreDB.Dismount
                strActionDone = "Dismounted"
        end if
        WScript.echo oFirstStorageGroup.Name + " - " + oMailboxStoreDB.Name + " " + strActionDone
else
        'Get an array of PF DBs
        oDBs = oFirstStorageGroup.PublicStoreDBs

        'Get count of DBs
        cDBs = UBound(oDBs) - LBound(oDBs)

        'Go through the PF DBs and look for the one the user wants to dis/mount
        for iDBs = 0 to cDBs
                oPublicStoreDB.DataSource.Open (oDBs(iDb))
                if oPublicStoreDB.Name = strMDB then
                        bFound = true
                        exit for
```

```
                        end if
                next
                if bFound then
                        if bMount then
                                oPublicStoreDB.Mount
                                strActionDone = "Mounted"
                        else
                                oPublicStoreDB.Dismount
                                strActionDone = "Dismounted"
                        end if
                WScript.echo oFirstStorageGroup.Name + " - " + oPublicStoreDB.Name + " " + strActionDone
                else
                        WScript.echo "The DB was not found in the Storage group"
                end if
        end if

sub Usage
        Wscript.echo "Usage: Mount <Server> <Storage Group> <MDB> <Mount or Dismount>"
        Wscript.echo "Usage: The parameters are case sensitive"
        Wscript.echo "Usage: Parameters with a space must be enclosed in double quotes."
 End sub
```

the second BC EVA job LAUNCH command

The second **BC EVA job LAUNCH** command runs a script that backs up the snapclone to tape. Rapid Recovery has integrated two backup applications:

- HP OpenView Storage Data Protector
- VERITAS NetBackup

hp OpenView storage data protector

HP OpenView Data Storage Protector integration with Microsoft Exchange Server allows you to perform online or offline backups of Exchange 2000 Server. Data Protector is cluster-aware. To integrate Data Protector and database applications so both are cluster-aware, install the Data Protector cell manager on the cluster and select the required Data Protector cell manager integration software components for Exchange 2000. A few configuration notes:

- Be sure you have an HP license.
- Install a cell manager. In this solution the cell manager was installed on the Exchange recovery server.
- Install the Data Protector client on the two Exchange cluster nodes.
- Add the <Exchange2000>\bin directory to the system path before any operation is performed.
- Install Exchange Server integration agent.

- If you have devices connected to the system, Media Agent automatically installs during initial setup of Data Protector.

During installation, select the integration software component for Exchange 2000. In this solution, only the Exchange Server option was selected to be installed on the virtual server.

Note: Normally when installing a backup application, the tape devices (SCSI drive and library) are already installed. In this solution the drivers on the HP MSL5026 library and the SDLT drives were installed before Data Protector. Even though Data Protector could recognize the drivers, it did not have complete control over them. These drivers had to be removed permanently before Data Protector could be installed.

Note: Another problem to watch for is the drive assignments on the library. The SCSI tape drives could be in the reverse order from what the library has configured and what Data Protector expects. A symptom of this problem is that when you try to assign tapes to a cell they can be seen by the inventory process but they can't be Data Protector formatted.

the Data Protector cell environment

The Data Protector cell is a network environment containing a cell manager, clients, and backup devices. The cell manager centrally manages the cell and administers the backup and restore operations. Systems that are to be backed up can be added to the cell and set up as Data Protector clients. When Data Protector performs a backup of data from these clients, it saves the data to media contained within the backup devices.

The Data Protector's internal database (IDB) keeps track of the files backed up, making it easy to browse and restore them, ether singly or collectively. This recording occurs whether the process was done using the GUI or command line interface.

The cell manager is the main control center for the cell and contains the IDB. It runs the core Data Protector software and the Session Manager, which starts and stops backup and restore sessions and writes session information to the IDB. Any system within a chosen cell environment can be set up as a Data Protector client. The role of the client depends on whether it has a Disk Agent or a Media Agent installed.

For the latest *HP OpenView Storage Data Protector Software Release Notes,* visit

http://www.openview.hp.com/products/data_protector/

BC EVA launches the scripts, but the Data Protector commands are run on the Exchange recovery server at the command line. The commands are customized for a static backup process and have to be changed if a volume name or drive letter is changed. For example, the **Omnibackup.cmd** file (see Script 5) is broken down as follows.

omnib launches Windows Data Protector commands with the following varables

- backupsrv (the Exchange recovery server)
- O: SG1, P: SG1LOG (the drive letters and volumes of the BCV)
- –device drive1 (the SCSI drive in the library)
- –mode full (a full backup)
- –protect (protects the tape from being written over)
- –pool (location of the tape that was picked)

script 5—Omnibackup.cmd

```
REM *** This file will issue commands to HP OpenView Storage Data Protector and run a full backup of the ***
REM *** below Volumes SG1, SG1LOG ***

omnib –winfs backupsrv:O: SG1 –device Drive1 –mode full –protect none –pool LAB4

omnib –winfs backupsrv:P: SG1LOG –device Drive1 –mode full –protect none –pool LAB4


REM *** These commands do the same as above and add a debug process, which directs an ***
REM *** output file to Omniback\log **
REM omnib –winfs backupsrv:O: SG1 –device Drive1 –mode full –protect none –pool LAB4 –debug 1-99 debuglog.txt

REM omnib –winfs backupsrv:P: SG1LOG –device Drive1 –mode full –protect none –pool LAB4 –debug 1-99 debuglog.txt
```

for more information

For a demo of BC EVA, click on the following link. This demo installs on any Windows 2000 or XP-based system. It simulates a basic SAN consisting of HSV and HSG controllers on two Windows hosts. Documentation is included.

Note: The file is 20MB. Load it to a PC and use locally.

http://storage.inet.cpqcorp.net/application/view/allibraries.asp?LIBID=&QLID=11&rdoType=filter&inpCriteria1=Category&inpCriteria2=nothing&inpCriteriaChild1=Tool&btnSubmit=Filter+It%21&inpCriteria1XML=&inpCriteria2XML=&inpCriteriaChild1XML=&inpCriteriaChild2XML=

additional links

For Microsoft Exchange 2000 Server information:
http://www.microsoft.com/Exchange/default.asp

For HP ActiveAnswers for Microsoft Exchange Server:
http://activeanswers.compaq.com/ActiveAnswers/Render/1,1027,2366-6-100-225-1,00.htm

For HP StorageWorks Enterprise Virtual Array information and documentation:
http://h18004.www1.hp.com/products/storageworks/enterprise/index.html

For Management Appliance information and documentation:
http:// h18004.www1.hp.com/products/sanworks/managementappliance/index.html

For HP StorageWorks Business Copy documentation:
http://h18004.www1.hp.com/products/sanworks/BC EVA/documentation.html

For HP StorageWorks Secure Path information and documentation:
http:// h18004.www1.hp.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 HP OpenView Storage Data Protector and backup guide using Exchange 2000:

Using HP Data Protector 5.0 to back up Exchange 2000 servers

For VERITAS NetBackup information and documentation:
http://www.veritas.com >> Products

trademark and date information

4/03