

Quick recovery for Microsoft® Exchange 2003 using VERITAS Storage Foundation for Microsoft Windows and HP storage — white paper



Executive summary.....	2
Test overview.....	2
Test environment.....	2
About the components.....	3
Test procedures and results	3
Solution configuration	4
Configuring the test environment.....	4
Configuring the Exchange Server.....	10
How the Snap commands work.....	16
Creating snapshots of the storage group.....	17
Verifying integrity.....	17
Restoring storage groups.....	19
Point-in-time recovery.....	19
Roll-forward recovery.....	20
Appendix A—Detailed test results	22
Appendix B—Test configuration BOM	26
Appendix C—Snap commands.....	28
For more information.....	29

Executive summary

This white paper describes how users can recover Microsoft® Exchange Server 2003 databases in just minutes. Microsoft Exchange Server 2003 leverages Microsoft's Volume Shadow Copy Service (VSS) framework to achieve extremely quick recovery of Microsoft Exchange Server 2003 environments. The test environment included three HP storage configuration options: the HP Smart Array 6402 controller, HP StorageWorks Modular Smart Array 1000 (MSA1000) storage area network (SAN), and HP StorageWorks Enterprise Virtual Array 5000 (EVA5000) SAN.

Creating VSS snapshots of the production Exchange volumes enables quick recovery. These snapshots (on-host persistent volumes), which are created without shutting down the Exchange application, can be used as a substitute for the current production volumes if a corruption or database loss occurs. This process provides a first line of defense that is significantly faster than restoration from standard backup media (typically from tape), although it does not replace the need for standard backup procedures as a disaster recovery best practice.

HP has provided a total solution that integrates HP storage with VERITAS Storage Foundation 4.1 for Microsoft Windows® and Microsoft Exchange Server 2003. This solution augments the backup process with snapshots of Microsoft Exchange Server 2003, which enables administrators to restore services within minutes, instead of hours.

Results: During all restoration testing, Exchange databases could be mounted in less than two minutes—in most cases, less than one minute. The time required to mount individual storage group databases for point-in-time restores was generally less than 30 seconds.

Test overview

HP has integrated and tested a quick recovery solution that customers can implement to recover their Exchange databases in minutes, instead of hours. The purpose of the testing was to show how HP and VERITAS Storage Foundation 4.1 for Windows could provide organizations running Microsoft Exchange Server 2003 with very fast recovery times.

Microsoft Windows Server 2003 includes the VSS framework, which enables customers to create snapshot copies of a volume. By using these snapshots, which can be created without shutting down the application, customers can achieve extremely quick recovery of Microsoft Exchange Server 2003.

Results: The test results show that high recovery rates can be achieved by following the configuration procedures and best practices. A recovery rate of just minutes is achievable by integrating HP storage along with VERITAS Storage Foundation 4.1 for Windows and Microsoft Exchange Server 2003.

Test environment

This project was tested and documented on the following three HP storage configurations:

- HP ProLiant Smart Array 6402 controller, direct attached storage, and 500 users—The users were distributed across two storage groups with a single mailbox store for each group.
- MSA1000 SAN—1,000 users were distributed across two storage groups with two mailbox stores for each group.
- EVA5000 SAN with a three-node Active/Active/Passive (A/A/P) Microsoft Windows Server 2003 cluster—This configuration supported 5,000 users, who were distributed across two Exchange 2003 virtual servers. The Exchange 2003 configuration included four storage groups per virtual server (eight total for the configuration), each with 625 users.

The user load was simulated using the standard Microsoft Messaging Benchmark MMB3 profile. More details are included in the Configuring the test environment section.

About the components

VERITAS Storage Foundation 4.1 for Windows, formerly known as VERITAS Volume Manager for Windows, extends the capabilities of dynamic disks in Windows, allowing snapshot images of Exchange data without taking services offline. These snapshot images can be used to recover Exchange quickly.

VERITAS Storage Foundation 4.1 for Windows takes advantage of the Microsoft VSS technology. The VSS framework includes three major components. Each has a particular role in snapshot creation:

- Requestor, which is typically a backup application, calls VSS to prepare for a snapshot. In this solution, the Storage Foundation FlashSnap VxSnap utility is the Requestor.
- Writer (VSS-enabled application) manages the data for the application. Microsoft Exchange Server 2003 functions as the Writer.
- Provider (vendor-specific software) creates the storage for the snapshot. Storage Foundation Provider integration is provided by FlashSnap.

By combining the management features of HP storage with the enhanced volume management of VERITAS Storage Foundation 4.1 for Windows, administrators can easily create VSS snapshots and achieve significantly reduced recovery times. The VERITAS Enterprise Administrator (VEA) user interface is used to manage and monitor all the storage objects in the system.

Test procedures and results

Testing was based on “split mirror snapshots,” which are separate persistent volumes containing an exact duplicate of all the data on the original volume at the time that the copy was created.

The test procedures were performed as follows:

- Test Run 1 established a baseline for Exchange performance. No mirrors were attached to volumes.
- In Test Run 2, mirrors were attached to all volumes for the duration of the test run to observe the performance impact of mirroring.
- In Test Run 3, VSS snapshots were created during a four-hour steady state period.

All subsequent test runs, during which mirrors or VSS snapshots were created, were compared against the baseline results. Performance Monitor Counters were tracked during all test runs. The tests demonstrated quick recovery times for Exchange 2003. The following activities were successfully performed:

- Volume creation
 - Created VSS snapshot volumes with Exchange online and without impact to the user
 - Successfully performed an integrity check of each volume to ensure that the volume can be used as a recovery source
- Database recovery
 - Demonstrated that volume is persistent and can be used to enable instant recovery
 - Demonstrated point-in-time recovery by utilizing the snapshot as the recovery source
 - Demonstrated roll-forward recovery by utilizing the database snapshot as the recovery source and applying the current logs

Testing was done using the Load Simulator (LoadSim) for the Microsoft Exchange Server 2003 default MMB3 profile. Server performance numbers are averages recorded during a four-hour steady state of each test run. The expected results were derived using standard performance points and generally accepted values. The Performance Monitor Counters reported are intended to provide a general indication of server performance.

Server processor utilization numbers were well below established guidelines. Memory and disk resources were within expected guidelines. Volume mirroring and snapshot creation did not result in significant performance degradation.

Results: During all restoration testing, Exchange databases could be mounted in less than two minutes—in most cases, less than one minute. The time required to mount individual storage group databases for point-in-time restores was generally less than 30 seconds.

A LoadSim 95th percentile score of less than 1,000 ms is generally considered a passing score.

Table 1.

Test Run		EVA cluster 5,000 users	MSA1000 1,000 users	SA6402 2,500 users
Test 1—Baseline	95% score	507	513	310
Test 2—Mirrored volumes	95% score	739	375	312
Test 3—VSS snapshot	95% score	432	382	401

The test results show that quick recovery rates can be achieved. Detailed results of the tests are described in Appendix A—Detailed test results.

Solution configuration

To configure the solution, the following tasks were performed:

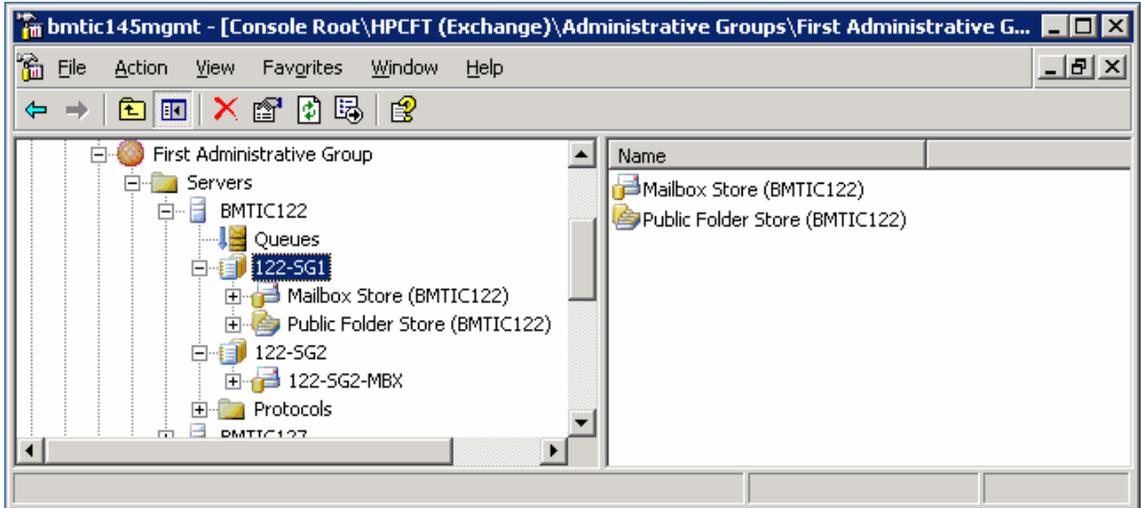
1. Configuring the test environment
2. Configuring the Exchange servers, which included creating dynamic disk groups, volumes, and mirrors
3. Creating snapshots of the storage groups
4. Verifying integrity by running `eseutil /k /l` against the databases and `eseutil /ml` against the logs
5. Restoring the storage groups

Configuring the test environment

Three types of arrays were set up for testing.

The first system consisted of a single server supporting 500 users. Users were distributed across two storage groups with a single mailbox store for each group. See Figure 1.

Figure 1.



User load was generated from a single HP ProLiant DL580 G1 server running a 500-user MMB3 profile.

- Server: HP ProLiant ML530 G2 system with 2-GB RAM (2x) and 3.0-GHz Intel® Xeon™ processor with Hyper-Threading enabled running Exchange 2003 and Windows Server 2003
- Storage: Smart Array 6402 controller, (12) 72-GB 10,000-rpm drives and eight 36-GB 10,000-rpm drives

Figure 2. HP ProLiant ML530 G2



Table 2. Database/LUN Mapping

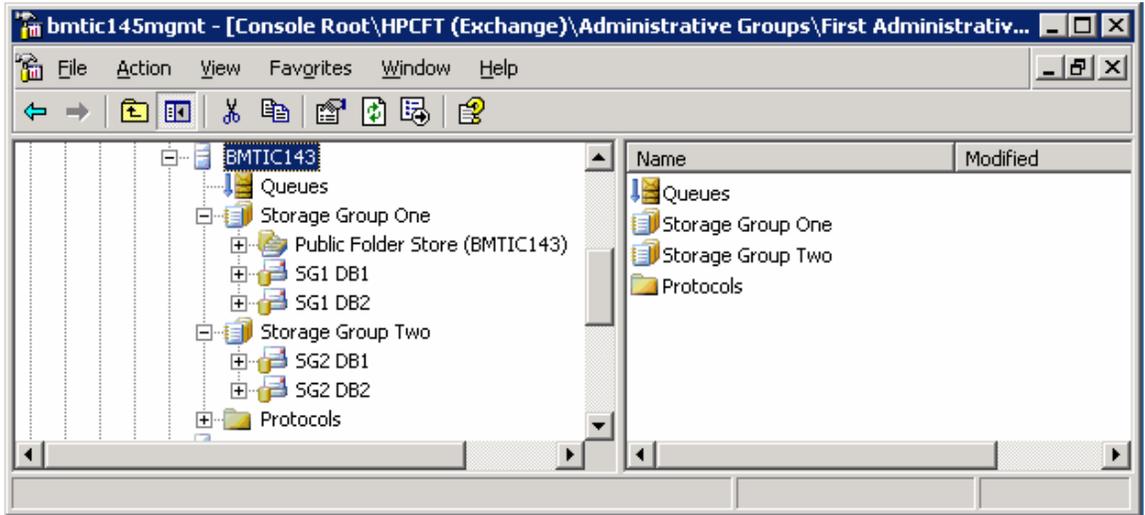
Storage Group 1	N:\	Array B LD 2	(4) 72-GB RAID 1+0
Storage Group 2	O:\	Array C LD 3	(4) 72-GB RAID 1+0
SG1 Logs	L:\	Array D LD 4	(2) 72-GB RAID 1+0
SG2 Logs	P:\	Array D LD 4	(2) 72-GB RAID 1+0
Snapshots	N/A	Array E LD 5	(8) 36-GB RAID 1+0

Note

Snapshots are not presented so they do not have drive letters assigned.

The second system consisted of a single server supporting 1,000 users. Users were distributed across two storage groups with two mailbox stores each. See Figure 3.

Figure 3.



User load was generated from a single ProLiant DL580 G1 server running a 1,000-user MMB3 profile.

- Server: ProLiant DL580 G2 system with 4-GB RAM and four 2.8-GHz Xeon processors with Hyper-Threading enabled, running Exchange 2003 and Windows Server 2003
- Storage: MSA1000, version 4.24 firmware, (12) 72-GB 10,000-rpm drives and (14) 36-GB 10,000-rpm drives

Figure 4. MSA1000

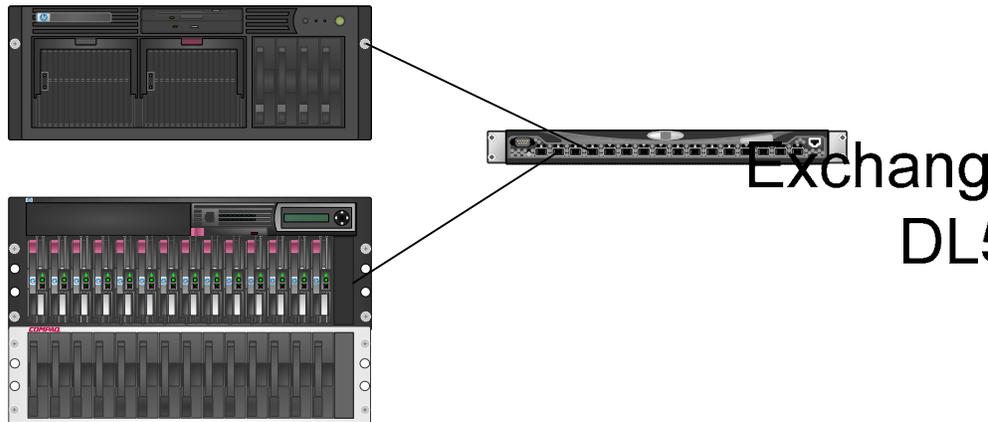


Table 3. Database/LUN mapping

Storage Group 1	N:\	Array B LD 2	(4) 72-GB RAID 1+0
Storage Group 2	O:\	Array C LD 3	(4) 72-GB RAID 1+0
SG1 Logs	L:\	Array A LD 1	(2) 72-GB RAID 1+0
SG2 Logs	P:\	Array A LD 1	(2) 72-GB RAID 1+0
Snapshots	N/A	Array D, E LD 4,5,6	(14) 36-GB RAID 1+0

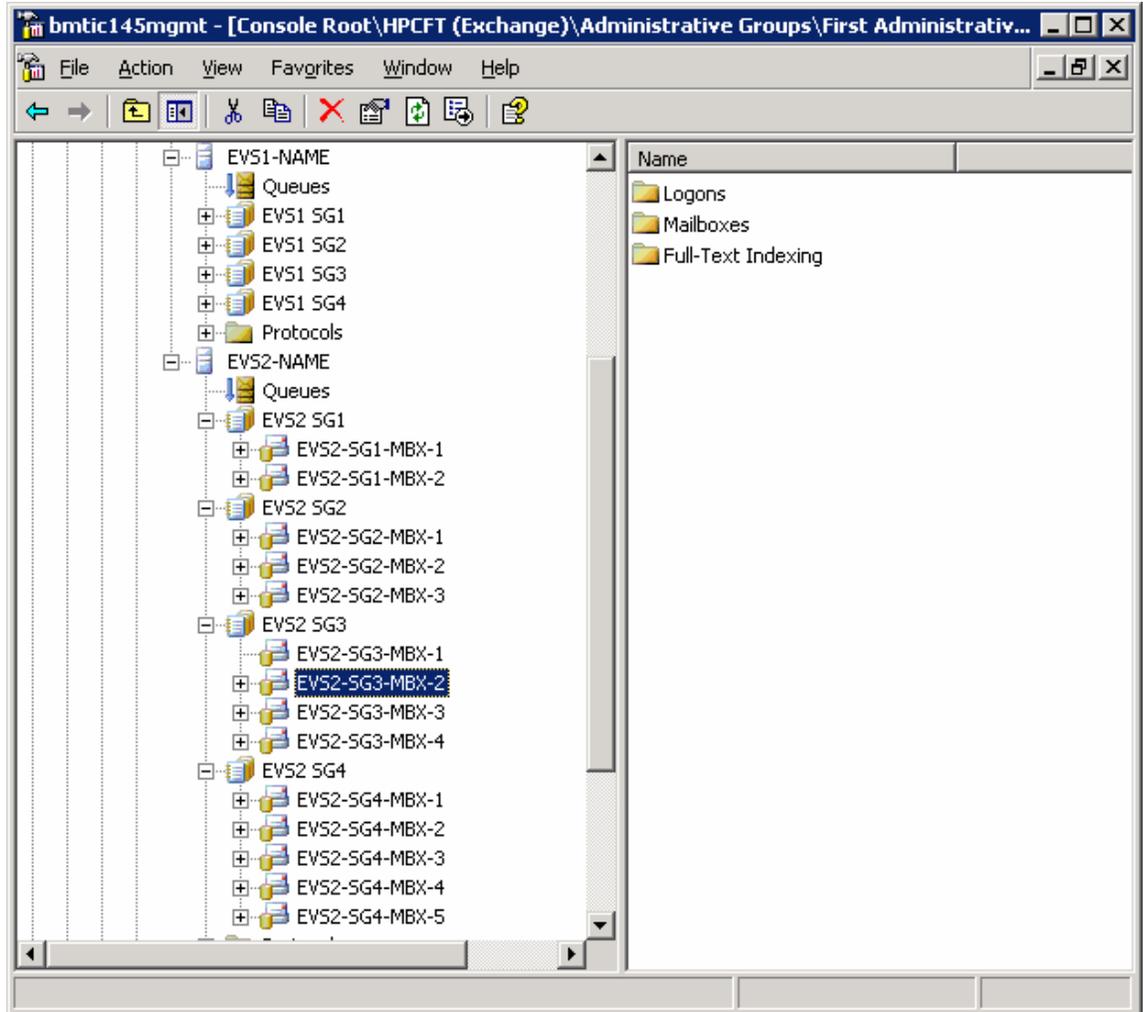
Note

Snapshots are not presented so they do not have file system drive letters assigned.

StorageWorks 42x4

The third system consisted of a three-node Windows Server 2003 cluster. This configuration supported 5,000 users distributed across two Exchange 2003 virtual servers. EVS1 had four storage groups, each with a single mailbox store. EVS2 had four storage groups with n+1 mailbox stores each. For example, storage group one had two stores, storage group two had three stores, and so on. Storage is a single EVA5000. See Figure 5.

Figure 5.



User load was generated from five ProLiant servers. Each server was configured to run a 1,000-user MMB3 profile.

- Servers: Three ProLiant DL580 G2 systems with 4-GB RAM and four 2.8-GHz Xeon processors with Hyper-Threading enabled, running Exchange 2003, Windows Server 2003, and Microsoft Cluster Services (MSCS) A/A/P cluster for 5,000 users
- Storage: EVA5000, VCS 3.010, (166) 72-GB 10,000-rpm SCSI drives

Figure 6. Cluster/EVA SAN configuration

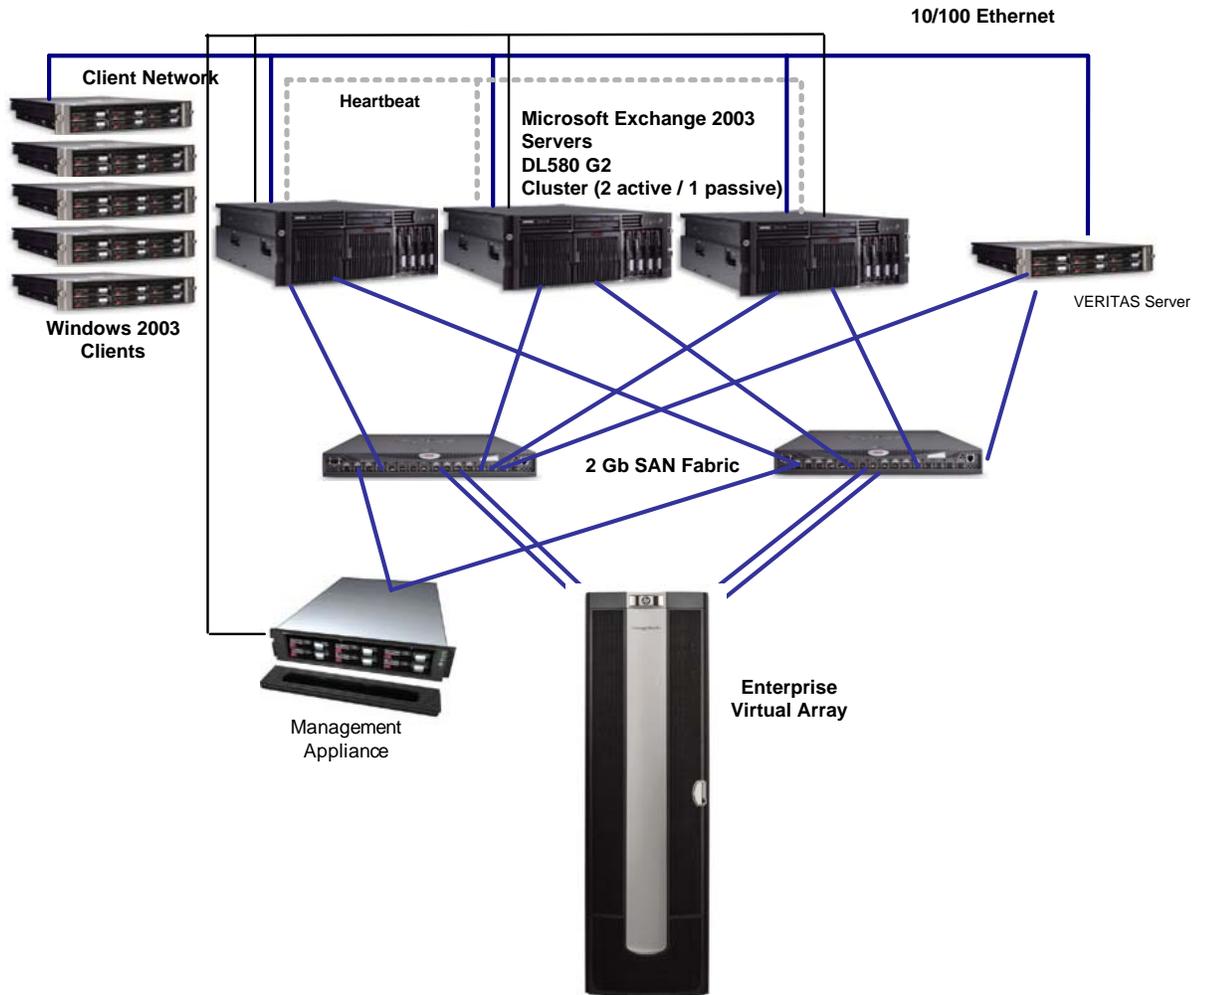


Table 4. EVA/database configuration

Disk group	Disks	Capacity	Contents
1	36	2,304 GB	EVS1 databases
2	36	2,304 GB	EVS2 databases
3	16	948 GB	Logs
4	52	3,388 GB	Snapshots
5	26	1,555 GB	Snapshots

Note

Snapshot space requirements can vary depending on deployment. The minimum additional storage requirement is equal to double the usable allocated volume size. Additional mirrors will require additional space equal to the volume being mirrored.

Table 5. LUN mapping

Server designation	Disk drive letter	Data type	Disk group	LUN size (GB)	RAID	LUN
Cluster	Q\$	Quorum	DG3	10	1	1
EVS1	E:\Exchsrvr (Mount Point)	SMTP	DG1	50	1	2
	E\$	SG1-Data	DG1	125	1	3
	F\$	SG2-Data	DG1	125	1	4
	G\$	SG3-Data	DG1	125	1	5
	H\$	SG4-Data	DG1	125	1	6
	E:\Log (Mount Point)	SG1-Log	DG3	50	1	7
	F:\Log (Mount Point)	SG2-Log	DG3	50	1	8
	G:\Log (Mount Point)	SG3-Log	DG3	50	1	9
	H:\Log (Mount Point)	SG4-Log	DG3	50	1	10
EVS2	I:\Exchsrvr (Mount Point)	SMTP	DG2	50	1	15
	I\$	SG5-Data	DG2	125	1	16
	J\$	SG6-Data	DG2	125	1	17
	K\$	SG7-Data	DG2	125	1	18
	L\$	SG8-Data	DG2	125	1	19
	I:\Log (Mount Point)	SG5-Log	DG3	50	1	20
	J:\Log (Mount Point)	SG6-Log	DG3	50	1	21
	K:\Log (Mount Point)	SG7-Log	DG3	50	1	22
	L:\Log (Mount Point)	SG8-Log	DG3	50	1	23

Note

Mount points allow for logical unit (LUN) scalability.

Configuring the Exchange server

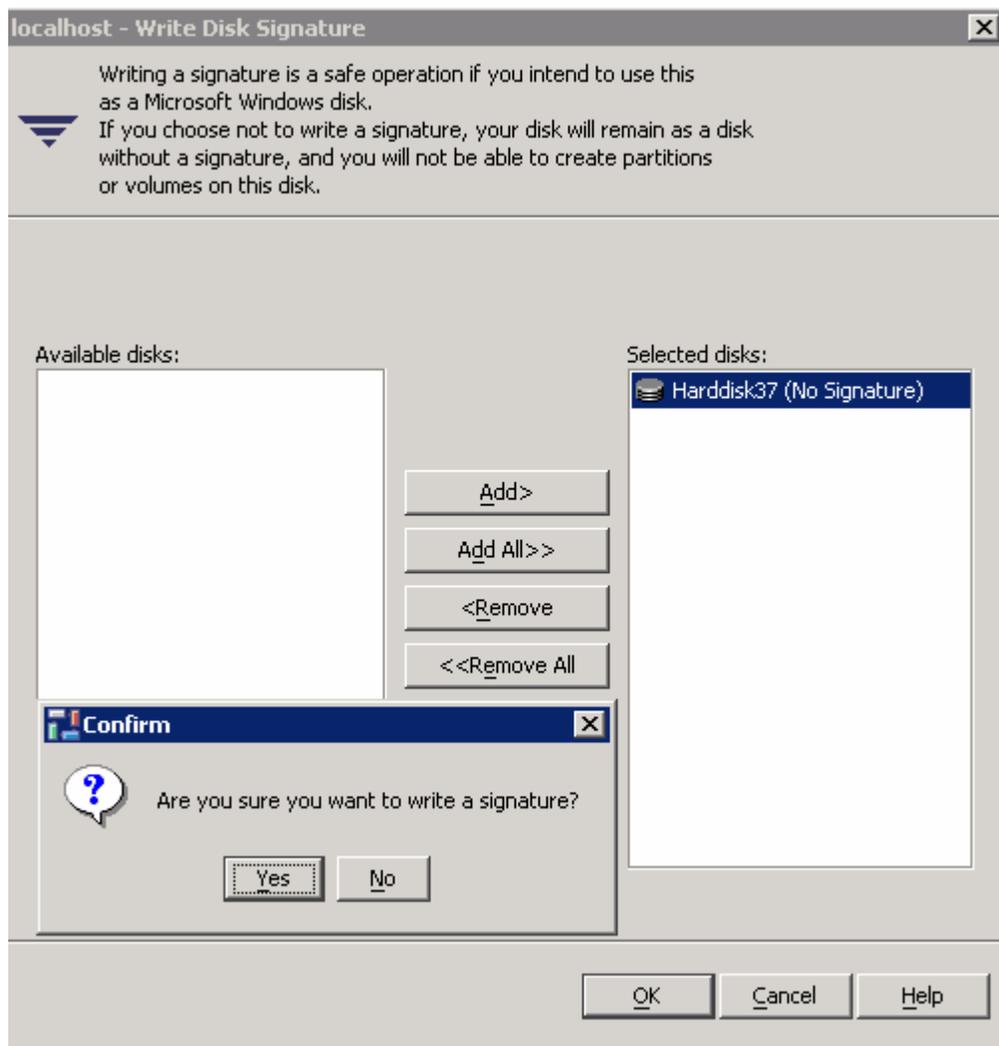
After the hardware was set up, the system was configured. To configure the Exchange server, the VEA interface was used to create dynamic disks and then create dynamic volumes to match the predefined storage layout for the Exchange server. For more information, refer to the VERITAS Storage Foundation 4.1 for Windows Solutions Guide for Microsoft Exchange Server at <http://www.veritas.com/Products/www?c=product&refId=31>.

Using the following procedure, a new dynamic disk group was created and all disks for Exchange storage were placed in the new disk group through the VEA user interface:

1. Write signatures to all disks.

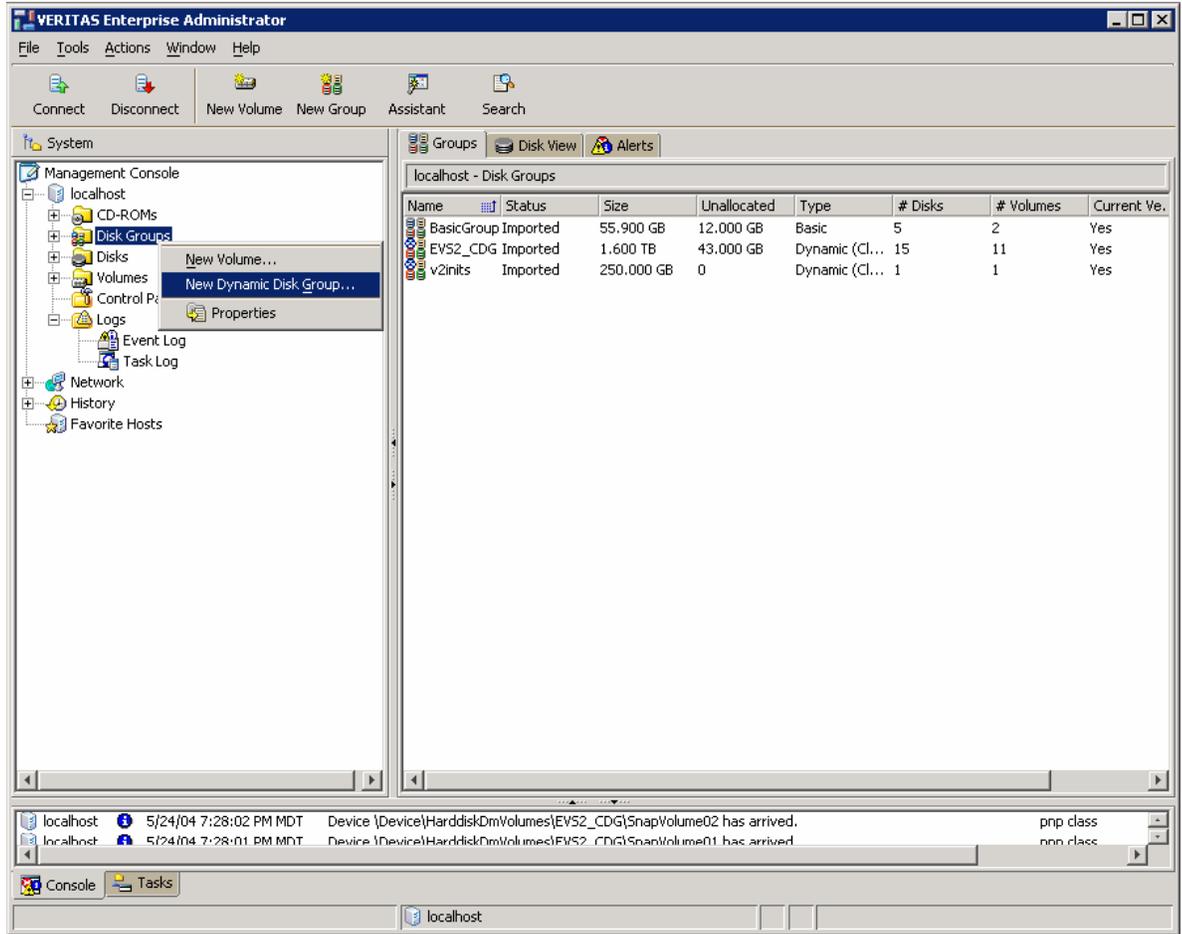
Right-click a disk, and select write signature. Add all disks, click OK, and then click Yes to write the signatures. See Figure 7.

Figure 7.



2. Right-click the Disk Groups folder, and then select New Dynamic Disk Group to start the disk group wizard. See Figure 8.

Figure 8.

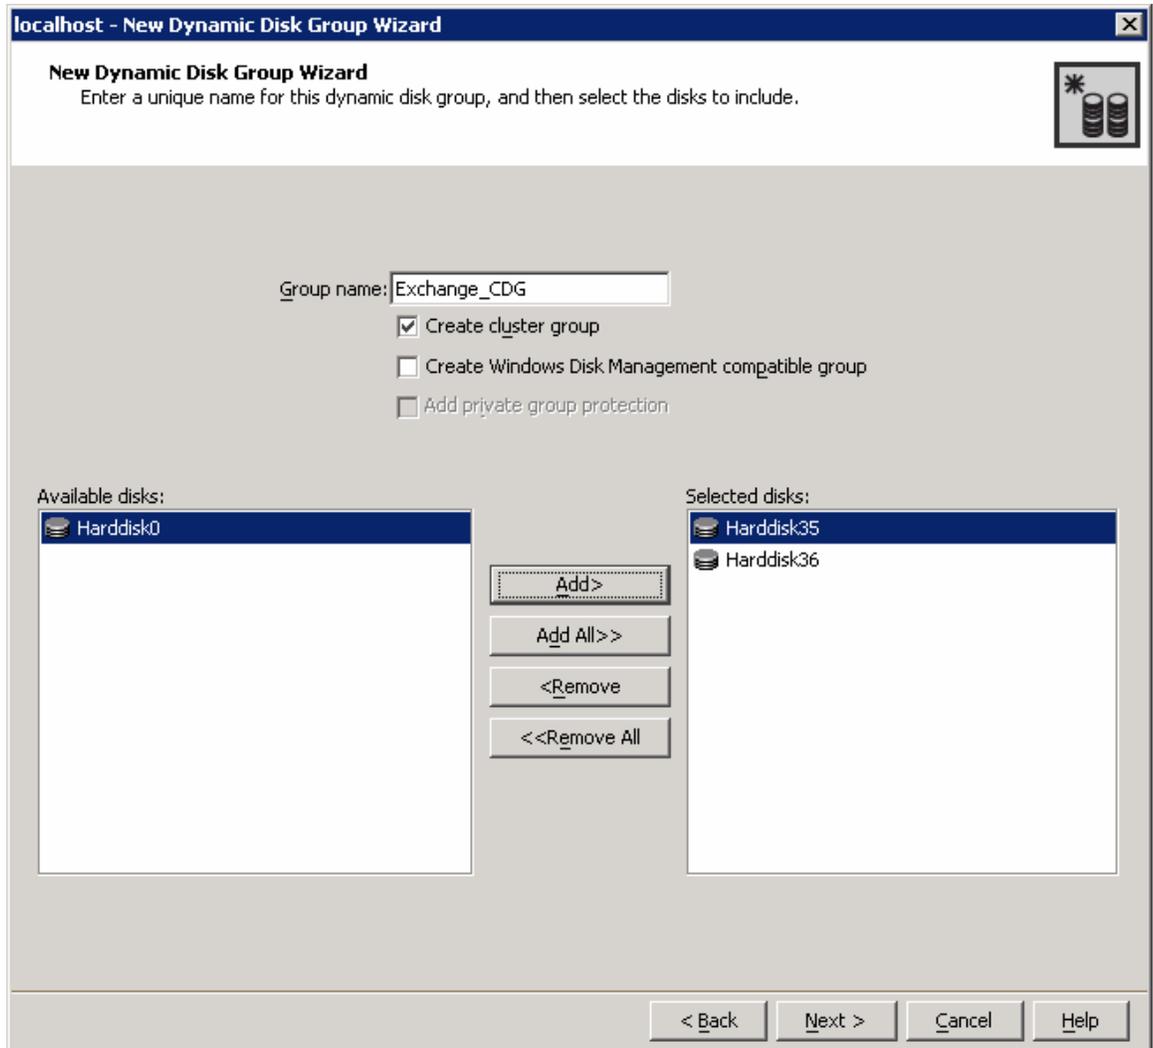


3. Add all disks that are part of the Exchange server. Do not add the system disk (C).

Note

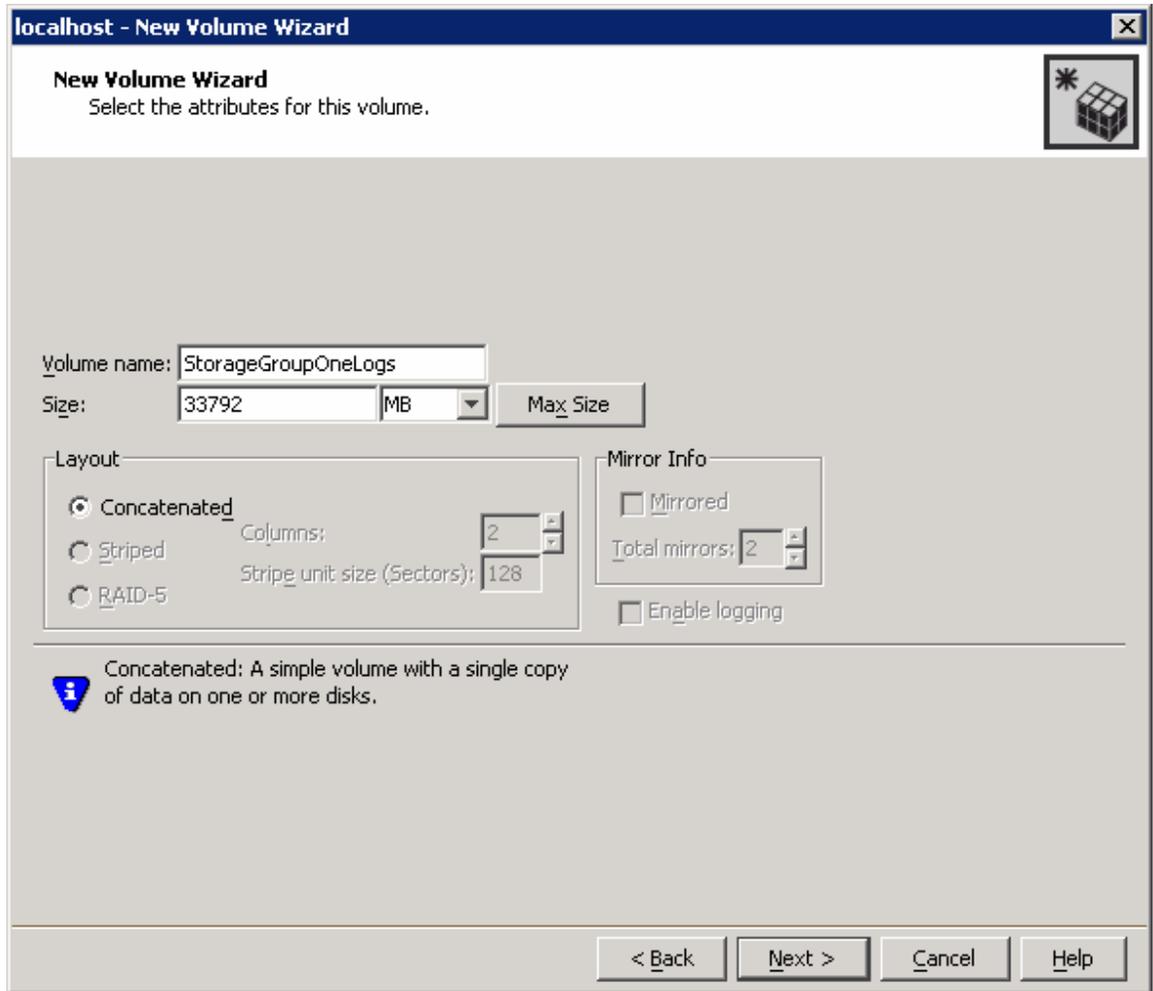
In a clustered environment, select cluster disk group, and add all disks for the Exchange Virtual Server. See Figure 9.

Figure 9.



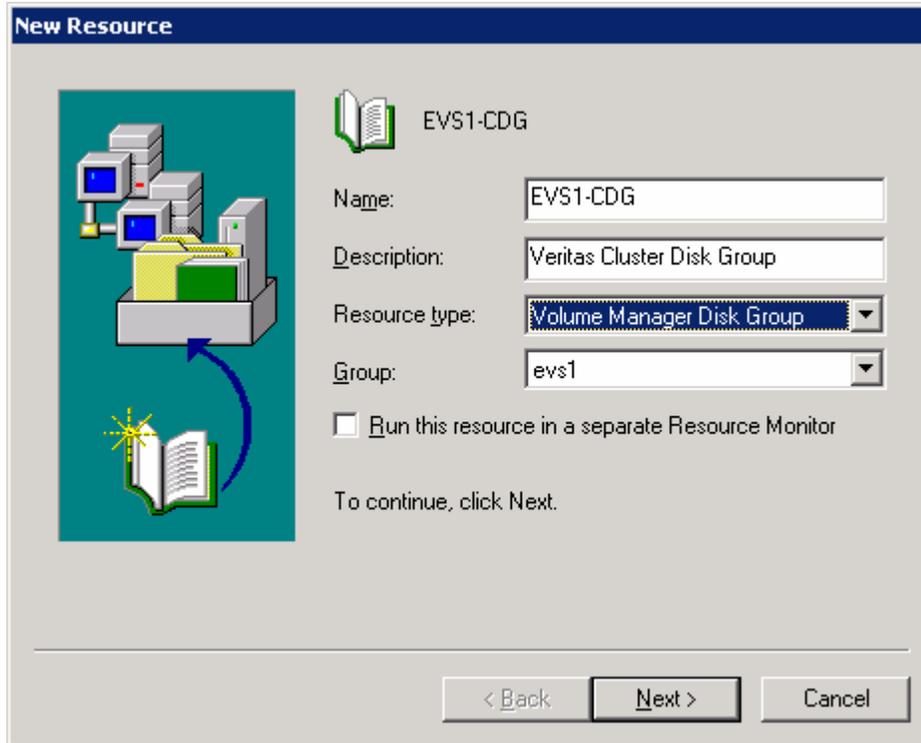
4. Complete Disk Group creation by following the wizard. After the Dynamic Disk Group Wizard completes, the disk group is created, and it contains the selected disks.
5. Use these disks to create dynamic volumes for the storage group database and transaction log files.
6. Right-click the Dynamic Disk group just created, and select new volume to start the new volume wizard. See Figure 10.
7. To create volumes that match the storage configuration desired, select each disk manually and create a separate volume for each disk.
8. Select an appropriate volume name, and enter the size of the volume. Do not select Max Size. Enter the correct size in GB of the volume. Leave space on the hard disk to allow for the Disk Change Object (DCO) volume (5 to 10 MB).

Figure 10.



9. In a clustered environment, the Disk Group for each Exchange Virtual Server must be added as a cluster resource. Using Cluster Admin, select new resource.
10. Select the resource type Volume Manager Disk Group, and then select the disk group created for the Exchange Virtual Server (see Figure 11). The Exchange System Attendant uses this resource as its physical disk dependency.

Figure 11.



11. For each volume, issue the Snap Start command to start creating the mirrors.

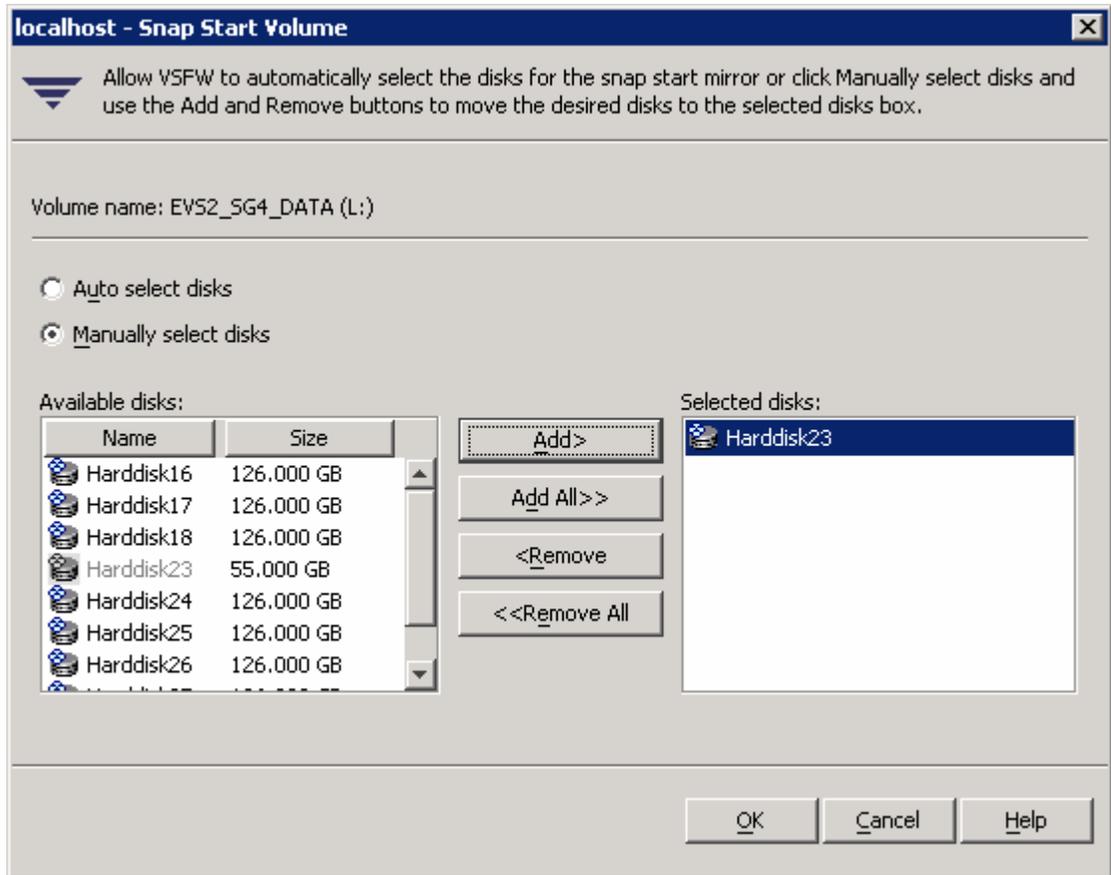
Best Practice: Because mirror synchronization can be a lengthy process, HP recommends creating the mirrors (known as Snap Start when snapshots are being performed) immediately after creating the volumes.

Use the command line interface (CLI) commands or the VEA user interface to create the mirrors, and perform one of the following:

- Using VEA, select each volume, and then highlight Snap and click Snap Start. Select the disk on which the snapshot is to be created. Click Add, and click OK. See Figure 12.
- Using the command line, enter the command `vxassist snapstart [drive letter of volume to be snapped] [disk snapshot is to reside on]`. For example, enter:

```
vxassist snapstart L: HardDisk23
```

Figure 12.



12. Using the Exchange System Manager, configure the storage groups by assigning the appropriate databases to the volumes. Mount all stores.

The Exchange server is now configured. The mirror will exist as a duplicate of the production volume until a snapshot is taken.

How the Snap commands work

The Snap Start command is used to create a mirror. After the command has completed, a mirror is established between the original volume and a secondary volume. This mirror will remain attached and contain the current contents of the original volume until the Snap Shot command is issued. (The Snap Shot command breaks the connection.) For more information, see Appendix C—Snap commands.

After the Snap Shot command has completed, the secondary volume becomes a separate persistent volume. By using the Snap Back command, the snapshot volume can be resynchronized with the original volume. Resynchronization can be bidirectional. A DCO volume that tracks changed blocks is maintained, which allows for a faster resynchronization because only changed blocks are written.

With the VxSnap CLI tool, VSS storage group snapshots can be restored. This tool integrates the metadata file created during the snapshot and ensures that all associated volumes are restored. During a VSS restore operation, the current volume is substituted with the snapshot. The original volume is then resynchronized using the snapshot.

Creating snapshots of the storage group

VSS-enabled snapshots can be generated for each volume using the VEA user interface or VxSnap CLI utility. However, to create VSS snapshots at the storage group level, you must use the VxSnap requestor.

Snapshots of the storage group are created to ensure that all the required volumes are snapped at the same time. A storage group can contain multiple databases, and all the databases in a storage group share the same set of transaction logs. Using the VxSnap VSS requestor to take snapshots of the storage group ensures that all required volumes are included. The metadata file contains the volume information for the snapshot of the specified storage group. During restoration, the metadata file is used to ensure the correct volumes are restored.

To create a snapshot of a storage group, VxSnap is used in the following form:

```
VxSnap -x metadatafilename.xml create writer="Microsoft Exchange Writer"  
component="Storage Group Name"
```

```
Vxsnap -x EVS1sg1snap.xml create writer="Microsoft Exchange Writer"  
component="EVS1 SG1"
```

Volume name and drive letter can be assigned to the snapshot volume at creation time. If none are assigned, the drive is not assigned a letter (presented), and a default volume name is given (SnapVolume_X).

```
VxSnap -x metadatafilename.xml create writer="Microsoft Exchange Writer"  
component="Storage Group Name" source=[parent drive  
letter]/driveletter=[DL for snapshot volume]/newvol=[name for newvolume]
```

```
Vxsnap -x EVS1sg4snap.xml create writer="Microsoft Exchange Writer"  
component="EVS1 SG4" source=H:/newvol=evs1sg4_datasnap1  
source=H:\log/newvol=evs1sg4_logsnap1
```

Snapshots can be created manually, but during testing, batch files were created to automate the process. These batch files were called using the Windows Task Scheduler but can also be launched with post-processing commands from a backup application.

Verifying integrity

VSS snapshots are copy backups of the database. For this reason, and because it is critical to ensure the validity of a volume that will be used as a restore target, it is important to verify the quality of the databases and logs in the snapshot set. The Exchange utility eseutil was used for integrity checking. After the snapshot is created using the VEA GUI, a drive letter must be added to present to the host. See Figure 13.

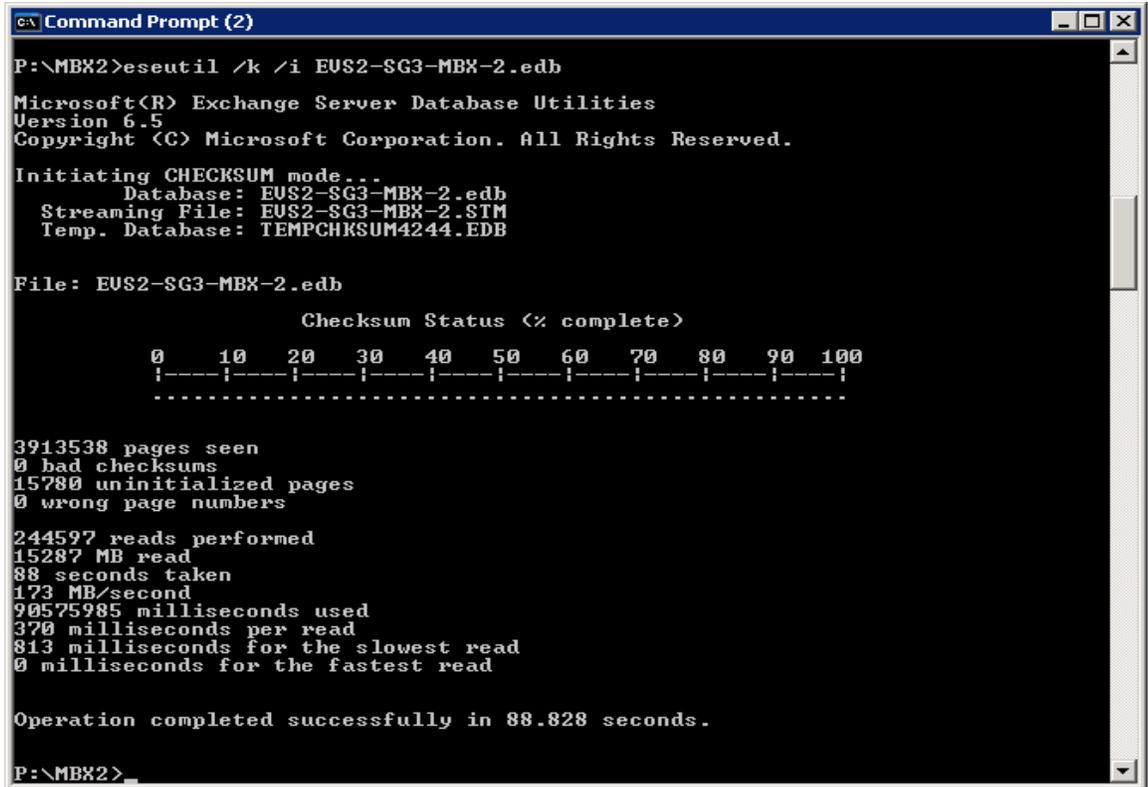
Results: Average completion time for the integrity check for 65-GB databases was 6 minutes. The average time for 15-GB database was 90 seconds. No errors were generated during testing.

To add a drive letter, the following procedure was performed:

1. Highlight and right-click the snapshot volume.
2. Select File System>Change Drive Letter and Path. See Figure 13.
3. Add a drive letter from the selection list, and click OK. The volume is now accessible using the selected drive letter.

During testing, integrity checking was performed on the host and had minor performance requirements. Because snapshot volumes were on separate spindles, verification did not impact the production volume. After the utility is completed, the drive letter can be removed using the VEA interface.

Figure 14.



```
Command Prompt (2)
P:\MBX2>eseutil /k /i EUS2-SG3-MBX-2.edb
Microsoft(R) Exchange Server Database Utilities
Version 6.5
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating CHECKSUM mode...
  Database: EUS2-SG3-MBX-2.edb
  Streaming File: EUS2-SG3-MBX-2.SIM
  Temp. Database: TEMPCHKSUM4244.EDB

File: EUS2-SG3-MBX-2.edb

          Checksum Status (% complete)
0      10     20     30     40     50     60     70     80     90    100
|-----|-----|-----|-----|-----|-----|-----|-----|-----|
.....

3913538 pages seen
0 bad checksums
15780 uninitialized pages
0 wrong page numbers

244597 reads performed
15287 MB read
88 seconds taken
173 MB/second
90575985 milliseconds used
370 milliseconds per read
813 milliseconds for the slowest read
0 milliseconds for the fastest read

Operation completed successfully in 88.828 seconds.

P:\MBX2>
```

Restoring storage groups

During the restoration process, the server file system is redirected to the snapshot volumes, which support the storage group. Because no data transfer from tape or disk is required for restoration, Exchange databases can be mounted immediately using the snapshot volume. When the databases are mounted, the snapshot volume is used to service the required application I/O while background resynchronization completes. When resynchronization has completed, the volumes are mirrored and snapshot-ready.

Results: During all restoration testing, Exchange databases could be mounted in less than two minutes—in most cases, less than one minute. The time required to mount individual storage group databases for point-in-time restores was generally less than 30 seconds. The time required to mount the storage groups during roll-forward recovery depends on the number of transaction logs.

Point-in-time recovery

By leveraging and managing snapshot technology to create point-in-time recovery volumes, you can fully recover database volumes in minutes. If log files become corrupted or damaged, you cannot assume that any of the stores are left in a consistent state. You must restore all the storage group stores, along with the corresponding log files.

The following procedure was used during testing to execute a point-in-time restore. Renaming or deleting the databases is not required but was done during testing to simulate a catastrophic failure.

1. Dismount all database stores in storage group.
2. Rename or delete database and log files. (This step was for test simulation purposes only.)
3. Using VxSnap, restore the storage group. From the command line or batch file, enter:

```
Vxsnap -x [metadatafilename].xml restore
```

Best Practice: When restoring more than one volume, the `-b` argument should be used for the restore command. This argument executes the resynchronization of the volumes as a background process.

4. Mount all stores.

After the stores are mounted, they are in the same state as they were during the last full VSS snapshot.

In a clustered environment, it is possible that the Exchange Virtual Server will move to another node. To allow for restoration, be sure to place the metadata file on a volume in the cluster disk group. The full path to the location must be specified in the VxSnap create and restore commands.

Roll-forward recovery

Roll forward occurs when there is a loss or corruption of databases, but the logs remain intact. In this scenario, the database files are restored, and the current logs are replayed, bringing the system to the current state. The following procedure was performed during this phase of testing. Renaming or deleting the databases is not required but was done during testing to simulate a catastrophic failure.

1. Dismount all database stores in storage group.
2. Rename or delete database files only.
3. Execute VxSnap restore, selecting database volumes only. Volume name or drive letter can be used to specify the volumes. From the command line or batch file, an example command is:

```
Vxsnap -x [metadatafilename].xml restore [volume names] -b
```

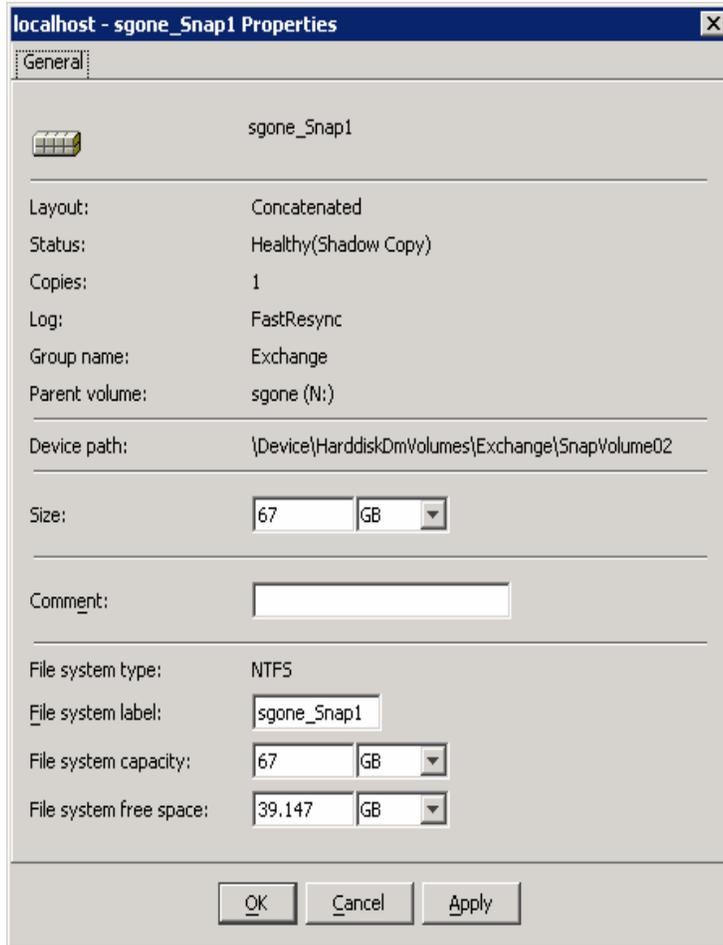
If storage groups are configured to use more than one volume for database files, each volume name must be specified.

```
Vxsnap -x [metadatafilename].xml restore -b
```

```
vol=\Device\HarddiskDmVolumes\[dynamic disk group name]\[snapshot  
volumename]
```

If default snapshot volume names are used, the name is displayed in the device path on the properties page of the snapshot. See Figure 15.

Figure 15.



In this example, the command for roll-forward recovery of Storage Group One would be:

```
Vxsnap -x sg1snap.xml restore vol=\\Device\\  
HarddiskDmVolumes\\Exchange\\SnapVolume02
```

Figure 16.

```
C:\scripts>restoresg1  
C:\scripts>vxsnap -x sg1snap.xml restore vol=\\Device\\HarddiskDmVolumes\\Exchange\\  
SnapVolume02  
Please wait while restore operation is going on.  
Successfully restored the volumes.  
C:\scripts>_
```

4. Mount all stores. Log replay begins, and completion depends on the number of logs present.

Appendix A—Detailed test results

The performance numbers listed are not benchmarking results. The systems were not configured for MMB3 benchmarks and do not adhere to the mandatory configuration parameters established by Microsoft for benchmark testing. These numbers are specific to these system implementations and are provided for reference only. The MMB3 profile was used to provide data that could be compared to expected performance in optimal configurations. The expected results for some counters reflect established pass/fail criteria and are shown for reference only.

Three test runs were performed. The first run was performed without mirrors attached to the volumes to establish a baseline. The second run had mirrors attached for the duration of the test to measure the impact of mirroring. The third run was started with mirrors attached, and snapshots were created during the four-hour steady state period. With mirrors attached, there is an expected performance degradation to disk writes. The impacts observed during testing were minimal, and overall performance was within acceptable levels.

The following tables contain the results gathered from Performance Monitor and reflect the averages during the four-hour steady state. Results are listed by storage option. The active nodes are reported for each Exchange Virtual Server (EVS) instance.

Table 6. Test run 1—Baseline configuration

	EVA5000 EVS1	EVA5000 EVS2	MSA1000	SA6402	Expected Results
Processor					
% processor time	20.6	21	3.4	12.2	<90
Memory					
Available MB	1,972	2,086	1,995	750	>500 MB
Pool non-paged bytes	42,163,545	35,355,999	19,516,702	31,076,627	<80,000,000
Pool paged bytes	161,229,121	96,037,784	35,388,284	68,825,433	<180,000,000
MS ExchangeIS Mailbox					
Receive queue size	0	0	0	0	<500
Send queue size	18	11	2	7	<500
Avg del time	20,550	54,009	13,362	4,276	
SMTP Server					
Cat. queue length	0	0	0	0	<10
Local queue length	37	21	2	5	<1,000
Remote queue length	38	2	4	680	<1,000
Logical Disk					
Avg queue length	20.7	15.9	26.2	13.9	

	EVA5000 EVS1	EVA5000 EVS2	MSA1000	SA6402	Expected Results
Current queue length	20	15	20	12	
Avg sec read	0.014	0.011	0.010	0.008	<100
Avg sec write	0.003	0.003	0.140	0.057	<100
Physical Disk					
Avg disk queue length	20.6	15.9	26.2	13.9	
Current disk queue length	20	15	20	12	

Table 7. LoadSim scores

	EVA5000	MSA1000	SA6402
95%	507	513	310

Table 8. Test run 2—All volumes mirrored

	EVA5000 EVS1	EVA5000 EVS1	MSA1000	SA6402	Expected Results
Processor					
% processor time	19.6	22.2	4.2	6.3	<90
Memory					
Available MB	1,982	2,053	2,244	821	>500 MB
Pool non-paged bytes	42,241,114	38,175,485	31,064,237	20,919,613	<80,000,000
Pool paged bytes	144,312,900	94,714,224	56,264,684	44,101,555	<180,000,000
MS ExchangeIS Mailbox					
Receive queue size	0	0	0	0	<500
Send queue size	26	26	1	1	<500
Avg del time	29,157	103,668	10,512	5,396	
SMTP Server					
Cat. queue length	0	0	0	0	<10
Local queue length	50	50	1	1	<1,000
Remote queue length	8	5	0	0	<1,000
System					

	EVA5000 EVS1	EVA5000 EVS1	MSA1000	SA6402	Expected Results
Processor					
Processor queue length	0	0	0	0	
Context					
Context switches/sec	9,315	11,248	3,406	2,735	
Logical Disk					
Avg queue length					
Avg queue length	29.5	25.5	4.9	6.6	
Current queue length					
Current queue length	27	25	5	5	
Avg sec read					
Avg sec read	0.016	0.015	0.007	0.006	<100
Avg sec write					
Avg sec write	0.016	0.012	0.013	0.051	<100
Physical Disk					
Avg disk queue length					
Avg disk queue length	37.8	31.8	6	9.7	
Current disk queue length					
Current disk queue length	35	29	6	7	

Table 9. LoadSim scores

	EVA5000	MSA1000	SA6402
95%	739	375	312

Table 10. Test Run 3—Included mirrors and VSS snapshots

	EVA5000 EVS1	EVA5000 EVS2	MSA1000	SA6402	Expected Results
Processor					
% processor time	11.7	10.4	3.9	6.9	<90
Memory					
Available MB					
Available MB	2,016	2,066	2,221	838	>500 MB
Pool non-paged bytes					
Pool non-paged bytes	37,975,711	33,364,093	34,467,446	21,007,959	<80,000,000
Pool paged bytes					
Pool paged bytes	123,080,815	58,056,121	62,512,145	68,706,684	<180,000,000
MS ExchangeIS Mailbox					
Receive queue size					
Receive queue size	0	0	0	0	<500
Send queue size					
Send queue size	2	30	1	2	<500
Avg del time					
Avg del time	12,778	0	11,454	4,676	
SMTP Server					

	EVA5000 EVS1	EVA5000 EVS2	MSA1000	SA6402	Expected Results
Cat. queue length	0	0	0	0	<10
Local queue length	2	9	1	0	<1,000
Remote queue length	0	0	0	2	<1,000
System					
Processor queue length	0	0	0	0	
Context switches/sec	6,900	7,630	3,312	3,088	
Logical Disk					
Avg disk queue length	8.4	5	4	6.3	
Current disk queue length	8	5	4	6	
Avg disk sec/read	0.009	0.007	0.008	0.006	<100
Avg disk sec/write	0.004	0.003	0.006	0.048	<100
Physical Disk					
Avg disk queue length	9.1	5	4	7	
Current disk queue length	9	5	4	8	
Table 11. LoadSim scores					
	EVA5000	MSA1000	SA6402		
95%	423	382	401		

Appendix B—Test configuration BOM

Note: Some part numbers listed may not represent actual tested configuration because of equipment availability.

Table 12.

	Quantity	Part number
ProLiant server with direct attached storage (500 users)	1	
ProLiant ML530 G2 Intel Xeon processor 3.0 GHz 1024-MB DDR SDRAM (1) NC7770 PCI-X Gigabit NIC	1	325246-001
NC7770 PCI-X Gigabit server adapter	1	244948-B21
Smart Array 6402/128ww	1	273915-B21
72.8-GB 10,000-rpm U320 Hot Pluggable hard drive WW	12	286714-B22
36-GB 10,000-rpm U320 Universal hard drive	8	286713-B22
ProLiant server with SAN-attached storage (1,000 users)		
ProLiant DL580 G2 Intel Xeon processor 2.4GHz 2048-MB DDR SDRAM (1) NC7770 PCI-X Gigabit NIC Intel Xeon MP X 2.4GHz, 2-M processor option kit	1	180285-001 325254-B21
2048-MB PCI1600 Registered ECC SDRAM memory kit (4x512 MB)	2 per server	202171-B21
NC7770 PCI-X Gigabit server adapter	1 per server	244948-B21
FCA 2104 2-Gb FC HBA	1 per server	245299-B21
MSA1000 4.24fw	1	201723-B21
72.8-GB 10,000-rpm U320 Pluggable hard drive WW	14	286714-B22
36-GB 10,000-rpm U320 Universal hard drive	14	176496-B22

	Quantity	Part number
ProLiant server with SAN-attached storage (5,000 users)		
Exchange back-end servers	3	
ProLiant DL580 G2	3	325134-001
Intel Xeon processor 3.0 GHz		
2048-MB DDR SDRAM		
(1) NC7770 PCI-X Gigabit NIC (embedded)		
Intel Xeon MP X2.0 GHz, 2-MB processor option kit		325254-b21
2048-MB PCI1600 Registered ECC SDRAM memory kit (4 x 512 MB)	1 per server	202171-B21
Compaq NC7770 PCI-X Gigabit server adapter	2 per server	244948-B21
36-GB 10,000-rpm U320 Universal hard drive (1 inch) for OS	4 per server	176496-B22
FCA 2104 2-Gb FC HBA	2 per server	245299-B21
HP StorageWorks Secure Path 4.0C for Windows (5 license)	1 per server	231292-B25
Microsoft Windows Server 2003 Enterprise Edition	1 per server	third-party
Microsoft Exchange Server 2003 Enterprise Edition	1 per server	third-party
VERITAS Storage Foundation 4.1 for Windows including the FlashSnap option	1	third-party
Storage—EVA	1	
EVA5000—2C12D-C: Controller assembly with two HSV110 controllers, 12 M5314 3U dual-redundant FC Loop 14-bay disk enclosures, four 12-port FC loop switches, 42U graphite storage cabinet	1	283199-B24
VCS PKGV3.010 dual HSV controller, base controller software	1 per array	250203-B25
72.8-GB 10,000-rpm dual-port 2-GB FC-AL 1-inch drive	168	238921-B22
Microsoft Windows NT®, Windows 2000, Windows 2003 Server Edition Kit 3.0b for EVA	1	333685-B21
HP StorageWorks SAN Switch 2/16	2	322118-B21
Short Wave Optical Transceivers	16 per switch	221470-B21
HP OpenView Storage Management Appliance III, including HP OpenView Storage Management Appliance 2.1	1	189715-003

Appendix C—Snap commands

Snap Abort—Permanently removes the snapshot volume

Snap Back—Re-attaches the snapshot to the original volume and resynchronizes the data contained on the volume

Snap Clear—Permanently removes the association between the original volume and the mirror

Snap Shot—Breaks the connection between the original volume and the mirrored volume, leaving the mirror as a complete point-in-time replica of the original

Snap Start—Creates the mirror, a volume created on the selected hard disk that is attached to the original volume

For more information

HP Storage Solutions for Microsoft Exchange Server 2003

<http://www.hp.com/solutions/microsoft/exchange/storage>

HP ActiveAnswers

<http://www.hp.com/solutions/activeanswers>

VERITAS

Technical Review Guide, datasheet, and other materials for Storage Foundation for Windows

<http://www.veritas.com/Products/www?c=product&refId=31>

Microsoft Exchange Server

<http://www.microsoft.com/exchange>

© 2004 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Intel and Xeon are trademarks or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

5982-7610EN, 09/2004

