

February 2000
11Y1-0200A-WWEN
Prepared by:
Storage Division
Compaq Computer Corporation

Contents

Executive Summary3
 Conclusion and
 Recommendation.....3
Uptime Booster Solution.....4
 Solution Components.....4
Configuration
Recommendations7
 Recommendations for BCV
 clones.....8
Expanded Exchange
Configurations.....12
Appendix A: Sample scripts.....14
 Script process.....14
 Exchange script fragments.....14
 Backup script fragments.....18
Appendix B: Overview of
SANworks Enterprise
Volume Manager.....19
 EVM graphical user interface...19
 EVM job interaction with
 external applications20
 EVM storage unit replication
 features20
 EVM general requirements.....21
 Operation of EVM in a
 Microsoft Exchange
 environment.....23
Appendix C: Overview of
Microsoft Exchange Server.....24
 Managing the Exchange
 Server 5.5 database.....24
 Common support issues for
 Microsoft Exchange Server.....25
Appendix D: Test disclosure.....27
 Workload methodology.....27
 Test Procedure28
 Test results.....30
 Descriptive Terms31
Appendix E: Related
documents.....32
 Compaq ActiveAnswers32
 Compaq StorageWorks32
 Compaq White Paper Index.....32
 Microsoft Exchange Server.....32
 Books32

Microsoft Exchange Backup/Restore Solution – Uptime Booster

Abstract: This Compaq solution for Microsoft Exchange 5.5 provides administrators with the ability to back up Microsoft Exchange data with minimal disruption to Exchange services, as well as recover Exchange data very rapidly in the event of corruption. This solution employs *Compaq SANworks Enterprise Volume Manager* software and customizable scripts to create and manage Business Continuance Volumes of the Exchange Information Store using cloning technology. *SANworks Enterprise Volume Manager* works in combination with industry tape backup solutions to provide a complete, end-to-end backup and restore solution for the Microsoft Exchange Server database. This white paper discusses the best practices for deploying the *SANworks Enterprise Volume Manager* in an Exchange Server environment. Performance data using the Microsoft Load Simulator on a reference platform is included. In addition, this paper describes how to configure Microsoft Exchange and deploy *SANworks Enterprise Volume Manager* as part of the overall disaster recovery plan.

Notice

The information in this publication is subject to change without notice and is provided “AS IS” WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

This solution is only valid for StorageWorks RA 8000 FC and ESA 12000 FC controllers. The RA 8000 FC tests were conducted with default settings and 128 MB cache. The Legato NetWorker, Computer Associates ArcServeIT, and VERITAS Backup Exec tests were run with default settings.

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

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

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

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

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

©2000 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Microsoft Exchange Backup/Restore Solution – Uptime Booster
White Paper prepared by Storage Division

First Edition (February 2000)
Document Number 11Y1-0200A-WWEN

Executive Summary

Enterprises of all sizes have come to depend on Microsoft Exchange for mission-critical messaging and collaboration. Therefore, these enterprises require reliability and continuous availability of Exchange data to serve customers and maintain internal productivity. The fact is, downtime resulting from lengthy backup and restore processes can be very costly in terms of lost revenue, poor customer satisfaction, or missed competitive opportunities. That makes reducing downtime a key component in developing a reliable messaging infrastructure.

While Microsoft Exchange is generally quite stable, Exchange Server databases (Information Stores) can occasionally become corrupted. Daily tape backups of Exchange databases are part of Microsoft's recommended disaster recovery strategy. But there is a direct correlation between the size of the Exchange Information Store and the amount of time needed to back up and restore it. As a consequence, when using an online tape backup solution, large Exchange Information Stores can degrade Exchange services for an extended period of time. Moreover, in the event the production database becomes corrupted, a recovery may require a restore from the last-known-good copy of the database file—again adversely affecting Exchange services. The amount of downtime required to restore a large Exchange database from tape would result in a long period of downtime, which would be unacceptable to many users.

The Microsoft Exchange Backup/Restore Solution – Uptime Booster minimizes Exchange downtime by using the *Compaq SANworks Enterprise Volume Manager* software and an Open Storage Area Network (SAN) configuration to facilitate a minimally disruptive backup and fast recovery of Exchange Information Stores. With *SANworks Enterprise Volume Manager*, users can create point-in-time copies of the Exchange Information Store called Business Continuance Volumes (BCVs) using cloning technology. BCV clones can be backed up to tape off-line across the Open SAN, avoiding lengthy impact to Exchange services. If corruption occurs within the Exchange Information Store, a BCV clone can then be used to restore the Exchange database in minutes, as opposed to hours typically required when restoring from backup tapes. *SANworks EVM* also supports the creation of snapshots with certain *StorageWorks* controller firmware. A BCV snapshot (if available) can be used to restore the Exchange database very rapidly.

With the Microsoft Exchange Backup/Restore Solution – Uptime Booster, Compaq provides customers with a complete set of hardware, software, and supporting utilities to achieve a highly available Exchange-based messaging environment. This unique backup and restore solution for Microsoft Exchange 5.5 supports the Compaq Enterprise Network Storage Architecture (ENSA) and delivers the high availability enterprises look for in a *Compaq NonStop™* eBusiness solution.

Conclusion and Recommendation

Administrators can use the Microsoft Exchange Backup/Restore Solution – Uptime Booster to enhance their backup strategy and avoid the downtime of restoring from a tape backup. BCV clones are not, however, a replacement for tape backups. Therefore, the Uptime Booster has been developed as a complete solution to facilitate backup of clones to tape, either directly or across an Open SAN. This solution is an effective and tested backup and restore option for administrators faced with extended degradation of Exchange services during backup or experiencing unacceptable periods of downtime while restoring Exchange data from tape.

Note: Log Management is important when backing up the Exchange Information Store. For more information, please refer to the following document for the recommended procedures:

<http://www.microsoft.com/exchange/55/whpprs/BackupRestore.htm>

Implementing the Microsoft Exchange Backup/Restore Solution – Uptime Booster requires additional capacity planning and proactive monitoring of the storage system. This paper contains recommendations for configuring the Exchange disk subsystem and procedures for using BCV clones to back up Exchange databases, as well as for using BCV clones to restore the Exchange database. It also includes a reference platform and provides the results of extensive testing, including the overhead and benefits of implementing BCV clones, as well as the impact and performance of three backup products for Exchange: Legato NetWorker, VERITAS Backup Exec, and Computer Associates ArcServe/IT.

IMPORTANT: Best practices recommend the use of BCV clones for backup and restore.

Uptime Booster Solution

Backup and recovery of an Exchange database is a critical process to ensure availability of Exchange services. Failure to back up Exchange data is one of the leading causes of downtime for sites that require a restoration from the last-known-good copy. Through the use of BCV clones, the Microsoft Exchange Backup/Restore Solution – Uptime Booster helps ensure that users' critical Exchange data is protected and readily available for recovery.

From our extensive experience in storage technology and Microsoft Exchange, Compaq has developed a set of best practices, which include recommendations and procedures for implementing the Microsoft Exchange Backup/Restore Solution – Uptime Booster. These best practices address the primary component of the solution—the *SANworks Enterprise Volume Manager*—and include customizable scripts for Microsoft Exchange 5.5, which can be used to automate backup and restore operations.

Solution Components

Figure 1 below indicates the Uptime Booster components and their relationships.




Figure 1

Table 1 below provides a Bill of Materials that includes some of the products used in the Microsoft Exchange Backup/Restore Solution – Uptime Booster. This table is an example only; you can choose the products best suited to your needs, including products not listed here.

Table 1. Uptime Booster Sample Bill of Materials

Item	Item Descriptions	Base SKU
1	EVM Software for Window NT-Intel	160094-B21
2	ACS 8.5F Controller Software	128697-B21
3	OS Platform Software Kit for NT (SWCC)	380551-001
4	RA8000 Rack with 2 HSG80 Controllers	380670-B21
5	Hard Disk Drives, 9 GB, 10k rpm, UWSE	380588-B21
6	ProLiant NT Backup Server (PL1850R - typical)	153554-001
7	FibreChannel HBA (switched fabric), KGPSA-BC	380574-001
	FibreChannel Host Adapter Kit / PCI (loop only)**	223180-B21
8	FibreChannel Storage Switch, 16 port	380578-B21
	FibreChannel SAN Switch / 16 with Fabric Software	158223-B21
	FibreChannel Storage Switch, 8 port	380591-B21
	FibreChannel SAN Switch / 8 with Fabric Software	158222-B21
9	FC Optical Short Wave GBIC	380561-B21
10	FibreChannel Optical Cable	234457-B22
11	Tape Library UW SCSI Host Bus Adapter	157299-B21
	FibreChannel Tape Controller II***	152975-001
12	Compaq Batch Scheduler v2.0 for Windows NT	388934-001
13	TL891 Tape Library	120876-B22
	TL895 Tape Library	349350-B24
	ESL9326D Tape Library	146205-B23
14	Legato NetWorker	
	VERITAS Backup Exec	3rd Party P/N
	CA ArcServe/IT	

* Optional (use any ProLiant NT server). Does not support HSG80 connectivity.

** Supported on Tape Library loop only. Comes with FibreChannel cable and GBIC for connectivity.

*** Required with FibreChannel Attached Tape Library

Table Explanation

Item 1: EVM Software – EVM software is necessary on every client and server on the SAN that needs to communicate or coordinate communications with hosts or storage devices. One instance of EVM is designated as the EVM server. When using EVM for Backup, this might well be the same machine that is running the backup application. For example, a configuration with nine application servers and a single backup server, 10 instances of EVM are loaded.

Item 2: Controller Firmware – One instance of ACS is necessary for each HSG80 controller. ACS 8.5F should be used for clones. ACS 8.5S is only supported on switches and offers both snapshots and clones.

Item 3: OS Platform Kit includes drivers and installation for the specific operating system and host involved.

Item 4: RAID Array – Use of either a StorageWorks RAID Array 8000 or Enterprise Storage Array 12000 does not affect the solution. Both must have HSG80 controllers for the solution to operate (the HSZ80 controller is not supported). While cache memory option of 512 MB is required for snapshots other options such as, power and packaging options, chassis, power supplies and environmental monitoring options have little influence on solution and therefore are not discussed in this paper. The choice of optional additional RA8000 expansion cabinets does not affect the functionality of this solution. More disks will increase the availability of spare disks from which to create BCV. It is possible to have a volume in a storage sub-system that is so large that it does not leave enough remaining spare disks to create a BCV. Remember that there must be exactly enough spare disks to match the size of every volume made into a BCV. If volumes are not all made into BCVs at once then the spares can be reused provided each BCV is returned to it's initial state prior to the next being created.

Item 5: Hard Drives – The choice of disk type does not affect the functionality of this solution. Two factors to consider for disk type are rpm and the number of spindles. For applications like mail the access pattern is largely random. In this case, the reduced rotational latency of the 10,000 over the 7,200-rpm disk can make a significant performance improvement. For this type of workload it is often desirable to use more smaller disks (to achieve more parallelism) rather than use fewer larger disks. In extreme cases, this can lead to a situation where, to get the required performance, many disks with little space of each actually in use. This strategy assumes the space required is less than 54 GB for a 24 disk RA8000 which allows for spare disks for BCV creation.

Item 6 – Backup Server - The backup server must be powerful enough to maintain performance levels during a backup/restore operation. For our testing we used a 2-processor 500MHz PL1600 with 1GB of RAM. No server sizing information is presented in this document.

Item 7: - FibreChannel HBA – use the KGPSA for all switched fabric solutions and loop solutions which contain the HSG80 on the SAN. For loop configurations which only have the backup server and the tape device, the 223180-B21 HBA must be used (NT support only).

Item 8: SAN Interconnect – Switches allow for significantly better performance, loops (hubs) present significantly limited performance and connectivity options. Since they are newer, switches have limited supported configurations whereas loops are supported by a variety of components. Storage arrays and tape can reside on the same switched fabric, making for the possibility of a simpler, more extensible installation. Disk and tape can not exist on the same loop; instead, Dual loops are required.

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

Item 10: Cables – order the length needed. The basis part number is 234457 and the suffix B2x denotes the length. Suffix B21 is 2, B22 is 5, B23 is 15, B24 is 30 and B25 is a 50 meter cable. Each controller comes with one five meter cable. All other cables will have to be ordered.

Item 11: Library Adapters – SCSI attached tape libraries require a special Host Bus Adapters (do not use standard server HBA adapters). FibreChannel-attached tape libraries (being SCSI devices) require a special FibreChannel to SCSI bridge (FCTC).

Item 12: Compaq Batch Scheduler – Allows running Minimally-Disruptive Backups automatically at off-hours at a given time each night. The entry usually consists of a single Windows NT console command line calling the EVM script. Alternatively, this functionality can be accomplished by using the WinAT Windows NT command scheduler included in the Windows NT Resource Kit.

Item 13: Library – pick the library based on the number of tape drives (performance) and number of slots (capacity). Only certain libraries are supported on specific ISV solutions, switches, others on loops, others via SCSI. Table 2 below provides a list of library and application support.

Item 14: ISV Backup Application Software – Different configurations will have different ISV support. The backup job is invoked by EVM with a single command line. Sample backup application-specific command lines are supplied with EVM documentation.

Table 2. Library and Application software support based on configuration

Infrastructure	Backup Application	Tape Library
Disk–FibreChannel Switch	Legato NetWorker V5.51	TL895
Tape–EBS Switch		
Disk–FibreChannel Switch	Legato NetWorker V5.51	ESL9326D
Tape–Direct SCSI		TL895
		TL891
	VERITAS Backup Exec V7.3	ESL9326D
		TL895
		TL891
	CA ArcServe/IT V6.61	TL895
		TL891
Disk–FCAL	Legato NetWorker V5.51	ESL9326D
Tape–Direct SCSI		TL895
		TL891
	VERITAS Backup Exec V7.3	ESL9326D
		TL895
		TL891
	CA ArcServe/IT V6.61	TL895
		TL891
Disk–FCAL	VERITAS Backup Exec V7.3	TL895
Tape–EBS FCAL		TL891
	CA ArcServe/IT V6.61	TL895
		TL891

Configuration Recommendations

The Microsoft Exchange Backup/Restore Solution – Uptime Booster has been validated in performance characterization tests using specific test configurations (See Appendix D). For those tested configurations, there are a number of recommendations you should follow to ensure the best results from this solution.

Using a BCV clone as part of a total backup and recovery solution requires additional capacity planning when designing your disk subsystem. For an Exchange server, the disk subsystem for the Exchange database should be designed using RAID fault-tolerant settings.

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

For best results, the Exchange Information Store and directory should be located on a separate volume from the transaction logs. The transaction logs must be on their own separate dedicated spindles using RAID 1 or RAID 0 +1. The DIR.EDB should also be separate from the PRIV.EDB and PUB.EDB. In addition, when using BCV clones, you must have free space on the RAID array equal to the size of the volume containing the Exchange databases.

Recommendations for BCV clones

Creating a clone

A clone should be created while Exchange is running. Therefore, be sure to create a clone when traffic is light because the I/O necessary to synchronize the clone with the other members of the mirror set will affect Exchange performance slightly. Depending on database size, it takes time for the clone creation to complete. Testing indicates that the clone will be built at about 60 GB/hr. Wait for the cloning operation to complete before splitting off the clone.

To split off the clone for backup, the Exchange database files must be in a consistent state. During normal operations, however, these files are in an inconsistent state. To bring them into a consistent state, a graceful shutdown of Exchange is necessary. The shutdown of Exchange is required only for the few moments when the clone is split off; once the clone is split off, Exchange services should be restarted immediately.

The specific steps and processes to create a BCV clone include the following:

- Use the *SANworks EVM* to create a BCV clone while Exchange is running.
- Stop the Exchange Information Store and dependencies
- Split off the BCV clone of the database files and map a drive letter
- Start the Exchange Information Store and dependencies
- Determine and store the last consistent transaction log

These steps can be easily automated using tested, documented, and customizable scripts provided as part of the Uptime Booster solution. Sample scripts are included in Appendix A of this paper, or they can be downloaded from the following URL:

http://www.compaq.com/products/storageworks/menu_solutions.html

Let's take a look at these steps in more detail.

Step 1: Use SANworks EVM to create a BCV clone

While Exchange is running, use the features of the *SANworks Enterprise Volume Manager* to create a BVC clone of the Exchange database files. Wait for this to complete before continuing on to Step 2.

Note: Typical creation time for BVC clones is approximately 60 GB/hour.

Step 2: Stop the Exchange Information Store and dependencies

Breaking off a BCV clone requires that the Exchange Information Store and any dependencies are stopped. Dependencies include the Event Service, Internet Mail Service, and the Key Management Service. Stopping these services will affect any users who are active on the Exchange Server.

By stopping Exchange services gracefully, the database pages in memory are committed to the Exchange database files on disk. The database file is now in a consistent state and a BCV clone can be created. Any service interruption should be a planned and scheduled activity. Do not schedule the splitting off of a BCV clone during periods when high activity is expected for Exchange.

If an Exchange service has been running for a long time, it may take a long time to stop. This is especially true of the Information Store and directory, which need to flush transactions to the database and close threads when the services are stopped. Again, it is very important to stop the Exchange Information Store gracefully. If the Information Store is not stopped gracefully, the database will be in an inconsistent state. The Information Store will then go through a built-in startup recovery, which will replay the log files that have not been committed to database file.

The following optional component services in Table 3 depend on the Information Store service running. If installed on the server, these services must also be restarted after the Information Store service is stopped.

Table 3: Optional Component Services

Service	Object
Exchange Event Service	MSEExchangeES
Internet Mail Service	MSEExchangeIMC
Key Management Server	MSEExchangeKM

Step 3: Split off the new BCV clone of the database files

With the Exchange Information Store gracefully stopped, use the *SANworks Enterprise Volume Manager* to split off the new BCV clone from the parent disk hosting the Exchange database. (Refer to the sample script provided in Appendix A)

Step 4: Start the Exchange Information Store and dependencies

Restart the Exchange Information Store service and any dependencies.

Note: Any estimates provided for the amount of time users may be offline from Exchange do not include the time required for users to reconnect to their Exchange Server. Estimates include only the time required to stop the Information Store and dependencies, split off a BCV clone, and restart the Information Store and dependencies. Testing has shown these steps can be accomplished in less than two minutes. Specific test results are provided in Appendix D at the end of this document.

Step 5: Determine and store the last consistent Exchange transaction log

Identify the last consistent Exchange transaction log. To facilitate this process, we recommend using the “Create BCV clone” script provided in Appendix A. This script executes a process that determines the state of the Exchange database and the last consistent transaction log for the database. The last consistent log generation is retained for the restore process.

Backing up a clone to tape

The SANworks Enterprise Volume Manager provides scripts to execute backup of the BCV clone to tape. Specific scripts are available for Legato NetWorker, VERITAS Backup Exec, and Computer Associates ArcServe/IT. Specific commands to execute these scripts are provided in an EVM white paper available at the following URL

Recovering from a clone

Administrators can use BCV clones to recover either the full Exchange database or individual Exchange mailboxes. In addition, the BCV clone can later be used to recreate a parent disk or become the parent disk.

Recovering a full Exchange database

The process of recovering the full Exchange database from a BCV clone include the following steps:

- Check the integrity of the BCV clone
- Verify the Information Store transaction log
- Stop the Exchange Information Store and dependencies
- Map the BCV clone to the parent disk drive letter
- Rename or Delete the MDBDATA checkpoint file
- Run ISINTEG –PATCH
- Start the Exchange Information Store and dependencies

While the decision to implement Exchange recovery requires human intervention, the actual recovery process can be scripted. Tested, documented, and customizable scripts enable scheduled BATCH/COM processes to run necessary commands and utilities within *SANworks Enterprise Volume Manager* to restore the database files from a BCV clone. Sample scripts are provided in Appendix A at the end of this document, or they can be downloaded from the following URL: http://www.compaq.com/products/storageworks/menu_solutions.html.

Now let's examine the steps to recover the full Exchange database from a BCV clone in more detail.

Step 1: Check the integrity of the BCV clone

It is critical to integrity check the BCV clone prior to using it for a restore. If the BCV clone fails the integrity check, than the restore must use the last-known-good tape backup.

The ESEFILE utility is available on Service Pack 3 for Microsoft Exchange Server 5.5. This utility will report any inconsistencies of the physical table structure of the database.

Step 2: Verify the Information Store Transaction Log

Verification of the Information Store transaction log ensures that all the required transaction log files are available for the recovery process and that the transaction logs are sequential. The verification process also verifies that the database signature in the transaction log files is consistent with the signature in the database header. The recovery process will not work if any of these checks fails.

Step 3: Stop the Exchange Information Store and dependencies

After the recovery process has verified the BCV clone and the transaction log files, the next step is to stop the Information Store service, if it is still running.

Note: Any users connected to the Exchange Server when the Information Store and dependencies are stopped will be required to reconnect after Exchange services are restarted.

Step 4: Map the BCV clone to the parent disk drive letter

The BCV clone is currently mounted to a drive letter that is not correct for the database path. This step involves unmapping both the BCV clone and the parent disk from their respective drive letters. The BCV clone can now be mapped to the drive letter that was previously used by its parent. Verify that the drive letter is the correct path for the database file.

Note: If a “disk is in use” message appears when attempting to unmap a drive letter, the server has a process running that needs to be disabled. A common situation is that the server may have Performance Monitor disk counters enabled (diskperf -y). Using the command line prompt, type diskperf -n. The server must then be rebooted.

Step 5: Rename or Delete the MDBDATA checkpoint file

Replaying transactions after a power outage or similar ungraceful shutdown is known as a “soft recovery.” This situation requires that the database files be updated from the transaction logs. Each time the Information Store service is started, the checkpoint file (EDB.CHK) is read to determine the starting point of uncommitted transactions. If EDB.CHK is missing, then the database service will read all the transactions beginning with the oldest.

Soft recovery is an automatic process that occurs without an administrator. At the completion of the operation, the database is consistent and the Information Store can start normally. To fully recover operations that are stored in the transaction log files, the EDB.CHK file must be renamed or deleted. If it is not renamed or deleted, then the recovery process would start recovery at the checkpoint and skip any previous transaction log files. This could result in lost data and possibly introduce logical corruption.

The EDB.CHK file will be recreated when the Information Store starts.

Step 6: Run ISINTEG -PATCH

The Information Store Integrity Checker (ISINTEG -PATCH) is an Exchange Server Utility that must be used when restoring an offline copy of the Information Store.

Exchange uses indexes that are referred to as Globally Unique Identifiers (GUIDs). The GUIDs use the timestamp of the database creation as a base. All objects created in the database are allocated a GUID that is derived from the base. The ISINTEG -PATCH utility will restore the GUIDs for the BCV snapshot by first validating the base GUID and then updating the store so that all new GUIDs are created correctly.

Note: The time required to run ISINTEG -PATCH is dependent on the number of transaction log files, CPU, disk subsystem, and memory.

Step 7: Start the Exchange Information Store and dependencies

Start the Exchange Information Store and any optional dependencies that may be installed. Dependencies include the Event Service, Internet Mail Service, and the Key Management Service.

Recovering an individual mailbox

A time-tested and reliable procedure for recovering individual Exchange mailboxes that have been deleted or corrupted is the use of a recovery server. With this procedure, the Information Store is restored from backup to a recovery server that is separate from the production mailbox server. The Exchange services are then started on the recovery server and the individual mailbox is accessed. The mailbox can be exported to a Personal Store (PST) file. This PST file can then be used to recover an individual mailbox on the production server.

The Uptime Booster solution simplifies and speeds this procedure by leveraging the Network Disk feature of the *SANworks Enterprise Volume Manager*. In this scenario, a production Exchange server running *SANworks Enterprise Volume Manager* can make BCV clones available to other servers on the network. For example, an Exchange server with a recent BCV clone of the Exchange Information Store would serve the BCV clone to another server. The recovery server would then mount the networked BCV clone as its Information Store. An individual mailbox or message could be recovered using this procedure, alleviating the need to restore from tape by making the BCV clone available immediately to the recovery server. For more information on the individual mailbox recovery procedures for Exchange Server 5.5, please see Microsoft TechNet and Exchange Server documentation.

Expanded Exchange Configurations

Beyond the typical configurations used in our performance characterization tests, there may be installations where some advantage may be taken by increasing the number of supported Exchange seats other than by implementing multiple instances of the Uptime Booster solution. In this section, we will review some of the relevant data that supports this expansion and the ramifications of expected performance.

General rules of Exchange configuration are that each Exchange seat consumes one disk IO/sec of capacity. The disks that we use in configuring the Uptime Booster solution are capable of about 140 IO/sec for the RAID 0+1 configuration and the Exchange load imposed on it. In this configuration, the *StorageWorks* RA 8000 holds 24 drives, the data volume is triple redundant (data plus mirror plus BCV), and there is a spare disk, as well as a mirror pair, for Exchange logging. As a result, the size of the data volume is seven disk drives. These seven disk drives, at 140 IO/sec, provide a capacity of 980 IO/sec, which supports approximately 1000 Exchange seats.

The *StorageWorks* RA 8000 and ESA 12000 support configurations of 24, 48, or 72 disk drives. Using the approach described above, these configurations would support 1000, 2000, or 3000 users respectively. The HSG controllers included in the *StorageWorks* RA 8000 and ESA 12000 are easily capable of handling the traffic intensity afforded by such an Exchange user load, but there will be a marked performance penalty during off-line backup operations using a BCV clone.

Since characterization of these extended configurations is beyond the scope of testing performed on the Uptime Booster, it is recommended that such expanded configurations be undertaken slowly with attention to acceptable Exchange seat performance at peak load times. Should a point be reached at which unacceptable performance is observed, an additional *StorageWorks* RA 8000 array should replace the ESA 12000 expansion cabinets. This will provide additional I/O capacity to alleviate the performance issue.

As in any performance characterization and estimation, the individual installation performance will vary largely based on the mix of Exchange users and their work habits. Many more light users can be supported by a particular Exchange configuration than heavy users. It is suggested that Compaq best management practices be used with Microsoft Exchange. White papers on this topic can be found at the following URL:

http://vcmpoapp02.compaq.com/ActiveAnswers/Global/en/solutions.1128/Solutions_Planning.1131/default.asp#tabbar

Appendix A: Sample scripts

EVM scripts are made up of at least four files. Each EVM operation is divided into “Do” and “Undo” sets of commands. Required with each of these is a “keyword” file. Example script fragments are shown in the *Best Practices for Exchange Database Management using Compaq StorageWorks Enterprise Volume Manager*. The document will be available at <http://www.compaq.com/products/storageworks/Storage-Management-Software/evmindex.html>. Due to the significant amount of individualization each EVM installation and operation requires, this document did not choose to include complete sample script listings.

EVM scripts are stored as ASCII files but they are created and modified using EVM, browser based GUI. Directly editing the ASCII scripts incurs the significant risk that they will not be read by the EVM GUI. Not only are these scripts unique to each host but they are sensitive to the sequence they are run.

Let us say the backup host is expecting to backup a BCV named e:. If another BCV named d: has not yet been mounted to the backup host then the new BCV will be called d: instead of e:. This will cause the backup to fail because it will be trying to backup e: which does not exist.

Similarly, if spare disks are needed in the controller to create a BCV, and another BCV is still using them then the job will fail.

Script process

Scripts must be authored and maintained using EVM. This process creates text files, but the text should not be edited as the resulting changes may not work with EVM.

The process requires the following decisions.

- Pick a name for this EVM script.
- Decide if the user application use a Clone or Snapshot BCV.
- Decide which of the existing LUNs on that controller are to be copied to the BCV.
- Create the application commands to pause and to restart.
- Indicate the remote backup host where the BCV will be mounted and the backup will be done.
- Type in the backup host command to start the backup job of the newly mounted BCV.

Exchange script fragments

The following script fragments are examples of how to pause Microsoft Exchange using EVM.

With Exchange Server 5.5, Exchange must be briefly shutdown to split off a BCV clone. Once a clone has been created, it can be backed up by any software that can do file backups.

The opportune time to shut down Exchange is when there is little or no activity on the server. If there is heavy I/O activity to the log files and Information Store, then Exchange will take a long time to shutdown. Therefore, the Exchange administrator will want to shutdown Exchange properly, which entails shutting down the required services in a particular order.

After Exchange is shutdown gracefully, the database is consistent since all logs have been committed to the Information Store. Once the clone operation is complete, which should be a matter of minutes for a particular Storage Set; the Exchange Server can be restarted again, in the proper service order.

The following is an example of a batch script that will stop Exchange.

```

:: (C) Compaq Computer Corp. - Sample script for StorageWorks Enterprise Volume Manager

REM Note: Sample script is expecting CPQLE.EXE, ESEUTIL.EXE, and CHECKDB.EXE to
REM be in a directory in the System PATH. CPQLEMESS.DLL is required in
REM the %systemroot%\system directory.
REM
REM Requires: Exchange 5.5 (Service Pack 2 or greater),
REM StorageWorks Enterprise Volume Manager,
REM CPQLE.EXE,CPQLEMESS.DLL,ESEUTIL.EXE,CHECKDB.EXE
REM
REM
REM
*****
REM Modify this section with drive:\path for edb.log
REM
*****

REM SetLocal variables

SET LogLocate=E:\exchsrvr\mdbdata

REM *** Modify this section with the variable EVM Job and EDB drive letter names from
REM your system.

REM *** EVMJOBClone is the variable name for the EVM Job created
REM The EVM Job is created from the Compaq StorageWorks EVM Manager
REM Additional PRE/POST processing can be scheduled from the EVM Manager
SET EVMJOBClone=stor1_CloneExchangePrivEdb

REM *** EDBDrive is the variable name for the volume that contains the EDB file
SET EDBDrive=G

REM *** EXCHANGEServer is the variable name for the Exchange Server
SET EXCHANGEServer=P6000x

REM
*****
REM End Modification section: Save the Clonecopy.cmd file
REM
*****

@echo OFF
REM SET echo=ON
@if not "%ECHO%"==" " echo %ECHO%

REM
*****

```

```

REM  STOP THE MICROSOFT EXCHANGE INFORMATION STORE AND
DEPENDENCIES
REM
*****

net stop MExchangeIS /y
REM IF %ERRORLEVEL%==1 GOTO EXITonESEStopErr

REM
*****
**
REM  DUMP edb.log header and store for input to log verification (CHECKLOG.EXE) utility
REM  Note: CHECKLOG.EXE is used in the Clonerec.cmd file
REM
*****

md c:\compaq\swevm\log\
eseutil /ml %LogLocate%\edb.log > c:\compaq\swevm\log\dumplog.out
dir %LogLocate%\edb*.log /OD /B >> c:\compaq\swevm\log\dumplog.out

REM
*****

REM  RUN EVM JOB TO UNDO THE OLD Clone
REM
*****

net time \\\%computername% >> c:\compaq\swevm\log\Clone.out
CD c:\compaq\swevm\bin
swevmcl %EXCHANGEServer% undo_%EVMJOBClone%
IF %ERRORLEVEL%==0 CPQLE 4 2000 "PROCESS:%0, Remove previous Clone Complete.
COMMAND:swevmcl %EXCHANGEServer% undo_%EVMJOBClone%"
REM CPQLE 4 2010 "PROCESS:%0, Clone:%EVMJOBClone% DID NOT DELETE. "

REM
*****

REM  NORMALIZE THE NEW Clone FROM THE PRIMARY DATABASE UNIT
REM
*****

swevmcl %EXCHANGEServer% %EVMJOBClone%
CPQLE 4 2001 "PROCESS:%0, Create new Clone COMPLETE. COMMAND:swevmcl
%EXCHANGEServer% %EVMJOBClone%"

REM
*****

REM  DETERMINE AND STORE LAST CONSISTENT TRANSACTION LOG
REM
*****

eseutil /mh %EDBDrive%:\exchsrvr\mdbdata\priv.edb > c:\compaq\swevm\log\PrivMH.out

```



```

IF %ERRORLEVEL%==1 CPQLE 2 2011 "PROCESS:%0, WARNING !!! eseutil /mh
%EDBDrive%:\exchsrvr\mdbdata\priv.edb NOT SUCCESSFUL "

eseutil /mh %EDBDrive%:\exchsrvr\mdbdata\pub.edb > c:\compaq\swevm\log\PubMH.out
IF %ERRORLEVEL%==1 CPQLE 2 2012 "PROCESS:%0, WARNING !!! eseutil /mh
%EDBDrive%:\exchsrvr\mdbdata\pub.edb NOT SUCCESSFUL "

checkdb.exe c:\compaq\swevm\log\PrivMH.out > c:\compaq\swevm\log\privedb.out
REM IF %ERRORLEVEL%==1 CPQLE 2 2013 "PROCESS:%0, WARNING !!!
CHECKDB.EXE Integrity check of %EDBDrive%:\exchsrvr\mdbdata\pub.edb NOT
SUCCESSFUL "
checkdb.exe c:\compaq\swevm\log\PrivMH.out >> c:\compaq\swevm\log\Clone.out

checkdb.exe c:\compaq\swevm\log\PubMH.out > c:\compaq\swevm\log\Pubedb.out
REM IF %ERRORLEVEL%==1 CPQLE 2 2014 "PROCESS:%0, WARNING !!!
CHECKDB.EXE Integrity check of %EDBDrive%:\exchsrvr\mdbdata\priv.edb NOT
SUCCESSFUL "
checkdb.exe c:\compaq\swevm\log\PubMH.out >> c:\compaq\swevm\log\Clone.out

REM
*****
REM   START THE MICROSOFT EXCHANGE INFORMATION STORE AND
DEPENDENCIES
REM   Note: The Microsoft Exchange Event Service, Key Manager Service, Internet
REM         Mail Service are dependencies of the Information Store
REM
*****

net start MExchangeIS
net start MExchangeES
net start MExchangeIMC
net start MExchangeKMS

net time \\\%computername% >> c:\compaq\swevm\log\Clone.out

EXIT

REM
*****
REM Error Section
REM
*****

:EXITonESEStopErr
CPQLE 2 2015 "PROCESS:%0, WARNING !!! The program %0 exited on an NET STOP
MExchange IS /y error. %EZStatement%"
GOTO END

:END
EXIT

```

Backup script fragments

The following script is an example of how to invoke a Computer Associates' ArcServe^{IT} backup job from EVM

```
-----
@echo off
rem ..... BUarc1.bat Single BU Job .....
rem ..... Sample .BAT file for EVM and CAI ARCServe .....
rem ..... Rev 1.0 .....
rem ..... Oct. 3, 1999 .....
rem .....
rem .....

rem ..Replace <Backup_Server_Name> with actual sever name.
rem ..Replace <job_name> with actual name of BU job from the EVM
rem ..planning Sheet. Omit the <>.
rem ..This batch file must be placed in the Compaq\SWEVM\BIN folder.

rem ..Refer to the EVM planning sheet for naming this file.

rem ..Use the Post command in Arcserve's Pre/Post Command Operation Page to
rem ..run the undo job command BParc1.exe and sCheck the "Immediate" box in the
rem ..Create Job page in EVM.

rem ..When you have finished, save this file as a .bat (batch) file.

c:
rem ..Go to the drive that the Backup Exec Utility is located on
cd \program files\CAI\
rem ..Go to the Arcserve directory
if exist archbatch.exe goto next
echo on
echo "ARCBATCH.EXE not found. Check Path!"
echo off
goto reset
:next
Arcbatch MODE=Execute /H=<Backup_Server_Name> S//00000001.qsd/<job_name>
if errorlevel 1 goto errornote
if errorlevel 0 goto end
:errornote echo on
echo "Error occurred during backup operation, refer to ARCServe Error Log."
echo off
:reset
cd \compaq\swevm\bin\
:end
```

Sample scripts for other operations used in the Uptime Booster performance characterization tests can be found in the file "Script Examples.ZIP" located at <ftp://ftp.compaq.com/pub/products/storageworks/solutions/ScriptExamples.zip>.

Further documentation on these command lines can be found in your backup application user manual.

Appendix B: Overview of *SANworks Enterprise Volume Manager*

Compaq SANworks Enterprise Volume Manager V1.1 is a key component of Compaq's Enterprise Network Storage Architecture strategy (ENSA). *SANworks Enterprise Volume Manager* provides advanced, centralized storage management capabilities in Windows NT computing environments. Using either the BCV snapshot or BVC clone feature can decrease the amount of time required to restore an Exchange mailbox server and provides an additional tool that is complementary to regular online backup techniques. The BCV snapshot or BCV clone of the database file allows administrators to perform integrity checks and offline copies of the clone without stopping the core services of the Exchange Server. Clones or Snapshots should not be viewed as a replacement for tape backups, or as a justification for unchecked growth of the Exchange Information Store.

EVM is a Web-enabled application program that allows you to establish an EVM Network of up to ten computers that can participate in automated storage unit replication jobs on Compaq StorageWorks Enterprise Storage Array 12000 systems or StorageWorks RAID Array 8000 FibreChannel systems.

Each EVM Network and its associated StorageWorks subsystem can be remotely managed from a central location using a Web browser. Access control features are included in EVM's HTTP server.

Figure B-1 below shows a typical EVM Network. The EVM Server and EVM Clients are noted with red and orange dots, respectively.




Figure B-1

EVM graphical user interface

Through an easy-to-use Web browser interface, EVM provides the means to create, manage and run storage replication jobs on the StorageWorks subsystem. Completing just four simple steps with the EVM Create Job wizard defines most jobs. The most complex jobs can be created in six steps using the wizard. EVM job pages provide you with the ability to create, run, validate, monitor, review and delete jobs. The EVM Network View and EVM Storage View pages provide graphical displays of the available resources on the EVM Network.

EVM job interaction with external applications

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

- EVM-database jobs that create nearline backup units on the *StorageWorks* subsystem.
- EVM-database-backup jobs that copy online databases and mount them on backup servers for automated tape backups.
- EVM-database jobs that copy databases and mount them on another server for data mining, or testing.

Creating high-availability solutions

By using batch files and scheduling applications, it is possible to create job sequences that can be run 24 hours a day, seven days a week.

EVM storage unit replication features

EVM provides up to two methods of copying storage units: cloning and snapshot. To select either method requires but a single click on the EVM graphical user interface. All of the underlying *StorageWorks* CLI commands that are required for these methods have been automated and incorporated into EVM. You can click on the blue hypertext links to find out more about each method.

Replicating storage units using cloning

Defining cloning

Cloning is a method of making a physical, point-in-time copy of a storage unit. EVM creates clones with a simple, single-click operation. All of the underlying *StorageWorks* CLI commands that are required for cloning have been automated and incorporated into EVM.

When you specify the clone method in a job, EVM will automatically make duplicates, or BCV units, and mount them on a computer with the drive letters that you specify. The BCV units are then available for a variety of purposes. You can also have the EVM job leave the BCV units unmounted, for example, as nearline backup.

ACS versions supporting cloning

Versions 8.5F and 8.5S.

Performance of clones

A clone can be built using one of two modes—Fast or Normal. Fast mode builds the clone at 20 MB/Sec, while Normal mode builds it at 7 MB/Sec. Normal mode is the default and recommended setting as it runs at lower priority than host I/O.

When replicating the Exchange Information Store, the most important factors are the capacity and number of disks being normalized. Assuming a moderate I/O load on the *StorageWorks* HSG80 controller, and an adequate configuration of disks in the array, a replication rate of 60 GB per hour can be achieved.

Cloning a storage unit using EVM

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

EVM applies slightly different procedures, depending upon the type of source unit. You should review the differences in procedures outlined below and consider them when planning EVM jobs and storage configurations.

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

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

Source Unit is Not Mirrored (RAID 0 and JBOD). EVM automatically creates a single member mirror set and then adds a second member. When the members are normalized, EVM splits the second member to create the new BCV unit.

Cloning RAID 3/5 units

RAID 3/5 units can NOT be cloned. Due to the presence of parity data, units that are configured as RAID 3/5 cannot be replicated by the clone method.

Cloning BCV units

BCV units can NOT be cloned. EVM tracks the BCV units that it creates with each job and will not allow you to select one as a source for another job.

EVM general requirements

System requirements summary

Host Operating System

- Windows NT Server 4.0 (Service Pack 5) or —
- Windows NT Server Enterprise Edition 1.0, with Cluster Hotfix

Host Adapters

- Windows NT compatible LAN adapter
- FibreChannel adapter - Emulex LP7000, Driver Version 4.3

FibreChannel

- Hub or switch configurations

StorageWorks Subsystem

- HSG80 controller (FibreChannel)
- ACS versions 8.5F or 8.5S
- HS Series Agent Version 2.2 (included on EVM CD-ROM)
- RAID Array 8000 or —
- Enterprise Storage Array 12000

Web Browser

- Microsoft Internet Explorer 4.0 or later or —
- Netscape Navigator 4.0 or later. Version 4.5 or later is recommended.

Compaq SANworks Enterprise Volume Manager can be deployed on Windows NT systems (server or workstation). It is complementary to industry backup solutions (e.g. Veritas Backup Exec™ for Windows NT, Legato NetWorker® for Windows NT, and Computer Associates ArcServeIT® for Windows NT), as well as NT's basic built-in backup capability. When layered on top of Microsoft Cluster Server (MSCS), it takes full advantage of MSCS's high availability features.

Note: Exchange Server requirements may exceed the minimum requirements for *SANworks Enterprise Volume Manager*. Consult your Compaq sales or support representatives for assistance in optimizing your specific configuration.

EVM supported software**Database Applications**

The following have been tested and qualified for use with EVM:

- Microsoft Exchange Server 5.5
- Oracle Corporation Oracle8

Backup Applications

The following have been tested and qualified for use with EVM:

- Computer Associates ArcServeIT for Windows NT Version 6.6
- Legato NetWorker for Windows NT Version 5.5
- VERITAS Backup Exec for Windows NT Version 7.3

Other Software Support Information

EVM supports the following:

- Windows NT, MSCS Clustering
- StorageWorks Secure Path Manager Version 2.2
- Compaq Batch Scheduler

Microsoft SQL Server

For the latest information on support of Microsoft SQL Server, visit the Compaq StorageWorks Web site that is listed in the EVM Toolbox.

Operation of EVM in a Microsoft Exchange environment

Please refer to the White Paper "Best practices for Exchange Database Management Using Compaq SANworks Enterprise Volume Manager" located at:

<http://www.compaq.com/support/techpubs/whitepapers/11ef-1199a-wwen.html>

Appendix C: Overview of Microsoft Exchange Server

Microsoft Exchange Server architecture provides reliable messaging, scheduling, and information sharing services. Each server has core and optional components. All server components are implemented as Windows NT Server services. The core components are installed during Setup and must be running at all times. In addition, they must all reside on the same server running Windows NT. The core and optional components maintain address and message databases, log transactions, and allow you to connect with other servers running Microsoft Exchange and with foreign systems.

Managing the Exchange Server 5.5 database

The Extensible Storage Engine (ESE) manages the Exchange Server 5.5 database. The Exchange database is stored in memory, log files, and database files. Exchange database files are in an inconsistent state during normal operations. When a graceful shutdown of the Exchange core services is complete, the database file is in a consistent state.

The Exchange database is the Information Store database where the data is saved in the PRIV.EDB and PUB.EDB files that are created by default. Table C-1 details the location and, where applicable, file names of ESE database files.

Table C-1. Typical Distribution of Exchange Database Files, Log Files, and checkpoint file.

Description	Location	Files
Directory database	\exchsrvr\dsadata	DIR.EDB
IS database	\exchsrvr\mdbdata	PRIV.EDB, PUB.EDB
Directory transaction log files	\exchsrvr\dsadata	EDB.LOG, EDBNNNNN.LOG
IS transaction log files	\ exchsrvr\mdbdata	EDB.LOG, EDBNNNNN.LOG
IS Database Checkpoint File	\EXCHSRVR\MDBDATA	EDB.CHK

The Exchange Information Store

The database files are the .edb files that are stored on disk. Each database is a single large file, with all the attributes of a normal operating system file. Internally, a database file is organized as a collection of 4 kilobyte (KB) pages which contain data definition, data, indexes, and so on.

Storing pages in memory

When a modification is made to the database, the page that needs to be changed is copied into memory, if it is not already loaded in memory. A database page is accessed in nanoseconds from memory. Modifying pages in memory increases the performance of ESE.

Saving transaction data in log files

Changes to the database are always written to transaction log files. These files store a sequential list of the changes that have been made to the database pages in memory. These log files provide a means for recovering changes that might be lost in case of a power failure.

Changing data

A transaction is a series of physical modifications to a database that leaves the database in a consistent state. These physical modifications are called operations.

An operation is the smallest change that can be made to a database, such as an Insert, Delete, Replace, or Commit operation. Several operations make up a single transaction. A single operation may leave the database in an inconsistent state.

Transactions meet ACID test

ESE supports full atomic, consistent, isolated and durable (ACID) transactions as defined:

Atomic: Operations performed in a transaction are either all completed or none of them completed.

Consistent: A transaction takes a database from one consistent state to another.

Isolated: Changes are not visible until the transactions are committed.

Durable: Committed transactions are preserved in the database even if the system crashes.

Common support issues for Microsoft Exchange Server

Unplanned downtime for Exchange Servers is a common support issue problem. Power outages, hardware failures, or other issues may cause Exchange to experience a failure that requires the system to attempt to get back to a consistent state.

- Long Restore Periods
- Physical Corruption of the Database
- Logical Corruption of the Database
- Operator, Process or Procedural Errors
- Lack of Proactive Management Techniques

Long restore periods

Today administrators face lengthy delays when restoring large databases from tape. *SANworks Enterprise Volume Manager* is a tool that in many cases will reduce the restore time.

Physical corruption of the database

At the lowest level, your data can become physically corrupt on the disk. This is usually a hardware-related problem and may require downtime for the users with mailboxes on the server. The downtime is a direct result of the time required to restore to a last-known-good backup. Physical corruption to the database is a serious problem. Online backups will halt when they encounter physical corruption. If the corruption occurs in a critical piece of table structure then the Information Store service may not start. The corruption may be limited to a single message in a mailbox or it may be in a page of the database that contains “white space.” If that’s the case the Information Store may continue to run without the ability to do online backups. Until the root cause of the corruption is determined the corruption may occur again. As a result, it’s very important that you detect physical corruption early and resolve the problems quickly. Currently to determine the root cause of the problem may require the system to be down. The root cause is not always easily identified either.

Logical corruption of the database

The most typical logical corruption occurs at the database level. For example, database engine failure may cause index entries in the database to point to missing values. Logical corruption can also happen at the application level, for example, in mailboxes, messages, folders, and attachments. An application level corruption might cause incorrect reference counts, incorrect ACLs, a message header without a message body, and so on.

Logical corruption can occur in the Information Store or the database engine. Because logical corruption can seriously damage your data, you shouldn’t ignore errors in the Information Store or the database engine. You can use ISINTEG to check problems with the Information Store, or ESEUTIL to check problems with the database engine.

IMPORTANT: You should use ISINTEG and ESEUTIL as a last resort only after you’ve tried to restore your system from backup.

Operator, process, or procedural errors

Logical corruption may manifest itself as random crashes of the database. It is crucial to have documented and tested procedures for all aspects of managing your servers. A large percentage of unplanned downtime is directly attributed to improper decisions made by administrators. Many administrators have wished for a good backup. Implementing *SANworks Enterprise Volume Manager* offers a recovery/backup tool for database files and logs.

Corruption of log files

A reason a database file becomes logically corrupt is that transaction log files are replayed incorrectly:

- The checkpoint is not removed before running isinteg –patch on a restored offline backup.
- Transaction logs are deleted or renamed, and therefore not replayed in order.
- Transaction logs are skipped during replay.

Lack of proactive management techniques

Unplanned downtime may be avoided by monitoring events in the Windows NT Event Viewer and having critical events notify support personnel.

Appendix D: Test disclosure

Performance characterization tests were conducted using the Microsoft Exchange Backup/Restore Solution – Uptime Booster under typical load conditions. *The results of these tests show that the use of a BCV clone has no significant impact on the performance of Microsoft Exchange and that backup to tape completes in nearly half the time of an online backup.*

Table D-1 below shows the results from tests using our typical configuration.

Table D-1. Load Simulator (Loadsim) results and backup rates

Operation	Backup Type	Loadsim Score 95% Weighted Avg. (ms)	Backup Rate
Baseline	None	120	N/A
Backup	Online	136	21 GB/Hr
Backup	From Clone	137	25-35 GB/hr
Restore	Direct to Volume	N/A	30 GB/hr

* Compaq's Best Practice for Exchange recommends using BCV clone for Backup.

Note: Both the *StorageWorks* RA8000 and ESA 12000 require 512 MB cache/controller for optimal BCV snapshot performance.

Workload methodology

Three Microsoft Exchange Servers were members of a site. Each of the tests was run on one of the servers. Each server was tested with 1000 medium Outlook clients. Different configurations of the hardware were tested to develop a baseline. The baseline did not use *SANworks Enterprise Volume Manager* BCVs. Additional tests were run to compare the overhead of the copy-out operation. The following configurations were tested:

- Baseline configuration
- Online Backup
- Clone
- Snapshot
- Backup Clone using Legato's NetWorker, CA's ArcServeIT and Veritas' Backup Exec
- Backup Snapshot using Legato's NetWorker, CA's ArcServeIT and Veritas' Backup Exec

These test configurations are further described in the section on "System Configurations."

The main tool used in generating the workload was the Microsoft Exchange Server Load Simulation (Loadsim) utility called "Load Simulator." Load Simulator is a tool for simulating a client user load on an Exchange Server. Its purpose is to enable a single Windows NT machine, called a "Loadsim client," to simulate multiple Microsoft Exchange client users.

For each test, the Load Simulator Medium Outlook Messaging Application Programming Interface (MAPI) canonical profile was chosen. The Medium MAPI profile reflects the task workload of a typical corporate e-mail user, including common daily mail tasks such as send, browse, read, and forward, as well as scheduling tasks and distribution list usage. The Response Time score represents a 95th-percentile score of the measured test run. The score is expressed in milliseconds (ms). A Response Time score of 1000 ms or less is considered an acceptable response time for e-mail users using Exchange Server’s MAPI protocol.

Multiple Intel-based Load Simulator client systems were used to drive the system under test to saturation levels. Up to a total of 500 users were simulated on each client system. The tests were run for eight hours to verify that the response times and all measured events and counters were within acceptable limits.

For detailed information on the workload used for testing, refer to the Microsoft Exchange Server White Paper written by Microsoft. This workload-sizing document is available from the Microsoft Web site at: <http://www.microsoft.com/Exchange>. The document provides a good description of the Load Simulator Utility.

Test Procedure

Setting up the simulation environment is the first task in a Loadsim test. Loadsim generates the user directory database and then initializes the test by populating the Exchange Information Store database. A test typically lasts for eight hours. During a test, the simulated users log onto 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, a command line utility packaged with Loadsim, can be used to parse the “loadsim.log” file and calculate the 95th percentile (95%) response time. A successful test must satisfy the Microsoft Exchange 5.5 UPS Policy Guidelines 1.01.

Tested system configurations

Tables D-2 through D-6 below describe the configurations that were used to test *SANworks Enterprise Volume Manager*.

Table D-2: Common system configuration 1

Configuration Compaq ProLiant 6000 (Common Exchange configuration for all tests)
(2) Pentium® II Xeon™ 400-MHz – 1 MB L2 cache per processor
512 MB RAM
KGPSA-BC – FibreChannel Adapter
10/100 Base-TX UTP
Windows NT Server Version 4.0 Enterprise Edition, Service Pack 5
Exchange Server Version 5.5 – Enterprise Edition, Service Pack 3
Smart 3100ES, 18 – 4.3GB drives
Compaq RA8000 FC Array Controller with 128MB Cache
Wide Ultra SCSI 9.1 GB 10k rpm drives (24 drives)

Table D-3: Common system configuration 2

Configuration Compaq ProLiant 1600 (Common Backup Processor configuration for all tests)

(2) Pentium® II Xeon™ 498-MHz – 1 MB L2 cache per processor

512 MB RAM

KGPSA-BC – FibreChannel Adapter

10/100 Base-TX UTP

Windows NT Server Version 4.0 Enterprise Edition, Service Pack 5

Exchange Server Version 5.5 – Enterprise Edition, Service Pack 3

Table D-4: Common backup processor configuration**Configuration Compaq ProLiant 1600**

(2) Pentium® II Xeon™ 498-MHz – 1 MB L2 cache per processor

512 MB RAM

KGPSA-BC – FibreChannel Adapter

10/100 Base-TX UTP

Windows NT Server Version 4.0 Enterprise Edition, Service Pack 5

Exchange Server Version 5.5 – Enterprise Edition, Service Pack 3

Table D-5: Exchange workload**Table 2. Click here and enter the Table Name****Exchange Workload**

1000 Medium Loadsims users

49GB Exchange Store – Raid 0+1

Table D-6: Tested configurations**Table 3. Click here and enter the Table Name**

Infrastructure	Backup Application	Tape Library	BCV Volume Type
Disk-FibreChannel Switch	Legato NetWorker V5.51	TL895	Snap
Tape - EBS			Clone
Disk-FibreChannel Switch	VERITAS Backup Exec V7.3	TL895	Snap
Tape - Direct SCSI			Clone
Disk - FibreChannel Switch	CA ArcServe/IT v6.61	TL895	Snap
Tape - Direct SCSI			Clone

Test results

Performance results were measured using Microsoft NT Performance Monitor. Measurements were obtained by measuring averages for the period of steady-state activity (after 1000 users were successfully logged on). Tests measure the messaging throughput of a multi-server, single-site topology.

IMPORTANT: There was no significant impact on the performance of Exchange using a BCV clone in this tested environment.

Actual Loadsim results can be viewed at

<ftp://ftp.compaq.com/pub/products/storageworks/solutions/LoadSim.zip>

For deployment-specific information, contact a Microsoft or Compaq representative. More information can be found at

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

Performance impact of restoring from BCVs

It is important to verify the integrity of your BCV prior to using it for an offline recovery. If using the ESEFILE utility, the integrity check should be scheduled during periods of low user activity. Testing has shown that the overhead of running ESEFILE immediately after starting the Exchange Information Store service will slow the logon of users until ESEFILE completes. With active users on the Exchange Server, a user's response time will be affected when ESEFILE runs. Table D-7 below shows that ESEFILE performs intensive reads on the BCV.

Table D-7: Performance impacts of running ESEFILE with active users

Performance Monitor Counter	Counter Value	Loadsim Medium MAPI Users
Disk Queue Length (IS)	42.85	3000
Disk Queue Length (Snapshot)	255.69	3000
Disk Read I/Os (IS)	95.41 per/sec	3000
Disk Read I/Os (Snapshot)	547.39 per/sec	3000

As shown in Table D-8, ESEUTIL exhibits similar read intensive processes.

Table D-8: Integrity check utility throughput

Table 4. Click here and enter the Table Name

Utility	Throughput for Utility	Load on Server during Integrity Check
ESEUTIL /G	17.34 GB/hr	No users connected
ESEUTIL /G	15.52 GB/hr	500 medium Loadsim
ESEFILE /S	197 GB/hr or 3.28 GB/min	No users connected
ESEFILE /S	193 GB/hr or 3.22 GB/min	500 medium Loadsim

Descriptive Terms

Average Time for Local Delivery

This counter is the average time, measured in tens of milliseconds, between submission of a message to the Information Store and delivery to all local recipients for the last ten (10) messages.

Messages Submitted

Submit calls made by clients. This equates to total message sends by users.

Messages Sent

Messages that the Information Store sends to the MTA (not messages sent by clients). Normally all messages submitted by the clients are sent to the MTA, except in the case where all recipients are local mailboxes. In that case, since all the deliveries can be performed locally, no message is sent to the MTA.

Message Recipients Delivered

Separate mailboxes that messages have been delivered to. Think of this as the number of Reads that are 'caused' by sending a message (one per recipient).

Message Opens/Sec

Messages accessed for reading per second.

Folder Opens/Sec

Folders opened for browsing per second.

RPC Read Bytes/Sec

RPC Bytes read from clients (i.e., submit calls).

RPC Write Bytes/Sec

RPC Bytes written to clients (i.e., message opens).

IS Send Queue Average Length

Send Queue Size is the number of messages in the private Information Store's send queue.

MTA Work Queue Average Length

Work Queue Length is the number of outstanding messages in the Work Queue, which indicates the number of messages not yet processed to completion by the MTA.

Appendix E: Related documents

Visit the Compaq and Microsoft Web sites for online tools and technical information to plan, deploy, and operate Exchange Server Solutions.

Compaq *ActiveAnswers*

<http://www.compaq.com/activeanswers>

- Performance and Tuning Guide: Microsoft Exchange 5.5 on Compaq Servers
- Boosting Performance on Microsoft Exchange Thorough Controller Caching
- Microsoft Exchange Server Backup and Restore Performance using Compaq 35/70 DLT Drives
- Backup and Restore Procedures for Microsoft Exchange Server
- Deploying and Operating Microsoft Exchange Server: How to Get It Right!
- Managing and Monitoring Microsoft Exchange Server

Compaq StorageWorks

<http://www.compaq.com/products/storageworks>

- Compaq SANworks Enterprise Volume Manager
- Compaq StorageWorks Management Software
- Compaq Batch Scheduler for Windows NT

Compaq White Paper Index

<http://www.compaq.com/support/techpubs/whitepapers>

Microsoft Exchange Server

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

Books

Microsoft Exchange Server V5.5 : Planning, Design, and Implementation
Tony Redmond; Digital Press; ISBN: 1555582133