NETAPP TECHNICAL REPORT

Solution Blueprint for Symantec Enterprise Vault Journal Archiving on NetApp Storage Systems

Nathan Walker - NetApp Technical Marketing Engineer
Rick Krieger – Symantec Senior Regional Product Manager

February 2010 | TR-3716

7,500-Seat Environment Execute Summary

This document is intended to serve as a blueprint for architecting and deploying Symantec™ Enterprise Vault™ in a 7,500-seat customer environment. We have evaluated the archiving throughput, storage growth, and costs based upon this scenario. As always, please refer to the latest technical publications on the NOW™ (NetApp® on the Web) site for updates on processes; Data ONTAP® command syntax; and the latest requirements, issues, and any relevant limitations. This document is intended for field personnel who require assistance in deploying and architecting an Enterprise Vault solution on NetApp storage systems.

1 INTRODUCTION

1.1 PURPOSE

This solutions blueprint is a guide for field engineers to use when implementing Enterprise Vault for journal mailbox archiving on NetApp storage systems. It describes the architectural requirements for Symantec Enterprise Vault journal archiving on NetApp storage solutions. This document provides a brief performance analysis of 7,500-journal mailbox archiving on a NetApp FAS3000 series storage device, protected by NetApp storage solutions. Now you can streamline your e-mail operations and minimize your risk with solutions from NetApp.

1.2 TARGET AUDIENCE

This document is intended for information technology professionals and storage architects who are responsible for corporate messaging infrastructure management. It assumes the reader has some technical experience installing, configuring, and administering the following technologies:

- Symantec Enterprise Vault
- NetApp Data ONTAP
- NetApp SnapManager® for SQL Server®
- NetApp SnapDrive® for Windows®
- Microsoft® SQL Server 2005
- Microsoft Exchange 2007
- Microsoft Windows Server 2003
- NetApp Operations Manager

1.3 SCOPE

This document only discusses architectural requirements for Journal Archiving with Symantec Enterprise Vault for Microsoft Exchange 8, NetApp FAS3000 series storage, Data ONTAP 7.3.1.1, NetApp Snapshot™ technology, SnapMirror®, and SnapManager for SQL Server. For installation processes, refer to the references section of this document for the Enterprise Vault Installing and Configuring Guide, the Data Protection Online Backup and Recovery Guide, and the Data ONTAP Software Setup Guide.

The focused objective of this paper will help archive and messaging administrators to consider the uniquely superior characteristics of unified NetApp storage system for archiving objectives. Included are helpful introductions to relevant technologies such as our SnapManager application suites and SnapMirror protection software.

As you plan to deploy Enterprise Vault on NetApp storage solutions, we encourage you to evaluate our other supporting technologies to make sure of complete project success. A qualified systems integrator or NetApp Professional Services will partner with you to affordably store, manage, protect, and retain one of your most precious assets—your data. Let NetApp help you unify and simplify your storage architecture. Our single-box solutions are more compatible, flexible, and secure. Independent third-party testing has validated our superior e-mail archive performance over those based on comparable EMC hardware.

2 NETAPP SOLUTIONS FOR SYMANTEC ENTERPRISE VAULT

2.1 ENTERPRISE STORAGE

Enterprise Vault stores and manages large volumes of structured e-mail data. Selecting and properly architecting the underlying storage system for an Enterprise Vault solution are critical to achieve superior performance. The three main data types Enterprise Vault repositories created are the vault store data, the indexes of archived content, and the vault store databases.

2.1.1 DISK TYPES

There are several data sets that can reside on different storage based on IOPs requirements. Enterprise Vault data can reside on different types of disks. The two discussed here are Fibre Channel (FC) and Serial ATA (SATA) disks.

FIBRE CHANNEL

Fibre Channel is a hard disk drive interface technology designed primarily for high-speed data throughput for high-capacity storage systems. They are a proven and reliable storage device, with high read/write speeds and the ability to handle high I/O loads. Fibre Channel disks are ideally suited for Enterprise Vault databases and indexes in high-performance environments.

SERIAL ATA

Lower-cost NetApp SATA storage solutions provide customers with an excellent opportunity to reduce storage costs or stretch their IT budget without incurring a noticeable application performance impact. SATA storage solutions can be leveraged in environments for storing and managing Enterprise Vault vault store data. Although SATA disks are not intended to replace higher-performance FC disk drives, knowing candidate applications that could perform satisfactorily on SATA disks helps achieve the best value and use of a given storage purchase. One possible modification of the architecture used in this paper would be to substitute SATA drives for the index and database Fibre Channel drives, with the addition of a Performance Acceleration Module (PAM) card. This would provide greater storage per dollar spent, with comparable performance.

SATA disk drives are becoming increasingly popular as a storage medium. SATA disks are a low-cost, high-capacity storage solution. Based on performance testing, SATA drives are suitable to host vault store data in several instances as well as indexes for customers with 7,500 users or more.

2.1.2 RAID-DP

RAID-DP® is a standard Data ONTAP feature that safeguards your data from double disk failure. Integrated with our WAFL® (Write Anywhere File Layout) file system, RAID-DP gives you data protection plus high performance. RAID-DP is a double-parity RAID 6 implementation that prevents data loss when two drives fail.

We integrated RAID-DP with the WAFL file system to make sure that the dedicated parity drives don't become a performance bottleneck. You get protection plus the performance you need for your most demanding applications.[1]

Additional details on RAID-DP are available in TR-3298.

2.2 DATA PROTECTION

Enterprise Vault has several data components that require proper planning for data protection. It is critical to maintain data consistency among these data sets at all times. The NetApp solution to back up an Enterprise Vault data set involves a series of procedures and takes advantage of the PowerShell capabilities in Enterprise Vault 8 to maintain data consistency among the distributed data components. The NetApp technologies described in this section can be used to back up Enterprise Vault data sets.

Many existing Enterprise Vault and NetApp customers have implemented one or more of the following technologies in their environment and are enjoying the benefits that they offer.

2.2.1 SNAPMANAGER FOR MICROSOFT EXCHANGE

We recommend using SnapManager for Microsoft Exchange to protect and manage your Microsoft Exchange environments. With its Microsoft Management Console GUI and PowerShell cmdlets you can centrally manage and automate the complex, manual, and time-consuming processes associated with the backup, recovery, and verification of Exchange Server databases.

SNAPMANAGER FOR MICROSOFT EXCHANGE BENEFITS

- Reduce backup times to seconds and restore times to just minutes.
- Add storage capacity and expand volumes without taking the Exchange Server or the NetApp storage system offline.

- Streamline management by automating common tasks to let your administrators spend less time on maintenance.
- Deploy into your existing infrastructure.
- Integrate SnapManager with your native Microsoft technology and frameworks.
- Integration with SnapMirror (synchronous and asynchronous) and SnapVault®.
- Provide disaster recovery and business continuity.
- Role-based access control.

2.2.2 SNAPMANAGER FOR MICROSOFT SQL SERVER

NetApp recommends using SnapManager for Microsoft SQL Server to protect the Enterprise Vault databases. SnapManager for Microsoft SQL Server is tightly integrated with Microsoft technologies to help you streamline database storage management while simplifying storage layout planning, backup, and restore operations for SQL Server databases. SnapManager for Microsoft SQL Server can save you time and money with our space-efficient backup capabilities and automated data management processes. With SnapManager, you can dramatically reduce SQL Server data recovery times from hours to minutes, making it one of the fastest backup and recovery solutions available. You can also use SnapManager for SQL Server with its PowerShell cmdlets to automate backup, recovery, and database cloning.

SNAPMANAGER FOR MICROSOFT SQL SERVER BENEFITS

- Near-instantaneous backup and fast restore of entire SQL Server databases and full text indexes using NetApp Snapshot technology
- Reduce storage costs with our space-saving backup technologies
- Increase productivity by automating routine database tasks
- Protect more data and increase your backup frequency without impacting performance
- Easy migration wizards to move databases to SAN and IP SAN environments
- Restore a failed database of any size to full production in minutes
- Easy-to-use, intuitive graphical user interface
- Rich backup scheduling and reporting
- Integration with SnapMirror for wide area data replication

2.2.3 SNAPSHOT COPIES

NetApp strongly recommends using Snapshot copies and SnapRestore® for Enterprise Vault vault store and index backup and restore operations. Snapshot provides a point-in-time copy of the entire vault store data and index in seconds without incurring any performance penalty, and SnapRestore can instantly restore an entire database to a point in time in the past.

For Snapshot copies to be effectively used with Enterprise Vault, they must be coordinated with the Enterprise Vault trigger file mechanism, which tells Enterprise Vault the backup process has been completed and it is now OK to remove safety copies or archived content. For this reason, NetApp recommends automatic Snapshot copies be turned off on volumes storing data files for the Enterprise Vault vault store data and the indexes. For additional details about the cascading Snapshot methodology, see TR-3709, Cascading Snapshot Copies for Enterprise Vault. For more information about the trigger file mechanism, see the Enterprise Vault administrator's guide.

2.2.4 SNAPMIRROR

There are several approaches to increasing data availability in the face of hardware, software, and even site failures. Backups provide a way to recover lost data from an archival medium (tape or disk). Hardware redundancy technologies also help mitigate the damage caused by hardware issues or failures. Data mirroring provides a third mechanism to enhance data availability and to minimize downtime. NetApp SnapMirror provides a fast and flexible enterprise solution for replicating data over local area, wide area, and

Fibre Channel (FC) networks. SnapMirror is a key component when implementing enterprise data protection and disaster recovery (DR) strategies. If a disaster occurs at a source site, businesses can access mission-critical data from a mirror on a remote NetApp system for uninterrupted operation.

SNAPMIRROR BENEFITS

- Block-level updates reduce bandwidth and time requirements.
- Data consistency can be maintained at a DR site.
- A DR plan can be tested without affecting production.
- Synchronization between source and destination sites is complete.
- Mission-critical data can be mirrored.
- A DR location can keep many Snapshot copies at once; data can be restored to a point in time before the data corruption occurrence.
- Data can be replicated between dissimilar NetApp storage systems.
- A standard IP or FC network can be used for replication.
- SnapMirror supports one-to-one, one-to-many, many-to-one, and many-to-many replication, referred to as *cascading* and *multihop.*

Note: To achieve a true disaster recovery plan, NetApp recommends the SnapMirror destination be to tape or to a geographically distant facility.

3 SCENARIO: JOURNAL MAILBOX ARCHIVING 7,500 MAILBOXES

3.1 CONFIGURATION

To design this journal archiving solution, the following environment was utilized.

Table 1) Environment configuration.

Quantity	Item	Version
One	Active Directory domain	2003
One	Enterprise Vault Microsoft SQL Server	8.0, SP2 Microsoft SQL Server 2005 Standard Edition (Version 9.00.1399.06 RTM)
One	Microsoft Exchange	2007, SP1
One	NetApp FAS3050	7.3.1.1
7,500 Users		

An in-house tool was used to generate the content to be archived: 750,000 message objects for the 7,500 mailboxes. This tool used a large dictionary file to create random message bodies. A Snapshot copy of a FlexVol® volume containing user home directories was used as a source repository for e-mail attachments. The tool simulated typical enterprise content by creating 62,500 messages with three mean recipients, 375,000 messages with 1.5 mean recipients, and the remainder with one recipient. A total of 175,096 distinct files were required to satisfy the attachment requirements.

4 SOLUTION BLUEPRINT

4.1 HIGH-LEVEL ARCHITECTURE

This section describes the detailed solution blueprint, including in-depth information about the Enterprise Vault architecture and the proposed storage configuration. It describes the way that Enterprise Vault should be configured to meet the customer business requirements and technical requirements discussed earlier in this document.

4.2 ENTERPRISE VAULT

4.2.1 RECOMMENDED ALLOCATION OF ENTERPRISE VAULT SITE SERVERS

The following figure depicts the local layout of the lab environment. Because the objective of this testing was not high availability, the database server and Enterprise Vault server were placed on the same hardware. The number of servers, server roles, and archiving policies will be determined by the messaging topology and particular business requirements of your organization. This diagram is meant to illustrate the structure of a typical Enterprise Vault layout. Your environment might not be arranged in this manner.

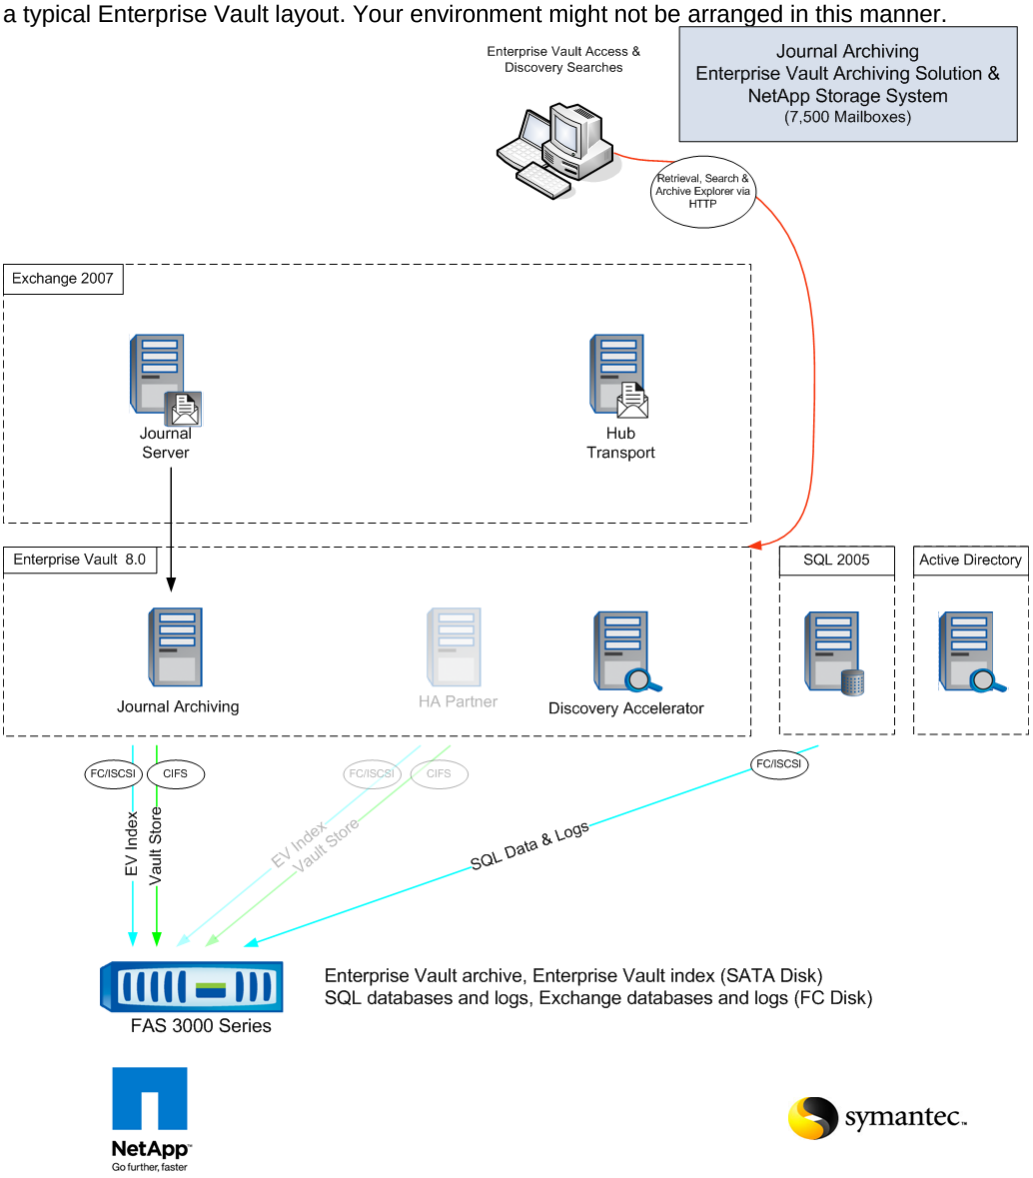

Figure 1) Enterprise Vault 8 reference architecture.

The following tables describe the server hardware setup used to evaluate the throughput performance of Enterprise Vault on NetApp storage solutions. Exchange 2007 was installed on one physical server. The second physical server contained Enterprise Vault 8.0 and Microsoft SQL Server. The Outlook client PC and Active Directory server are not detailed herein. These computers were virtualized and stored on an aggregate within the same NetApp storage system.

Table 1) Lab server hardware and software details.

Component	Details
Model	IBM System x 3250 M2
CPU	Intel® Xeon® 3060, 2.4Ghz, dual core
RAM	5GB
Internal disk	1 x 74GB SAS
Used network interfaces	1 onboard Gigabit Ethernet, for both iSCSI and data communication
OS	Microsoft Windows Server 2003 Enterprise Edition R2 (64 bit)
iSCSI initiator	Microsoft iSCSI Initiator 2.08
SnapDrive for Windows	6.1
SnapManager for Exchange	5.0 (e-mail server only)
SnapManager for Microsoft SQL	5.0 (database server only)

Table 2) Storage layout, Symantec Enterprise Vault 8 server.

Drive Letter	Description	Type
C:\	Operating system	Local disk
D:\	Application install directory	Local disk
I:\	Index, shopping, MSMQ	NetApp LUN
S:\	Microsoft SQL binaries, databases, logs	NetApp LUN
UNC	Vault store data	NetApp CIFS share

Table 3) Storage layout, Microsoft Exchange 2007 server.

Drive Letter	Description	Type
C:\	Operating system	Local disk
M:\	Application install directory, Exchange databases and logs	NetApp LUN

4.3 STORAGE REQUIREMENTS

Enterprise Vault requires space to store archived items and metadata that describes those items, as well as space to be used for steady state processing. The following list outlines the main items to consider when sizing an Enterprise Vault environment.

- SQL databases
- Index locations
- Index level (brief, medium, full)
- Vault store partitions
- Shopping locations
- MSMQ storage locations
- PST holding temp folder
- PST migration temp folder

The one-year storage projections that follow in this section focus on the following items:

- SQL databases
- Index locations
- Vault store partitions

The initial storage requirements and growth rate of the index, archive, and database depend upon the unique characteristics of your message environment. The next two tables list some common parameters that can be used to provide a reasonably accurate estimate of required storage for Enterprise Vault archiving.

Table 4) Year one Enterprise Vault 8 storage parameters.

Parameter	Metric
Number of journal messages per day	750,000
Journaling fan-out factor, due to multiple journal mailboxes	1.75
Number of original messages sent/received daily	428,571
Average size of messages (kB)	130
Estimated percentage of messages with attachments (%)	20
Average attachment size (kB)	200
Average attachments per message with attachments	1.5
Number of work days per year	260
Attachment SIS ratio	2.0
Average compression percentage of attachments based on mix of file types (%)	60
Index size (%)	12

Table 5) E-mail growth parameters.

Assumption	
Annual growth in number of mailboxes being archived (%)	10
Annual growth in number of messages sent/received (%)	10
Annual growth in size of average message/attachment (%)	10

The following table contains storage estimates after one year of journal mailbox archiving based on the metrics of the two previous tables. The archive, indexes, and databases will grow an average of 28% each year. This knowledge will help storage and archive administrators to plan IT capital expenditures. The storage numbers listed in this table do not include the processing of an e-mail backlog or ingestion of PST files. The total storage requirement for Enterprise Vault can be reduced through FAS deduplication. Deduplication removes redundant blocks of data within a FlexVol volume. Deduplication provides different amounts of storage savings depending on the type of data. Review TR-3765 to learn more about deduplication of Enterprise Vault data. A thorough analysis of your particular archive needs can be understood by engaging with Symantec professional services or with a qualified systems integrator.

Table 6) Year one Enterprise Vault 8 storage estimates.

Year One Storage Estimate	Archiving from Journal Mailbox After One Year
Number of messages to archive over year one	**215,475,000**
Estimated number of attachments over year one	**64,642,500**
Total size of messages to archived in year (GB)	**28,050**
Vault store NTFS (GB)	**5,655**
SQL vault store database (GB)	**100.34**
SQL fingerprint database (GB)	**8.60**
Total SQL size (GB)	**108.94**
Index size (GB)	**3365.97**
Total EV storage required for year 1 steady state (GB)	**9,130**

These items grow as content is archived; they represent the majority of the storage requirements for an Enterprise Vault environment. The system shown in Table 7 was used to hold Enterprise Vault index, archive, and Microsoft SQL database. In addition, one aggregate was dedicated for ESX guests. This meets the requirements for a fully functional Enterprise Vault site.

Additional disk shelves were contained in the storage frame, but were not used for this evaluation. The author of this document recognizes NetApp best practices were not followed for RAID group sizing of the database aggregate. This was due to a limited supply of disk drives. For RAID-DP best practices, consult TR-3298.

Your archiving server arrangement should be chosen according to your specific needs and data center best practices. For example, most production servers will have dedicated and redundant network connections for data, monitoring, and iSCSI traffic. Our storage controller and servers each had a single Gigabit network connection for iSCSI, CIFS, monitoring, and management traffic.

Table 7) Lab FAS3050 storage system details.

Component	Details
Controller	FAS3000 series
Data ONTAP	7.3.1.1
RAM	3GB (factory default is 4GB)
CPU	2 x dual core 2.8Ghz Intel Xeon
Archive aggregate (RAID-DP)	16 x 206GB 7200 RPM SATA disks, 1 RAID group
Index aggregate (RAID-DP)	16 x 206GB 7200 RPM SATA disks, 1 RAID group
Exchange aggregate (RAID-DP)	16 x 134GB 10,000 RPM FC disks, 1 RAID group
DB aggregate (RAID-DP)	15 x 134GB 15,000 RPM FC disks, 1 RAID group
VMware® aggregate (RAID-DP)	16 x 134GB 10,000 RPM FC disks, 1 RAID group

Table 8 represents a standard disk configuration for three aggregates to be used in an environment that implements this reference architecture. The total storage for year one should be procured at the inception of the project to have the proper number of disks in each RAID group. The FAS3050 storage controller used for these tests did not contain a PAM card. As noted earlier, SATA drives and a PAM card might have been substituted for Fibre Channel drives, with comparable archiving performance.

Table 8) Recommended year one storage layout.

Server	Aggregate Name	Raid RAID Group	Disk Capacity/Type	RAW Aggregate Capacity	Usable Aggregate
EV	EV_SQL	14+2	144GB10K FC	2.04TB	1.54TB
EV	EV_INDEX	14+2	250GB 7.2K SATA	3.5TB	2.4TB
SQL	EV_ARCH	14+2	250GB 7.2K SATA	3.5TB	2.4TB

4.4 STORAGE LAYOUT

4.4.1 DISK UTILIZATION AND APPLICATION THROUGHPUT

This section documents the archive characteristics and data processing rates for the Enterprise Vault indexes, vault stores, and SQL databases on the target storage system. These data points are based on the actual metrics during an archiving window of 750,000 messages. Table 9 shows the characteristics based upon data obtained from Enterprise Vault reports and the relevant environment settings behind these estimates.

Table 9) Lab measurements of Enterprise Vault 8 archiving on NetApp 3050.

Assumptions	
Archiving rate (average items per hour)	28,514
Average attachment size (kB)	104
Single instance ratio	1.1
Compression percent (%)	60.1
Indexing level	Full

To initiate this evaluation, the journal archive task was allowed to run until all messages had been archived from the target journal mailbox. The tests were executed four times with the identical content in the Exchange message store to make sure of consistent results. The results in tables 9 and 10 are from the final test cycle. All four runs had approximately the same performance results. SnapManager for Exchange, SnapManager for SQL, and SnapDrive for Windows were used to quickly revert the environment to its original state before starting the next test iteration. In this way identical content was archived each time. There were no tuning switches to skew performance. Only the Enterprise Vault wizard and administrator console were used to manage the application and database. Every effort was made to make sure the archive server was a close approximation of what would be used for a 7,500-user Enterprise Vault environment. One exception was to enable circular database logging within Exchange. Because the index LUN also had the MSMQ directories, the IOPs were about four times higher than the test iteration with these queues on a local disk. In a large environment, it may be advisable to move these queues to a dedicated LUN.

Table 10) Lab IOPs and throughput of Enterprise Vault 8 archiving on NetApp 3050.

	Disk Read/Write Rate (GB Read/Written per Hour)	Disk Usage (IOPS)	Data Profile
Vault store	2.43	*227.0	Sequential writes
Indexes	8.84	147.3	Random reads and writes
SQL DB and logs	1.13	37.86	Random reads and writes

* CIFS operations per second, reported by NetApp Performance Advisor.

While the journal mailbox content was being archived the NetApp storage controller, CPU utilization averaged less than 11%, even with the VMware guests, Exchange databases, and logs on the same storage frame, whereas the average CPU utilization on the Enterprise Vault server was 85%, within the same time period. The Enterprise Vault 8 performance guide estimates a dual core server can ingest 25,000 items per hour with an average size of 70kB. Our performance throughput of 28,514 items per hour was 14% better, even with a larger average message size. A faster Enterprise Vault server with more cores and dedicated network connections would have significantly improved the performance. Refer to the Enterprise Vault performance guide for further details.

4.5 CASCADING SNAPSHOT COPIES METHODOLOGY

Symantec Enterprise Vault deployments require a well-designed storage and network architecture to address data availability, performance, and reliability. Enterprise Vault has numerous moving parts in reference to its underlying data sets. These various components might be distributed across a number of servers within a data center. Enterprise Vault operates with Microsoft SQL Server and various client services to provide archiving, management, and data lifecycle functions. Without a solution such as cascading Snapshot copies, it is a challenging task to back up the distributed components of Enterprise Vault data sets while maintaining end-to-end data consistency. Each component of an Enterprise Vault data set plays an important role. For example, following a disaster recovery all databases, indexes, and vault stores must be recovered following the procedures described in the recovery section of the Enterprise Vault administrator's guide. If the database was recovered from a different state than the indexes or archives, there will be an inconsistency that might result in lost data and extended production downtime. Likewise, if indexes are lost or corrupt, it might take days to complete the rebuild cycle. This rebuild activity on the server will likely interrupt certain user operations and compete for resources with other Enterprise Vault activities. The loss of user productivity and service interruption will create a lack of confidence in the technology and will prevent IT operation staff from focusing on value-added tasks.

This methodology provides a solution to address such issues by creating backups using cascading Snapshot copies. Creation of these Snapshot copies occurs almost instantly. With this strategy, the Enterprise Vault server is placed in a read-only mode for a very brief period while the Snapshot operation is triggered on the database and data files. Once the operation is complete, the server is returned to a normal read-write mode. Then SnapMirror will be used to transfer minimal data to the DR or backup site. By following the example of this paper, an effective archive disaster recovery solution can be facilitated on NetApp storage solutions.

The backup procedure to maintain data consistency without the cascading Snapshot copies involves a complex and time-consuming alternative process that might not capture a point-in-time image of the entire distributed Enterprise Vault environment. By using the cascading Snapshot copy methodology, an Enterprise Vault administrator can configure an elegant solution for backup and recovery requirements. The example script found in TR-3709 can be easily extended to handle multiple Enterprise Vault and SQL Servers within a site.

The following image illustrates how the point-in-time Snapshot copies of the SQL database, Enterprise Vault indexes, and the Enterprise Vault archive are captured with the solution we propose.

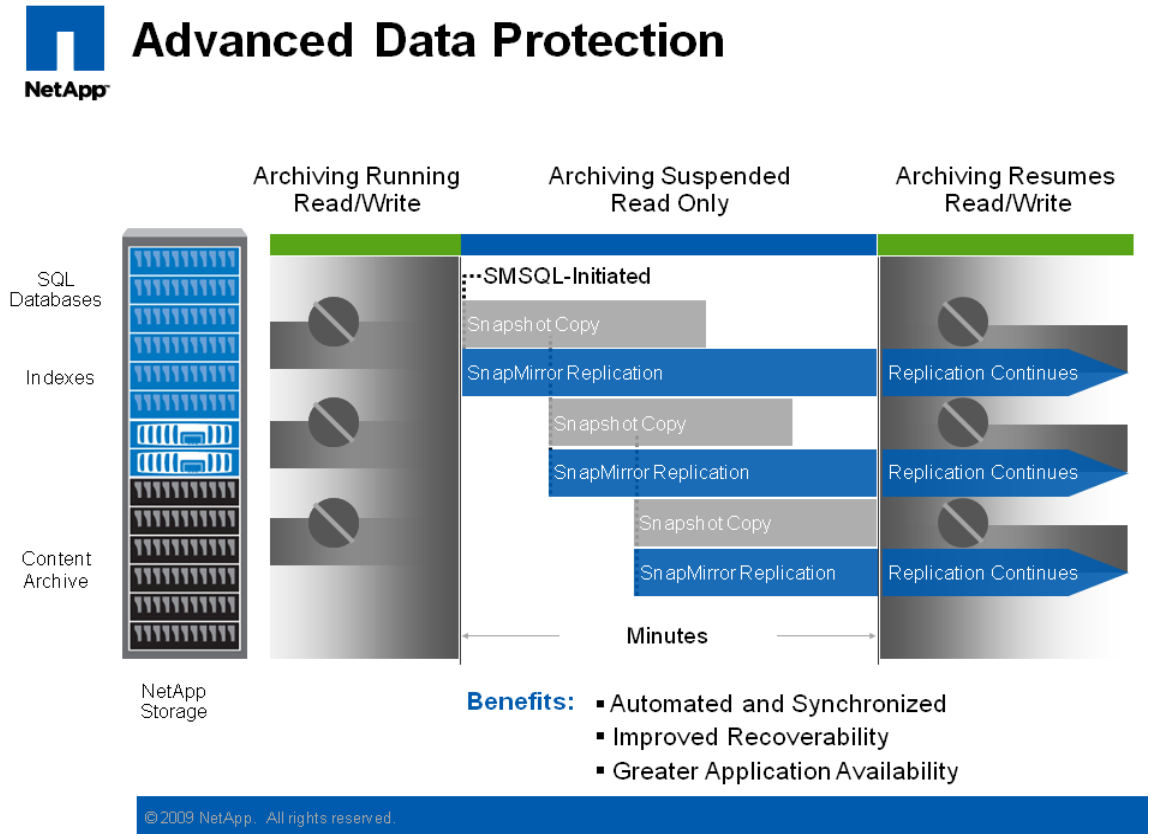

Figure 2) NetApp cascading Snapshot copies.

4.6 HIGH AVAILABILITY

To provide mission-critical and high-availability solutions to customers, Symantec offers a planned strategy in case of system downtime. High-availability server configurations for Enterprise Vault use clustering or Enterprise Vault Update Service Location (USL)/building blocks configuration. With these technologies, Enterprise Vault can be configured in active-active or active-passive (N+1) mode.

With the USL active-active mode, Enterprise Vault services can run on both servers simultaneously (using USL only). If one server fails, the second server takes over the additional services.

With the USL clustered active-passive mode, all Enterprise Vault services run on one server. The passive server simply waits in standby until the production server fails.

Here are several possible Enterprise Vault high-availability solutions in use at customer sites:

- SAN or NAS boot
- Enterprise Vault Update Service Location (USL)
- Active-passive pair
- Enterprise Vault Warm Standby (N+1)
- Clustering with Veritas™ Cluster Server (VCS/SFW-HA)
- Clustering with Microsoft Server clusters: MSCS

NetApp recommends using volume management software to handle movement of SAN disk resources on the Enterprise Vault servers. This facilitates faster and easier failover with USL. For example, SnapDrive for Windows, SnapMirror, or Protection Manager can be used to accomplish these activities.

4.6.1 UPDATE SERVICE LOCATION

To use the Update Service Location function, Enterprise Vault must be installed with DNS aliases for all the physical computers. This abstraction layer creates a virtualization of Enterprise Vault computers and the services that run on them. Essentially, the hardware is decoupled from the identities defined in Enterprise Vault. When a failure takes place, the computer DNS alias can be directed to either another server running Enterprise Vault (active-active) or a hot spare (N+1). The USL command is run, and Enterprise Vault checks which services should be configured on the server to which the alias is pointing. If new services are needed, they are created automatically; if there are more services than required, they are removed. As long as the underlying Enterprise Vault database is still available, user downtime can be calculated by the length of time it takes to update a DNS alias and run the Enterprise Vault Update Service Location command. If Enterprise Vault is to be configured in an active-active solution, you must calculate the workload for two servers and make sure that a single server can handle the extra requests. When running in a failed state, Enterprise Vault should be configured to run in read-only mode. Running the SQL Server on the same hardware as the archive application will make the USL failover a little more complicated. You should consider running the SQL databases for Enterprise Vault on dedicated hardware.

Figure 3 depicts a recovery of the application by employing the Enterprise Vault USL capabilities after having failed over the database to the recovery server.

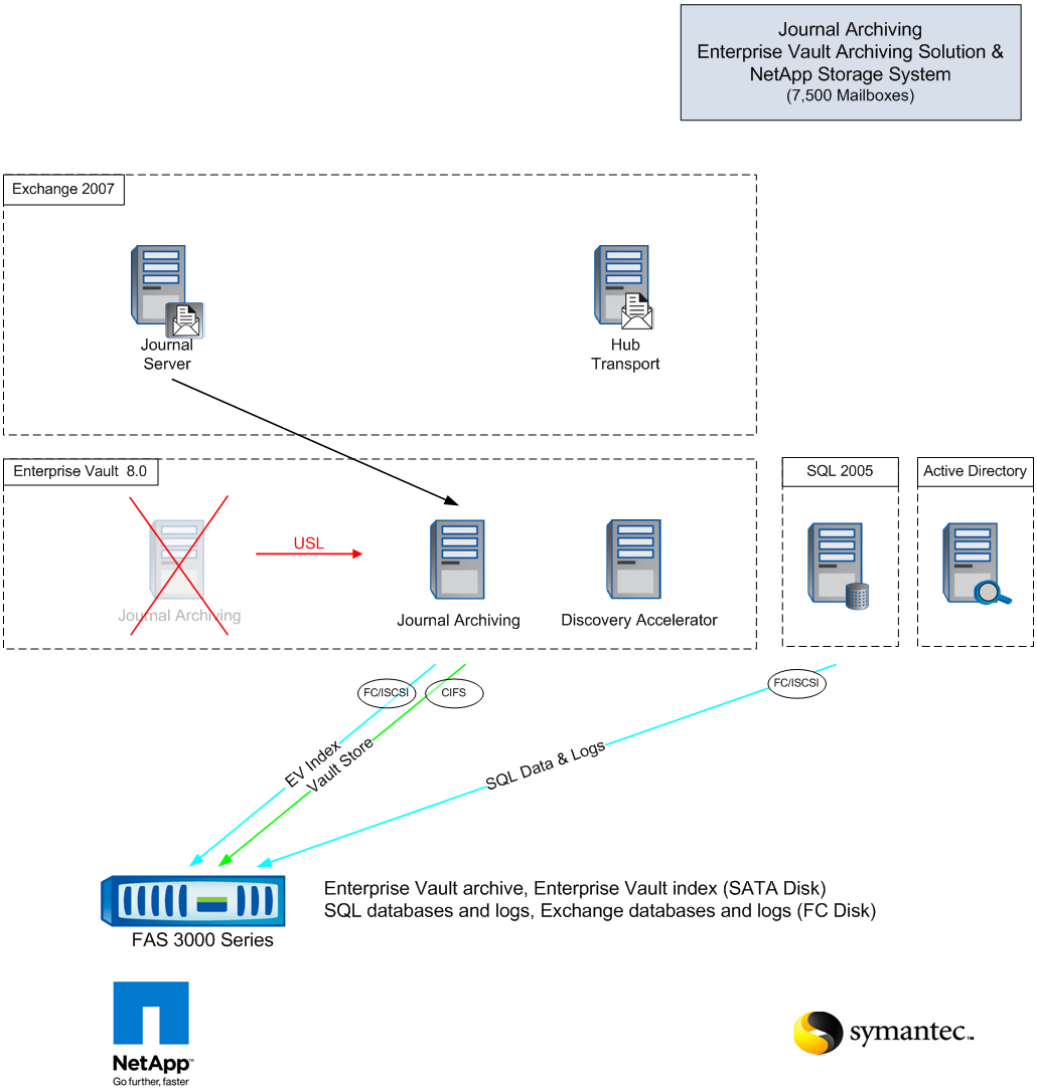

Figure 3) Enterprise Vault 8 recommended architecture for high availability.

4.7 HIGH-LEVEL RESTORE SCENARIOS

The following table lists some typical disaster scenarios and possible options to remediate the issue. You should test these scenarios within a lab environment to build the specific processes for your company.

Table 11) Disaster recovery options.

Scenario	Impact	Restore Action
A single Enterprise Vault server failure	All access is unavailable.	Perform disaster recovery process, described in Chapter 23, Enterprise Vault Administrator's Guide.
Loss of SQL Server	All access is unavailable.	Perform disaster recovery process, described in Chapter 23, Enterprise Vault Administrator's Guide. Restore Microsoft SQL databases by using SnapManager for SQL.
Loss of a single NetApp storage controller	All access is unavailable, if not clustered. No service impact if clustered.	If clustered, the surviving storage controller at the same site automatically takes the name of the failed controller and services all iSCSI disk requests from all local servers.
Loss or corruption of index data	User Web search and archive explorer might be unavailable.	Recover corrupted or lost folder by using SnapRestore. Rebuilding the index is an alternative, although time-consuming, option.
Loss of directory SQL database	All archiving, retrieval, and search are unavailable.	Stop all Enterprise Vault services and restore Microsoft SQL databases using SnapManager for SQL Server.
Loss of vault store SQL database	Archiving and retrieval of the lost database are unavailable.	Stop all Enterprise Vault services and restore Microsoft SQL databases using SnapManager for SQL Server.
Loss of fingerprint SQL database	Archiving and retrieval of the lost database are unavailable.	Stop all Enterprise Vault services and restore Microsoft SQL databases using SnapManager for SQL Server.
Loss of vault store data	Archiving to and retrieving from the lost vault store are unavailable.	Recover lost folder by using SnapRestore.
Full disaster, such as loss of all servers	All access is unavailable.	Perform disaster recovery process described in Chapter 23 of the Enterprise Vault Administrator's Guide.

5 BEST PRACTICE RECOMMENDED CONFIGURATIONS

5.1 STORAGE LAYOUT BEST PRACTICES

NetApp recommends putting the MSMQ storage area on a separate volume or separate spindles if possible to enhance performance as Enterprise Vault uses MSMQ extensively.

If necessary, shopping locations can be located on the same volume as the index or vault store data.

The Enterprise Vault configuration process now creates eight index locations per indexing service. These can be placed on separate LUNs, placed in folders inside one larger LUN, or in separate CIFS shares. The LUN and folder sizes should be designed to accommodate efficient and regular backups. Your index

locations should be placed on NetApp storage and protected by SnapDrive for Windows. SnapDrive helps host administrators provision storage and manage it directly from the host, enabling them to flexibly define backup policies and resize storage on the fly without any disruption of application service. SnapDrive understands the operating system, volume manager, and file system details necessary to coordinate Snapshot activities between the host and NetApp storage system. Refer to TR-3197 for a more complete technical discussion of SnapDrive.

5.2 I/O REQUIREMENTS AND DISK SELECTION

It is important to strike a balance between storage capacity and I/O throughput. Larger drives such as 1TB SATA drives provide tremendous storage capacity and value, but with these larger drives the amount of I/O that can be handled per given storage unit decreases. Yet, for sequential reads and writes NetApp SATA drives perform as well as FC drives since WAFL keeps the head stationary for constant streams, and the higher data density compensates for the slower disk RPM. Care should be taken to understand the amount of I/O that a given disk subsystem or volume can provide.

If indexes are to be placed on SATA disks in an environment of this size, then the volumes used by the Enterprise Vault indexes should be sized to provide 2,000 IOPs for adequate end-user search, discovery search, archiving, and reindexing performance.

5.3 OPPORTUNISTIC LOCKING ON THE NETAPP STORAGE SYSTEM

The CIFS protocol allows a client to request the ability to cache locally the contents and attributes of an open file. This usually results in a dramatic performance gain. However, NetApp recommends disabling opportunistic locking (oplock) for Enterprise Vault indexes because the indexes usually contain extremely important data in large files, so if oplock is enabled, a lot of important data being cached on the client could get lost if the network or the power fails. Opportunistic locking is relevant only for CIFS shares; it is not needed for a LUN.

Opportunistic locking can be disabled on the NetApp storage by using the command option `cifs.oplocks.enable off` or by using FilerView®.

5.4 OPPORTUNISTIC LOCKING ON ENTERPRISE VAULT ARCHIVING SERVER

It is a best practice to disable opportunistic locking on the Enterprise Vault server as well, so that opportunistic locking is not used even if it is enabled on the NetApp storage solution. This can be done by configuring the OpslocksDisabled registry key on the Enterprise Vault server.

For more information, see the following technote: http://seer.entsupport.symantec.com/docs/280922.htm.

5.5 HARDWARE INITIATOR FOR ISCSI CONFIGURATIONS

There is some evidence that a software iSCSI initiator can affect CPU performance. Therefore NetApp recommends attaching the server to the storage by using a hardware initiator card for iSCSI.

A hardware initiator uses dedicated hardware, typically in combination with software (firmware) running on that hardware, to implement iSCSI. A hardware initiator mitigates the overhead of iSCSI, TCP/IP processing, and Ethernet interrupts and therefore might improve the performance of servers that use iSCSI. This solution blueprint was not developed with an iSCSI hardware accelerator.

Refer to the Symantec Enterprise Vault Indexing Best Practice Guide for further guidelines when using iSCSI for the index.

6 STORAGE EFFICIENCY COMMENTARY

We have developed a survey-based methodology to evaluate a storage environment and then calculate the cost and capacity savings of our efficiency technologies. The calculator has approximately 50 input fields. An analysis of our 7,500-user environment was used and compared against a generic, centralized storage environment without the benefit of our storage efficiencies. We applied the values and growth projections for this e-mail archiving scenario and filled in remaining values based upon IDC's "Storage Workloads" model. The resulting analysis shows a 20% reduction in the total cost of ownership over a three-year period.

We estimate it would cost U.S.$605,000 at the primary data center over three years for storage augmentation, storage maintenance, power, cooling, and data center floor space costs to support the Exchange and Enterprise Vault services. The same considerations applied to NetApp storage solutions cost only U.S.$395,000. This resulted in a savings of U.S.$209,000, or 65%. The kilowatt per hour costs were those for May 2009 in San Francisco, California. When using NetApp storage solutions the power and cooling costs alone were 24% lower. Using NetApp requires 59% less storage, so augmenting storage capacity can be diverted to a later date.

Because of the complexities and markdowns during the typical sales cycle, it was assumed the year one storage was already in place. The calculations did not account for any soft cost savings such as IT staff or end user productivity. The calculations assumed all data would be replicated to a disaster recovery environment, and therefore Snapshot copies are required. As discussed in TR-3635, the use of cascading Snapshot copies creates a consistent point-in-time copy of the Enterprise Vault index, archive, and databases within a matter of minutes, versus traditional backups, which require a backup window proportional to the size of the target data.

These figures represent a scenario of 7,500-user journal archiving on our storage technologies. Your return on investment might be greater or less than this. For further details on this particular analysis, or to have one created for your storage environment, contact your local NetApp sales representative.

7 CONCLUSION

Our storage systems allow you to efficiently consolidate SAN, NAS, primary, and secondary storage on a single platform. You get the ultimate in scalability, versatility, and availability for your e-mail archives when you use our technologies.

This technical report highlights the importance of the message that a joint solution based on Symantec Enterprise Vault and NetApp storage solutions provides the best-in-class archiving technology for your company. A NetApp storage solution complements Enterprise Vault capabilities in a simplified architecture with support for backup and restore functions.

The model described in this paper gives an overview of the Enterprise Vault architecture for 7,500-user mailbox journal archiving. This paper serves as a starting guide for designing and deploying a high performance e-mail archive based on Symantec Enterprise Vault and NetApp storage solutions. During the design phase, it is important to collaborate with Microsoft Exchange and SQL Server specialists in addition to your local Enterprise Vault architects. The partnership you build with these experts will guide you through the analysis, planning, and successful deployment of your archiving solution based on NetApp and Symantec technologies. We invite you to contact your local sales representative for a thorough review of your e-mail archive needs.

Data growth is inevitable. Reducing data center power, cooling, and space costs while storing the maximum amount of data for the lowest possible cost and for longer periods of time should not come with the penalty of reduced performance or increased administration. "Storage efficiency without compromise" is NetApp's promise: We will reduce your storage burden while improving manageability and performance.

8 REFERENCES

8.1 TECHNICAL REPORTS

TR-3525: Storage Performance Management

TR-3635: Symantec Enterprise Vault Data Protection with Network Appliance Storage System

TR-3487: SnapVault Best Practice

TR-3446: SnapMirror Best Practices

TR-3500: Installing Enterprise Vault with NetApp Storage Systems

TR-3298: NetApp Implementation of RAID Double Parity for Data Protection

TR-3765: Enterprise Vault 8.0 Storage Efficiency on NetApp Storage

TR-3719: Technical Overview of NetApp SnapDrive

8.2 BEST PRACTICES GUIDES

Symantec Enterprise Vault 8 Performance Guide

Symantec Enterprise Vault Indexing Best Practice Guide

8.3 COMPATIBILITY MATRIX

Symantec Enterprise Vault 6.0, 7.0, 2007, and 8.0 Compatibility List