Technical report:

Microsoft Exchange and IBM System Storage N series with NearStore and Symantec Enterprise Vault

Integrating for Data Archival

Document NS3267-0

October 9, 2007

Table of contents

Abstract

Regulatory compliance and other data storage requirements demand efficient and reliable data archival solutions. IBM System Storage N series with NearStore and SnapLock features in combination with Symantec Enterprise Vault offers a simple, cost-effective, and best-of-breed solution to archive e-mail and other messages in a Microsoft Exchange environment.

Introduction

Securities and Exchange Commission (SEC) regulations require financial broker-dealers to retain e-mails and messages for a specified period of time after their origination. In addition, internal policies are increasingly being established by companies across various industries requiring e-mails be retained for future reference and/or audit. As a result, retaining e-mail messages for longer periods of time and being able to quickly search and retrieve specific records are becoming critical capabilities for many businesses. According to the SEC, financial broker-dealers must also exclusively preserve key business records such as e-mail as non-erasable, non-rewritable (including as WORM—write once, read many—media volumes) that is fully indexed and easily searchable for two years from origination. Given this demanding set of requirements, how can businesses easily and effectively ensure compliance with these regulations?

Symantec has a suite of products available to assist businesses with achieving and maintaining compliance for e-mail and message archival regulations. Enterprise Vault leads the way in archival functionality. In combination, non-erasable, non-rewritable WORM volume functionality (via SnapLock® software) is now available on magnetic disk drive-based storage (leveraging optical and tape media storage) through the IBM® System Storage™ N series with NearStore® feature. The combined solution introduces critical new technology performance to further assist businesses with achieving regulatory compliance. The combination of Enterprise Vault using SnapLock on the IBM N series with NearStore feature storage product offers a simple, cost-effective, and best-of-breed solution to businesses required to archive e-mail and messages.

Purpose and Scope

This paper will provide a complete and tested solution using Enterprise Vault to archive e-mails and messages on an IBM N series with NearStore feature. All steps required to implement the solution outlined in this paper will be provided. This paper will address relevant background information on required infrastructure and IBM N series with NearStore technology including SnapLock.

The solution covered in this paper is limited to Enterprise Vault retention of Microsoft® Exchange e-mails and messages to an IBM N series NearStore volume. Symantec offers additional tools that can further assist with both initial archival and ongoing management of archived e-mails and messages. Uses for the NearStore feature or NearStore SnapLock technology besides for e-mail and message archival are outside the scope of this paper.

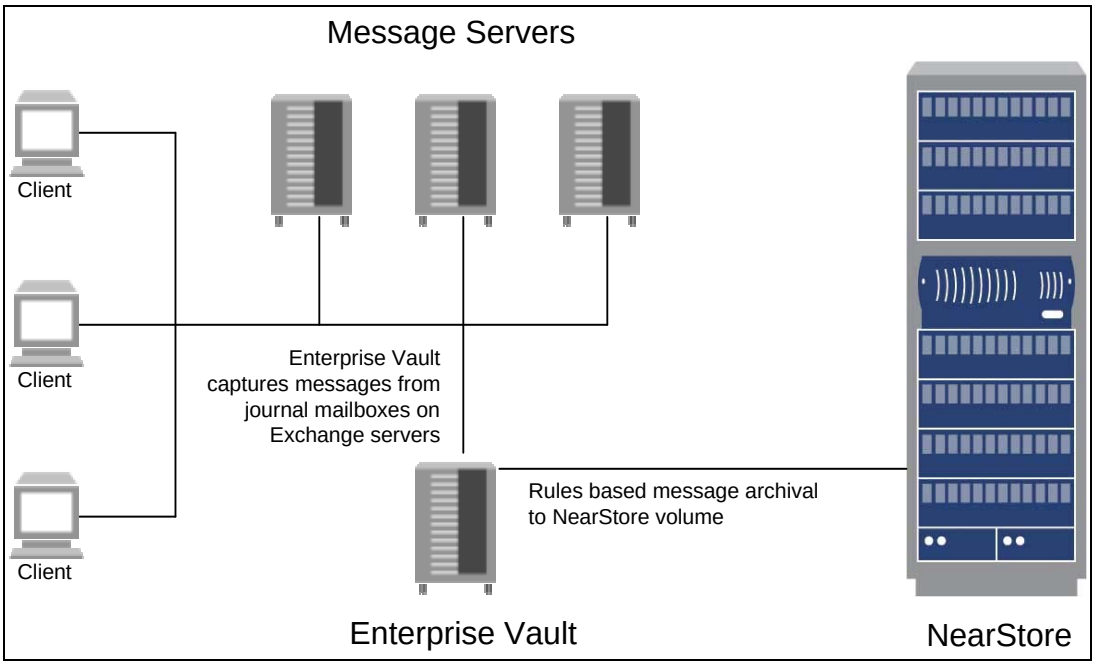

Figure 1. IBM N series with NearStore Feature and Enterprise Vault solution architecture.

In a typical deployment scenario, the system running Enterprise Vault is logically situated between the Exchange servers and the IBM N series with NearStore. Enterprise Vault monitors message activity on Exchange servers and takes actions depending on how its rules are configured. In larger deployments, additional servers may be required to effectively run Enterprise Vault, Exchange, and Active Directories Domain Controllers (ADDCs). Other options for larger deployments would be a private network for Enterprise Vault connectivity to Exchange servers, ADDCs, and the IBM N series with NearStore.

Enterprise Vault archival and retrieval of data between the Microsoft Exchange server and IBM N series with NearStore are transparent to the user. IBM N series with NearStore is a common internet file system (CIFS) share that the Enterprise Vault server maps as a network drive. After network and systems are configured, policies are created within Enterprise Vault to automate the entire message archival process.

Infrastructure

Network

Connectivity

Most institutions required to comply with message retention standards are large enough to warrant having private gigabit network connections between the Enterprise Vault server, the Exchange server, the ADDC server, and the IBM N series with NearStore data store. Either setting up a separate switch or creating a virtual LAN (VLAN) on existing switches would equally suffice. Client connectivity to the Exchange server can continue over current network infrastructure.

Isolating the message archival network in the manner just described affords two important benefits. Potential latency and contention between systems used in this solution white paper are eliminated by minimizing the number of overall servers attached to the network. This proposed network implementation also improves security of the e-mail archival. Having this archival available on a shared network drastically increases the risk of unwanted or illegal data access.

Existing Microsoft Environment

Exchange servers and ADDCs should already exist on the network. Other standard network services, such as a domain name server (DNS) or a timeserver, should also be present.

Enterprise Vault

Hardware

The actual hardware requirements for Enterprise Vault can widely vary, depending on the number of Exchange servers, amount of data being archived, bandwidth, and storage medium. Contact your account representative for specific requirements to fit your environment. For a test or pilot install, we suggest the Enterprise Vault run on an independent, enterprise-class server.

Software

At a minimum:

- Enterprise Vault requires Microsoft Windows 2000 Server with SP 2 and IIS and MSMQ enabled.
- SQL Server 2000 with Service Pack 2 is required (on the same server or on the network).
- Outlook 2000 Service Pack 2 with CDO should be installed on the same designated Vault server prior to Vault installation.

There are additional software requirements for Exchange servers utilizing Outlook Web Access (OWA). Please contact appropriate representatives for more information.

IBM N series with NearStore

System Requirements

Support for IBM N series with NearStore and SnapLock volumes is provided in IBM System Storage N series with Data ONTAP® 7.1 or later. In addition, please note that Data ONTAP 7.1 and later deliver additional capabilities (e.g., flexible volumes in aggregates) outside of the scope of this report.

Licensing Requirements

SnapLock and CIFS are both license-based features on IBM N series appliances.

IBM N series with NearStore/SnapLock Configuration

Enterprise Vault integrates with an IBM N series with NearStore whether its volumes were created as non-WORM, as SnapLock (non-erasable, non-rewritable WORM volumes), or in combinations of each. The solution in this paper uses IBM N series with NearStore with SnapLock storage for e-mail and message archival. This flexibility in an Enterprise Vault and IBM N series with NearStore environment gives businesses configuration options to meet various needs while maximizing their investment in an archival infrastructure.

SnapLock technology provides the same non-erasable, non-rewritable guarantees as traditional optical or tape WORM volumes, substantial performance improvements over traditional WORM media, and adherence to open protocols for data access. IBM N series with NearStore implementation of SnapLock on magnetic disk drives utilizes CIFS and network file system (NFS) open protocols to store and access archived data. The open protocol aspect of this solution provides a very natural and flexible way to manage, store, and retrieve WORM volumes via regular CIFS and NFS clients. Other disk-based WORM volume storage solutions require the use of proprietary application programming interfaces (APIs) and complicated file-addressing schemes to store and access data.

Creation of a SnapLock aggregate and SnapLock volume on IBM N series with NearStore

Creating a SnapLock volume in Data ONTAP version 7.1 and later must be done from either the console or a telnet or ssh session to the IBM N series with NearStore and not through the IBM System Storage N series with FilerView® GUI. After a session is established, use the following command:

```
nearstore> aggr create snap_aggr –L 3
nearstore> vol create snap_vol snap_aggr 100g
nearstore> cifs shares -add snap_vol /vol/snap_vol
```

Provided SnapLock Compliance was licensed, using the –L switch in the above `aggr` command creates a SnapLock Compliance aggregate called `snap_aggr1` comprised of 3 disks. If SnapLock Enterprise was licensed, then that would be the resulting SnapLock aggregate type. If both SnapLock Compliance and SnapLock Enterprise are licensed simultaneously, then either Compliance or Enterprise would follow the –L switch to select the proper SnapLock aggregate type. The second command creates a 100GB SnapLock Compliance flexible volume called `snap_vol` in the `snap_aggr` aggregate. The third command allows Microsoft servers and clients to access the volume as a CIFS share.

Once the above steps are completed, verify the SnapLock Compliance volume and CIFS share status by using `vol status` and `cifs shares`, respectively.

With the permanence of SnapLock, prior consideration should go into how SnapLock volumes and their underlying RAID groups are created and maintained. Before creating a SnapLock volume, please see separate technical documents concerning SnapLock functionality and best-practice guidelines.

IBM N series with NearStore CIFS share

IBM N series with NearStore CIFS share of SnapLock volumes supports all standard authentication and security features implemented in a Microsoft AD environment. This capability allows for seamless integration with existing corporate security policies for data protection.

Enterprise Vault Configuration

Enterprise Vault stores archived messages as individual files in the IBM N series with NearStore. Vaults are collected together in vault stores, which contain vault store partitions. A vault store partition resides on a Windows NT® File System (NTFS) volume set or other type of storage device. When you set up Enterprise Vault, you specify on which device each partition is created.

A single vault store can be divided into a number of partitions, which can be on different disks or media. As a vault store grows, you can add partitions to extend the space available. A vault store can have only one active partition, which is the partition into which all new items are archived. You can change which partition is active at any time. As you enable each mailbox, you specify which vaults are assigned to which vault stores.

Once a drive or pathway has been created and the Vault service account has been given proper access rights to it as outlined in the previous section, start the New Vault Store wizard from the Administration Console. To do this from the Administration Console—a Microsoft Management Console (MMC) snap in—right-click Vault Store container, point to New and click Vault Store. Alternatively, click the Add New Vault Store icon on the toolbar.

When the store is named and the databases are selected, you are prompted to create a partition (Figure 2).

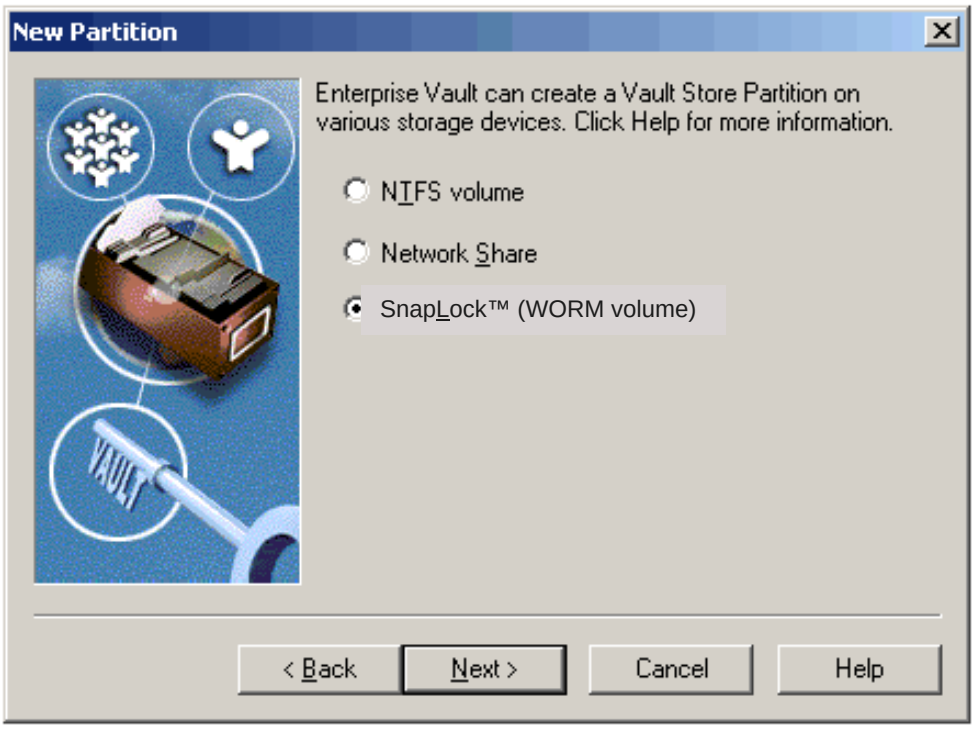

Figure 2. Creating a partition with Enterprise Vault using SnapLock.

Here, you will be asked for the type of storage solution. For non-WORM storage on an IBM N series with NearStore, the second option Network Share is used. For non-erasable, non-rewritable SnapLock WORM volume storage, the highlighted option (as seen above) is used. Either choice leads the administrator to a wizard page where he or she is asked to select a folder on a network share (Figure 3). Again, make sure that the Service Account for the Enterprise Vault has been given read/write permissions to this folder.

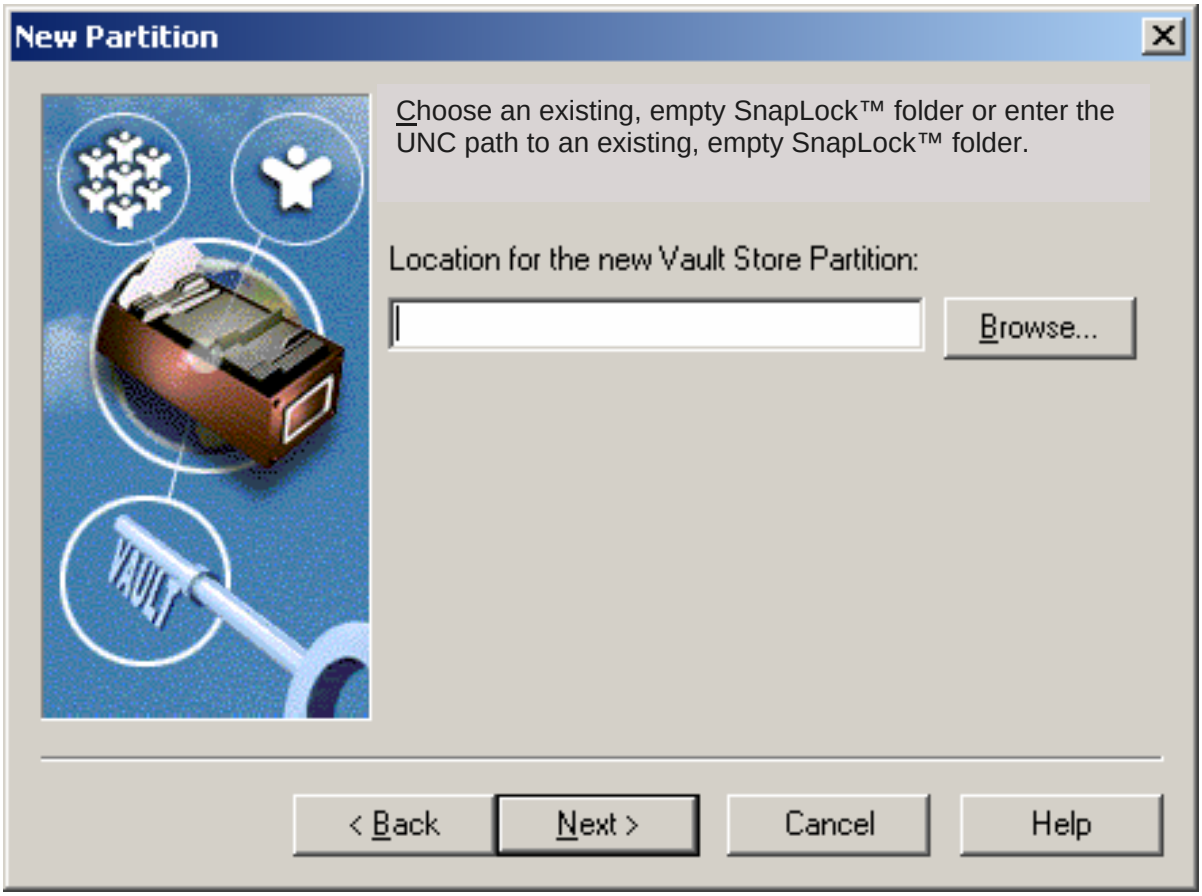

Figure 3. Locating the new partition.

Share archived items: Select this to make Enterprise Vault optimize the storage of shared items by archiving to a single file on the storage media rather than archiving multiple copies. For example, if a message that has been sent to a large distribution list is archived, the message is stored on disk just once, rather than once for each recipient. There is a slight performance cost in sharing the storage, but the space savings normally outweigh this cost.

On a SnapLock share, you may want to uncheck this option to prevent Enterprise Vault from attempting to update files since they are already stored uniquely and cannot be modified once committed to non-erasable, non-rewritable WORM volume status.

Caveats

IBM N series with NearStore SnapLock considerations

Files on an IBM N series with NearStore SnapLock volume get their SnapLock state flag set when their status is changed to read-only. Once this trigger event has occurred, attempts to modify or delete the file will fail. This is true no matter which user triggered the transaction to SnapLock state and which user is trying to modify or delete the file (i.e., administrator). Deletion of directories is allowed if no SnapLock state files exist within their hierarchy.

Data protection

IBM N series with NearStore utilizes RAID-DP to provide fault tolerance against possible hard drive failures. In the event of a single or dual disk drive failure, a copy of all the data still exists via the two parity drives and the IBM N series with NearStore will immediately begin rebuilding data on an available spare drive. Given this scenario, traditional IBM N series best practices guidelines should be followed so two or more hot standby drives is available in the event of a single or dual drive disk failure. Critical business considerations like disaster recovery and business continuance in the event of a IBM N series with NearStore outage are outside the scope of this paper but do warrant further contingency planning.

Trademarks and Special Notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP, FilerView, Snapshot and WAFL are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.