Technical report:

Symantec Enterprise Vault and IBM System Storage N series with NearStore

Deployment Guide

Document NS3382-0

September 24, 2007

Table of contents

Abstract

Symantec has a compelling product to meet e-mail and file-system storing, archiving, and management challenges and increase the customer experience. This paper describes the method to integrate Symantec Enterprise Vault with an IBM System Storage N series with NearStore feature configuration to take advantage of file system archival as well as the backup, recovery, and compliance solutions from IBM.

Executive Summary

Messaging and collaboration servicing combined with an efficient archival solution in a configuration where a storage solution can improve the total customer experience is the secret to addressing the customer problems in that space. Symantec adds value for customers taking advantage of both technologies. Symantec recently added the file system archival (FSA) feature in addition to e-mail archival in its Enterprise Vault product. This paper discusses the procedure to integrate the Symantec product and the required components and other products with an IBM® System Storage™ N series with NearStore® feature solution. This paper specifically will discuss the FSA feature in such an IBM N series environment.

Background

In the age of information, e-mail has become the mode of communication in the business community. Attachments to e-mail have served a business purpose in the communication process, and it is now a norm in day-to-day business processes. The increased acceptance of communication systems has led to documents and messages being stored electronically. Once the information is stored electronically, need for an efficient method of storing, archiving, and managing it is a challenge. In addition to e-mail archival, FSA offers a new feature to archive needed files and save storage space. Symantec has a compelling product to meet these challenges and increase the customer experience. In addition to e-mail archival, end users are looking for ways to manage file systems and to intelligently archive and manage files.

A number of compliance regulations have been enacted mandating archiving the content and being able to produce the information when needed. Although not all businesses are required to follow these regulations, it is interesting to see that businesses are looking for ways to protect the contents. Regulations such as SEC Rule 17a-4, the Healthcare Insurance Portability and Accountability Act (HIPAA), and CFR 21 are forcing certain businesses such as those in financial, insurance, and healthcare sectors to protect the content.

In addition to the compliance regulations, a large number of customers are adopting the policy to protect the data. In addition to archiving and managing e-mail communication, FSA and managing communication have become important requirements. To solve the customer's business needs, Symantec introduced the FSA feature with its Enterprise Vault 5.0 Service Pack 3 product. Businesses require an efficient backup and restore method in the file archival configuration. IBM N series adds value in providing performance-improving, backup, data replication, and data recovery features that come with its own operating system, called IBM System Storage N series with Data ONTAP®.

Introduction

Businesses need a plan to store and archive content (such as e-mail, faxes, and files) that enables them to search quickly, yet provides data security. E-mail content has to be stored and managed, and users must be able to search it and obtain the right content as needed. Enterprise Vault has these capabilities. It is also important to understand the fundamentals of unstructured information lifecycle management to know the business needs. Enteprise Vault provides functions and features to achieve business success. This paper explains the basic procedure and steps involved to integrate Enterprise Vault software with IBM N series storage solutions. It also discusses the FSA feature of Enterprise Vault in IBM N series configurations. FSA allows archiving and managing the files on the file system, including private and public folders.

The benefits of file archival, performance and backup, recovery, and configurations to include disaster recovery by deploying the data replication features of IBM N series will be discussed in a separate paper.

Symantec E-mail Archival

Users have continued to have issues with managing e-mail content because of the rate of increase in content growth. In addition, the increase in the number of attachments to e-mail messages has been the primary reason for the increase in e-mail content handled by e-mail servers.

Developing software to archive e-mail by optimizing the use of storage, yet provide management simplification was a challenge. Symantec has a suite of products to address this challenge and provide a method for quick implementation, a system that is transparent to users with an open application programmable interfaces (API). Symantec offers Enterprise Vault, which provides store, index search, and index retrieve capabilities in Microsoft® Windows® Exchange environments. FSA in Enterprise Vault enables file archival and search and retrieval capabilities. Symantec provides open APIs to store, manage, and discover any content in a customer's environment. The system provides archiving services to (minimum versions listed) Microsoft Exchange 5.5, Windows 2000 and 2003, and PST files; SharePoint Portal V1 and V2; Microsoft File System (NTFS/CIFS); instant messaging, and Internet mail (SMTP).

Enterprise Vault works in Windows Exchange environments, and it is helpful in understanding the need for Enterprise Vault where the mailboxes store the data. As data grows, it is hard to maintain stability and yet provide performance in the Exchange Server environment. Attempts to solve this issue included creating the auto-archive with PSTs, setting quotas, and even adding Exchange Server. The above workaround may not solve all the issues, and cost of deployment becomes a major issue. Exchange Server provides the content for Enterprise Vault to be stored in NTFS configuration. Instead of placing the messages in a database, as in Exchange, Enterprise Vault stores the messages as standard files, one file per message in one or more NTFS partitions. This approach will offer extensive benefits such as the following:

- No additional maintenance is required
- It is easy to recover a corrupted item as compared to the entire database
- More data can be stored without performance degradation
- Central management of storage for efficient usage of storage
- Support of multiple Exchange Servers.

Figure 1 shows a simple configuration of Enterprise Vault managing e-mail and FSA.

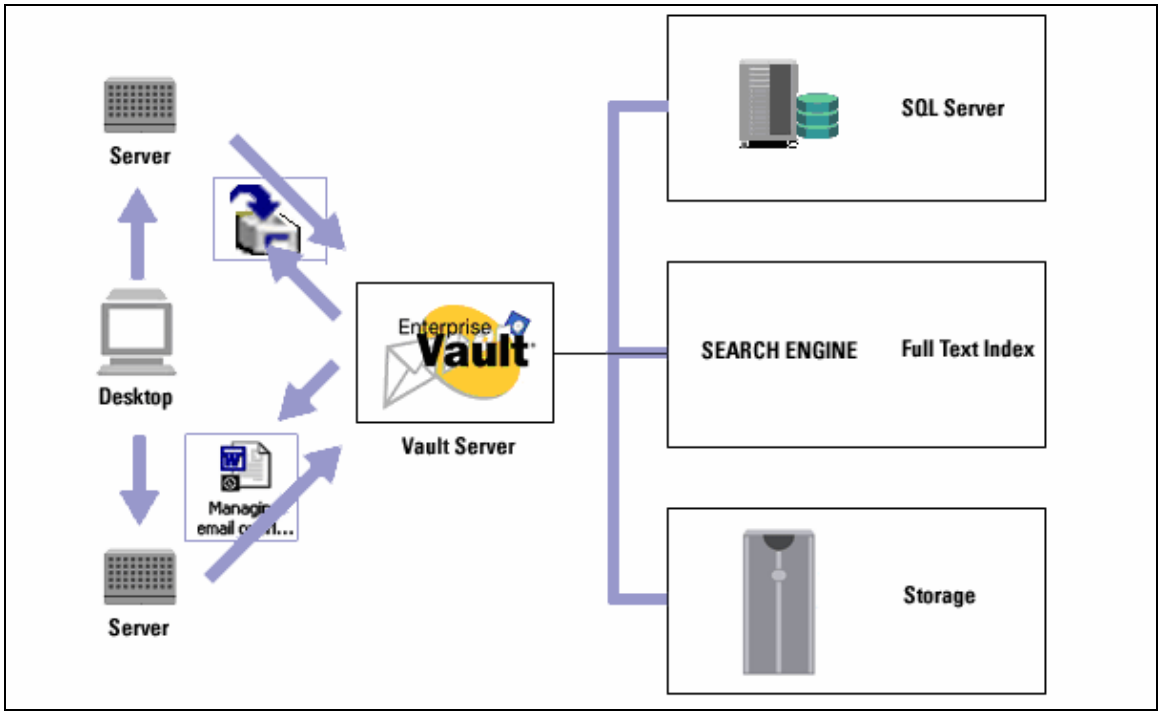

Figure 1) Enterprise Vault environment configuration.

Symantec File System Archival

Enterprise Vault has functionality that extends beyond e-mail to file-based data. It includes the capability to archive the system files. It is also possible to configure Enterprise Vault to use the FSA component. FSA functionality is limited to file systems that can be presented as NTFS. Files that reside on network share, corporate data, and documents are typical examples for archiving into an Enterprise Vault system. Note that it is dangerous to archive Windows system files, and this paper strongly recommends excluding the archival of system files and other files that are critical to running the operating system. If all the files in the operating system are archived, the operating system server will not be able to be restarted. In such scenarios, the system has to be recovered using the system backup, and Enterprise Vault data has to be recovered from its backup.

The creator and migrator tools in Enterprise Vault allow moving the archived data to a different device. By using FSA, files can be stored where it is easier to manage the data in terms of backup/recovery and data replication features. Archived files provide a centralized location for information storage and allow data mining. FSA provides flexible policy control and quick recovery of data.

It is important to note that Enterprise Vault allows backup of the content that is required to be stored with lifecycle management and provides end users with access to the data. The Enterprise Vault file system adds value by addressing data management and maintains end-user access. Flexible policies can be set to archive and manage a specified type of files. A single archive for all data is allowed without the restriction of files, electronic content, or other content type.

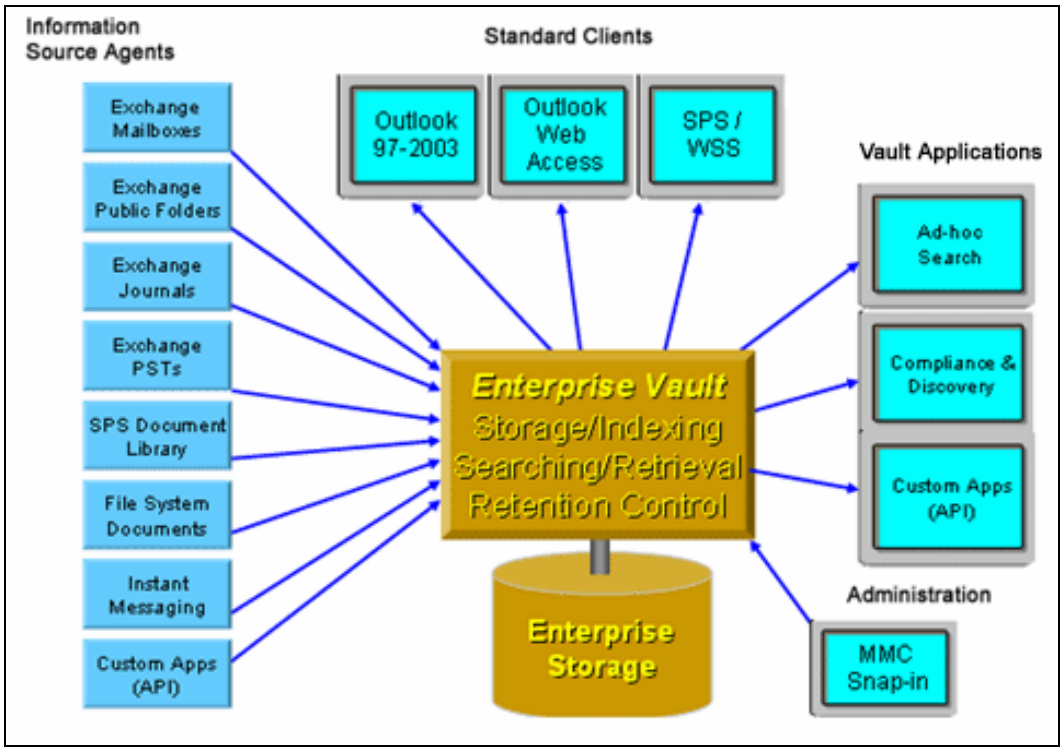

Figure 2) Enterprise Vault overview.

Figure 3 shows a framework for Enterprise Vault highlighting all the components present in an archiving environment. Note that the file system component is one level below the universal access layer, along with Exchange and SharePoint components. This shows that FSA can be configured with or without other components such as Exchange Server or SharePoint product features.

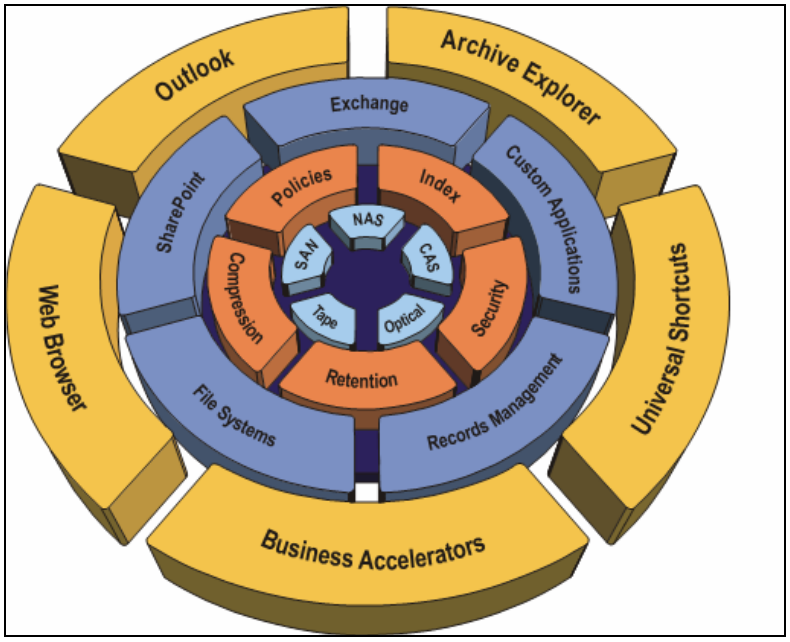

Figure 3) Enterprise Vault: showing the file system archival component.

IBM N series Storage Solution

The IBM N series storage appliance contains many redundant hardware features. Built-in RAID (redundant array of independent disks) protects against downtime due to disk failures. In the event of a disk failure, automatic reconstruction takes place on a hot spare disk with notification sent to the system administrator. The Data ONTAP software's storage health monitor proactively monitors the disk drives and storage connections for any potential problems. Redundant power supplies and cooling fans are included in the system unit and disk shelves. Cooling fan speed and system temperature are also monitored, and notification is sent if there is a problem. Disk drives, power supplies, and cooling fans are all hot swappable. Additional hardware redundancy can be achieved by deploying the following configurations: virtual network interfaces, Fibre Channel multipath, and IBM N series storage cluster configurations.

IBM N series with Data ONTAP 7G allows storage administrators to create more flexible and scalable storage configurations by using features such as IBM System Storage
N series with FlexVol™, where the volume can be grown or shrunk as needed on IBM N series storage.

Network Connectivity

To install and configure Enterprise Vault, the Microsoft Exchange Server SQL Server database requires different types of network connections. Network connection between Exchange Server and IBM N series storage should use either the Fibre Channel Protocol (FCP) or the Internet Protocol (IP)–based Small Computer System Interface (iSCSI) protocol. The storage configured with either iSCSI or FCP can be maintained much more easily with a software feature called IBM System Storage N series with SnapDrive®. SnapDrive software will ease storage device maintenance with data management for backup/recovery or scaling the storage devices as data grows. Providing continuous access to storage is one element of high data availability. The other elements are the integrity and recoverability of the data. The IBM N series storage appliance has several built-in features and optional software for data integrity and protection.

Enterprise Vault also requires a network connection between the Windows Servers and the IBM N series storage devices. This connection must use Gigabit Ethernet.

It has become standard to use a private Gigabit Ethernet network connection between the Enterprise Vault Server, Exchange Server, the Active Directory domain controller server, and the IBM N series data devices. Either setting up a separate switch or creating a VLAN on existing switches would suffice equally. Client connectivity to Exchange Server can continue over current network infrastructure.

FSA allows setting a policy for private and public folders and hence using the network shares. Having this archival available on a shared network drastically increases data storage usage. FSA helps to archive the files and index their content, and hence the data mining becomes an important benefit with FSA.

Microsoft System Environment

Enterprise Vault works on the Windows platform and supports Microsoft Exchange Server and SQL Server relational databases on Windows 2000, Windows XP, and Windows 2003 Server platforms. Microsoft Exchange requires the local disk configuration, and hence the use of iSCSI and/or FCP local disk configurations is required on IBM N series devices. It is suggested that SnapDrive be installed and configured on both the Exchange and SQL database servers to ease the storage management issues. Enterprise Vault supports IBM N series storage as long as the storage can be presented to the system as an NTFS file system. Configuring the NTFS system using IBM N series storage can be achieved using the Common Internet File System (CIFS) and/or iSCSI and/or FCP.

Prerequisites

In order to install and successfully configure Enterprise Vault, certain prerequisites must be met. It is important to note that Enterprise Vault supports Microsoft Exchange Server in a SQL Server database environment. Even though Enterprise Vault supports Microsoft Exchange 5.5 and 2000 Servers, FSA is supported on Exchange 2000 or 2003 Servers. The Enterprise Vault 5.0 SP3 offers more features using Exchange 2000 and 2003 Servers than Exchange 5.5. This paper will not discuss Exchange 5.5 Server configuration.

Enterprise Vault supports only a SQL Server environment, and SQL Server 2000 is required on a Windows Server. In a production environment, Exchange Server and SQL Server must be installed on separate Windows Servers. This will provide much-required performance when dealing with large amounts of data.

Enterprise Vault requires network share at the storage front to enable file searching and archiving. It is recommended that the network share use the Unified Network Connectivity (UNC) path to maintain the same network path across multiple client machines and servers. If the storage is mapped using network share and assigned a drive letter, it may create issues for data visibility across different client machines, and the drive letter may differ and cause some issues.

Regarding the Windows Server requirement, it is assumed that a separate server is available for each Exchange Server and database server and all the required Windows service packs are installed. If the FSA component is configured as a standalone feature, Exchange Server is not required. However, most customers use an Exchange Server configuration, and hence this paper assumes that Exchange Server is used.

IBM N series storage requirements depend on the configuration of Exchange Server, SQL Server, and the Enterprise Vault software and the storage requirement of vault stores. Exchange Server, Enterprise Vault software, and SQL Server may be configured on high-performance systems, while IBM System Storage N series with NearStore feature can be deployed for storing the data on the vault stores. System configuration requires iSCSI and/or FCP to complete the local storage requirements.

Design Configuration

Sections 3.1 and 3.2 gave a brief product overview of Enterprise Vault. An Enterprise Vault system can be grouped into four sections: information source agents, user clients, and vault applications together with Enterprise Vault to store and administer the vault. Information agents include Exchange mailboxes, public folders, file systems, SharePoint, etc., and vault administration will be done with Microsoft Management Console (MMC).

In a Symantec environment, several message servers and desktop clients may exist. Enterprise Vault supports multiple Exchange Servers, and a typical limit is eight Exchange Servers per Enterprise Vault. Each Exchange Server is configured to be connected with a single SQL Server. IBM N series systems are configured with a SAN or an IP-based SAN configuration. Storage is used to configure Exchange Server and SQL Server data. Note that the same system can be configured on both Exchange Server and a database server as local disks to install and configure the Exchange Server and SQL Server database.

To archive the e-mail and files using FSA, a second IBM N series with NearStore was configured using network path. For this purpose, we used the CIFS protocol to map the drive using an UNC path. Providing an UNC path provides the same path name across the network.

In Figure 4, Enterprise Vault has access to both Exchange Server and SQL Server database data. The system is configured to have a local storage configuration, and the IBM N series with NearStore provides a network share to be able to archive the files and e-mail.

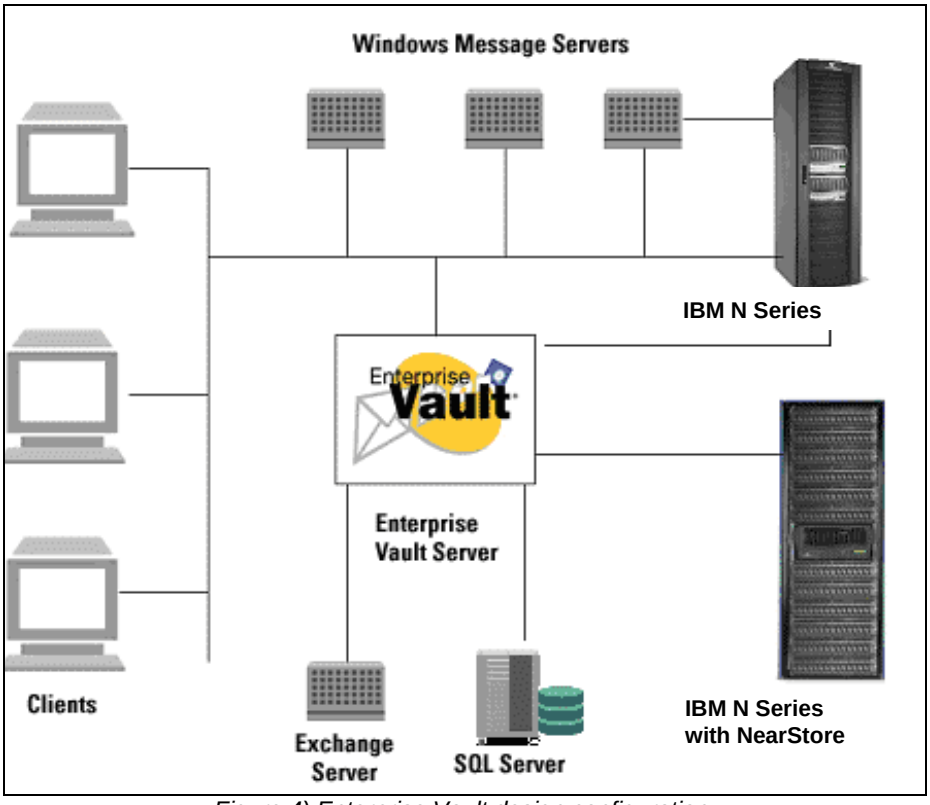

Figure 4) Enterprise Vault design configuration.

Configuration of Local Disks and SnapDrive

SnapDrive software integrates with the Windows Volume Manager so that the IBM N series can serve as virtual storage devices for application data in Windows Server environments.

SnapDrive manages virtual disk logical unit numbers (LUNs) in an IBM N series, making these virtual disks available as local disks on Windows hosts. This allows Windows hosts to interact with the virtual disks just as if they belonged to a directly attached RAID array.

- SnapDrive enables online storage configuration, virtual disk expansion, and streamlined management
- It integrates IBM System Storage N series with Snapshot™ technology, which creates point-in-time, read-only images of data stored on virtual disks
- It works in conjunction with IBM System Storage N series with SnapMirror® software to facilitate disaster recovery from asynchronously mirrored destination volumes.

SnapDrive supports both iSCSI and FCP, and using either an iSCSI software initiator or host attach kits, SAN or IP-based SAN configurations can be configured on IBM N series storage. To install SnapDrive software on IBM N series storage devices, refer to the SnapDrive installation guide.

It is recommended that any system connected to a host reside in the same broadcast domain as that host, so that virtual disk I/O commands do not need to traverse router hops. For Windows cluster configurations, do not permit internal cluster traffic on a Gigabit Ethernet (GbE) network used for host system data transfer. Instead, use a Fast Ethernet connection for all cluster traffic. This practice ensures that a single network error cannot affect both the connection for internal cluster traffic and the connection to the quorum disk.

Mapping the Network Share on NearStore

In order to provide the archival destination, network connectivity between the Enterprise Vault Server and IBM N series with NearStore is configured. To complete the configuration, verify that the NearStore server name is entered in the Windows domain and the network connectivity is established. Once the network connectivity is established, create a volume of desired size. Starting with Data ONTAP 7.1, the IBM N series provides a great flexibility in defining and configuring volume sizes. Depending on the need and growth of data, the volumes can either be expanded or shrunk, provided the right type of volume is created, such as FlexVol, which allows growing or shrinking the disk volume size.

Configuring Write Once, Read Many Storage Using SnapLock

Before configuring the volume and folder archival policy and rules to configure the Enterprise Vault Server, a volume with IBM System Storage N series with SnapLock® enabled has to be configured.

Currently IBM N series storage devices support creating a flexible volume depending on the storage requirement by configuring a larger aggregate to allow the creation of flexible volumes to be created on top of the aggregate. This will enable storage administrators to exploit the benefits of storage provisioning.

IBM N series offers a more robust solution with a new type of RAID protection named IBM System Storage N series with RAID-DP™. RAID-DP stands for RAID Double Parity, and it significantly

increases the fault tolerance from failed disk drives over traditional RAID. At the most basic layer, RAID-DP adds a second parity disk to each RAID group in a volume. Whereas the parity disk in a RAID4 volume stores row parity across the disks in a RAID4 group, the additional RAID-DP parity disk stores diagonal parity across the disks in a RAID-DP group. With these two parity stripes in RAID-DP, one horizontal and the other diagonal, data protection is obtained even in the event of two disk drives failing in the same RAID group.

To create a SnapLock volume, follow these steps:

1. Set up and configure the IBM N series storage devices.
2. Once these devices are available, create or use an existing volume with SnapLock enabled. To create a SnapLock volume, log in to the IBM N series or IBM N series with NearStore system and issue a volume create command using the "-L" option.

Here is the checklist to create the volume with SnapLock enabled:

3. Set up and configure the IBM N series storage devices.
4. Verify that the IBM N series device has the necessary SnapLock licenses for CIFS and/or NFS.
5. Create a SnapLock aggregate by issuing a command with the "–L" option (DataONTAP 7G).
6. Create the SnapLock volume by issuing a command with the "–L" option.
7. Create necessary qtrees and CIFS shares.

Current Configurations

In completing this report, we used a sample configuration with Exchange Server running Enterprise Vault and Exchange Server on a Windows Server. SQL Server is running on a separate Windows Server. The system is configured with SnapDrive to set up the local disks to be able to install Exchange Server and SQL Server data. NearStore is configured as a destination for archival of files and e-mail content. Several desktop clients use the Exchange mailboxes.

Enterprise Vault: Preinstallation

It is important to understand the different components that are installed with Enterprise Vault. This section will provide the information required for preinstallation requirements, including the software required and the tasks that are to be performed before installing Enterprise Vault.

Enterprise Vault has the following components:

- Windows services
- Vault administration console
- Web-based components to provide access to archives
- User extensions to allow clients to access archived items
- Outlook Web access extensions
- Microsoft Exchange forms.

Preinstallation task include analyzing the requirements of performance and high availability, and the existing infrastructure architecture helps to implement a right solution. If FSA is installed without Microsoft Exchange Server, configuration details regarding Exchange Server may be skipped. Enterprise Vault has several services, and installing these services will enable the vault administrator to configure and run services on that server. The service components can be installed on any computer on which the services are run. Some of the services included in Enterprise Vault are mentioned below.

- Admin service. One per server, and this service is installed automatically with the installation of any of the service components of Enterprise Vault.
- Directory service. One per server, and it is installed initially per directory. It requires access to SQL Server for an Enterprise Vault directory database.
- Indexing service. One per server, and it must have a connectivity to a physical storage location to store the index data.
- Storage service. One per server, and connectivity to a physical storage configuration is required. It also requires access to Microsoft SQL Server for the vault store databases and IIS and MSMQ services.
- Shopping service. This also requires a physical connectivity to storage configurations and requires IIS.
- Archiving services agent. This is required for each Microsoft Exchange Server computer. More than one archiving service may be configured, and the service requires MSMQ and Outlook with collaboration data objects (CDOs).
- Public folder service. This is required for each public folder root directory.
- Journaling service. This is required for each journal mailbox. It requires MSMQ and CDO to run.
- Retrieval service. This is required for each Exchange Server, and it should be installed on the same Enterprise Vault Server as the archiving or journaling services.

Verify that active server pages, IIS, Microsoft .NET, and MAC components are installed and registered as a virtual directory in IIS called “Enterprise Vault.” Exchange forms are installed within the Exchange organizational forms library with ownership rights to that library. The vault administration console is a snap-in to the MMC. The administration console may be installed on any computer from which Enterprise Vault is to be managed.

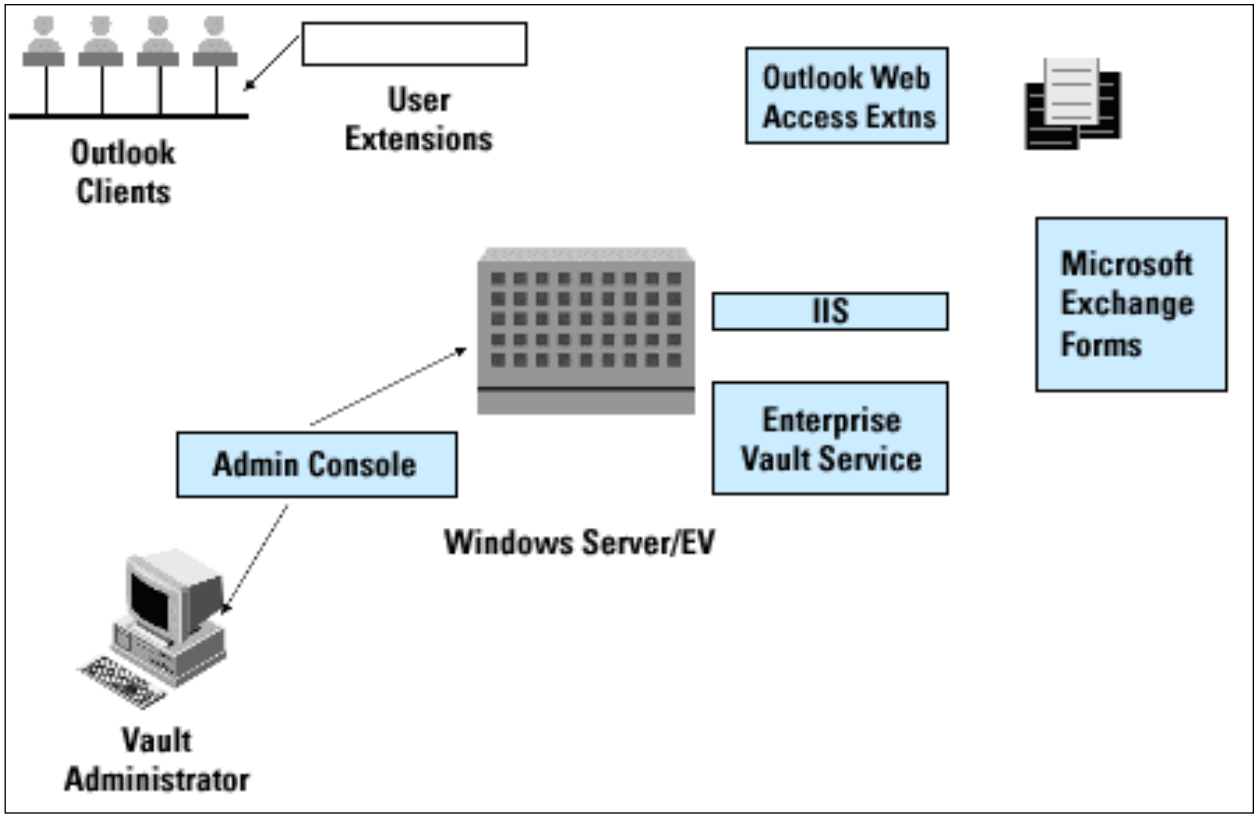

Figure 5) Preinstallation of Enterprise Vault configuration.

Installing the Prerequisites before Enterprise Vault

In order to have a successful installation and configuration of Enterprise Vault, follow the prerequisite software sequence to avoid issues with dynamic library loads (DLLs). Here is the sequence to complete the preinstallation tasks. It is important to note the platform configuration where the installation and configuration needs to be completed. If you are installing Enterprise Vault on Windows Server 2003 and Windows 2000 communicating with Exchange Server 2003 or Exchange 2000 Server, follow these tasks in the order they are given.

Obtain all prerequisite software and note the requirements for each service that will be installed on the computer. Some services require a physical connectivity and running services such as MSMQ and IIS.

1. Step 1: Windows 2003 or Windows 2000 with Service Pack 3—Windows 2003 Standard Edition or Enterprise Edition, Windows 2000 Advanced Server, or Windows DataCenter Server may be used.
2. Step 2: Outlook needs to have CDO components if Exchange Server is not running on the computer.
3. Step 3: Install SQL Server. SQL Server is recommended to be installed on a separate server. Enterprise Vault works with Windows authentication mode and mixed-mode authentication, and SQL Server must be case insensitive.
4. Step 4: On Windows 2000 Server, verify that Service Pack 3 is installed.
5. Step 5: Server Manager for Exchange Server—Enterprise Vault and Exchange Server are to be installed on the same server. Do not install the Server Manager for Exchange on an Enterprise Vault Server running Windows 2003.
6. Step 6: MSMXL that comes with the redistributable software folder on the Enterprise Vault software media. Alternatively install Internet Explorer V6.
7. Step 7: Microsoft data access component (MDAC) V2.6 or later, and the software comes with the Enterprise Vault media.
8. Step 8: Microsoft .NET Framework V1.1 software that comes with the redistributable software folder on the Enterprise Vault media.

Enterprise Vault services need access to the network with appropriate access permissions. This is accomplished with one service account. Enterprise Vault services run under this account, and it is shared by all Enterprise Vault computers in all Enterprise Vault sites. This account should be allowed to log on as a service and act as a part of the operating system and debug program user rights**.**

Exchange Permissions

If the Microsoft Exchange Server is being installed and configured, the vault service accounts need the full control access to each Exchange Server processed by Enterprise Vault. In case of Exchange 2003, it must also be delegated Exchange full administrator in the appropriate administrative group.

Configuring the SQL Login

To create the directory and vault databases, Enterprise Vault needs to access SQL Server. This means that you should verify the network connection between the Enterprise Vault Server and SQL Server machines before installing Enterprise Vault. Then create a SQL login.

Configuring the Microsoft Message Queue

Microsoft Message Queue (MSMQ) Server needs to be configured on the Enterprise Vault Server. It is also required to use a DNS alias. A simple version of the alias such as vaultserver.mydomain.com is recommended to be configured to point to a vault.

Completing the Preinstallation Tasks

Setting the Service Account Permissions

Setting the appropriate permissions is required to continue with Enterprise Vault software installation. Assign “full control” permissions to the vault service account on Exchange Server. On our test install, we used Exchange Server. To set the permission, open the Exchange System Manager, expand the servers and properties, and then select the “security” tab.

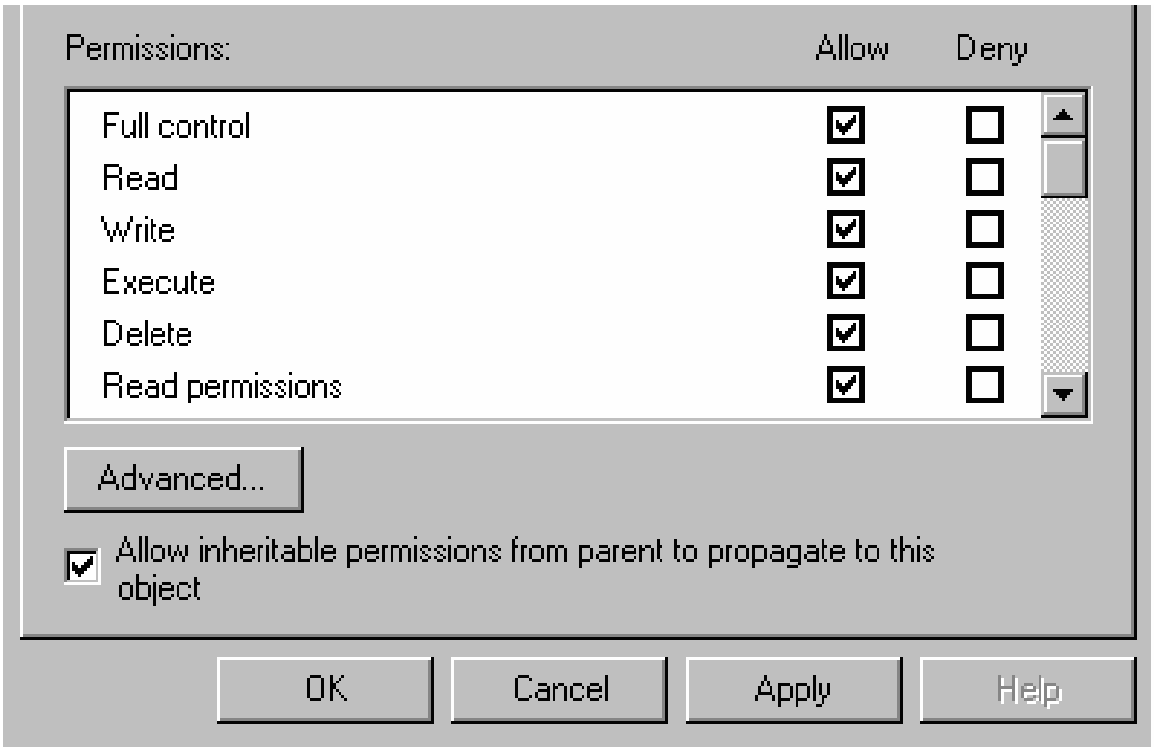

Figure 6) Vault admin service account permissions.

Setting the DNS Alias Name

Create a DNS alias. On our configuration, evault1 was used as the alias name after setting the permissions.

Creating SQL Login

Using SQL Enterprise Manager, create a SQL login for the vault service account. If a separate group manages SQL Server, contact the SQL database system administrator to perform the task. To create the required SQL login, use SQL Enterprise Manager. On our system, we added a new login and grant access with the server role as a database creator was configured.

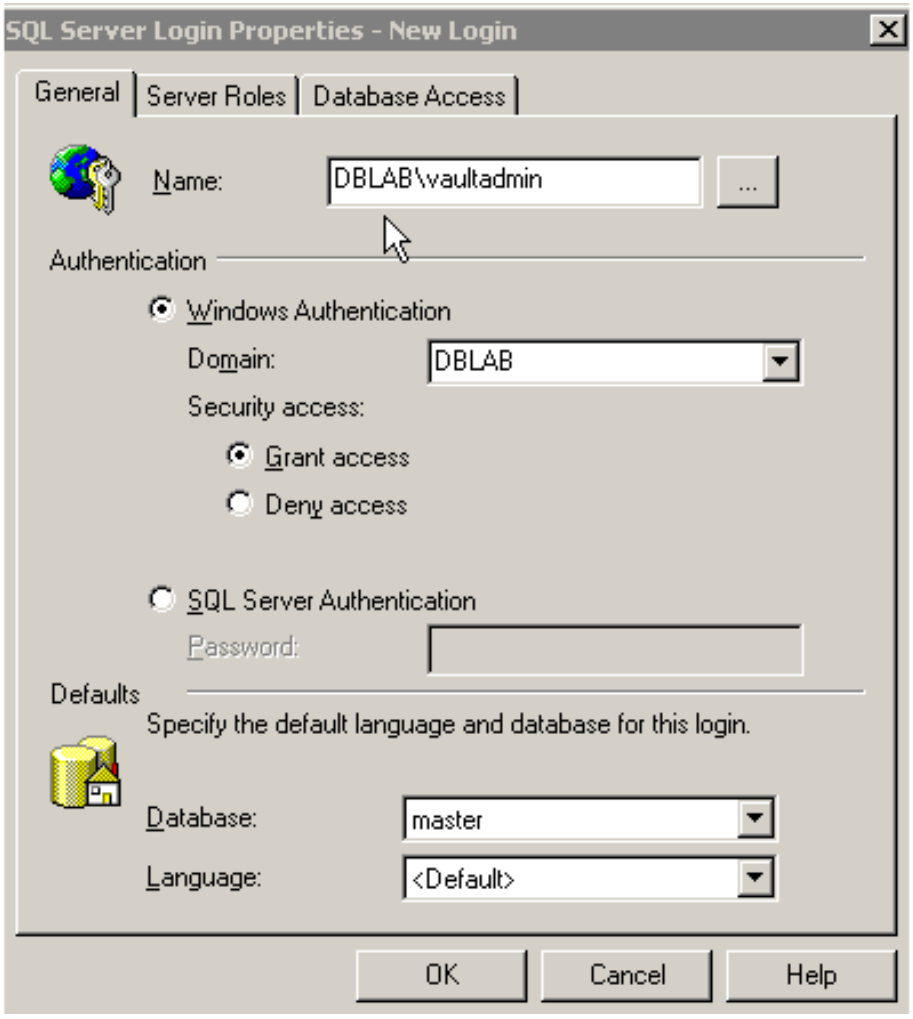

Figure 7) Creating a SQL login.

After creating the SQL login, verify that administrator privileges for the vault service account are set properly by opening MyComputer and managing local users and groups. Now the configuration is ready to install Enterprise Vault.

Installing Enterprise Vault

In the previous section, we discussed the preinstallation requirements. Once the preinstallation requirements are met, Enterprise Vault software can be completed. The Enterprise Vault installation process provides the available choices for installing the required components. Note that administration console service is installed as part of Enterprise Vault installation. A virtual directory is created and registered in IIS called “Enterprise Vault.” Before installing Enterprise Vault, stop the IIS admin to stop the dependent services and continue with the setup wizard, which will guide you through the installation by selecting Enterprise Vault and the administration console as the required components. Before installing Enterprise Vault, stop the Internet services (IIS) and follow the instructions to complete the installation. Note that FSA will be installed after the Enterprise Vault Server is installed and configured, and hence the file placeholder services component may be unchecked.

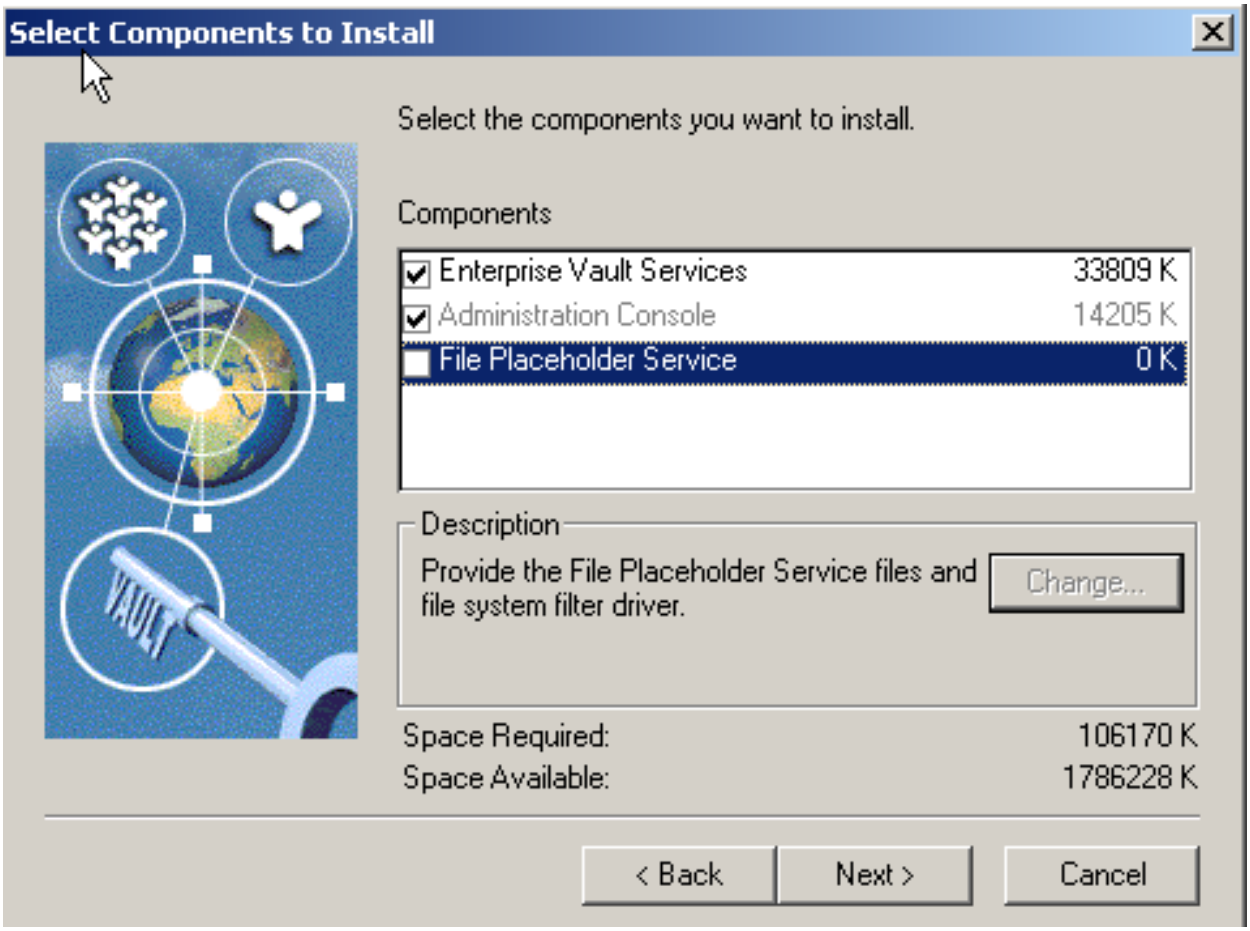

Figure 8) Enterprise Vault components.

Postinstallation Tasks

After installing Enterprise Vault, certain postinstallation tasks need to be completed. Setting security for Enterprise Vault Web access application, creating an Outlook profile, and distributing the Microsoft Exchange Server forms are some of the tasks to be configured as a part of postinstallation configuration. To set up the security for Web access on Windows, this paper suggests using Internet services (IIS) on the Enterprise Vault Server. The next step is to create a folder in the organization forms library on Exchange Server. In this section, the vault service account (VSA) on the Enterprise Vault Server can install the forms from Microsoft Outlook using any mailbox with an account that has owner permission on the specified folder in the organizational forms library. The Enterprise Vault Server is enabled with the license key file to the Enterprise Vault installation directory.

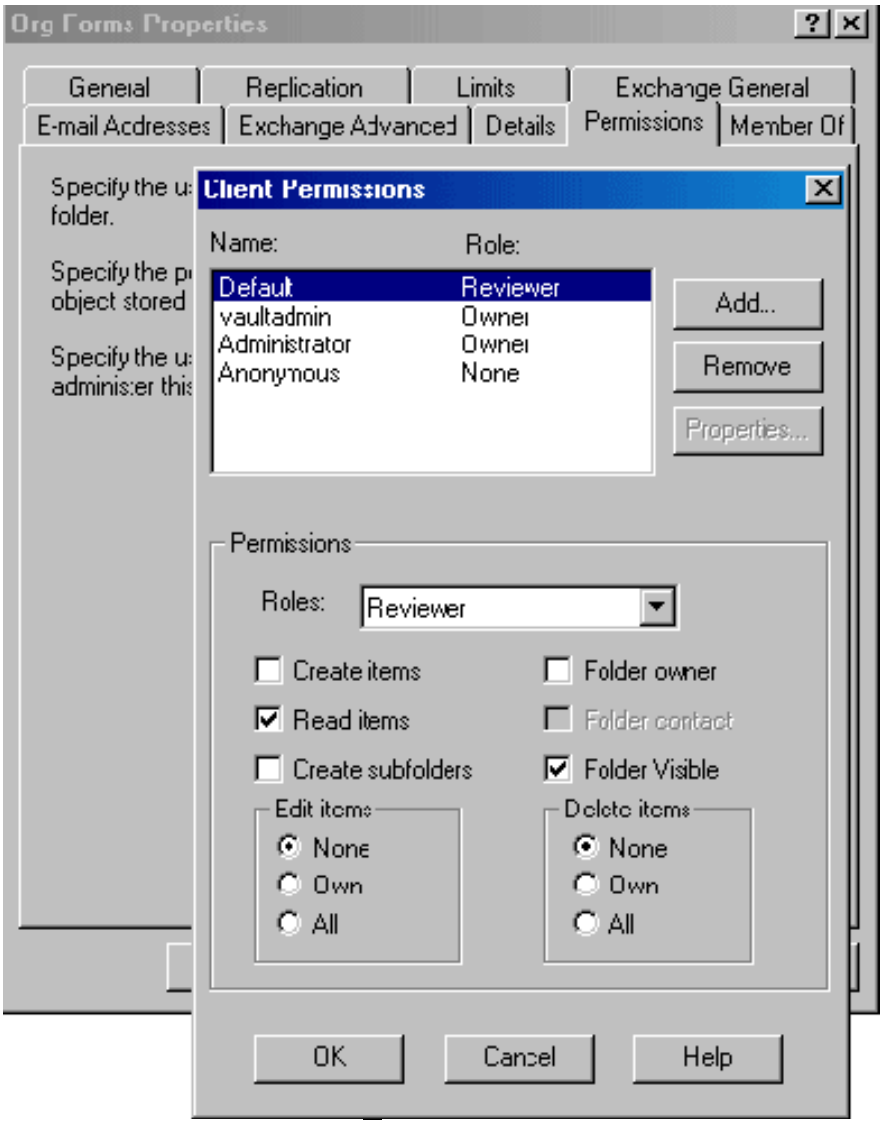

Figure 9) Configuring Exchange organizational forms.

Enterprise Vault WORM Storage Configuration

Using the Enterprise Vault configuration utility, complete the following activities:

- Create a vault directory
- Create a vault site
- Add services to the vault computer.

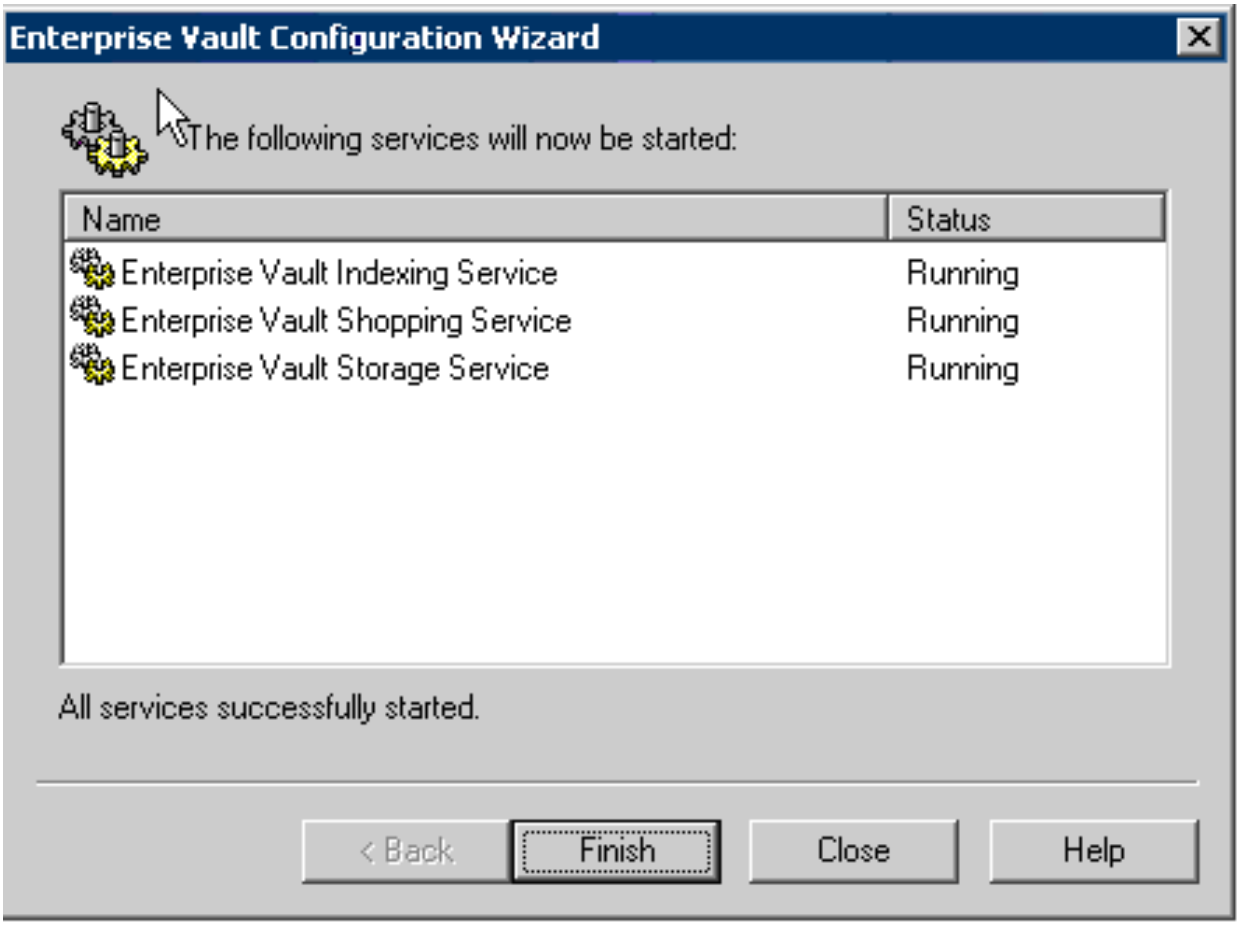

Figure 10) Enterprise Vault services.

Enterprise Vault can create a vault store partition on storage devices such as an IBM N series SnapLock volume to take advantage of write once, read many (WORM) capability with the retention period features. In the test setup, SnapLock. Enterprise Vault required the storage to present it as an NTFS file system.

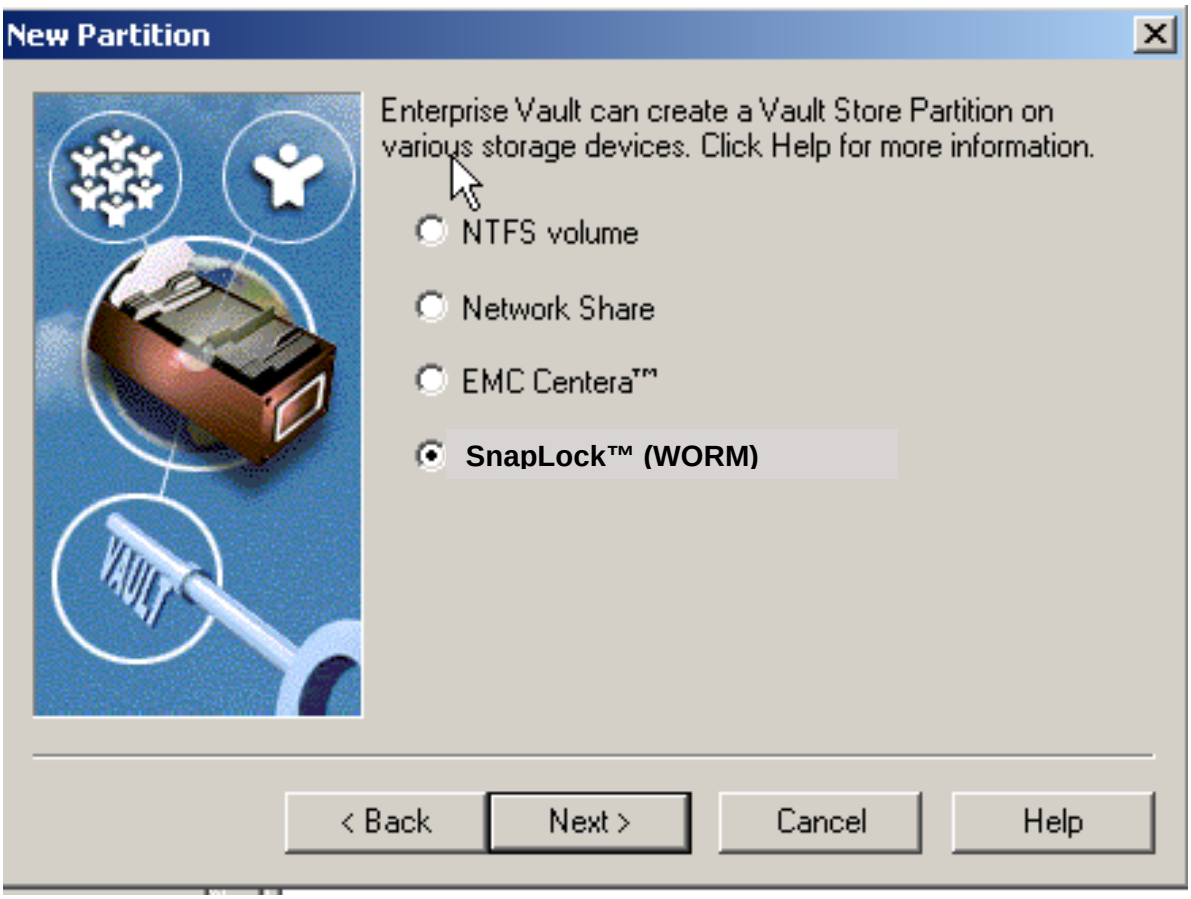

Figure 11) Vault store partition with SnapLock WORM storage.

File System Archival

An Enterprise Vault site computer runs one or more Enterprise Vault services by sharing the same configuration. The Enterprise Vault Server establishes a relationship with Microsoft Exchange Server. Enterprise Vault archivea items from Microsoft Exchange Server mailboxes and public folders. If the existing Microsoft Exchange Server topology is unsuitable for Enterprise Vault, a new collection of Microsoft Exchange Server computers to be served by an Enterprise Vault site can be configured. When the vault site is spread across multiple Exchange Servers, users require the following configuration.

- An Enterprise Vault site cannot serve Exchange Server computers outside that organization. Exchange Server computers must be in the same Exchange Server organization.
- Under the control of one Enterprise Vault site, the information from a single Exchange Server computer must be archived. This means the different Enterprise Vault sites cannot manage archiving from the same Exchange Server.

Before installing the FSA component, analyze the method that can be related to the Exchange Server organization. Several possible installation strategies are available, including the following common ones:

One Enterprise Vault site for each Exchange Server site

- One Enterprise Vault site for a part of a single Exchange Server site
- One Enterprise Vault site for parts of multiple Exchange Server sites
- One Enterprise Vault site for many Exchange Server sites
- Many Enterprise Vault sites for one Microsoft Exchange Server site

Design and Sizing Requirements

In order to design a proper FSA system, a reasonably good estimate of the server and storage requirements is required. The main design points should be based on whether the solution is centralized, decentralized, or multiple vault. It is also important to check what components of Enterprise Vault need to be archived. The following list shows some of the examples:

- Mailbox archiving
- Public folder archiving
- Journal archiving
- PST migration
- Office vault
- Compliance acceleration
- Discovery acceleration.

During the design analysis phase, consider whether the high-availability feature is required. Note that the “high-availability” solution with Enterprise Vault is designed considering Enterprise Vault Server availability rather than increasing Enterprise Vault Server performance. A high-availability solution may in fact degrade performance while improving server availability in a case of the failure of the primary Enterprise Vault Server.

Lastly, consider how many vault servers need to be configured with an estimated required two years of storage. By following these requirements and analysis, a better system may be configured.

Topology Selection

In this section, different topologies that can be used in designing a system are discussed. In a centralized topology, the vault server is located in one site; servicing Exchange Servers are also located in that site and possibly other remote areas. In a decentralized topology, vault servers are located in a central site, and service and local Exchange Servers provide back-end services for remote services. Vault servers at each remote location run “archiving” services. A hybrid topology is a combination of centralized and decentralized design. In a multiple solutions environment, a separate, independent vault server for each location is involved. Selecting a particular topology depends on Exchange organizational structure and location of the server. It also depends on the amount of storage that will be archived.

The advantages of a centralized topology are that it is simple, costs less, and is easy to set up and manage. It could affect WAN links and increase the access time for remote users to view the archived items. This topology uses MAPI connections to remote Exchange Servers over WAN, and a single vault server is located at one site.

In a decentralized design, vault servers are located at each site and use MSMQ services to connect remote vault servers over WAN. This leads to an advantage of MSMQ in reducing the traffic of MAPI requests. The topology adds complexity to implementing the system design, and it may increase the access time to view the archived items.

A hybrid solution uses both centralized and decentralized topologies by running the vault servers at only selected remote sites. Other sites are serviced by centralized sites. Multiple independent solutions offer a better solution; however, it is an expensive, complex design. It offers no single search across an organization’s archives.

The number of archiving services on a vault server depends on the limitations of Windows memory management that is available for Windows services under the local system account. With archiving and retrieval services configured, the same vault server can archive up to 14 Exchange Servers. This leads to the requirement of having multiple vault servers even when the archiving throughput could be achieved with a single vault server.

Vault Server Specification

Any Windows Server running Windows 2000, Windows 2003, and Windows XP is supported, and it is recommended to have dual Pentium or XEON processors with at least 2GB of memory to be able to archive up to 40,000 messages per hour. This assumes that each message item is based on 50KB average size. IBM N series storage provides the required RAID to protect the data at the storage level. Configure a sufficient storage requirement for SQL Server database data and logs. It pays well to design a simple configuration. Use of multiple servers and a topology other than centralized may not be a requirement in most scenarios.

Database Storage Configuration

While using IBM N series storage, local NTFS storage can be configured using FCP or iSCSI, and the storage can be easily managed by SnapDrive. Using SnapDrive and FlexVol, the storage growth can be scaled according to the needs of data growth. IBM N series storage provides RAID protection and allows creating an instant backup and quick recovery using Snapshot and IBM System Storage N series with SnapRestore®. Both backup and recovery can be managed by the SnapDrive solution. SnapDrive is integrated seamlessly with MMC. This meets the requirement of storage area management for placing the database data and log devices.

Mailbox Archiving Policy

The next steps are to set up Enterprise Vault to archive from mailboxes and establish a strategy for implementation. In order to achieve this, it has to be decided if the policy affects across the site, which mailboxes it will affect, and archiving and retention policies defined. It also allows configuring whether to keep a safe copy after the item is archived.

Similarly, a public folder archiving strategy has to be developed by deciding folders that are to be archived and whether the archived items are to be deleted from Exchange Server. The journalizing service allows a few options while configuring to know the entities that need to be tracked.

Installing File System Archival Component

FSA is installed using the installation utility of Enterprise Vault software; uncheck the Enterprise Vault services and the administration console, leaving the file placeholder services, and follow instructions to add a new file server for archiving.

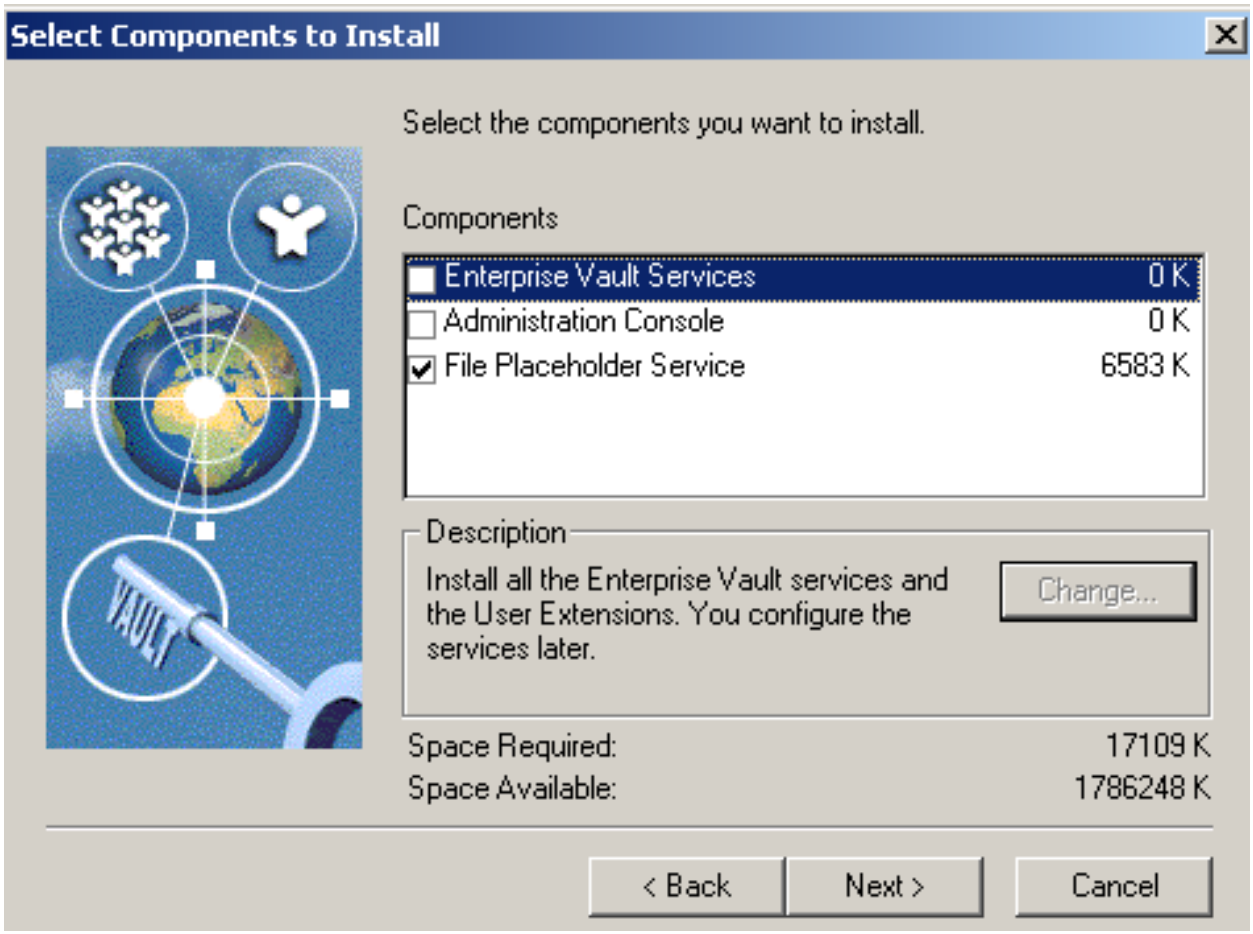

Figure 12) Installing file system archival components.

After the new file server for archiving is configured and added, Enterprise Vault displays the services component installed. On our test setup, the components shown in Figure 13 were installed.

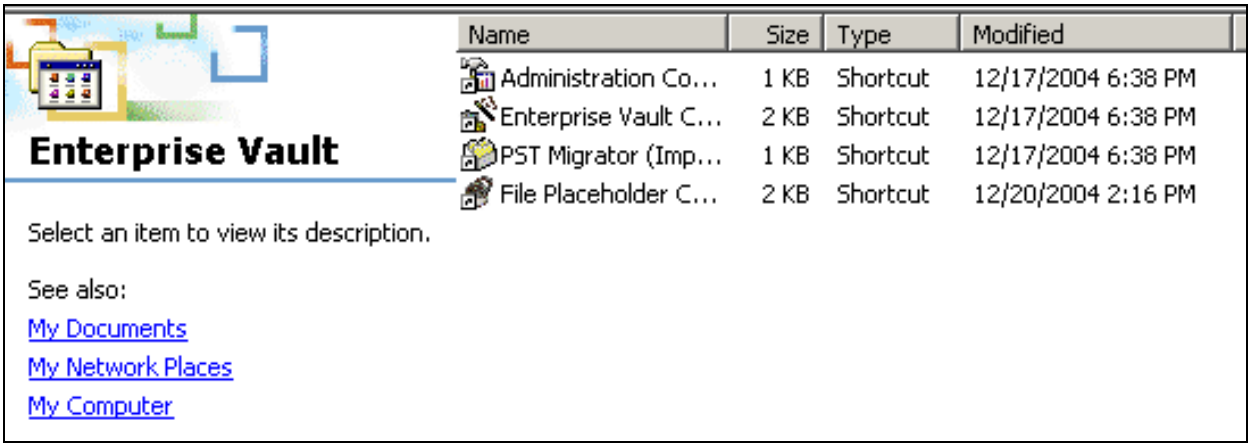

Figure 13) Enterprise Vault components.

Adding a File Server

This approach will offer several benefits. Some of these benefits include: in order to configure the file system archival successfully, a place holder service has to be installed and add inform the Enterprise Vault Server about the file servers with the volume and folders that are to be archived. In this section, the procedure to add a file server and set the archival policies is discussed. Note that each file server can be serviced by only one Enterprise Vault Server. A single Vault Server can also archive data from several file servers. As a pre-requirement, the storage services must be running within the Enterprise Vault Server. This configuration allows the administrator to simply specify the files to be archived and leave the process of creating the new required services to the Enterprise Vault Server. Placeholder shortcuts can only be used on devices hosted by a Windows operating system. Note that File placeholder services on all file servers and the Enterprise Vault Server.

Summary

With Enterprise Vault, messages are extracted from Exchange Server and stored on storage. This storage has to be presented as an NTFS file system. Instead of placing messages in a database, messages are stored as standard files. This approach offers several benefits, including:

- No complex maintenance is required, as opposed to databases
- A particular item can be recovered in case of a corruption, as opposed to the entire data set
- Without affecting performance, more data can be stored
- Multiple Exchange Servers can be handled by one or a few Enterprise Vault Servers.

The Enterprise Vault Server has several drawbacks in terms of data availability and dependability. To access the data of archived files or to access the files, SQL Server must always be up and running. In the case of database corruption, it has to be recovered from the backup copy, losing all the recently archived items (files). Enterprise Vault works only in Microsoft Exchange environments. Space savings due to FSA may be offset by creating a secondary copy in the form of HTML. The data replication could take a significant amount of time and resources. Creation of a second copy of an HTML file after the file is archived reduces the space savings from archiving and compressing. Restoring the corrupted database could be disastrous in an enterprise environment.

The disadvantages listed above can be easily addressed by exploiting the advantages of IBM N series storage solutions, including with the NearStore feature. The Enterprise Vault and IBM N series product integration design will take advantage of both Enterprise Vault and the IBM N series storage solution to offer an efficient and highly available data solution.

IBM N series and Symantec are committed to providing Enterprise Vault users with superior solutions designed to meet their business needs. IBM N series systems and data management solutions ensure that Enterprise Vault data is protected and available 24x7. With IBM N series, you get solutions that are easy to use, deploy, and manage, with high availability and exceptional performance at the lowest total cost of ownership in the industry.

IBM System Storage N series with SnapManager for Microsoft SQL Server is a complete data management solution that provides backup and restore features using Snapshot technology. By reducing backup and restore times, minimizing application outages, and consolidating database storage, SMSQL delivers a cost-effective solution for managing critical SQL Server databases.

In conclusion, the recommendations made in this paper are intended to be an overview of best practices for *most environments*. This paper should be used as a set of guidelines when designing and deploying Enterprise Vault. To ensure a supported and stable environment, familiarize yourself with the resources provided in this paper and involve an Exchange specialist if necessary.

Caveats

All possible combinations of hardware platforms and storage architecture and software options have not been tested. If you use a different Windows Server OS or a different version of Enterprise Vault, then significant differences in your configurations could exist that might alter the procedures necessary to achieve the objectives outlined in this document.

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, FlexVol, NearStore, RAID-DP, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapRestore and Snapshot 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.

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.

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.