

Deployment Guide for ZANTAZ EAS with Lotus Notes: NetApp Storage Solution

Network Appliance Inc.
TR-3528

Abstract

Electronic communication, largely in the form of e-mail, has become one of the most critical tools for success in many of today's business environments. Yet an organization's increased reliance on e-mail also brings with it a fair share of challenges: IT must constantly ensure that users have ongoing access to e-mail while at the same time designing the back-end infrastructure to handle the growing volume of e-mail documents and messages it is now forced to store and protect. Enterprises have begun to deploy two types of solutions to help them handle the increasing burden of e-mail: larger, more centralized and reliable storage systems from companies like NetApp; and a new breed of powerful e-mail and file management and archiving solutions, such as that found in ZANTAZ Enterprise Archiving Solution (EAS). This paper describes some high-level deployment practices when using NetApp storage with ZANTAZ EAS in a Lotus Notes environment. These practices are meant to facilitate Lotus Notes-related e-mail and file archiving, search, and retrieval.

Table of Contents

1. Executive Summary	4
2. Introduction	4
3. Overview of ZANTAZ EAS.....	4
3.1 EAS Server.....	5
3.2 IIS Server.....	5
3.3 Document Store	5
3.4 EAS Database	5
3.5 Index Server	5
3.6 Storage Manager Server	5
3.7 EAS Client Software.....	5
4. Overview of NetApp Storage	6
5. Overview of Archiving in Lotus Notes Environments.....	7
6. General Considerations Prior to Deployment	8
6.1 Network Considerations	9
6.2 Lotus Notes-Specific Considerations	9
6.3 System Sizing Considerations	10
6.4 Recommended Use of Database Partitions.....	12
6.5 Considerations for High Availability.....	12
7. Microsoft Windows and Lotus Notes Prerequisites	13
8. General Deployment Guidelines for Zantaz EAS with Lotus Notes	14
8.1 Pre-Installation Checklist.....	14
8.2 Installation Overview	16
9. General Deployment Guidelines for NetApp Storage.....	17
9.1 Setting up Local Disks with NetApp SnapDrive	17
9.2 Mapping the Network Share on NearStore	18
9.3. Configuring Write Once, Read Many Storage Using SnapLock	18
10. Summary and Caveats	18
Appendix A: Typical Enterprise Configuration of EAS Deployment	19

1. Executive Summary

This paper discusses high-level information on deploying ZANTAZ EAS components along with NetApp storage solutions in Lotus Notes environments. It is intended to give the reader an overview of the prerequisites and general steps required to successfully deploy this type of software and storage architecture in today's enterprise organization.

2. Introduction

E-mail and file archiving is a powerful strategy in today's enterprise to contain and manage the exponential growth in e-mail systems and file shares. Software-based archiving to a centralized data store, like that provided by NetApp storage, has also grown in popularity as organizations find themselves facing a number of trends:

- The need to comply with mounting legislative requirements to account for their data
- The growing cost of potential e-discovery requests by the courts
- Initiatives that stress increased corporate accountability and stricter corporate governance

In effect, these trends all require a new focus on tracking data, storing and protecting it efficiently, and managing it for easy searching and retrieval. The use of software like ZANTAZ EAS for file and e-mail archiving in Lotus Notes environments, in conjunction with NetApp storage, addresses these needs effectively. This paper explains the high-level steps involved in deploying ZANTAZ EAS with NetApp storage.

3. Overview of ZANTAZ EAS

Managing the ever-increasing size of IBM Lotus Domino and Lotus Notes environments can be a daunting task for any administrator. ZANTAZ EAS offers a targeted solution that can effectively archive both e-mail and files, as well as offload much of the burden of Lotus Notes mailbox management. Uniquely suited to large, geographically dispersed environments, EAS offers a distributed architecture that provides centralized management and administration of overall e-mail policies, while still allowing branch or remote offices the option to search, discover, and administer their own local archival document and e-mail stores.

Some high-level benefits to using EAS with Lotus Domino server and Lotus Notes e-mail clients include:

- Improved end-user productivity by eliminating mailbox quotas and providing easy search and retrieval of historical e-mails and files
- The ability to move or consolidate e-mail attachments, documents, and public folders to a centralized NetApp secondary storage system while maintaining ultrafast search and retrieval from the archive
- 50% to 80% reduction of overall e-mail storage costs with single-instance storage, eliminating duplicate copies of e-mails and files stored in the NetApp archive
- A unified architecture for backup, mirroring, archiving, and compliance from a single disk platform
- A rich policy engine for e-mail retention and deletion, ensuring data permanence, accuracy, and security for compliance with industry regulations and corporate policies
- Fast response times and dramatically lower costs for electronic discovery requests

ZANTAZ EAS is able to offer such benefits largely through its distributed client/server architecture, which consists of the following key components:

- The EAS server
- The IIS server (retrieval Web server)
- Document store

- EAS database
- Index server
- Storage Manager server
- EAS client software

The following sections explain the general function of each of these components.

3.1 EAS Server

The EAS server, also known as the EAS Archive Server, performs the common tasks of archiving (writing to the document store), replacing objects with stubs, restoring objects, etc. Depending on the size of the installation, the EAS server may be composed of two components:

- a parent server
- One or more child servers

3.2 IIS Server

The IIS server, also referred to as the retrieval Web server, responds to user requests to retrieve e-mail from the document store. It is the primary retrieval server.

3.3 Document Store

The document store, also known as the document archive repository, is the enterprise storage source location where archived e-mail or file content resides.

3.4 EAS Database

The EAS Database, also known as the *Database server*, is responsible for storing metadata about the archive. This metadata includes:

- Information about all EAS processes
- The location and permissions to access each archived e-mail or file
- Information such as message headers, folder structure, and mailbox profiles

At the customer's discretion, the EAS database server can be based on either Microsoft® SQL Server or Oracle®.

3.5 Index Server

The index server performs indexing and the creation of search catalogs for both e-mail and files. Its unique design in the ZANTAZ EAS architecture offers powerful scalability and high-performance indexing functionality, with total flexibility when it comes to configuring an archive solution.

3.6 Storage Manager Server

The Storage Manager server implements automatic, policy-driven archive management that optimizes the performance and functionality of ZANTAZ EAS archive storage.

3.7 EAS Client Software

EAS client software resides on the user's PC and allows users to manually specify which e-mails they would like to be archived, during the next scheduled archival run. It also works with the IIS server to let users send retrieval requests for viewing e-mails or files that have already been archived.

4. Overview of NetApp Storage

The majority of tasks performed by the ZANTAZ EAS archival system relies on a series of rapid read or write requests made to or from one or more back-end document stores. A successful EAS archival implementation relies on a number of factors, not the least of which is speed of I/O operations to and from the document store. Speed is critical to a number of core archival functions, including:

- The speed at which e-mail or files can be archived
- The speed at which the document store, filled with archived e-mail or files, can be searched for the purposes of discovery
- The speed at which archived e-mails or files can be retrieved by a user request

In the areas of speed, affordability, reliability, and sheer ease of management, NetApp storage delivers the characteristics best suited for use in ZANTAZ EAS environments.

Speed and Top Performance. NetApp storage provides some of the fastest search and retrieval platforms for EAS archives.

A Centralized, Secure Repository at a Low TCO. NetApp storage solutions, such as NetApp NearStore®, enable centralized storage of large volumes of EAS archive data onto very low-cost storage tiers. Instead of adding more dispersed Lotus Domino servers to handle the growth of Notes NSF files, NetApp secondary storage allows Notes administrators to archive older data off their primary production servers, reduce the number of servers, and even consolidate the remaining primary production e-mail files onto a smaller number of centralized network repositories, such as highly available NetApp FAS3000 storage.

Archives Safe from Modification or Premature Deletion. NetApp SnapLock® write-once, read-many (WORM) functionality allows archived data to be stored securely as an unalterable record, following the strictest criteria for regulatory compliance and litigation support. E-mail or file data is sent to a designated SnapLock volume with industry approved WORM functionality.

Scalable Archives Can Be Sized Up or Down On-the-Fly. NetApp storage is based on its own unique operating system, Data ONTAP® 7G, which effectively aggregates system resources such as disk spindles and capacity through the use of features such as FlexVol™. FlexVol allows users to dynamically create or resize virtual volumes that span any number of underlying NetApp disk storage. Each ZANTAZ EAS document store can be pointed to this type of virtual volume, which the administrator can increase or decrease in size by following a few steps. There's no need to reconfigure LUNs. The process of provisioning or adjusting the size of document stores for your Lotus Domino archive becomes much simpler, thanks to features like FlexVol.

Multiple Layers of Redundancy to Ensure Uptime. Any successful archive solution depends on the document store being available for searches and retrieval. Here, the levels of availability offered by the underlying storage system are a critical factor. NetApp storage offers a number of hardware redundant features that protect archival data from unnecessary interruptions and downtime events. These include dynamic RAID-DP™ functionality that guards against the prospect of two disk drives failing in the same RAID group at the same time. RAID-DP stripes data diagonally across each RAID group and adds a second parity drive to any given RAID group. This implementation of built-in RAID functionality with double-parity support is offers added protection for environments using lower-cost ATA disk, such as that employed with NetApp NearStore—all while maintaining the high performance required from demanding archival, search, and retrieval environments. (Other hardware-related features that support uptime include auto-rebuilds of failed disks, the use of hot spares or hot-swappable components such as disk drives, power supplies and cooling fans, automated storage health monitoring, and the use of redundant components.)

5. Overview of Archiving in Lotus Notes Environments

ZANTAZ EAS is designed to carry out archiving functions for both IBM Domino mail and Windows® file servers. The following diagrams demonstrate how the archival and retrieval functions occur on the network path between the Domino server and one or more related EAS document stores.

In general, the EAS archiving server uses predefined policy settings and archive schedules to send designated Lotus Notes NSF e-mails to the document store repository for archiving. As part of the write process to the document store, the archiving server leaves a stub behind in the users' mail files reflecting each e-mail that has been successfully archived. (As mentioned in section 3, the EAS archiving server is typically composed of either a parent server by itself or a combination of parent and child servers when deployed in more distributed environments.)

Metadata information about each archived e-mail is also stored in the database server. Clients wanting to search or retrieve previously archived data send requests to the retrieval Web server, or IIS server. Sites may also use the EAS Storage Manager for other administrative functions associated with archived data stored in the document store.

Figure 1 shows EAS components used to archive e-mail from IBM Lotus Domino servers.

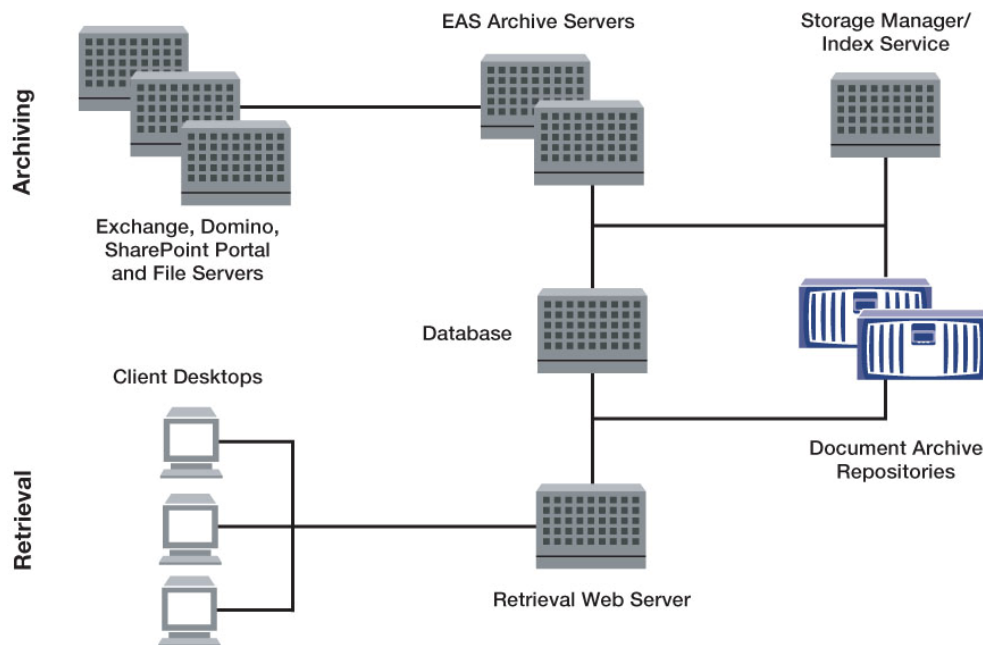


Figure 1) Archiving with IBM Lotus Domino servers and ZANTAZ EAS.

Figure 2 shows an example of multisite archiving of IBM Lotus Notes e-mail and files with ZANTAZ EAS and NetApp storage.




Figure 2) A geographically distributed archiving architecture.

6. General Considerations Prior to Deployment

This document describes some high-level practices to follow in planning to deploy a ZANTAZ EAS archiving solution with NetApp storage for Lotus Domino environments. It assumes that most Lotus Domino installations are based on the Microsoft Windows operating system. (For installations using AS/400 platforms, refer to the ZANTAZ reference documentation or get further implementation details from your ZANTAZ professional services contacts.)

Before deploying this combined software and hardware solution, it pays to do some groundwork in the following areas:

- Network considerations
- Lotus Notes-specific considerations
- System sizing based on e-mail volume, not mailbox size
- Setup of database partitions for optimal growth of the system
- Considerations for high availability

6.1 Network Considerations

Gigabit Ethernet Connections. For optimal archival “ingest” and retrieval performance, a Gigabit Ethernet network is highly recommended for use between EAS server components. Each EAS server should be configured with dual NICs for redundancy. Ideally, each NIC should also go through a separate set of switches and routers to ensure the existence of a secondary path in case of failure on the hub, switch, or router. Also, if network-attached storage (NAS) is in use as a document store, it should also be set up on a separate GigE switch (segment) to ensure optimal performance. In a data center environment, a load balancing switch for the IIS server is also recommended to enable load balancing of all user-driven archive and retrieval requests.

Network Placement of Child Servers. EAS child servers can be located remotely, across a WAN connection. This topology is appropriate in environments where Lotus Domino servers are already geographically distributed. In this case, the act of archive processing and compression can still be performed locally by a child server at the remote locale. The child server then communicates with the parent server over the WAN. Using a local EAS child server in these cases can reduce network traffic by up to 50%. A WAN child server can also be used as a redundant server if the parent server goes down.

6.2 Lotus Notes-Specific Considerations

Before deploying a ZANTAZ EAS solution to archive Lotus Notes and Domino-specific data, it often makes sense to answer the following questions in relation to the current Lotus Notes deployment.

Support for iNotes Web Access. If the Domino environment uses iNotes for Web access to e-mail (also known as Domino Web Access or DWA in Lotus Domino 7), some customization steps are required to reflect the unique iNotes settings in the EAS archive. For details, see the *EAS for Lotus Notes Installation Guide*.

Support for Offline Synchronization. If you want certain users or groups to be able to access and search the archive while they are offline, you must follow these steps to enable this functionality:

1. During initial EAS installation, select the option to install the ZANTAZ custom template. This corresponds to the EAS for Domino Server MSI package, which will be installed on one of your Domino servers.
2. Using your Domino Administrator account, log in to the Domino server where this package resides and locate the database file, EASLocalDS.NSF. This file represents the EAS Local Doc Store. Sign this file using the Domino Server ID (make sure that client workstation ECLs allow this signature). Then open the file in the Notes client. On the Database tab, click the “Allow Local Doc Store” button and identify the users for whom you want to enable offline access.
3. Under File > Preferences > User Preferences, select “Enable scheduled local agents.” (For the background synchronization process to run properly, all users’ Notes clients must be set to allow the execution of these agents.)
4. Have each user open the server copy of their mail file. If they do not already have a local replica of their mail file, create one. In either case, make sure that the local replica is up to date by running a replication task. When complete, click the EAS button and select “Synchronize the Offline EAS Archive.”

Depending on the size of the archive, the initial synchronization may take some time. It should create a local user database file, EAS_<username>.nsf, to be found in the database mail directory. Before going offline, users must also select the Synchronize option if they want access to the latest archived data.

Support for Journaling with EAS. Lotus Notes Mail Journal database archiving is supported starting with ZANTAZ EAS 5.0 and supports both types of Domino journaling, the local repository as well as the mail-in database. Before you can archive Journal mailbox information to an EAS document store, you must register a new user to link to the Notes Journal mailbox. This is accomplished through the Domino Administrator interface. Once the user is registered, open document properties on the Person Document and change the security setting for the new user file from "All readers and above" to "LocalDomainAdmins" and "LocalDomainServers" (or similar groups according to your company's standards). The EASAdmin account should already be a member of one of these administration groups. If you are using the local repository format of journaling, you need to update the Journal Template design, mailjrn.ntf. (If you are using the mail-in database journaling method, you must use the ZANTAZ mail template, then apply this to the journal file as you would for any regular user).

To update the Journal Template design (for a local repository journal):

1. Delete a series of views and copy new views from the EAS_50_Journal.ntf file.
 2. Update the overall template by inheriting the design from the master template. In this case, the master template is the name of the updated mailbox template (EASMailTemplate.ntf).
 3. Sign the template database (mailjrn.ntf) with the server ID or appropriate administrator ID.
 4. For every Journal database that you plan to archive with EAS, do the following: Open the specific Journal database (journal.nsf or a set that appears in the Domino directory). Then edit Database/Properties so that you inherit the new design from master template "EAS Journal Design." Right-click Journal database, select Database/Refresh Design, and locate the server name where the mailjrn.ntf file resides.
1. Edit the design of the journal template file (mailjrn.ntf). Add (or replace) a select number of views identified in the EAS_50_Journal.ntf file.
 2. Sign the template database (mailjrn.ntf) with the server ID or appropriate administrator ID.
 3. For every Journal database that you plan to archive with EAS, do the following:

Run a Load Convert task from the Domino Administrator Console on each Journal mailbox, pointing to the updated mailjrn.ntf template for the new design.

6.3 System Sizing Considerations

When determining the type of underlying server and storage hardware required, it is important to correctly calculate the potential size of the system on a few fronts.

Message Volume vs. Number of Users. When determining the type of server hardware required to house the EAS server components, it's important to first calculate the typical volume of messages going in and out of your Domino environment. (Instead of sizing the infrastructure based on the current number of Notes users, ZANTAZ believes that sizing based on volume of messages is a more accurate method to ensure optimal performance.) See Table 1 for ZANTAZ general guidelines. These guidelines do not account for redundant or DR/BC configurations. In those architectures, an EAS customer might choose to split up one dual-processor server with 4GB of RAM into two dual-processor servers for redundancy and load balancing, where each smaller server is equipped with 2GB of RAM. Ultimately, sizing guidelines depend less on the amount of RAM available and more on the volume of traffic that is anticipated on each server.

Suggested Systems to host EAS Solution Components				
Message Volume (per day)	EAS Server	Index Server	IIS Server	SQL Server
1,000	Single Processor + 2GB RAM			
5,000	Single Processor + 2GB RAM			Single Processor + 2GB RAM
25,000	Single Processor + 2GB RAM			Single Processor + 2GB RAM
50,000	Single Processor + 2GB RAM	Single Processor + 2GB RAM		Single Processor + 2GB RAM
100,000	Dual Processor + 4GB RAM	Dual Processor + 4GB RAM		Dual Processor + 4GB RAM
500,000	Dual Processor + 4GB RAM	Dual Processor + 4GB RAM	Single Processor + 2GB RAM	Dual Processor + 4GB RAM
1,000,000	Quad Processor + 8GB RAM	Quad Processor + 8GB RAM	Single Processor + 2GB RAM	Quad Processor + 8GB RAM

Table 1) Server sizing guidelines.

Sizing for Capacity and Future Growth. For the contents of each NSF database file included in the archive, plan for the archive to take up roughly 50% of the current file size. You should also plan for the archive to grow an average of 30% to 40% per year. This sizing information is based on general assumptions and may vary significantly for each customer environment. ZANTAZ recommends that you consult with a Professional Services representative to get a site-specific sizing recommendation.

Sizing and Architecture of the Database Server. For easier management of the Domino archive environment, some sites may choose to partition Lotus Notes mailboxes, based on the business segment or location with which they are associated. In that case, multiple database servers can be used. If you are using more than one database server, make sure that each one has a related EAS parent server associated with it. In this way, archiving is specific to each designated partition. Search and retrieval, however, is globally applied across all database servers and all partitions. The IIS server may just be automatically redirected to another EAS database server. Additionally, read-only database instances can also be replicated elsewhere in order to speed retrieval and search time from the IIS server when used in distributed environments.

Two sizing approaches can be employed for the database server. The first is to size the server using transactions per minute. The second is to use a clustering approach for the database server such as Oracle® Real Time Application Clusters (RACs) or Microsoft® SQL Server Cluster. Clustering enables a single, unpartitioned database view for multiple database servers, and the ability to add servers as required to meet performance demands. The architecture of these clusters still limits the practical number of servers to around 16 to 20. The clustering software also supports a transaction-based form of load balancing. Sequencing of overlapping transactions is also maintained.

(These general sizing recommendations are based on the assumption that policy-based archiving will be applied to all mailboxes in the organization. This presupposes a high ingest rate of messages per user mailbox and its associated <username>.nsf file.)

Other factors should also be considered when determining the capacity size of the database server to be used. These include the number of messages to be archived, the types of policies in use for archiving and retention, and whether or not you apply the ZANTAZ EAS software space-saving functionality of single-instance storage, which prevents duplicate e-mails or files from being archived to the document store. Each e-mail's metadata (message header, etc.) resides in the database server and requires an average of 750 bytes of database server capacity per message. Statistics show that two-thirds of all messages are redundant. Given that number, you might assume that 750 bytes of capacity is sufficient to store the metadata on the database server associated with every three messages sent to the archive.

Therefore, if you estimate that the environment will send as many as 30 million messages a day that implies a single-instance reduction rate of 66%. This equates to a daily requirement of 7.5GB of capacity required on the database server, which works out to about 1.8TB per year. In this example, however, you might estimate that 2TB of usable storage (4TB mirrored or 3TB RAID 5) would allow for growth while monitoring the actual utilization and not need to be increased for the first year.

Note: These are general guidelines. A ZANTAZ architect would typically define the specific amount of storage required for each installation.

6.4 Recommended Use of Database Partitions

As mentioned in section 6.3, some Domino sites may choose to use database partitions to speed database maintenance tasks, including the ability to back up, restore, or run operations on multiple database partitions concurrently.

Over time, the EAS database for large institutions can grow to extremely large sizes. ZANTAZ recommends setting up a vertical partitioning scheme, with database tables sectioned by Delivery Date. This scheme would employ read-only segments for simple backup and administration, and for improved archiving and retrieval performance.

The solution would probably be partitioned on a maximum retention period and/or delivery date (for example, 36 months for SEC regulations). This helps determine when a partition can be dropped. EAS Storage Manager would then be used to delete data as it expires (partition 37, for example). Deletions would occur in the rollover partition. The advantage of this solution is that it allows for possible consolidation of remaining data into smaller partitions, and thereby reclaims unused space. Only modified partitions need to be backed up after an EAS Storage Manager purge.

Following this prescription, all message-related tables should be partitioned using the same method: range partition by month. EAS uses Delivery Date for partitioning. For more recommendations on how best to configure partitioning, see your ZANTAZ support representative.

6.5 Considerations for High Availability

If it's important to your Domino site to keep archiving and retrieval of the archived data available at all times, then it pays to devote some time to ensuring that the ZANTAZ EAS system has no single point of failure. Some recommendations that help ensure high availability have already been mentioned in other sections. These include the existence of dual NICs on each EAS server component. Other availability considerations are discussed in this section.

Child Server Availability Practices. To ensure fast archiving and retrieval functions at all times, some sites may choose to deploy multiple child servers. Some common practices include:

Designating a minimum of two child servers with the same underlying hardware and settings as the parent server. (This allows one of the child servers to be promoted to the parent server function, if the current parent server becomes unavailable.)

Adding a buffer, or reserve, of an extra 20% of child servers. This allows the EAS environment to continue the archiving process, even if up to 20% of the child servers fail at the same time.

Database Server Availability Practices. The database server is a critical component in the ZANTAZ EAS environment. It is touched by every other server. Thus it makes sense to spend some extra time planning for how to ensure that the database server remains available, or can be quickly restored to working order in the event of a failure. ZANTAZ recommends that customers use a log-based, real-time replication approach to ensure that the archived message is stored correctly in the database.

Due to a concern related to potential data loss or potentially significant performance degradation, ZANTAZ does *not* recommend other methods for preserving availability of the database server. For instance,

ZANTAZ discourages the use of a standby database server with mirrored storage as well as the use of a hot standby database server with storage replication (typically in the form of volume or file snapshots). The only exception to this rule is the use of a read-only, standby database server that offers read-only access to the EAS archive for the purpose of search and retrieval.

IIS Server Availability Practices. To ensure the best levels of reliability and performance, customers should set up each document store or repository to be accessed by a minimum of three instances of IIS.

7. Microsoft Windows and Lotus Notes Prerequisites

Before deploying ZANTAZ EAS in a Lotus Notes environment, you should check to make sure that you have already installed the following operating system and application software components throughout the environment.

Minimum EAS Server Software Requirements. The following software must be installed before installing the EAS server:

- Microsoft Windows 2000 or 2003 Server/Advanced Server with all service packages

- MSI 2.0

- MSXML 4 SP2

Other Software Required for Installation. Install the following software on either the local or networked computers before beginning the EAS installation:

- Microsoft SQL Server 2000 or 2005 or Oracle server (Oracle 9iTM or later), with the latest service packs.

- Oracle database client, including Oracle ODBC drivers (for Oracle installations only)

- Lotus Domino server 6.5 or later with LDAP configured on one of the Domino servers

- Microsoft IIS 5.0 or later

- Any file service accessible from Windows (will be used for the EAS document store and search catalogs)

- A Windows domain is required to provide an environment ID for an EAS license

- Lotus Notes client version 6.5 preinstalled on the EAS server machine

The following software is also required for use by EAS e-mail and Web clients:

- Lotus Notes mail client 6.5.x or later

The Notes client must be the same as or later than the version of the client on the EAS server. Confirm that all clients have the same version of the DTD file as the version that appears on the EAS server machine that will be extracting data from NOTES mail files to EAS. (The DTD file is located in the NotesXMLSCHEMAS subdirectory.)

- Internet Explorer 5.x or later or Netscape Navigator (required for EAS Web clients)

8. General Deployment Guidelines for ZANTAZ EAS with Lotus Notes

This section provides an overview of the preinstallation and general installation steps required when deploying ZANTAZ EAS in a Lotus Notes environment. For more details, refer to the *EAS for Lotus Notes Installation Guide* and the *EAS for Lotus Notes Administrator Guide*.

8.1 Preinstallation Checklist

Before installing the EAS components, it is important to set up the proper Domino-related accounts, permissions, and directory structures, as well as the proper Oracle/SQL Server database environment. Use the following checklist to help keep track of these types of groundwork activities:

- Use Active Directory to create and configure the following EAS-related user accounts:
 - **EASAdmin:** To be used for administration. This account should be configured with Domain User rights and a minimum of Local Administrator rights for the EAS server.
 - **EASIIS:** To be used by the EAS IIS server to retrieve archived content from the document store. This account requires both a local batch job right and a local login directly to the IIS server.
 - **SQLsrv or ORCLsrv:** To be used for database administration. This user account should be created only if the database will reside on a different network server than the one hosting the EAS service.
- Configure Domino accounts:
 - Using the Domino Administrator client, register a new user account, **EASAdmin**, under the same administrator name as above. Then add this account to the Domino domain's Administrator group (such as **LocalDomainAdmins**). This group needs to have Manager access to all mail files. Finally, save the ID file to disk. This file will be copied later to the EAS server.
- Create one or more Document Store directories.

The Document Store directory is a shared directory on the network where archived e-mail and files are stored. At least one document store must be created in a ZANTAZ EAS environment. Distributed geographic environments may also choose to create a document store at each location. The Document Store directory serves the following purposes:

- The EAS server places archived content at this location. (Requires both read and write access to the directory.)
- The EAS IIS server processes search and retrieval requests and sends them to and from the Document Store directory. (Requires read-only access to the directory.)
- Backup operators access this location to perform any related backups of the repository. (Requires read-only access to the directory.)

These high-level steps are required to create a Document Store directory:

1. Make the underlying storage system visible from within Windows Explorer as a valid network share.

When using NetApp storage as the document store, you must first ensure that the NetApp system, such as NetApp NearStore, is visible as a network share from within Microsoft Windows Explorer. To follow the steps required to mount the storage system and present it as a network share, refer to section 9.

2. Using Explorer, locate the file server associated with the NetApp system.
 3. Change the file server's properties to make it a hidden share, named something like **DocStore01\$**. Grant share permissions as follows to the following network user accounts:
EASAdmin: Select the **Full Control** option from the permissions list.
EASIIS: Select the **Read** option from the permissions list.
[Backup Operator Account Name]: Select the **Read** option from the permissions list.
 4. Duplicate the new share permissions using the Security tab and remember to remove all other accounts.
- Configure the Oracle or SQL Server database.

A series of custom settings and steps are required when setting up the database that you plan to use for EAS. Detailed configuration steps can be found in the *EAS for Lotus Notes Installation Guide*. The following list offers a few of the high-level steps and sizing guidelines:

- **For Oracle Installations:** Sites using Oracle require installation of the Oracle client, configuration of the Net Naming Service, and installation of the Windows ODBC driver for Oracle. ZANTAZ recommends configuring the User Table Space setting to a minimum of 300MB; configuring the Index Size to a minimum of 30MB; and configuring the Temporary Table Space to a minimum of 50MB. When planning the database size, a good rule of thumb is to allow a minimum of 400 bytes per item. That means 400MB of storage capacity will be needed per one million items.
 - **For SQL Server Installations:** Create an EAS database in SQL Server with the following default options associated with the data file:
 - Set File Growth to a minimum of 10%.
 - Set Maximum File Size to "Unrestricted File Growth."
- Configure Domino and Lotus Notes.

You must perform the following tasks prior to EAS system installation.

- Confirm that the same DTD file version exists on the Notes Client machine that will host the EAS server, as well as the other Notes Client machines. See section 7 for details.
- Make sure that all database templates are signed with an ID that has appropriate ECLs for user desktops. This is typically the server ID.

If any users will need offline access to the EAS archive, select the **Enable Scheduled Agents** option in User Preferences for each Lotus Notes client. For details about offline support, refer to section 6.2.

- Make sure that Domino is running the LDAP task. In the LDAP configuration, publish another two attributed for EAS: Under the DominoPerson class, add the **MailFile** and **MailServer** attributes.

8.2 Installation Overview

This section offers a very high-level overview of the order in which various EAS components are installed. For specific installation steps, refer to the *ZANTAZ EAS for Lotus Notes Installation Guide*.

In general, once customers have completed the preinstallation steps, they proceed with the installation as follows:

1. Install the EAS software.
2. Configure the EAS database.
3. Configure Domino for use with EAS.
(This includes pointing to the correct EAS IIS server, with related EAS Web server DLL and search index DLL. It also involves using the default Domino mail templates, or preparing to use customized Domino mail templates for use with EAS.)
4. Configure the EAS Administrator.
(This includes connecting to the database, identifying the LDAP settings, and identifying the location of the EAS parent server.)
5. Configure the document store.
*(This is where you point to the location of the EAS archive repository. If you are using NetApp storage, this is where you specify the appropriate UNC path to the NetApp storage system, and where you select the **NetAppStorageTek** option. You also set up the directory structure for the archive, where each directory name reflects either a given **day** [YYYY/MM/DD] or a given **hour** [YYYY/MM/DD/HH].)*
6. Configure the EAS server ID and connect to the EAS server.
7. Configure licenses.
(You must have as many user licenses as there are individual mailboxes to be archived.)
8. Install and configure EAS retrieval components.
(This includes installing EAS IIS server instances, as well as testing for connectivity of the IIS server, the database server, and the document store.)
9. Install the EAS Administrator client and all other EAS end-user clients.
(This is where you can install both the EAS for Files client and the EAS for Notes client.)
10. If your organization uses customized Lotus Notes mail templates, merge the custom mail templates with the existing design elements found in the EAS templates. Custom mail templates are typically used to support Lotus iNotes (Domino Web Access) functionality, as well as Domino journaling functionality.
11. Set up Lotus Notes journaling to work with EAS archiving.

9. General Deployment Guidelines for NetApp Storage

To prepare NetApp storage for use as a ZANTAZ EAS document store, a few key steps are necessary. This section covers the following activities:

- Setting up local disks with SnapDrive®
- Mapping the network share on a NetApp NearStore system
- Setting up WORM with SnapLock

9.1 Setting up Local Disks with NetApp SnapDrive

NetApp SnapDrive software integrates with the Microsoft Windows Volume Manager so that NetApp storage systems can serve as virtual storage devices for application data in Windows server environments.

SnapDrive manages virtual disk logical unit numbers (LUNs) in a NetApp storage system, and makes these virtual disks appear as local disks on Microsoft Windows servers. This allows Windows servers to interact with the virtual disks just as if they belonged to a directly attached RAID array.

Benefits of SnapDrive include:

- Easy configuration of online storage and on-the-fly expansion of virtual disks.
- Integration with NetApp Snapshot™ technology, which creates point-in-time images of data stored on virtual disks
- The ability to work in conjunction with SnapMirror® software to facilitate disaster recovery from asynchronously mirrored destination volumes that often reside at an off-site DR location

SnapDrive supports both iSCSI and FCP protocols ; using either an iSCSI software initiator or host-attach kits, SAN or IP-based SAN configurations can be configured on NetApp storage. To install SnapDrive software on NetApp storage devices, refer to the *SnapDrive 4.0 Installation and Administration Guide*.

Practices that Network Appliance recommends include:

- Making the storage subsystem associated with the document store use the same broadcast domain as other EAS server components. Where possible, this can prevent the virtual disk I/O commands from having to traverse router hops.
- When using Lotus Domino server cluster configurations, do not permit internal cluster traffic on a Gigabit Ethernet (GbE) network used for host/NetApp storage 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. To here
- Use NetApp with the DataONTAP 7G operating system. This allows the creation of aggregates (virtual volumes that can span more than one underlying, physical disk) via NetApp FlexVol technology. It also allows the use of more advanced RAID-DP (RAID Double Parity) protection.
- When dealing with environments like a larger Lotus Notes installation with lots of data I/O activity involved in archive and retrieval, configure one or two large aggregates (or virtual volumes) as opposed to a number of smaller aggregates. This maximizes the number of available disk spindles significantly. Improving on access to disk spindles can also speed up the ZANTAZ EAS archive environment's read/write performance.
- To determine the best SnapDrive configuration for your environment, consult your NetApp sales engineer.

9.2 Mapping the Network Share on NearStore

In order to provide the document store destination for the EAS archive, verify that the NearStore server name is entered in the Windows domain and that the network connectivity has been verified.

Use a UNC path and the CIFS protocol to map the NetApp NearStore drive. Using a UNC path makes the same path name available across the ZANTAZ EAS archiving network

9.3. Configuring Write Once, Read Many Storage Using SnapLock

If a Lotus Notes environment requires extra assurance that the EAS archive repository (and document store) will not be altered, a NetApp FlexVol aggregate, or flexible volume, can be set up that uses NetApp SnapLock write-once, read-many (WORM) technology. Once this SnapLock volume is created, the EAS document store can be configured to send archives to it, much as it would send them to any other archive.

To create a SnapLock volume, follow these steps:

1. Set up and configure the NetApp storage devices, which may include fabric-attached storage devices (FAS series filers) and/or NearStore devices.
2. Verify that the NetApp device has the necessary SnapLock licenses for CIFS and/or NFS.
3. Create the SnapLock volume by issuing a command with the `-L` option. You can also turn an existing volume into a SnapLock-enabled volume in much the same manner.
4. Create necessary qtrees and CIFS shares.

10. Summary and Caveats

This deployment guide strives to highlight some of the key planning and deployment tasks that will lead to success using ZANTAZ EAS in conjunction with NetApp storage for archiving e-mail and files in Lotus Notes environments. Many referenced topics offer general practices that may differ from those needed at a specific customer installation. It is meant to be used as a source of information and a possible starting point for those looking into what may be involved to deploy such an archiving system.

Other combinations of Lotus Notes with ZANTAZ EAS and NetApp storage may require different activities than those described here. Customers are encouraged to clear any platform combinations with the ZANTAZ current list of certified and supported platforms. Customers interested in deployment recommendations specific to their unique environment should also contact ZANTAZ Professional Services for assistance.

Appendix A: Typical Enterprise Configuration of EAS Deployment

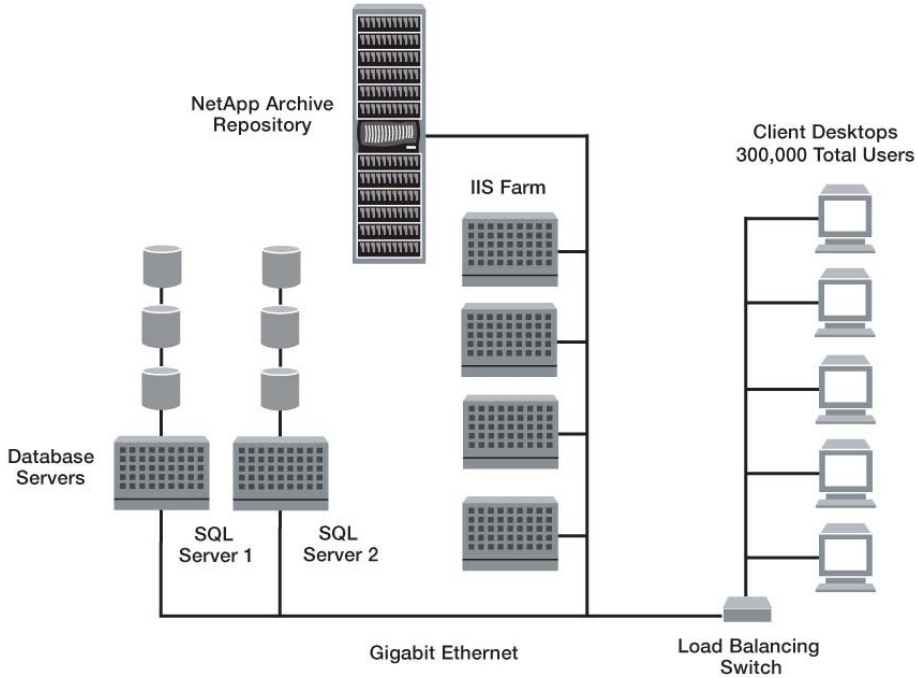


Figure 3) Example of a large-scale data center architecture with ZANTAZ EAS – retrieval.

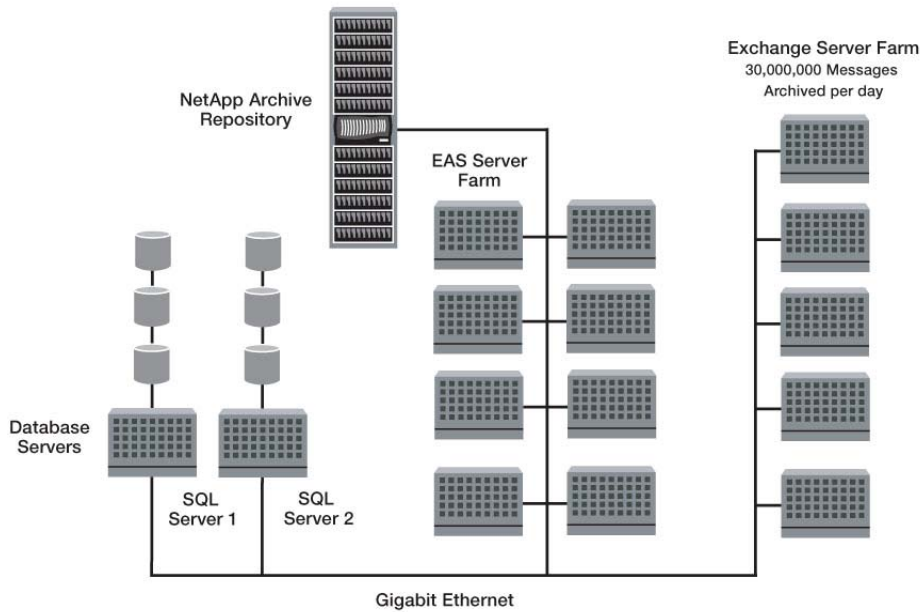


Figure 4) Example of a large-scale data center architecture with ZANTAZ EAS – archiving.

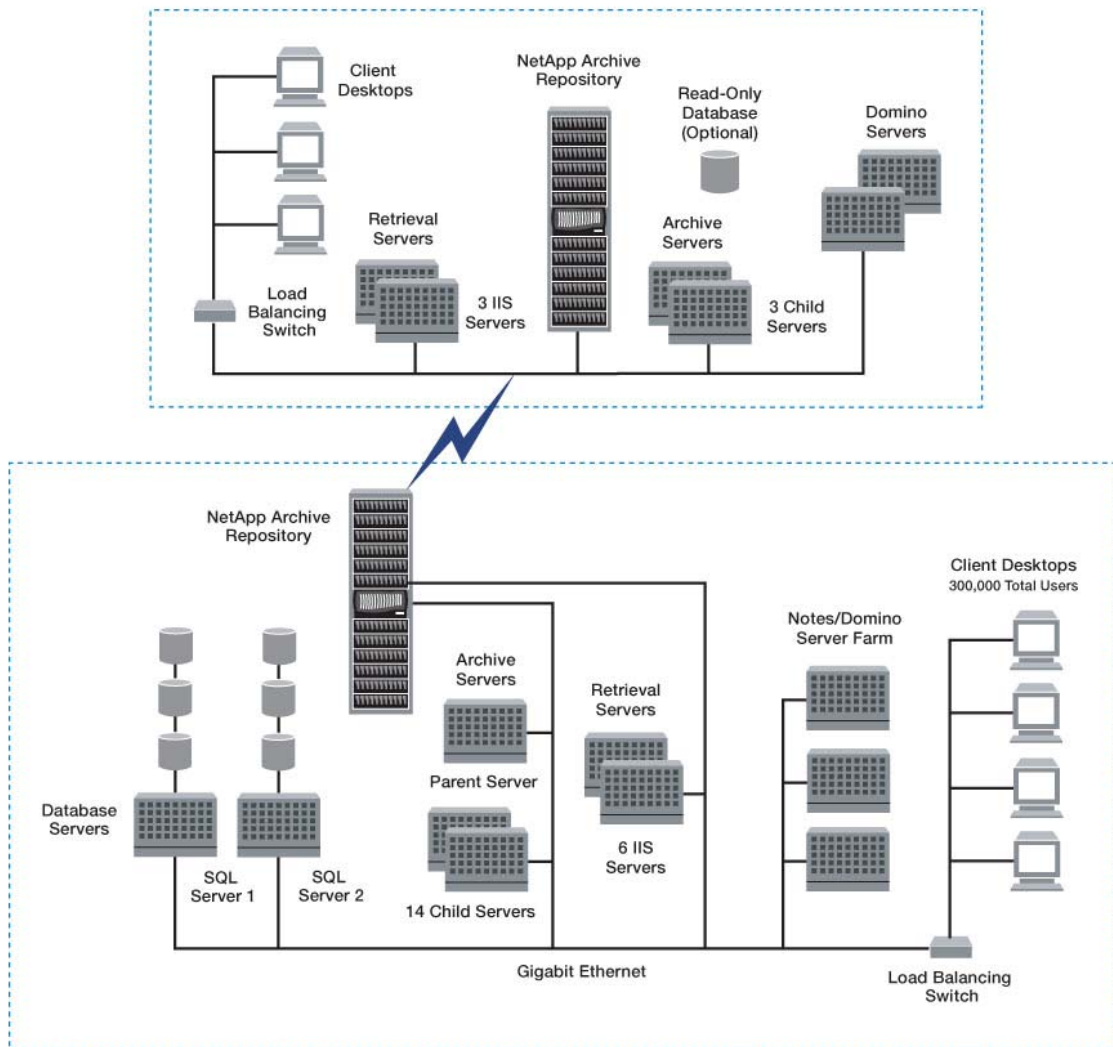


Figure 5) Example of EAS archiving and retrieval in a large-scale enterprise data center architecture.