Willkommen

Bienvenidos

Welkom

Bienvenue

Welcome

mailArchiva
email archiving & compliance

MailArchiva Enterprise Edition V2.6

Installation and Administration Guide

For Windows / Linux

stimulus software

1. INDEX

2. IMPORTANT NOTICE

This Administration Guide covers the installation and configurationGoogle of MailArchiva Enterprise Edition on both Linux and Windows platforms.

It is essential to read this Administration Guide prior to install.

3. CONTACT INFORMATION

Contact Method	Contact Information
Sales/Support Line	(USA) +1-(713)-343-8824 (EU) +44-20-80991035
FAX	(USA) +1-877-350-2328
E-mail Queries	info@mailarchiva.com
Enterprise Support	support@mailarchiva.com
Knowledge Base	http://www.mailarchiva.com/knowledge

4. TECHNICAL REQUIREMENTS

Category	Technical Requirement
Operating Systems	Windows XP Professional, Windows 7, Windows Server 2003, Windows Server 2008 Ubuntu or Redhat Linux Sun Microsystem Solaris or Open Solaris Mac OS X
Disk Storage	Compatible with most Storage Area Networks (SANs) and Network Attached Storage (NAS) devices
Hardware	Hardware varies widely according to no. mailboxes Please refer to the Hardware Sizing Spreadsheet
Mail Servers	Microsoft Exchange 5.5 / 2000 / 2003 / 2007 / 2010 Microsoft Small Business Server Google Apps / Gmail IPSwitch IMail Lotus Domino AXIGEN Postfix Sendmail Qmail Exim Zimbra Neon Insight Communigate Pro Scalix Kerio

5. OVERVIEW

MailArchiva is the email archiving and discovery server software of choice for companies ranging from twenty five users all the way up to sixty thousand users and beyond. It works in conjunction with popular mail servers to ensure that company emails are archived and kept over several years. It thereby helps companies retain valuable knowledge and comply with email archiving legislation such as the Sarbanes Oxley Act. Using MailArchiva's powerful search capabilities, employees and auditors are able to search across millions of archived emails and attachments at the click of a button. The product helps them find the information they are looking for quickly and easily.

In many jurisdictions around the world, the law requires that company emails are kept for up to seven years. MailArchiva is designed to help your company comply with US and EU legislation such as the Sarbanes Oxley Act (SOX), Gramm-Leach Bliley Act (GLBA) and the Freedom of Information Act (FOIA).

The product is compatible with a wide range of mail servers: including Microsoft Exchange 2003, 2007 and 2010 (with full support for envelope journaling), Lotus Notes, Zimbra, IpsSwitch IMail, Axigen, Sendmail, Postfix, Neon Insight, Qmail, Communigate Pro, Java Messaging Server, Kerio and others. Companies have complete flexibility in choosing the hardware and operating system configuration best suited to their needs. For instance, the product can be deployed on a variety of operating systems including Windows, Linux, Solaris, Mac OS X and others. In addition, it is capable of archiving emails to a wide range of network storage devices and storage arrays.

In contrast to many other email archiving systems, MailArchiva stores emails directly on the file system. This design allows one to avoid the pitfalls associated with storing information in a database; namely: high maintenance costs, size restrictions, backup complexity and increased potential for total data loss.

Emails are stored encrypted in standard Internet mail format (RFC822). RFC822 is the industry standard format for storing and transporting email messages. Thus, MailArchiva ensures that your information will remain accessible over the long-run.

The business benefits of MailArchiva include:

Preserve vital company knowledge. MailArchiva ensures that your company email archives remain intact over the long-term.

Access email information with ease. MailArchiva's ultra fast search engine enables you to accurately and efficiently search through years of email data.

Improve user productivity. There is no need for employees to waste time looking for old emails. When using MailArchiva, users can find the emails they are looking for at the click of a button.

Monitor and audit employee communications. Using MailArchiva, department heads can ensure that their employees are communicating effectively with one another and their customers.

Comply with email compliance legislation. MailArchiva provides all the tools needed to ensure that your organization complies with US and EU email archiving legislation (e.g. Sarbanes Oxley Act).

Enhance mail server performance. When company emails are stored in a long-term archive, older emails can be safely deleted from the mail server, thereby freeing up precious mail server resources.

Reduce legal exposure. Most firms will eventually be involved in a legal action of some sort. A technical inability to produce emails that the court has required is no longer an acceptable defence.

Lower storage costs. MailArchiva's de-duplication and compression technology can reduce your storage costs by up to 60%.

Low cost solution. There is no requirement to purchase any additional third-party software and the product is priced competitively.

Highly scalable. MailArchiva is capable of archiving over 150 emails per second on a standard server and is designed to deal with terabytes of data.

6. HIGH-LEVEL FEATURES

Category	Feature	
Internationalization	Multi-language	√
	English, German, Russian, Japanese translations	√
Email Server Support	Exchange 5.5/2003/2007/2010	√
	IpSwitch Imail	√
	Sendmail	√
	Postfix	√
	Qmail	√
	Communigate Pro	√
	Lotus Domino	√
	Zimbra	√
	Neon Insight	√
	Sun Messaging Server	√
	Kerio	√
	Scalix	√
	See Knowledge Base For More	√
Email Archiving	Internal/Outgoing/Incoming	√
	Define Archiving Rules	√
	Define Retention Policies	√
	Message De-duplication	√
	Attachment De-duplication	√
	ZLib Compression	√
	3DES Encryption	√
	Multiple Disk/Volume Support	√
	Failover and Recovery	√
	Auto Volume Creation	√
	WORM Drive Support	√
High Speed Search	Inside Word, Power Point, Excel, PDF, RTF, ZIP, tar, gz, Open Office	√

	Complex Search Criteria	√
Integrity Checking	Digital Signing of Archives	√
	Auto Verification of Archives	√
	Advanced XML Signatures	√
Web Console	Access to Archive From Web Console	√
	Active Directory, LDAP, Basic, Google Authentication	√
	NTLM v2 Windows SSO Authentication	√
	Bulk Export, Print, View, Delete, Restore	√
	Flexible Role Based Security	√
	Save Search Results	√
System Interfaces	SMTP Server	√
	Sendmail milter Server	√
	POP / IMAP Client	√
Monitoring	Search through Audit Logs	√
	System Status Reporting	√
	Real-time Status Info	√
	Real-time Status Chart	
Exchange Server	Message Stubbing	√
	Flexible Stubbing Rules	√
	Multiple Exchange Servers	√
	Multiple Exchange Stores	√
Import	Import From PST	√
	Import From .EML	√
	Import From MBOX	√
	Import From Exchange 2003 / 2007 / 2010	√
	Import From Google Apps	√
Export	Export To /EML Format	√
Outlook Integration	Access Archive From Within Outlook	√
Exchange Stubbing	Stub Attachments in Ms Exchange	√
	Access Stubbed Attachments From With Outlook	√
Ease-Of-Use	Self-Install	√
Google Apps	Archiving Using Google Audit API	√

	Bulk Message Import	√
	Login To Web Console Using Google Credentials	√
Shell	In built command line Shell (Accessible via SSH)	√

7. ARCHITECTURE

The MailArchiva server archives emails from external mail systems such as Microsoft Exchange, Postfix, Sendmail and others. It can either accept SMTP or Sendmail milter traffic from these external mail systems or it can fetch mail from them using IMAP or POP.

The MailArchiva Server can run on any server on your network provided it has TCP/IP connectivity to your mail server. Also, if you intend to authenticate users logging into the server console using Active Directory, it follows there must be TCP/IP connectivity between the MailArchiva server and the Active Directory server. For optimal performance, and to minimize changes to the server hosting your mail system, it is recommended that the MailArchiva server run on a dedicated server platform.

In addition to archiving e-mails, the server provides a web interface that is used to administer the product. This interface, referred to as the “server console”, also provides the capability for users to search and retrieve e-mails. Access to the server console is restricted to authenticated users only. An authenticated user may assume an administrator, auditor, user or custom-defined role. Each of these roles implies a different set of entitlements, which are discussed later in this guide.

For simplicity’s sake, the server may be configured to authenticate users using credentials contained in a simple XML configuration file (Basic Authentication). Alternatively, the server may be setup to authenticate users using Microsoft Active Directory (Active Directory Authentication) or using basic LDAP authentication. The benefit of authenticating with Active Directory or an LDAP server is that you can manage all your user accounts centrally, using standard administration tools.

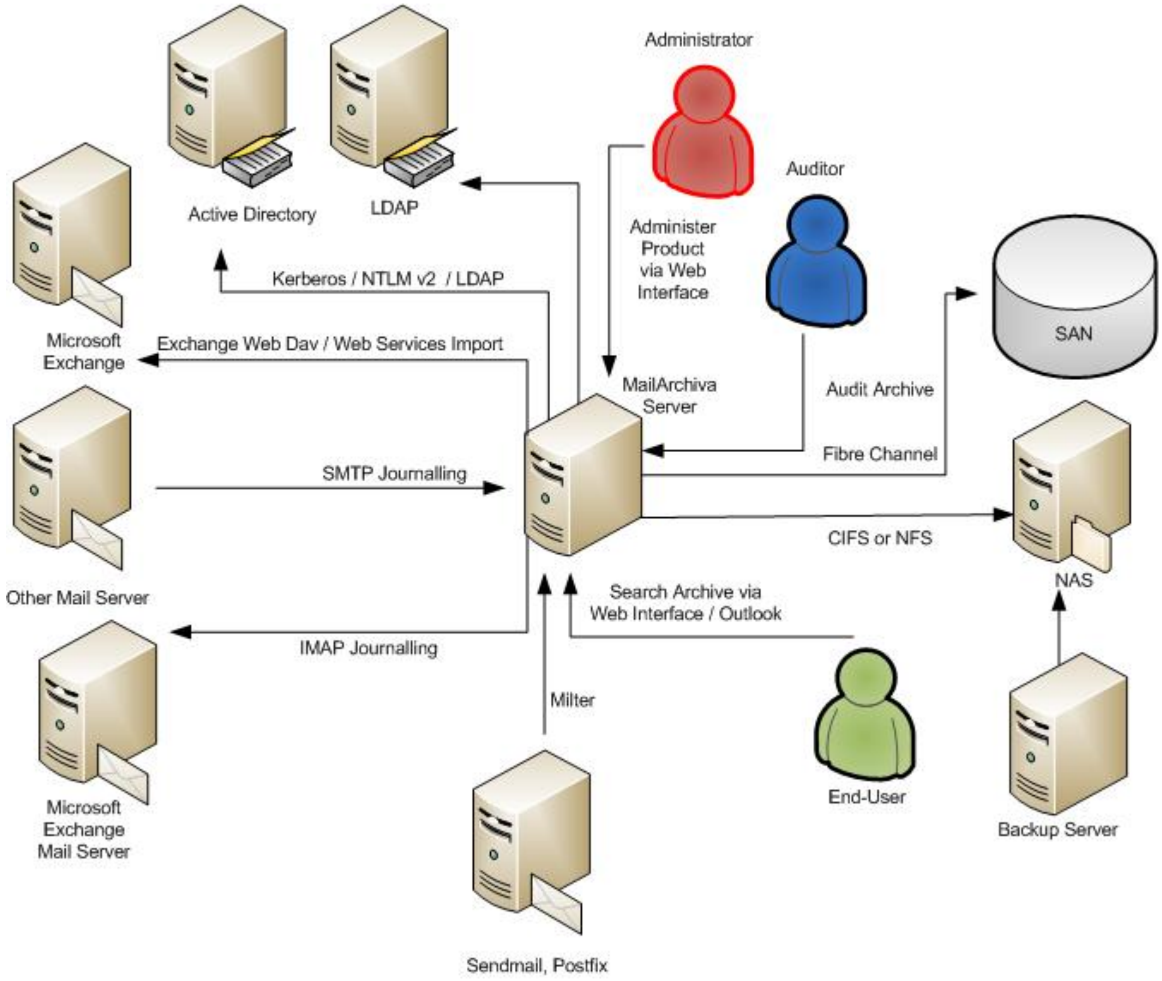

Figure 1 MailArchiva Architecture

If there is a firewall running between any of the components in the architecture, you will need to refer to the communications ports as described in Table 1.

Source	Destination	Protocol	Ports
MailArchiva Server	Active Directory	NTLM, LDAP	8090*, 389
MailArchiva Server	Microsoft Exchange	IMAP	143, 993
Sendmail/Postfix	MailArchiva Server	Sendmail milter	8092*
Mail Server	MailArchiva Server	SMTP	8091*
MailArchiva Server	MailArchiva Server	SSH	2022

Table 1 Communication Ports

* by default, you can change this port

The performance of MailArchiva is largely dependent on the performance characteristics of the hardware environment within which it runs. When planning your hardware configuration it is important to consider factors such as motherboard/chip architecture, CPU speed, number of cores, amount of memory, Ethernet speed, and storage configuration. In larger sites, the server may require more CPU power and larger amounts of memory. If there is any doubt about the needed hardware configuration for your environment, please refer to the MailArchiva Sizing Requirements Spreadsheet, MailArchiva Knowledge Base, or contact us.

Considering the large volume of traffic that will be passing between your mail server and MailArchiva, it is a good idea (especially in larger sites) to install a 1 GB or higher Ethernet link between them. This is especially important if you plan on connecting MailArchiva to an Ethernet based networked storage device since the same pipe may be used for both the retrieval and storage of emails.

The choice of storage hardware and configuration varies greatly depending on the volume of emails the archiving server is expected to handle. In small environments (0-100 mailboxes), two in built SATA drives organized in a RAID configuration is often sufficient. At larger sites, since high speed searching across the large indexes requires low latency disk access times, it is advisable to keep the search engine index and email archive store information on separate drives. While, MailArchiva is capable of, and is indeed optimized for, archiving to remote Network Attached Storage (NAS) Devices, it is never a good idea to store the index data remotely as this will adversely affect the performance of searches. In addition to archiving to NAS devices, enterprise level customers can comfortably configure MailArchiva to archive emails to Storage Area Networks (SAN).

8. INSTALLATION

Before you begin the installation procedure, ensure that you have met the technical requirements of the product stated earlier. Please take note of all IP addresses, usernames and passwords entered during the installation process.

8.1. Server Installation on Windows

Step 1. Install MailArchiva

Run the MailArchiva Server Setup and follow the instructions on screen. Please install both the MailArchiva Server and Application Server components.

Step 2. Check Availability of Port 8090

By default, MailArchiva uses port 8090. Before starting the server, ensure that port 8090 is not being used by another application. You can do this by typing “netstat -abn” from the console. If port 8090 is in use, edit the file C:\Program Files\MailArchiva\Server\conf \server.xml and change all references from “8090” to the desired port.

Step 3. Start MailArchiva Server

The MailArchiva application appears in the Windows task tray. Double click the MailArchiva task tray and click Start. Verify that the server is started correctly by clicking Start->Program Files->MailArchiva Console Login. If you see a login box in the browser window, the MailArchiva server is installed correctly. You can control the MailArchiva service directly from the Windows Services applet in the Control Panel.

8.2. Server Installation on Linux

The MailArchiva Server can be installed on a variety of Linux distributions and operating systems. The instructions in this section illustrate the steps required to install the server on Fedora specifically.

The below procedure may vary slightly on different Linux distributions. However, armed with sufficient knowledge of your distribution, you should be able to setup MailArchiva on your preferred Linux system with relative ease.

Step 1. Install/Upgrade MailArchiva Server

To install the server, type the following:

```
tar –xvzf mailarchiva_enterprise_edition_server_v2_5_0_linux.tar.gz
cd mailarchiva_dist
./install
```

Following the above, the server executables will be installed /usr/local/mailarchiva/server.

```
cd /usr/local/mailarchiva/server
```

♦

Step 2. Check Availability of Port 8090 and Port 8091.

By default, MailArchiva uses port 8090 and port 8091. Before starting the server, ensure that these ports are not being used by another application. You can do this by typing “netstat -vatn” from the console. If port 8090 is in use, edit the file /usr/local/mailarchiva/server/conf/server.xml and change all references from “8090” to the desired port.

Step 3. Start MailArchiva Server

To start the MailArchiva Server from the commandline type:

```
sudo /etc/init.d/mailarchiva start
```

Type: “http://localhost:8090” in a web browser to access the web console. If you cannot access the console, check that port 8090 is open on your firewall and examine the log files in /usr/local/mailarchiva/server/logs.

To stop the server, you would type:

```
sudo /etc/init.d/mailarchiva stop
```

8.3. Basic MailArchiva Server Configuration

The instructions below cover the minimum configuration steps needed to ensure that MailArchiva is ready for archiving. Refer to Section 9 for more detailed information on how to configure the product.

There are at minimum five configuration tasks that need to be performed before the server is ready to start archiving emails. These tasks are performed in the MailArchiva Console Configuration.

Step 1. Login to the MailArchiva Console

Login to the Admin Console using the default master username and password (username “admin”, password: “admin”). Click the Configuration button in the top menu to enter the area where the product is configured.

Step 2. Enter the Master Login Password

Click the Logins section and enter a suitable master admin login password. The master account is an all-powerful account that has access to all the features of the product. In future, to access the MailArchiva console as the master account, simply login as “admin” and the password.

Step 3. Enter an Archive Encryption Password

In the Volumes section, enter a suitable encryption password. Care should be taken to remember this password, as if it is forgotten, all archived information will be lost permanently! There is no back door.

Step 5. Add a Volume

In the Volumes section, add a new volume. A volume specifies where your archived emails and corresponding index data is stored. The store path specifies the location where the actual emails are to be stored, while the index path refers to the location where the search engine index data is to be stored. For performance reasons, it is important to keep index data on a local hard drive. The store path may refer to an external network drive, if desired. To specify a network location, on Windows, a UNC path may be specified provided the MailArchiva server service is running under an account with sufficient privileges to read and write to the external device. On Linux, UNC paths cannot be directly specified and must be mounted separately from the operating system (if MailArchiva is not in appliance mode).

Step 4. Add Local Mail Domains

In the Domains section, enter one or more local mail domains. For instance, if your users send emails to user@company.com or user@abc.com, both these domains should be added.

8.4. Mail Server Configuration

The MailArchiva email archiving system can be configured to inter-operate with a wide array of mail servers. As such, the configuration steps vary depending on your particular choice of mail systems.

Mail Server	Sections
MS Exchange 2003	8.4.1, 8.4.3, 11.2
MS Exchange 2007 / 2010	8.4.2, 8.4.3 or 8.4.4,11.2
Google Apps	8.4.3, 8.4.5, 11.1
IpSwitch IMail	Refer to IMail Docs
Sendmail	8.4.6
Postfix	8.4.7
Qmail/Exim	8.4.8
Neon Insight	Refer to Knowledge Base
Lotus Domino	Refer to Knowledge Base
Zimbra	Refer to Knowledge Base
Sun Messaging Server	Refer to Knowledge Base
AXIGen	Refer to Knowledge Base
Communigate Pro	Refer to Knowledge Base
Kerio	Refer to Knowledge Base
Scalix	Refer to Knowledge Base
Alternate	Refer to Knowledge Base

Table 2 Installation Sections to Complete

*Knowledge Base (see: http://www.mailarchiva.com/knowledge)

For additional mail server support, please refer to the MailArchiva Knowledge Base at http://www.mailarchiva.com/knowledge

There are two methods by which MailArchiva is typically integrated with Microsoft Exchange:

- **IMAP Archiving** – MS Exchange is configured to forward a copy of every email processed by the mail server to a temporary "journal mailbox". MailArchiva's inbuilt IMAP client periodically connects to Exchange's IMAP server, retrieves and archives the mail and then deletes it from the mail server.

- **SMTP Archiving** – MS Exchange is configured to forward copies every email processed by the mail server directly to MailArchiva's inbuilt SMTP server.

Both of these methods require that Microsoft Exchange's message journaling feature is enabled. Microsoft Exchange supports three different types of message journaling: standard journaling, BCC journaling and envelope journaling. In standard journaling, when an e-mail message is copied to the journaling mailbox, that message does not include BCC or alternative recipient information. Furthermore, if the message is addressed to a distribution group, the addressing information does not contain the individual recipients comprised of the distribution group. BCC journaling is similar to standard journaling except that the BCC field is included with all archived messages. In envelope journaling, all available RFC2821 and RFC2822 recipients are captured. Thus, an archived message includes all available header information, including BCC fields and the full expansion of distribution groups.

Please refer to Table 3 for an overview of Microsoft Exchange related features supported by MailArchiva Enterprise Edition.

Microsoft Exchange	MailArchiva EE
Standard journaling	√
BCC journaling	√
Envelope journaling	√
Multiple mail stores	√
Multiple exchange servers	√

Table 3 MailArchiva Exchange Features

8.4.1. Microsoft Exchange 2003 IMAP Journaling

Step 1. Create a Journal Account

On the server running Microsoft Exchange, using the Active Directory Users and Computers browser, create a Windows user account where all incoming and outgoing mail will be temporarily archived. This account must reside on your company's domain (i.e. not a local machine account).

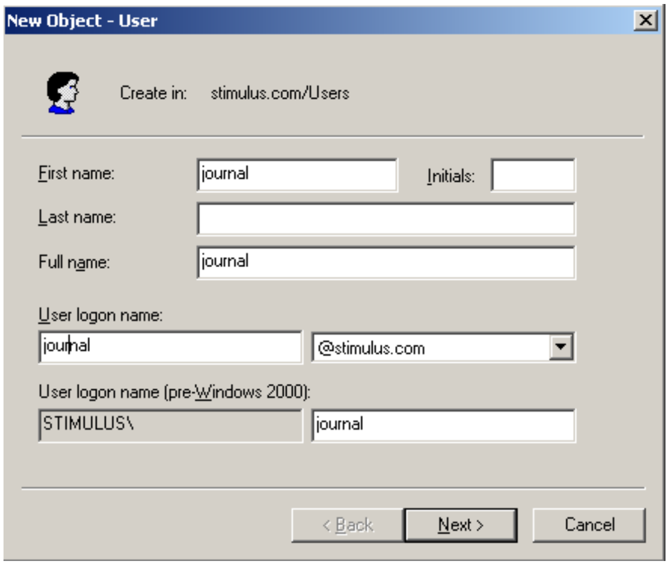

Create in: stimulus.com/Users

First name: journal

Initials:

Last name:

Full name: journal

User logon name: journal @stimulus.com

User logon name (pre-Windows 2000): STIMULUS\ journal

Figure 2 Journaling Account Creation

Step 2. Enable Journaling on Microsoft Exchange

On the same server, run the System Manager Application included with Microsoft Exchange. Locate the Mailbox Store node in the tree view on the left. It is in Servers->First Storage Group->Mailbox Store. Right click the Mailbox Store object and click Properties. A dialog will appear as in Figure 3. Click Browse and enter "journal" for the object name. Click OK. Journaling is now enabled for the Mailbox Store

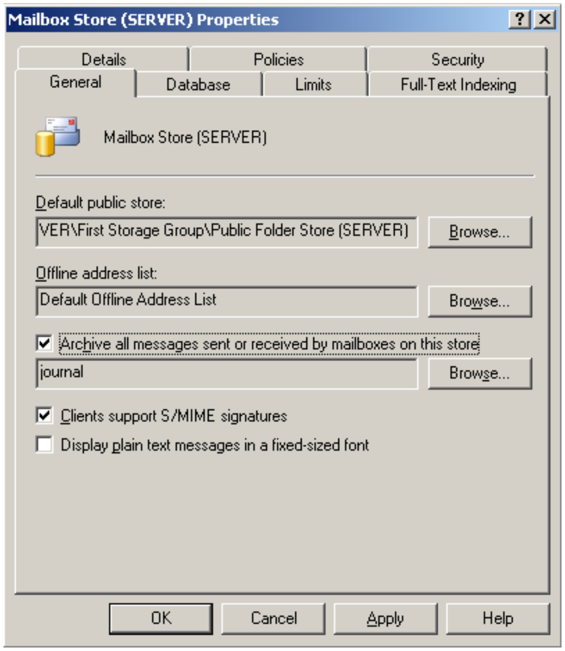

Mailbox Store (SERVER)

Default public store: VER\First Storage Group\Public Folder Store (SERVER)

Offline address list: Default Offline Address List

☑ Archive all messages sent or received by mailboxes on this store

journal

☑ Clients support S/MIME signatures

☐ Display plain text messages in a fixed-sized font

Figure 3 Enable Journaling

Step 3. Enable Envelope Journaling

1. Install the latest Service Pack
2. Download the Exejcfg.exe utility from Microsoft's Download Center

To enable envelope journaling, from command prompt, type:

```
Exejcfg -e
```

As of writing, the Exejcfg.Exe utility can be downloaded from http://www.microsoft.com/downloads/details.aspx?familyid=E7F73F10-7933-40F3-B07E-EBF38DF3400D&displaylang=en

Step 4. Start IMAP Service

Start the Microsoft Exchange IMAP Service in Windows Services.

Please perform the configuration steps in section 8.4.3 to configure the MailArchiva product to connect to the journal account setup in Microsoft Exchange.

Step 5. Configure Your IMAP Server Setings

If necessary, configure the Microsoft Exchange IMAP server to accept plaintext authentication. In addition, MailArchiva expects the journal envelopes it receives to be in plaintext. Care should be taken to ensure that Microsoft Exchange's IMAP server will not convert the journal messages to TNEF, rich text or HTML.

To verify that message archived in MailArchiva are in fact envelope journal messages, select Configuration->Search tab->Export Method->journal Message. Thereafter, in Search, export a message and view a text editor. The message should contain a journal envelope as described in Chapter 15. There should be no HTML or special character codes in the journal envelope. If there are special characters, then

TNEF or HTML conversion is not disabled as it should be.

8.4.2. Microsoft Exchange 2007 / 2010 IMAP Journaling

Step 1. Create a Journal Account

On the server running Microsoft Exchange, open the Active Directory Users and Computers console, right-click on the Users container and choose New -> Contact from the menus. When prompted, enter the first name, last name, full name and display name of the contact you're creating and click OK.

Step 2. Enable Journaling Agent

Go to the Hub Transport server, open the Exchange Management Shell, and execute the following command:

```
Get-TransportAgent
```

The Get-TransportAgent command will return a status of either True or False, indicating whether or not the journaling agent is enabled.

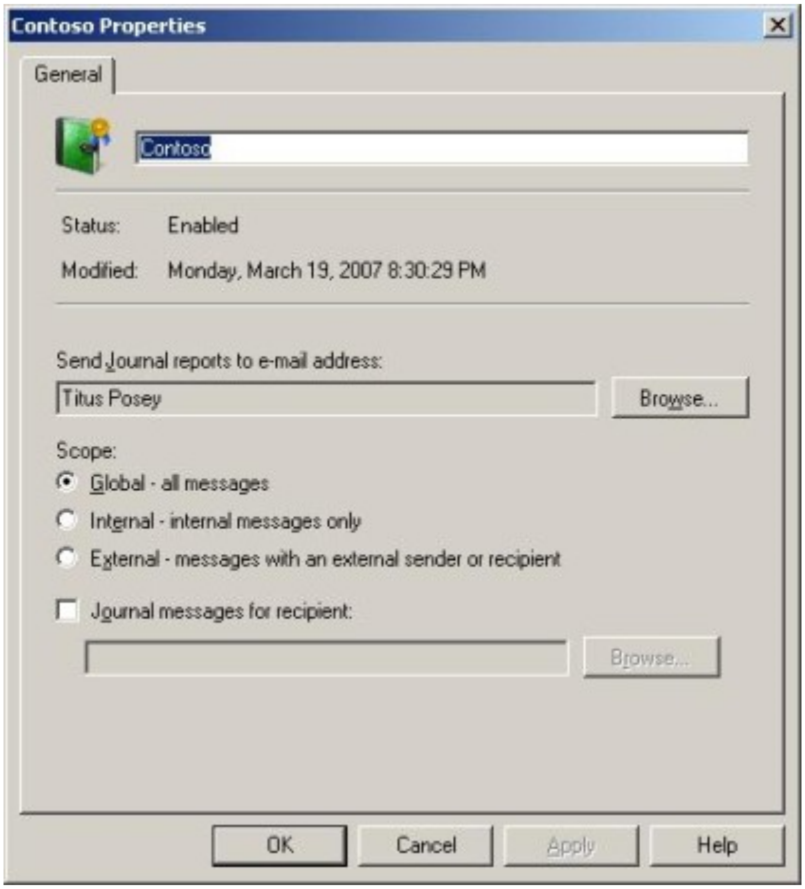

Contoso Properties

General

Contoso

Status: Enabled

Modified: Monday, March 19, 2007 8:30:29 PM

Send Journal reports to e-mail address: Titus Posey

Scope:

● Global - all messages

○ Internal - internal messages only

○ External - messages with an external sender or recipient

□ Journal messages for recipient:

Figure 4 Select Journal Recipient

If the journaling agent's enabled status is False, then you will have to enable the journaling agent before continuing. To do so, enter the following command:

```
Enable-transportagent
```

The command prompts you to enter additional information.

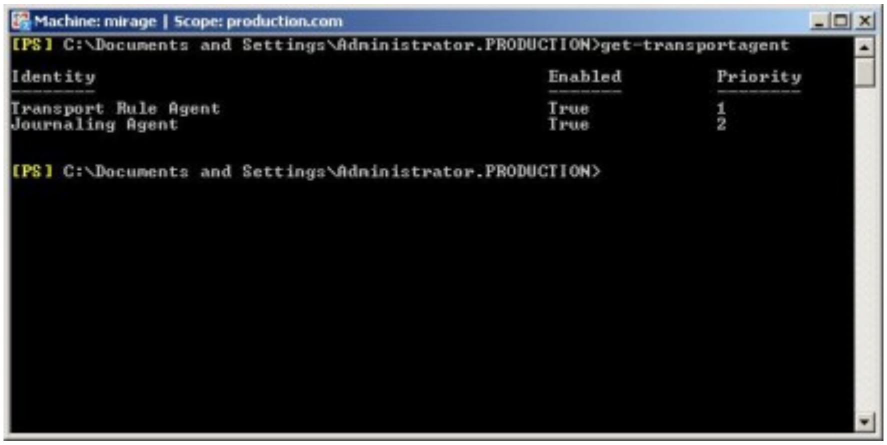

```
[PS] C:\Documents and Settings\Administrator.PRODUCTION>get-transportagent

Identity                                                Enabled         Priority
--------                                                -------         --------
Transport Rule Agent                                    True            1
Journaling Agent                                        True            2


[PS] C:\Documents and Settings\Administrator.PRODUCTION>
```

To enable the journaling agent, enter the words "Journaling Agent."

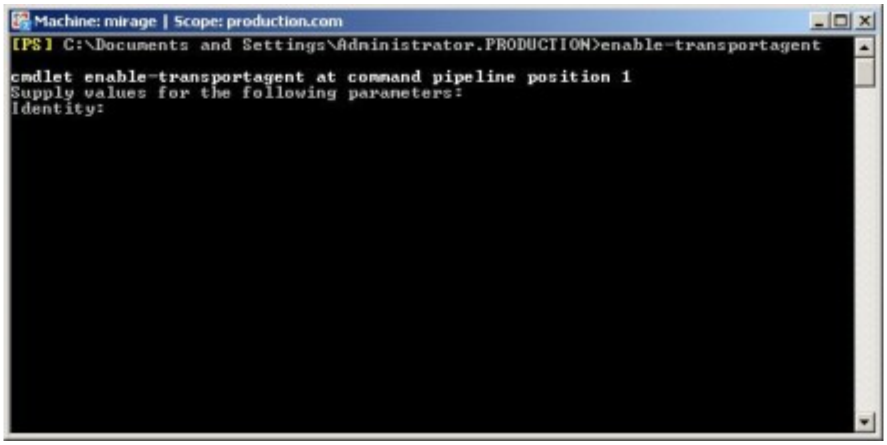

```
[PS] C:\Documents and Settings\Administrator.PRODUCTION>enable-transportagent

cmdlet enable-transportagent at command pipeline position 1
Supply values for the following parameters:
Identity:
```

The Exchange Management Shell doesn't provide you with any confirmation that the journaling agent has been enabled. Enter the Get-TransportAgent command once again to confirm that the journaling agent has been activated.

Step 3. Create a Journaling Rule

The process of creating an Exchange 2007/2010 journaling rule is fairly simple:

1. Open the Exchange Management Console and navigate through the console tree to Organization Configuration -> Hub Transport.

2. Select the Journaling tab, and then click the New Journal Rule link found in the Actions pane to open the New Journal Rule wizard. Enter a descriptive name for the Exchange 2007 journaling rule that you are creating.

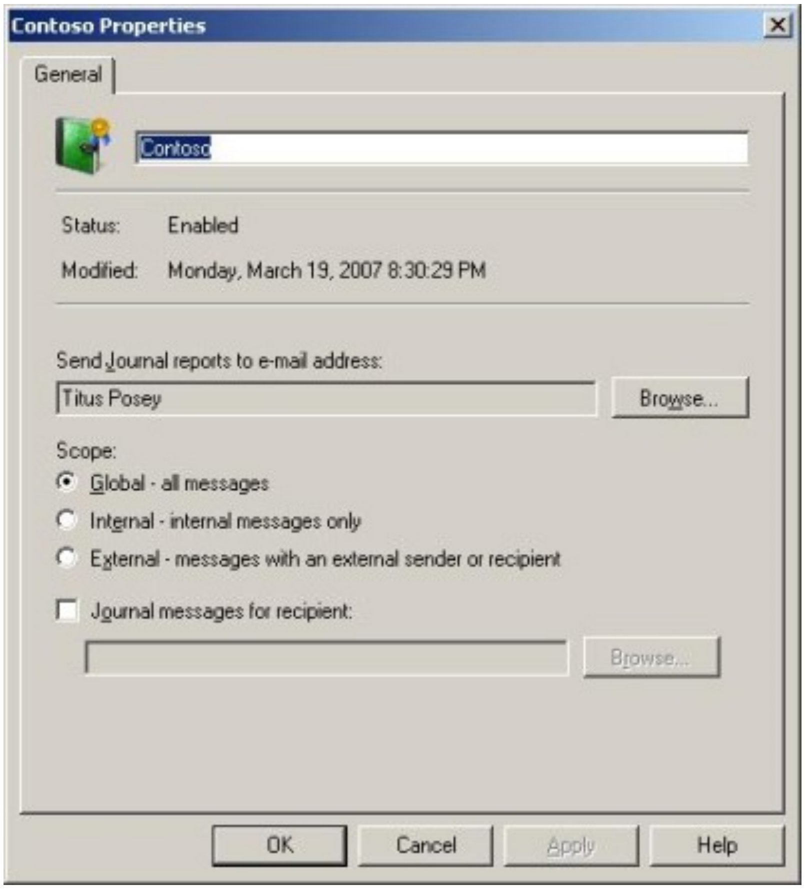

Contoso

Status: Enabled

Modified: Monday, March 19, 2007 8:30:29 PM

Send Journal reports to e-mail address: Titus Posey

Scope:

● Global - all messages
○ Internal - internal messages only
○ External - messages with an external sender or recipient

☐ Journal messages for recipient:

3. Just beneath the Rule Name field is the Send Journal Reports to E-Mail Addresses field. This is where you supply the name of the journal mailbox.

4. In the scope field, select all messages.

5. In the next field, specify the journal mailbox as the recipient.

6. Finally, the wizard contains a checkbox that you can use to enable the rule upon creation. This is checked by default, unless you choose to deselect it prior to creating the journaling rule.

7. After filling in the wizard's various parameters, click the New button and the Exchange 2007 journaling rule will be created.

Step 5. Configure Your Exchange IMAP Server Setings

Configure the Microsoft Exchange IMAP server to accept plaintext authentication. In Exchange 2007, the IMAP server settings are accessible from the Exchange Management Console in Server Configuration > Client Access > IMAP4 > Properties -> Authenticaton tab.

MailArchiva expects the journal envelopes it receives to be in plaintext. Care should be taken to ensure that Microsoft Exchange’s IMAP server will not convert the journal messages to TNEF, rich text or HTML. In Exchange 2007, this option is accesible from the Exchange Management Console in Server Configuration -> Client Access -> IMAP4 -> Properties -> Retrieval Settings tab. Select the Message MIME format to “Best Body Format”.

To verify that message archived in MailArchiva are in fact envelope journal messages, select Configuration->Search tab->Export Method->journal Message. Thereafter, in Search, export a message and view a text editor. The message should contain a journal envelope as described in Chapter 15. There should be no HTML or special character codes in the journal envelope. If there are special characters, then TNEF or HTML conversion is not disabled as it should be.

Step 4. Start IMAP Service

Start the Microsoft Exchange IMAP Service in Windows Services.

Please perform the configuration steps in section 8.4.3 below to configure the MailArchiva product to connect to the journal account setup in Microsoft Exchange.

8.4.3. MailArchiva Server IMAP Configuration

On a fresh install of the Exchange product, the IMAP connector is switched on and ready for action, however, it is possible that it has since been disabled. Before continuing, ensure that Microsoft Exchange's IMAP service is switched on.

Step 1. Add a Journal Account Connection

In the Journal Accounts tab of the MailArchiva server console configuration screen, click the "Add Journal Account" button and edit the following:

- Select IMAP as the preferred protocol
- Enter the server address of your Exchange server
- Enter the Microsoft Exchange journal account username and password
- For "Connection Mode", select TLS when available
- Select "No certificate authentication".
- Set "Listen for message arrival notifications from server" to checked.

Depending on how Exchange's IMAP server is configured, you may able to authenticate using a plain username such as "john" or an NTLM username may be required.

The NTLM login name is somewhat different in MailArchiva. Assuming the journal user's email address is "J.User@server.com", its Windows NT login name is "juser", its NT domain name is "dom", and its Exchange mailbox name is "Joe User", then a username of "dom\juser\J.User" should be used (Note: dom\juser will not suffice!).

Thus, if the email address on the journal account is journal@company.com, its NT login name is 'journal', your NT domain name is "company", then you would specify "company\journal\journal".

The IMAP Idle feature allows MailArchiva to receive notifications when new mail arrives in the journal account, rather than having to continuously poll the mail server. Your mail server may or may not provide support for IMAP Idle. If you experience problems, please disable it.

If you are having trouble authenticating with your mail server, telnet to port 143 on your Exchange server to establish whether Exchange's IMAP server is reachable.

If so, it is a good idea to experiment with all the different authentication methods (e.g. SSL, TLS, no auth or TLS when available).

Some versions of Exchange 2007 have an IMAP server that reports support for plain authentication even when it is disabled. The "disable plain authentication" option is a workaround to this bug that wil cause MailArchiva's IMAP client to always disable plain authentication and choose another.

Due to a problem in the Microsoft Exchange IMAP server, insufficient number of bytes may be retrieved when reading big messages. There are two ways to workaround this Exchange bug:

(a) The Exchange IMAP server provides a configuration option called
"fast message retrieval" to the UI. Simply go to the site, server
or recipient, click on the "IMAP4" tab, and one of the check boxes
is "enable fast message retrieval". Turn it off and the octet
counts will be exact. This is fully described at

> **http://support.microsoft.com/default.aspx?scid=kb;EN-US;Q191504**
>
> **(b) Set the "mail.imap.partialfetch" property to false. You'll have to set this property in the Properties object that you provide to your Session.**

Step 2. Test Journal Account Connection

Click the ‘Test Journal Account Connection’ button to determine if the connection is established. If the test is successful, save your configuration settings and emails should start appearing in the search results in a matter of a few seconds. If MailArchiva cannot establish a connection to Microsoft Exchange’s IMAP server, verify that you entered the correct information and that Microsoft Exchange’s IMap connector is listening. You could also try using both the full journal account name (e.g. journal@company.com) and the short name (e.g. journal).

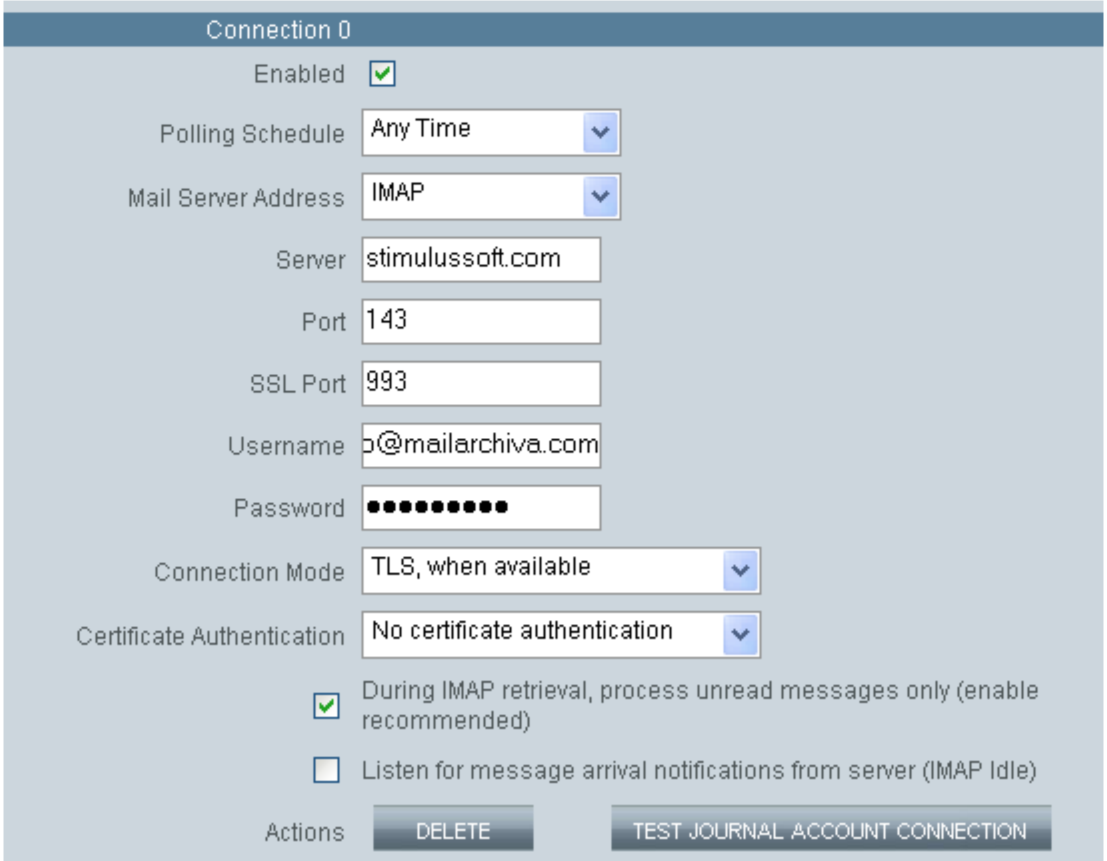

Connection 0

Enabled	☑
Polling Schedule	Any Time
Mail Server Address	IMAP
Server	stimulussoft.com
Port	143
SSL Port	993
Username	@mailarchiva.com
Password	●●●●●●●●●●
Connection Mode	TLS, when available
Certificate Authentication	No certificate authentication

☑ During IMAP retrieval, process unread messages only (enable recommended)

☐ Listen for message arrival notifications from server (IMAP Idle)

Actions

Figure 5 Journal Account Connection To Microsoft Exchange

8.4.4. Microsoft Exchange 2007 / 2010 SMTP Journaling

The SMTP journaling approach is typically used when high performance archiving is needed (it is much faster than using the IMAP approach). The use of SMTP journaling is generally discouraged for the simple reason that if MailArchiva goes down, Exchange could give up sending emails to it for a period and important mail could be lost.

To achieve journaling via the SMTP approach, a contact in Microsoft Exchange must first be created as follows:

1. Open Active Directory Users and Computers
2. Right-click an organizational unit in which you want to create the contact, point to New, and then click Contact
3. Enter the following: First Name: MailArchiva Last Name: Archive Display Name: MailArchiva Archive. Click OK
4. Open the Exchange Management Console
5. Expand Recipient Configuration, right-click Mail Contact, and then click New Mail Contact
6. Click Existing Contact, browse to and select the recently created MailArchiva Archive contact, then click OK
7. Click Next
8. For the External Email Address field, click Edit, enter the SMTP MailArchiva archive email address (archive@mailarchiva.server.com) then click OK. The FQDN of the MailArchiva server should be used, it must ping from Exchange server, and the SMTP connector must be setup to listen on port 25.
9. Click Next, then click New

To configure the message format settings for the SMTP contact in Exchange 2007 SP1:

1. Open the Exchange Management Console
2. Expand Recipient Configuration, then select Mail Contact
3. In the result pane, select the SMTP contact
4. In the action pane, under the SMTP contact, click Properties
5. In the General tab, from the Use MAPI rich text format list, select Never.

With this setting, journal reports are sent in MIME rather than S/TNEF.

To enable standard journaling:

1. Open the Exchange Management Console
2. Expand Server Configuration, then click Mailbox
3. In the result pane: select the server for the mailbox database for which you want to enable journaling
4. In the work pane, right-click the mailbox database, then click Properties
5. In the General tab, select Journal Recipient
6. For Journal Recipient, click Browse, select the MailArchiva archive recipient, then click OK
7. Click OK

The SMTP Journaling requires that MailArchiva's inbuilt SMTP server is listening at port 25. In the MailArchiva console configuration, select the Listener's tab and change the listening port from 8091 to 25.

8.4.5. Google Apps / Gmail

MailArchiva is capable of archiving emails from all Google Apps accounts belonging to an entire organization. It achieves this by instructing Google to continously forward copies of all mails to a journal mailbox that you create. MailArchiva then retrieves the mails using its embedded IMAP client. The mails that MailArchiva receives are specially crafted Google Journal Messages that are designed to preserve header information. Once archiving is setup correctly, MailArchiva will parse, index, archive, and display these journal messages in the usual manner.

Following the steps below to setup Google archiving:

1. Login using Administrator account to Google Apps
2. Click on Domain settings
3. Under Domain Settings Click User Settings
4. Tick 'Enable provisioning API'

5. Save changes
6. Create a new mailbox in Google Apps called “journal”
   (This mailbox will be used as a temporary copy-all destination for all mailflow in your organization)
7. Login to the Gmail "journal" account
8. Click Setting->Forwarding and POP/IMAP->IMAP Access->Enable IMAP
9. Save changes
10. Login to MailArchiva as Admin
11. In MailArchiva->Configuration->Archive, enter your Google administrator account email address and password
12. Specify the name of the journal account. I.e. “journal” in the Journal mailbox field
13. Check interpret Google journal messages
14. Check Enable Google Archiving
15. In MailArchiva->Configuration->Journal Accounts, click New Journal Account
16. Select IMAP as the Protocol
17. Specify “imap.gmail.com” as the server name
18. Enter the equivalent of “journal@company.com” and the associated password in the username and password fields
19. Select “SSL” as the Connection Mode
20. Enable “Listen for message arrival notifications from server”
21. Click Save
22. Test the Connection to ensure it is authenticating correctly

It should take 5-10 minutes before messages will be begin appearing in the journal mailbox. After some time, send a test message to confirm that archiving is functioning correctly. If no emails appear after 10 minutes, click System Status->System Alerts to see there are any error messages outputted. For example, MailArchiva may not have sufficient permissions to set monitors on all user accounts. With troubleshoot logging enabled, a detailed explanation of the problem will be outputted in the debug log.

Should MailArchiva be unable to able to keep up with the mailflow in your organisation, you may need to disable “Listen for message arrival notifications from server” and set a low (e.g. 10 msec) polling interval.

The steps necessary to setup Google login to the Web Console and for importing old emails from Google directly are discussed later in this Guide.

8.4.6. Sendmail

The MailArchiva server incorporates a Sendmail milter server and thus is able to integrate with Sendmail and postfix directly.

(1) Add the following to Sendmail’s sendmail.mc file:

```
INPUT_MAIL_FILTER(`mailarchiva', `S=inet:8092@127.0.0.1')dnl
```

(2) Compile the sendmail.mc file

```
sudo m4 /etc/mail/sendmail.mc >/etc/mail/sendmail.cf
```

(3) Restart send mail

```
sudo /etc/init.d/sendmail restart
```

8.4.7. Postfix

(1) Add the following to Postfix’s main.cf file:

```
milter_default_action = tempfail
smtpd_milters = inet:127.0.0.1:8092
```

(2) Restart Sendmail

```
sudo /etc/init.d/postfix restart
```

8.4.8. Qmail / Exim

Consult the MailArchiva Knowledge Base (http://knowledge.mailarchiva.com) as it contains a patch that provides interoperability with Qmail.

9. MAILARCHIVA SERVER CONFIGURATION

The MailArchiva server configuration settings are accessible from the Configuration section of the server console. Only users with the appropriate administrator rights can view or modify configuration settings.

To access the web console from Windows, click Start->Program Files->MailArchiva-> MailArchiva Console Login. On Linux, type: “http://localhost:8090” in a web browser.

If the server is installed correctly, you should see the MailArchiva login screen. If this is the first time you are configuring the product, login to the console using username “admin” and password “admin”.

❗ **The default web console login username is “admin” and password is “admin”.**

9.1. Local Mail Domains

When configuring MailArchiva for the first time, it is necessary to add one or more mail domains associated with your company. To do this, click “Add Domain” in the domain section of the configuration screen. An example domain is “company.com” or “company.org”. The entered domains are used by the MailArchiva server to assess whether the origin and destination of emails are internal or external to your organization. For example, when applying archive rules, the server will match the domain of a given email address with all of the configured domains.

9.2. Archive Encryption Password

All emails stored are encrypted using triple DES password-based encryption. Before using the server to archive emails, you need to choose and enter an Archive Encryption Password in the Volumes tab of the Configuration screen.

Bear in mind, the password you enter is irrecoverable, so it is very important that you remember it. Furthermore, since the password holds the key to your archived emails, you need to ensure that the password is kept highly confidential and secret. It is also important to bear mind that you cannot change the password once the server has begun to archive emails.

Once you have set the encryption password, it is essential to backup the file server.conf located in mailarchiva\server\webapps\MailArchiva\WEB-INF\conf from the root of your MailArchiva installation directory. This file contains your password and a specific salt value used for email encryption purposes. If you lose either of these, you will be unable to access the emails archived by the server in perpetuity!

It is of paramount importance that a backup of the server.conf configuration file is made and that is stored in a secure location.

9.3. Volumes

Archived emails are organised into one or more logical volumes. Each volume consists of an index and a store. The index is used to enable auditors to perform efficient search queries on the archived data. The store consists of multiple sub-directories where the archived information is kept.

When creating a volume, the index path and store path can refer to any location on disk. Furthermore, volumes are defined in terms of their order of preference. When a volume has reached its size limit, the server will automatically switch over to the next available volume on the list. This mechanism allows one to archive information on multiple hard disks, without necessitating manual intervention.

❗ **Never store the index data on a remote drive such as NAS. MailArchiva's search engine requires very low latency when accessing the index.**

❗ **Archive data may be stored on a remote drive since this data is accessed infrequently.**

To create a volume, click the "New Volume" button in the Configuration screen. Enter a path for the store and index (e.g. "c:\store1" and "c:\index2"). If you've created more than one volume, click the "Up" and "Down" buttons to organise them according to your order of preference.

It is often desirable to archive emails to a remote NAS or SAN disk. In the Windows version of MailArchiva, it is possible to specify a UNC path of the remote drive (e.g. \\server\store\store0) as the store path location. Before doing this, ensure that the MailArchiva service is running under an administrator account. From the Services Control Panel applet (not the MailArchiva task tray icon configuration!), select the service, right click, select Properties, select Logon Tab, enter your Administrator account login account details, save and restart the MailArchiva server.

❗ **If the volume store path is set to a remote location and the volume is shown as EJECTED, it is likely that you have run into a permissions issue. Ensure that MailArchiva is running under an account with sufficient privileges to read and write to the remote drive.**

Specifying UNC paths are not supported in the Unix versions of MailArchiva. Rather, a mount point must be defined in your /etc/fstab file and set to the base location of your NAS or SAN disk. For example, the following is added to the /etc/fstab file:

a. smbfs

```
//192.14.12.122/archive mnt/archive smbfs username=admin,password=pw
```

b. nfs

```
192.14.12.122:/archive /mnt/archive nfs
```

Thereafter, when specifying the store path, simply enter /mnt/archive/store0.

Ocassionally, due to network errors, it is possible that your mount point may become dismounted. In this case, during archiving, instead of writing to the remote mount point, your data may be written to local directory called /mnt/archive/store0. This is undesirable as the emails will be stored in a location outside MailArchiva's control. To prevent this from happening, it is a good idea to manually create the directory /mnt/archive/store0 and set it to read-only.

> ❗ **If you are running MailArchiva in 'Appliance Mode', the external mount points can be defined from within the Volumes section of the web console configuration.**

Once you've created a volume, you'll notice that it is assigned the "NEW" status, as described in Table 4 below.

Volume Status	Description
NEW	The volume has just been created and has not been saved.
UNUSED	The volume has been saved but it does not contain any information.
ACTIVE	The volume is currently being used for archiving purposes.
CLOSED	The volume is searchable, however, no further information can be written to it.
UNMOUNTED	The volume is not searchable, nor can it be made active.
EJECTED	Volume was removed without explicitly unmounting it.

Table 4 Volume Status

Once the configuration is saved, the volume status will be assigned the “UNUSED” status. When the first email is archived, the server will automatically switch over to the first unused volume on the list and set its status to “ACTIVE”. This volume will stay active until such time as its maximum size is exceeded, the disk is full, or the volume is explicitly closed.

Only one volume can be active at a time and once the active volume is closed, no further data can be written to it and it cannot be reopened from the server console. The purpose of this behaviour is to ensure that archive data is stored chronologically across multiple volumes. If for any reason, an active volume is accidentally closed, the volume status can be edited in a file called volumeinfo in the volume store path (if WORM drive support is disabled) or in the volume index path (if WORM drive support is enabled).

If at any stage during the archiving process, the server finds that an active volume is not available, it will always activate the next unused volume on its list. Assuming there are no remaining unused volumes available, the server will stop the archiving process until such time as a new volume is added.

When using removable disks, it is not recommended to remove the disk containing the active volume data without closing/unmounting the volume first. Any physical disk containing a closed volume may be removed provided that the volume whose store path refers to it is unmounted first.

When users search for emails, the search is conducted across both active and closed volumes. In the unlikely event that a volume's search index is corrupted, it can be regenerated. Re-indexing is a time consuming process and is only reco/mmended in the event of data loss. To re-index a volume, you need to close it first, and click on the "Re-Index" button.

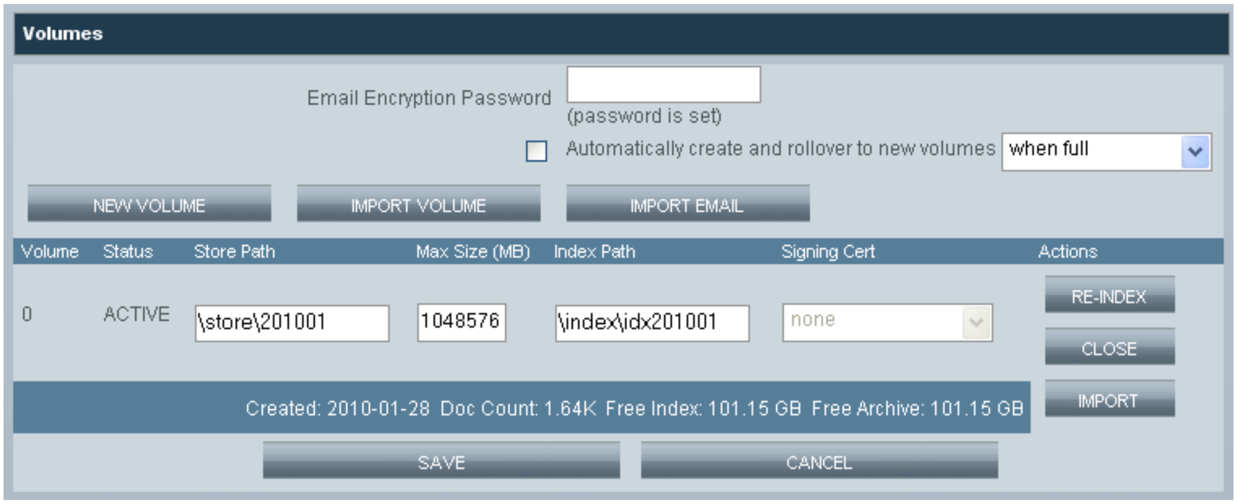

Volumes

Email Encryption Password
(password is set)

□ Automatically create and rollover to new volumes: when full

Volume	Status	Store Path	Max Size (MB)	Index Path	Signing Cert	Actions
0	ACTIVE	\store\201001	1048576	\index\idx201001	none	RE-INDEX CLOSE IMPORT

Created: 2010-01-28 Doc Count: 1.64K Free Index: 101.15 GB Free Archive: 101.15 GB

Figure 6 Volume Configuration

In addition to defining volumes manually, one can configure MailArchiva to create and rollover to new volumes based on certain conditions, such as when the volume is full or when a certain time period has elapsed. This feature is useful for two reasons: (1) it allows one to keep volumes to a defined size so that they can be backed up on DVD media (2) it allows one to store archive information on a monthly, quarterly, annual basis so that the information can be organized chronologically.

When using the auto volume feature, if a date in the format YYYYMM is appended to the store and index path, when a new volume is created, the store and index path of the last volume created will be updated with the current date. Assuming an existing volume has a store path of C:\store\201001 and an index path of C:\index\201001. When a new volume is automatically created in the following month, it should have a path of C:\store\201002 and C:\index20102. This feature enables you to record archived emails in chronological order.

In addition, if a numeric value of less than six digits is added to the end of the volume store and index path, when a new volume is created, the number appended to the previous volume's index and store path will be automatically incremented by one.

9.4. Console Access

Access to the MailArchiva web console and its features can be restricted to authorized users only. There are three types of authentication mechanisms supported:

- Basic – users are authenticated using credentials stored in a configuration file
- Active Directory – users are authenticated using their Windows credentials
- LDAP – users are authenticated against users in a directory server such as OpenLDAP

In all modes, if authentication is successful, a role is assigned to the authenticated user. The user's role determines what the user can and cannot do within the application. A user's role can also restrict which emails the user can see. For instance, user's assigned the User role can be given to permission to see their own emails only. Refer to Section 9.5 to learn more about MailArchiva's role mechanism.

9.4.1. Master Password

Before selecting an authentication method, it is prudent to specify the master password in the Logins section. The master password is essentially the "root" password for the system. It serves as a convenient back door in the event the system is unable to authenticate with Active Directory or LDAP servers. To login as the master user, use the username "admin" and your chosen master password. The master user has access to all privileges in the system.

Note: If you forget the master password, edit the file server.conf located in

/usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/conf (Linux)
C:\Program Files\MailArchiva\Server\Webapps\ROOT\WEB-INF\conf (Windows)

Remove the security.login.master.password and security.login.master.username fields.

9.4.2. Basic Authentication

In the Basic Authentication mode, the server authenticates users from credentials stored in an XML configuration file. The users.conf configuration file is located in mailarchiva\server \webapps\MailArchiva\WEB-INF\conf from the root of your MailArchiva installation directory.

You can either add users directly using the MailArchiva server console configuration screen or by editing the server.conf directly. The server.conf file, as illustrated in Figure 7, contains a list of users, each of which has a username, role and a password. The users listed in users.conf will login using their username and any of the domains in the Domains section appended to it.

Once a user is authenticated, the user will be assigned the specified role.

```
<User§s version="1.0">
  <User username="admin@company.local" role="administrator" password="123"/>
  <User username="user@company.local" role="user" password="abc"/>
  <User username="auditor@company.local" role="auditor" password="xyz"/>
</Users>
```

Figure 7 Users.conf

When Basic authentication mode is enabled, users are able to change their own passwords in the Preferences section.

9.4.3. Active Directory Authentication

In Active Directory (AD) authentication mode, the server uses NTLM v2 and LDAP protocols to authenticate users residing in Active Directory. The login procedure is a five step process:

1. MailArchiva authenticates with Active Directory user using a service computer account (you'll see later how this computer account is created)
2. MailArchiva searches for the login user in Active Directory using the login name
3. MailArchiva binds (authenticates) with the login user using the supplied password
4. MailArchiva assigns a role to the user based on the defined role assignments
5. MailArchiva extracts the user's email addresses from the mail LDAP attribute for use in search filtering

Note: If you are upgrading from earlier versions of MailArchiva, you should know that MailArchiva's authentication mechanism has changed from Kerberos to NTLM v2 authentication.

NTLM v2 authentication requires that the service account is a computer account, not a normal user account. Thus, to upgrade, you will need to create a computer account in AD, set a password on the computer password using the scripts provided, and change the service account to service$@business.local. Note the dollar ($) sign in the service acount UPN is used to denote a computer account (as opposed to a user account in Active Directory).

Field	Description	Example
DNS IP Address	IP address of your DNS server. It is used to resolve fully qualified domain names	192.168.0.1
Active Directory Adddress	The fully qualified domain name of Active Directory	active.business.local

Base DN	The distinguished name of the location in AD where MailArchiva should start searching for end-user entries.	dc=company,dc=com
Service Account Login	The FQDN of the service computer in AD.	service$@business.local
Service Account Password	The service computer password	
Mail Attribute	The mail attribute where the user's email addresses are obtained	ProxyAddresses
Email Value	The regular expression used to extract the email value from the mail attribute.	SMTP:(*.)
Bind Attribute	The attribute used to search for the user using login username in AD's LDAP. Leave this as is, unless you want users to be able to login using email address, or some other attribute.	SAMAccountName
NTLM Authentication	When NTLM authentication is enabled, MailArchiva will perform single-sign-on authentication with the users session.	Disabled

Figure 8 Active Directory Configuration

In order to authenticate with Active Directory, MailArchiva requires that a new computer account is created in Active Directory and that a password to the account is set. While it is possible to create a new Computer using Active Directory Users And Computers, there is currently no way from the GUI to set passwords on Computer accounts. For this purpose, a VBS script called ADSetupWizard.vbs is included with the server distributable. The script, when executed with Domain Administrator privileges, will automatically create a Computer in Active Directory and set a pasword on the Computer account. It will also output the AD configuration settings that are appropriate for your setup.

The procedure for configuring Active Directory authentication is as follows:

1. Included with the MailArchiva server distributable is VBS script called ADSetupWizard.vbs. This script can be found in the following location:

```
C:\Program Files\MailArchiva\Server\ADSetupWizard.vbs (Windows)
/usr/local/mailarchiva/server/ADSetupWizard.vbs (Linux)
```

2. Login to any computer nearby to (and including..) the MailArchiva server as a Domain Administrator. Copy the ADSetupWizard.vbs script from the above location to the local machine and run it.
3. Follow the Wizard instructions to create new “service” Computer account in Active Directory and a set a password on the service account.
4. When the Wizard completes, take note of the settings needed to define the AD settings in MailArchiva.
5. Open the MailArchiva Configuration console, select the Logins menu on the left. Choose Active Directory authentication and enter the settings outputted by the AD Wizard.
6. Next, click the New Role Assignment button to create a mapping between a role in MailArchiva and an Active Directory attribute.

Note: If the ADSetupWizard.vbs script generates the error "AccessDenied 80070005", it may be necessary to temporarily disable Windows UAC on the machine where the script is executed.

Note: If you experience problems running the ADSetupWizard.vbs script, as an alternative, you can create a computer manually in using Active Directory Users and Computers. Thereafter, run the SetComputerPassword.vbs script (located in the same location as ADSetupWizard script) to set the computer password.

Note: Microsoft requires that the user assigned the impersonation rights should not also have administrator rights assigned.

When assigning roles to Active Directory users, it is necessary to select a role, select an LDAP attribute and enter a match criterion.

Field	Description
Role	Role to be assigned
LDAP Attribute	LDAP attribute to use for the role assignment
Match Criterion	A value that is compared against a corresponding LDAP attribute in Active Directory for an authenticating user.

Table 5 Role Assignment Fields

To complete the attribute and match criterion fields, it is useful to understand how roles are assigned to users during console authentication.

A user in Active Directory has a set of LDAP attributes associated with it. These attributes are essentially properties about the user (e.g. account name, user group, etc.).

During console authentication, once the user has been identified, the value of the attribute selection is retrieved from Active Directory. This value is compared against the value entered in the match criterion field. If there is a match, the selected role is assigned to the user.

To assign a role to a Windows user, select "SAMAccountName" as the LDAP attribute and enter the user's name in the match criterion field.

To assign a role to all users within a user group, select "memberOf" in the attribute field and enter the distinguished name of the user group in Active Directory (e.g. "CN=Enterprise Admins, CN=Users, DC=company, DC=com").

Note: The match criterion field also accepts regular expressions for complex pattern matching requirements.

LDAP Attribute	Match Criterion Value
memberOf	Active Directory user group CN=Enterprise Admins,CN=Users,DC=company,DC=com
userPrincipalName	jdoe@company.com
SAMaccountName	Jdoe
distinguishedName	CN=John Doe,CN=Users,DC=company,DC=com

Table 6 Match Criterion Sample Values

In specifying the match criterion field, it is useful to lookup the LDAP attribute name and values associated with a user. This is done by clicking the Lookup button and entering a user's username (e.g. admin@company.com) and a password. A simple way to assign a role to an individual user is to copy one of the values of any of the attributes described in Table 5 and paste them into the match criterion field. There is likely to be an error in your configuration if the Lookup dialog does not return any LDAP attribute values.

Once all role assignments are configured, execute a Test Login to ensure that your Kereberos settings, LDAP settings and user roles have been configured correctly. If problems are encountered, enable server debugging as described in Section 27.1.1 to determine the source of the problem.

If you are unable to get AD authentication working in your environment, it is possible to authenticate with AD using password-based LDAP authentication instead. To do this, select LDAP authentication, enter the mail attribute to be “proxyAddresses” and “SAMAccountName” to be the bind attribute. You will also need to clear out the default login name suffix in the Logins section. See section 9.4.4 for more information.

Multi-Domain Authentication Tip: if your organization has multiple domains, MailArchiva must be configured to connect to AD’s Global Catalog Server running on port 3268. To do this, change your Active Directory server FQDN to the equivalent of company.com:3268. Set the base DN to be empty.

9.4.3.1. NTLM Authentication

The MailArchiva console supports single-sign-on authentication with Windows using NTLM authentication. With NTLM authentication enabled, there is no need for a user to manually log in to MailArchiva. Users will be logged to the MailArchiva console automatically using their Windows credentials. Before enabling NTLM authentication, be sure that standard AD authentication (without NTLM authentication) is working correctly.

There are few necessary measures to ensure the correct functioning of the NTLM authentication feature:

1. NTLM authentication must be enabled in AD authentication settings
2. There must be a matching role assignment in MailArchiva for each Windows domain user where access to MailArchiva is intended
3. From the connecting user’s client computer, the MailArchiva server must be addressable by fully qualified domain name (FQDN)

4. The address containing the FQDN of the MailArchiva server must be added as a trusted site in Internet Explorer's Local Intranet security zone. To do this, click Tools->Internet Options->Security->Local Intranet->Sites->Advanced. Type in the address of the MailArchiva server (e.g. http://mailarchiva.smallbusiness.local). Do not use the IP address of the server – it will not work! For test purposes, the MailArchiva server's FQDN can be added to hosts file of the client computer.

On condition that all four of the above conditions are met, when entering the MailArchiva console URL, users will be logged in automatically.

> ❗ **When NTLM authentication is enabled, to explicitly login as the master user or another user, it is necessary to specify the URL equivalent of http://mailarchiva.smallbusiness.local/signonform.do**

The easiest way is to implement NTLM authentication on every workstation in the company is by adding the fully qualified domain name of the MailArchiva server to the following registry key in Microsoft's Group Policy Editor:

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet
Settings\ZoneMap\Domains
```

9.4.4. LDAP Authentication

When LDAP authentication is enabled, MailArchiva authenticates to a directory service such as OpenLDAP using pure password-based credentials.

The following process occurs during LDAP console login:

- MailArchiva authenticates with the directory using a service account DN and a password
- MailArchiva searches for the user, starting from the Base DN, by matching the supplied username with the Bind Attribute (normally, UID)
- MailArchiva retrieves the DN of the located user
- MailArchiva uses the retrieved user DN and user password to login into the directory
- Once logged in, MailArchiva looks for a matching role and retrieves the user's email address from the Email Attribute field (usually, email or mail).

Since directory structures tend be to unique across different organizations, care must be taken to ensure that the base DN, service account login DN, bind attribute and email attribute is correct for the target directory. For example, some companies use "mail" as the location where user email addresses are stored, while others use "email".

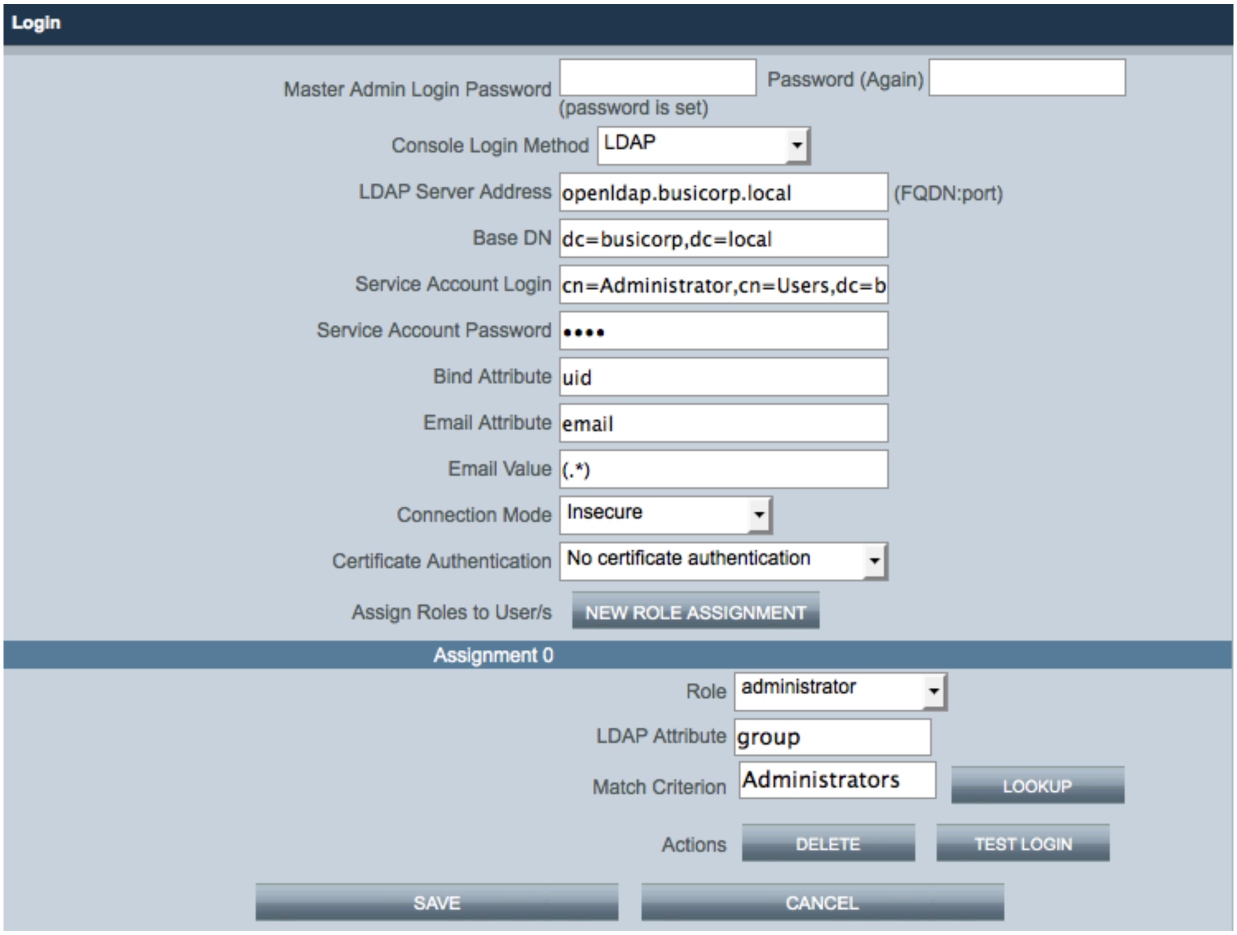

Login

Master Admin Login Password	(password is set)	Password (Again)	
Console Login Method	LDAP		
LDAP Server Address	openldap.busicorp.local	(FQDN:port)	
Base DN	dc=busicorp,dc=local		
Service Account Login	cn=Administrator,cn=Users,dc=b		
Service Account Password	••••		
Bind Attribute	uid		
Email Attribute	email		
Email Value	(.*)		
Connection Mode	Insecure		
Certificate Authentication	No certificate authentication		
Assign Roles to User/s	NEW ROLE ASSIGNMENT		

Assignment 0

Role	administrator
LDAP Attribute	group
Match Criterion	Administrators

Figure 9 LDAP Login Configuration

To determine the structure of a directory, it may be useful to connect to it using a the Linux command line utility ldapsearch or one of the many LDAP browsers available.

Once the correct LDAP settings have been entered, it is necessary to create one or more nrole assignments for purpose of assigning MailArchiva roles to the users residing in the directory. Refer to section 9.4.3 to develop an understanding of the role mapping mechanism.

Field	Description	Example
LDAP Server Address	Fully qualified domain name of LDAP server	openldap.company.com:389
Base DN	The distinguished name of the location in AD where MailArchiva should start searching for end-user entries.	dc=company,dc=com
Service Account Login DN	The distinguished name of an admin user in LDAP	cn=Administrator,cn=Users, dc=company, dc=com
Service Account Password	The service account password	
Mail Attribute	The mail attribute where the user's email addresses are obtained	mail
Mail Value	The regular expression used to extract the user's email address from the mail attribute	(*.)
Bind Attribute	The field in LDAP that contains the username or login name of the user.	uid

Figure 10 LDAP Configuration

9.4.1. Google Apps Authentication

With Google Apps login enabled, users are authenticated using their Google credentials. The following process occurs during Google console login:

- MailArchiva authenticates using Google's ClientLogin API with the supplied login credentials
- The user's group information is retrieved from Google
- A matching role is assigned to the user based on the defined role mappings

Field	Description	Example
Admin Email	Email address of Google Administrator	admin@company.com
Admin Password	Google Administrator password	passwords

Figure 11 Google Authentication Configuration

In Configuration->Login, select 'Google' as the Console Login Method. Enter the Google administrator email address and password in the fields provided. Click the New Role Assignment button to create one or more role mappings. Refer to section 9.4.3 to improve your understanding of the role mapping mechanism.

It is imperative to create at least one role mapping. The purpose of the mapping is to assign a user or group of users a role. For instance, to map all Google Administrators to the Administrator role, select 'Administrator' as the Role and in the Selection field choose 'adminsitrators'. Once all role mappings have been defined, click Save and perform a Test Login to establish whether the login settings are configured correctly.

Field	Description	Matches
Administrator	Google Administrators	[leave empty]
Normal Users	Google Users	[leave empty]
Specific User	Specific Google User	[username of specific user]
Specific Group	Specific Google Group	[name of specific group]

Figure 12 Role Mapping Selections

9.5. Roles

During the console login process, the user is assigned a security role. The security role determines what the user can do and which emails the user can see. There are two main aspects to role definition:

- Permissions – what the user can do (e.g. delete email)
- View filters – which emails the user can see (e.g. only emails within a domain)

There are three built-in roles in the system: administrator, auditor and user. The default permissions and view filters associated with these roles are described in Table 7 and Table 8, respectively.

Role	Allow Delete	Allow View	Allow Print	Allow Export	Allow Save	Allow Send	Allow Configure
User	No	Yes	Yes	Yes	Yes	Yes	No
Audit	No	Yes	Yes	Yes	Yes	Yes	Yes
Admin	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Master	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 7 Built-In Role Permissions

Role	View Filter
User	Can only view own emails (all addresses must match user's email address)
Audit	Can view any email
Admin	Can only view own emails (all addresses must match user's email address)
Master	Can view any email

Table 8 Built-In Role Email Filters

9.5.1. User Role

The User Role is used for the purposes of enabling employees to access to their own emails in the archive. The macro "%email%" in the view filter restricts users who are assigned the role to viewing emails that concern themselves only. In addition, users assigned the User Role are, by default, unable to alter the configuration of the system.

When defining a view filter, the macro %email% will be replaced with the email address of the logged-in user. Thus, by selecting “any address” and entering “%email%” as a value, you will effectively limit the user to seeing their own emails only.

9.5.2. Audit Role

The main different between the built-in Audit Role and the User Role is that its view filter is empty, meaning auditors are able to access the emails of every person in the company.

9.5.3. Administrator Role

Administrators are capable of modifying the configuration of MailArchiva, excluding role definition and login configuration. For security reasons, they are not permitted to view all emails of all users in the archive.

9.5.4. Master Role

The master role is assigned when a user is logged in as the “admin” user. The master user is all powerful and has the ability to access all functions of the system, including view all emails in the archive.

9.5.5. Built-In Roles

If the built-in roles are not suitable, you can define one or more custom roles. To define a custom role:

- Click on the “Add Role” button in the Roles section of the server console configuration screen
- Enter an appropriate name for the role
- Select the permissions associated with the role
- Add a view filter clause to limit which emails users assigned the role can view

It is possible to define a role that would allow all those who are assigned to it, the permission to view emails of all users belonging to a specific active directory group. This capability supports the scenario where the director of a department would like permission to view all emails sent and received by members of the department. To do this, simply select the User Group field in the role's view filter and enter the name/DN of the corresponding group in Active Directory.

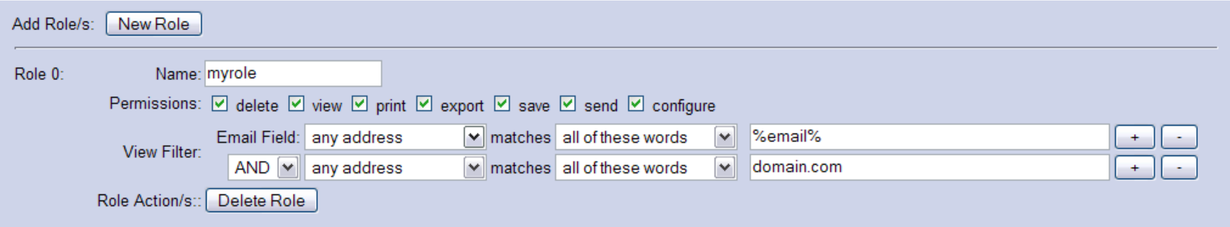

Role 0: Name: myrole

Permissions: ☑ delete ☑ view ☑ print ☑ export ☑ save ☑ send ☑ configure

View Filter: Email Field: any address matches all of these words %email%

AND any address matches all of these words domain.com

Role Action/s::

Figure 13 Custom Role Definition

9.6. Enabling HTTPS Access

To protect against the sniffing of login passwords and confidential emails on the network, it is advised (though optional) to enable HTTPS (HTTPS / TLS). Follow the steps below to achieve this:

Step 1. Generate & Install Certificates

Follow the instructions outlined in Section 9.7 to generate a server certificate and install trusted CA certificates.

Step 2. Obtain the KeyStore Password

All certificates and keys are stored in an encrypted key store. To obtain the password, for the key store, run the script getkeystoresecret to obtain the password.

```
C:\Program Files\MailArchiva\server\getkeystoresecret.bat (Windows)

/usr/local/mailarchiva/server/getkeystoresecret (Linux)
```

When running the script, you will be asked to enter the salt and encryption password. Accept the default value for the salt and enter your volume encryption password in the password prompt.

Step 3. Add Entry to Tomcat server.xml

MailArchiva runs under the Apache Tomcat application server. Edit the following file:

```
C:\Program Files\MailArchiva\server\conf\server.xml (Windows)

/usr/local/mailarchiva/server/conf/server.xml (Linux)
```

Uncomment and modify the following entry in the Tomcat server.xml file.

```
<Connector port=”443″
maxThreads=”150″ minSpareThreads=”25″ maxSpareThreads=”75″
enableLookups=”true” disableUploadTimeout=”true”
acceptCount=”100″ debug=”0″ scheme=”https” secure=”true”
clientAuth=”false” sslProtocol=”TLS”
keystoreFile=”/usr/local/mailarchiva/server/webapps/ROOT/conf/mailarchivacerts”
keystorePass=”[output of getkeystoresecret]” />
```

Set the keystoreFile above to the location of your key store.

```
C:\Program Files\MailArchiva\server\webapps\ROOT\conf\mailarchivacerts (Windows)

/usr/local/mailarchiva/server/webapps/ROOT/conf/mailarchivacerts (Linux)
```

Set the keystorePass to the output of the getkeystoresecret utility. After the server restarts, the console should be accessible using https://localhost.

Alternative method: Follow the normal procedures for enabling TLS in Tomcat 6, namely: http://tomcat.apache.org/tomcat-6.0-doc/ssl-howto.html

You should wish disable plain-text HTTP access and enforce the use of TLS, edit the MailArchiva applications, web.xml file located as follows:

```
C:\Program Files\MailArchiva\server\webapps\ROOT\web.xml (Windows)

/usr/local/mailarchiva/server/conf/webapps/ROOT/conf/web.xml (Linux)
```

Towards the bottom of the web.xml file, immediately above the closing </web-app> tag, add the following security-contraint entry:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>Automatic SLL Forwarding</web-resource-name>
    <url-pattern>/*</url-pattern>
  </web-resource-collection>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

9.7. Digital Certificates

This section describes how to generate and install server and trusted certificates. Certificates are used throughout MailArchiva for the purposes of secure communications and digital signing. Their use is entirely optional in MailArchiva.

Step 1. Generate a Certificate Signing Request (CSR)

1. In the Certificates tab of Configuration, click “New Server Cert”.

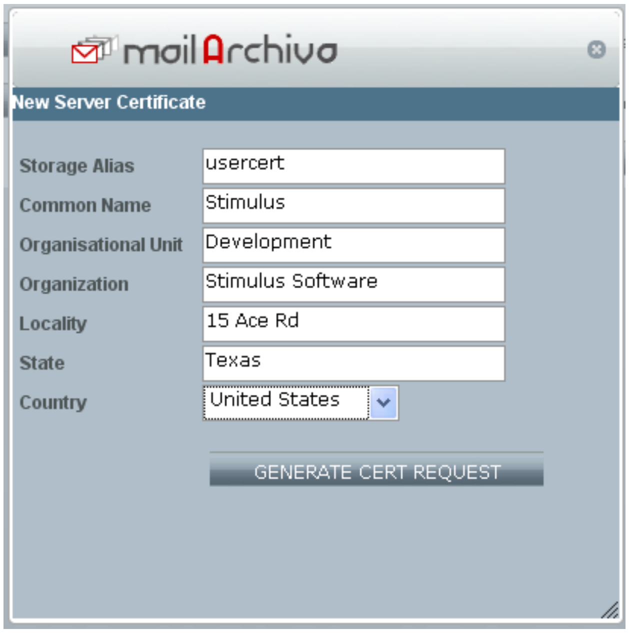

New Server Certificate

Storage Alias	usercert
Common Name	Stimulus
Organisational Unit	Development
Organization	Stimulus Software
Locality	15 Ace Rd
State	Texas
Country	United States

Figure 14 New Certificate Request

2. Choose and enter a storage alias. This alias is an arbitrary name (chosen by you) that you will later use to refer to the certificate
3. Complete all the relevant demographic information
4. Click the “Generate New Cert Request” to generate the CSR
5. Select and copy the certificate signing request to the clipboard.

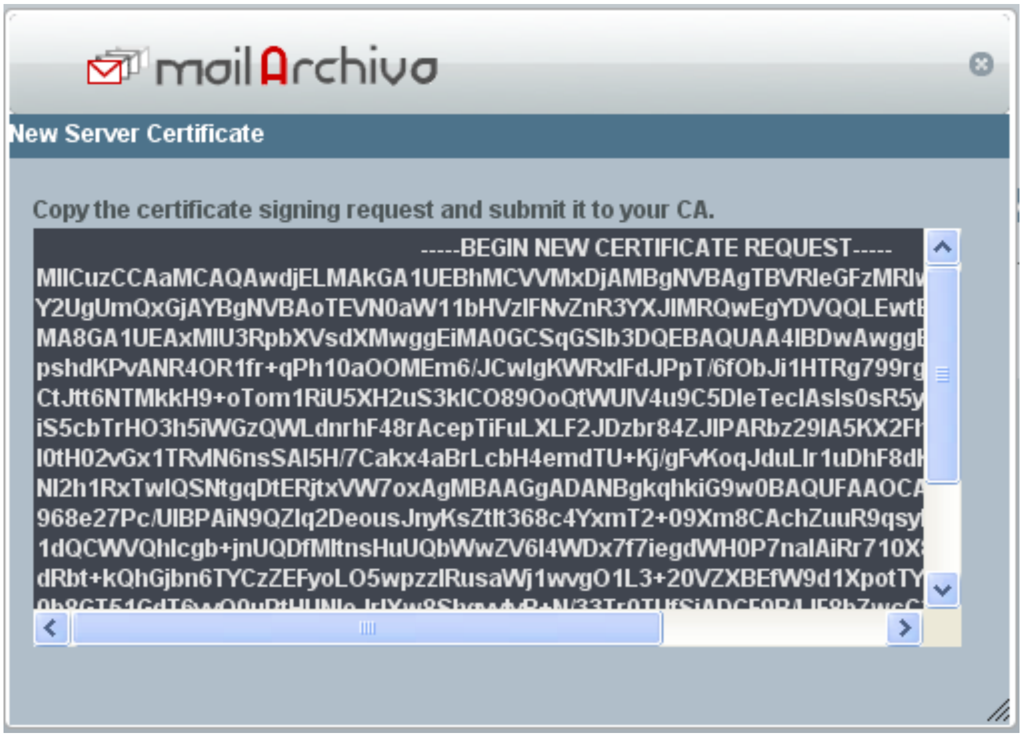

New Server Certificate

Copy the certificate signing request and submit it to your CA.

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIICuzCCAaMCAQAwdjELMAkGA1UEBhMCVVMxDjAMBgNVBAgTBVRleGFzMRIv
Y2UgUmQxGjAYBgNVBAoTEVN0aW11bHVzIFNvZnR3YXJlMRQwEgYDVQQLEwtE
MA8GA1UEAxMIU3RpbXVsdXMwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggE
pshdKPvANR4OR1fr+qPh10aOOMEm6/JCwlgKWRxlFdJPpT/6fObJi1HTRg799rg
CtJtt6NTMkkH9+oTom1RiU5XH2uS3kICO89OoQtWUlV4u9C5DleTeclAsls0sR5y
iS5cbTrHO3h5iWGzQWLdnrhF48rAcepTiFuLXLF2JDzbr84ZJIPARbz29IA5KX2Fh
I0tH02vGx1TRvN6nsSAI5H/7Cakx4aBrLcbH4emdTU+Kj/gFvKoqJduLIr1uDhF8dK
NI2h1RxTwIQSNtgqDtERjtxVW7oxAgMBAAGgADANBgkqhkiG9w0BAQUFAAOCA
968e27Pc/UIBPAiN9QZIq2DeousJnyKsZtIt368c4YxmT2+09Xm8CAchZuuR9qsy
1dQCWVQhlcgb+jnUQDfMltnsHuUQbWwZV6I4WDx7f7iegdWH0P7naIAiRr710X
dRbt+kQhGjbn6TYCzZEFyoLO5wpzzlRusaWj1wvgO1L3+20VZXBEfW9d1XpotTY
```

Figure 15 Copy CSR to Clipboard

6. Close the New Server Cert dialog. You should see a place holder for your signing certificate listed in the Server Certificates list in the Certificates tab.
7. **Important!** Click Save to save the changes to your certificate store

Step 2. Obtain Certificates from the Certificate Authority (CA)

Purchase certificate from a Certifying Authority such as Verisign. While obtaining the certificate you will be required to paste in the Certificate Signing Request (CSR) generated earlier.

> **❗ For test purposes, you can obtain a free 15 day trial SSL certificate from Verisign (http://www.verisign.com/)**

In the case of Verisign, your certificate will be emailed to you. Copy the server certificate to a file and save it to a file called user.cert on your Desktop.
Similarly, copy the intermediate CA certificate to a text file called intermediate.cert.
Finally, copy the CA root certificate to a text file called root.cert.

When using Verisign, the links to download the intermediate and CA certificates are included in the email containing your server certificate.

Step 3. Import the Certificates

1. Click "Import CA Cert", select the root.cert file, enter "cacert" as the storage alias and click Import.
2. Click "Import CA Cert", select the intermediate.cert file, enter "intermediatecert" as the storage storage and click Import
3. Click "Import Server Cert", select the user.cert file. Selected the same alias as used when you generated the CSR above.

Note: The order in which you import the certificates is important. You must first import the root CA certificate, then the intermediate CA certificate and finally the server signing certificate.

After following the above, the signing certificate and CA certificates should be visible in the Certificates list.

9.8. Archive Rules

In some circumstances, it may not be desirable to archive all incoming emails but only a select few. In such cases, it is necessary to define one or more archive rules. Archive rules are used to specify the conditions that determine whether or not an email should be archived at the time it is received by the system.

Each rule consists of one or more clauses. Each clause consists of an email field, an operator and a value. When processing a clause, the value of the selected email field is retrieved from the email and compared against the value specified in the clause. If they match, the action, either “ignore”, “archive” or “do not archive”, is applied. For example, to ensure all emails addressed to john@company.com are archived, you would simply select the field “to”, select the “contains” operator and enter “john@company.com”.

The ordering in which archiving rules are processed is significant. The rules are processed sequentially from top to bottom. A rule that appears before another will always be processed first. If none of the rules match, the default rule applies. In the definition of the default rule, there is the option to archive incoming, outgoing and/or internal emails.

9.9. Retention Policy

The retention policy settings are used to either delete or retain emails residing in the archive. Each retention rule specifies the conditions under which a set of emails is either retained or deleted, and the time period that applies.

A retention rule consists of an action, a time period and a set of conditions. The action can be either “hold” or “delete”. The time period is specified in number of days since the email was sent. The match conditions are expressed by adding one or more filter clauses. They define the set of emails to which the rule applies.

> ❗ **The definition of filter clauses can have catastrophic implications for the integrity of the archive if they are defined incorrectly. Thus, it is strongly advised to perform equivalent searches in the Search screen in order to verify and test that the defined clauses will produce the desired results.**

Like archive rules, the ordering is significant. A preceding rule will always be processed before a subsequent one. For example, a matching hold action will apply irrespective of the matching delete actions below it.

If none of the defined retention rules match, the default rule applies. In the default rule, one can specify the number days after which time the emails will be deleted from the store.

> ❗ **It is possible to define different retention periods for different departments, by selecting the User Group email field and entering the DN of the department's memberOf group in Active Directory.**

9.1. Configuring Outgoing SMTP Settings

The SMTP client settings can be found in Configuration->General->Outgoing SMTP Settings. These settings are used by MailArchiva for the purposes of sending messages to users and system administrators. In the first instance, when a user sends messages back to the server in the search interface, the messages are sent via the SMTP protocol. In the second instance, if status reporting and alerts are enabled, MailArchiva will send these as normal email messages again via the customer's mail server.

When users elect to send messages back to the server, MailArchiva tries to preserve all message header information as far as possible. Thus, header detail such as the from, to, bcc, cc, sent date, message ID, and so forth is preserved. The consequence of this is that there must be high degree of trust between the mail server and MailArchiva. The mail server must allow MailArchiva to send emails on behalf of all users in the system.

To achieve this, in most cases, it is imperative to enter a valid username and password in the SMTP client settings. Without proper authentication, the mail server is highly unlikely to deliver the messages sent by MailArchiva, especially when header information is preserved.

For Microsoft Exchange in particular, it is necessary to either specify the credentials for a Domain Administrator or assign the Send As permission to the authenticating user.

❗ in some instances, Exchange wont deliver a message if exactly the same message is already present in a user's mailbox.

9.2. Archive Settings

Setting	Default Value	Description
Duplication mode	Deduplication on	Where or not to deduplicate messages
Combine RCPTTO info for identical messages	Yes	When the same message is received multiple times via SMTP / milter protocol, combine all RCPTO info fom each session session into that message.
Store attachments larger than..separately	32768 bytes	Attachments will be de-duplicated when larger than specified value.
Do not archive messages larger than .. megabytes	100	Archive messages less than the amount specified (only applies to SMTP archiving)
Maximum no. Archiving threads	10	No. Simultaneous threads performing archiving. More threads take advantage of the parallelism of multicore processors, but consume more memory resources.
Perform disk space checks	Yes	MailArchiva will track disk space usage in order to roll over volumes if space is exceeded. This activity has I/O overhead as MailArchiva has to step through every email in the archive to determine whether or not the disk space limit has been reached. Disable if volume size limits

		can be safely ignored. This will result in reduced I/O overhead and increased performance.
Recover emails		Click this button to force MailArchiva to process all emails in the no archive queue.
Retain malformed messages	No	If MailArchiva receives a message that is significantly different from the RFC2822 spec, it will still be retained. This results in lower I/O performance since the message has to be cached to disk temporarily.
Noarchive location		The location where messages in the no archive queue are stored
Quarantine location		The location where messages in the quarantine are stored.
WORM Drive Support	No	Whether or not a WORM drive is being used. In this case,MailArchiva will be sure to write data only once to the store path. The volumeinfo file used to maintain volume status will be stored in the volume index path and not the store path.

9.3. Index Settings

Setting	Default Value	Description
Index message body	Yes	Whether or not to index the body of messages
Index attachments	Yes	Whether or not to index the contents of attachments
Auto detect email index	No	When enabled, MailArchiva will attempt to detect the language used in an email and index it using appropriate analyzers

Default indexing charset	UTF-8	When indexing an email, if the character set is not specified, MailArchiva will default to indexing using the selected character set
Default zip file name charset	UTF-8	Select a default character set for Zip file name usage
Mac chars to index per field	5000	Maximum no. of characters in the body of a message or attachment content permissible to index
Indexing Method	Precise	Whether or not to apply porter stemming to search terms during indexing. With stemming enabled, MailArchiva will index "run" and not "running". Use stemming if greater usability is desired at the expense of accuracy. Note: search accuracy is important if retention rules are specified, as these rules rely on MailArchiva's search engine to select the emails to retain or delete. After changing this setting, a re-index is necessary.

9.4. Search Settings

Setting	Default Value	Description
Max Results	10,000	The default maximum no. of results returned in a search query. Higher values will require more memory.
Initial sort order	Descending	Whether or not the search results should be sorted in descending/ascending order or left unsorted. If searches are too slow, consider setting this value to unsorted.

Initial sort field	Archive Date	The initial field used for sorting purposes.
Default date type	Sent Date	Whether or not sent date, receive date or archive date should be used for date selection during searching.
Allow searching using archive date only	Yes	The use of archive date is preferable from a performance standpoint as it allows MailArchiva to ignore searching in volumes that don't fall between specified date ranges.
Show message preview in search results	Yes	Show search highlighting in search results.
Export	10,000	The maximum no. of emails that can be exported at a time
Delete	10,000	The maximum no. of emails that can be deleted at a time
Send	10,000	The maximum no. of emails that can be sent at a time
Export method	Original Message	Whether or not exported messages should contain the journal envelope or not (applicable for Exchange journaling only)

9.5. General Settings

Setting	Default Value	Description
System Locale	OS Locale	The default language settings for system level messages
Override client locale in user's browser	Yes	Do not obtain language settings from user's browser. Force to a specified client locale
Client locale		The chosen client locale to use when overridden

		as above
Default System Theme	Cool Breeze	The default theme used for users logging into the system
Outgoing SMTP Connection Settings		In addition to including its own inbuilt SMTP server for receiving messages, MailArchiva has an inbuilt SMTP client that it uses to replay emails back to the mail server and send status reports. Specify these settings if users will have the option to restore emails from the archive or if the use of system status reports is desirable.

10. EXCHANGE CONNECTION

This section outlines the preparation steps that are necessary for direct import and stubbing of messages in Microsoft Exchange.

10.1.1. Microsoft Exchange 2003 Preparation Steps

MailArchiva integrates with Microsoft Exchange 2003 using its WebDav API. It is thus necessary to ensure that WebDav is installed and configured as described in the below steps.

Step 1. Install and Enable WebDav IIS Component

From Add *Remove* Programs in the Control Panel, click Server -> Internet Information Services -> World Wide Web Service -> WebDAV Publishing, and ensure that WebDav Publishing is installed.

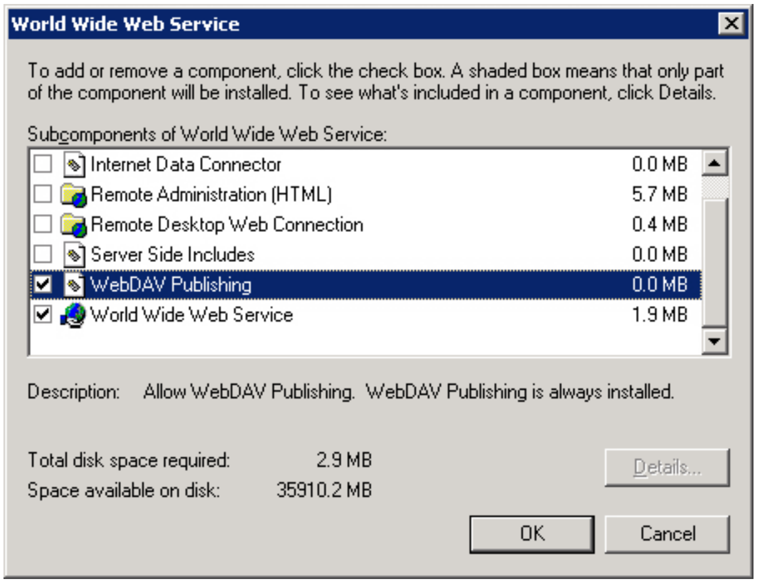

World Wide Web Service

To add or remove a component, click the check box. A shaded box means that only part of the component will be installed. To see what's included in a component, click Details.

Subcomponents of World Wide Web Service:

☐	Internet Data Connector	0.0 MB
☐	Remote Administration (HTML)	5.7 MB
☐	Remote Desktop Web Connection	0.4 MB
☐	Server Side Includes	0.0 MB
☑	WebDAV Publishing	0.0 MB
☑	World Wide Web Service	1.9 MB

Description: Allow WebDAV Publishing. WebDAV Publishing is always installed.

Total disk space required: 2.9 MB
Space available on disk: 35910.2 MB

Once the Web Dav component is installed, ensure that it is allowed in the IIS Manager.

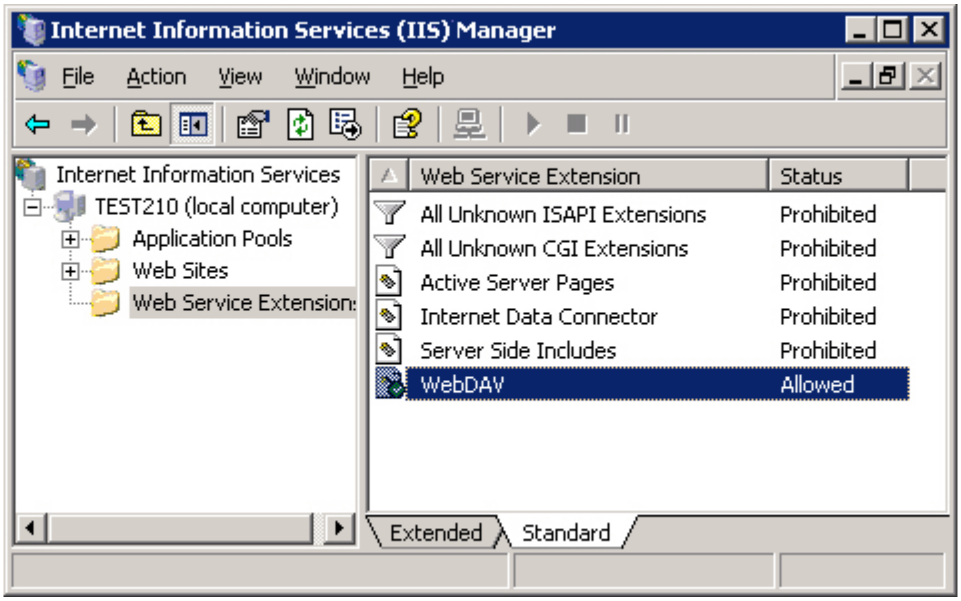

Web Service Extension	Status
All Unknown ISAPI Extensions	Prohibited
All Unknown CGI Extensions	Prohibited
Active Server Pages	Prohibited
Internet Data Connector	Prohibited
Server Side Includes	Prohibited
WebDAV	Allowed

Step 2. Grant A User Full Mailbox Rights

MailArchiva requires an account with sufficient privileges to read from all mailboxes in Microsoft Exchange.

1. Create a new mailbox called “journal” in Microsoft Exchange or use an existing one that does not already have any Admin privileges assigned.

2. On the server running Microsoft Exchange 2003, Start 'Exchange System Manager'.

3. Open the server object within the appropriate Administrative Group. Expand the server object. Expand the appropriate 'Storage Group'. Locate the required mailbox store, right-click and choose the 'Properties' option.

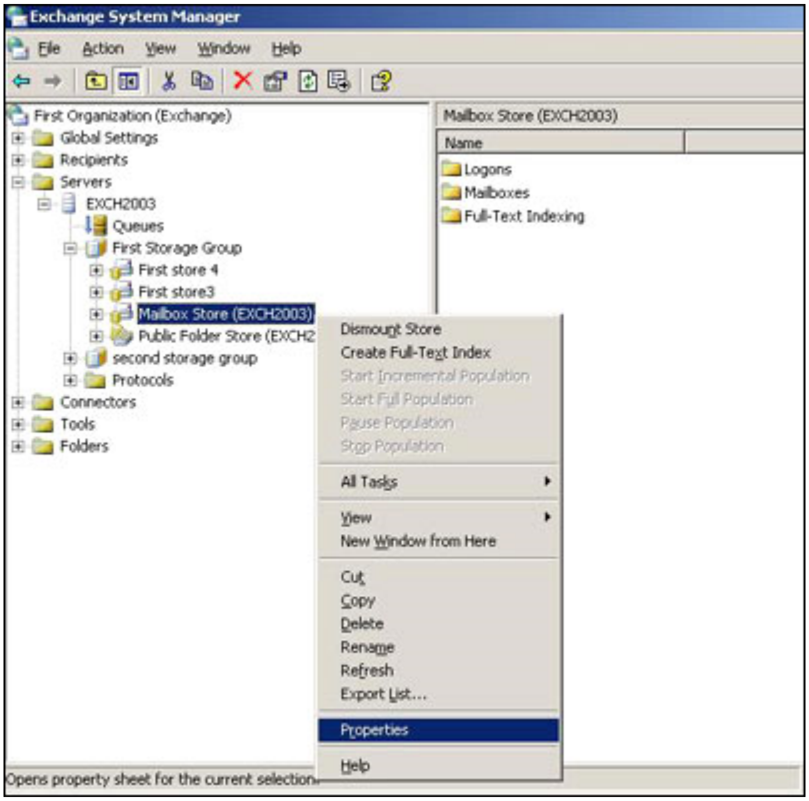

3. On the 'Properties' window click the 'Security' tab.
4. Click 'Add' and then click on the Active Directory service account click 'OK'.
5. Ensure that the ‘journal’ account is selected in the 'Name' box.
6. On the 'Permissions' list, click 'Allow' next to 'Full Control' and then click 'OK'.

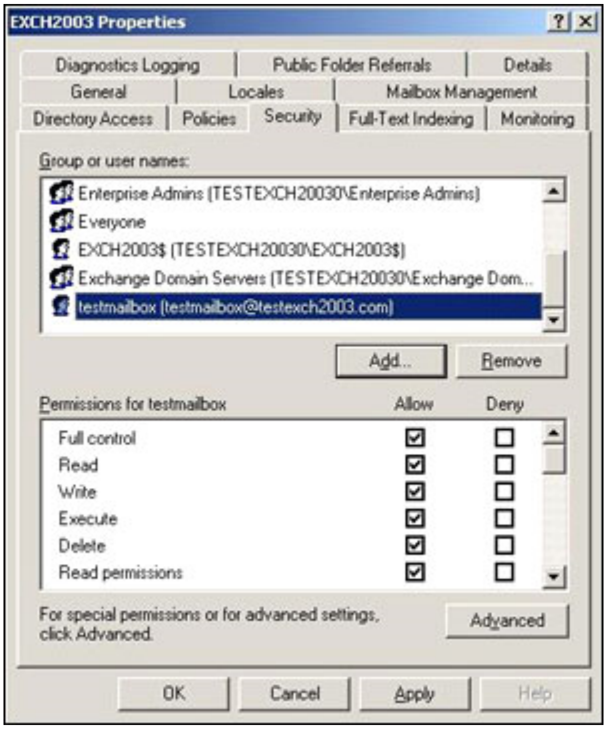

Group or user names:

- Enterprise Admins (TESTEXCH20030\Enterprise Admins)
- Everyone
- EXCH2003$ (TESTEXCH20030\EXCH2003$)
- Exchange Domain Servers (TESTEXCH20030\Exchange Dom...
- testmailbox (testmailbox@testexch2003.com)

Permissions for testmailbox	Allow	Deny
Full control	☑	□
Read	☑	□
Write	☑	□
Execute	☑	□
Delete	☑	□
Read permissions	☑	□

For special permissions or for advanced settings, click Advanced.

4. Click 'Ok' to finish

10.1.2. Microsoft Exchange 2007 Preparation Steps

Integration with Microsoft Exchange 2007 occurs by way of Exchange’s inbuilt web services API. This API is enabled and accessible by default. However, to perform a successful import, one needs to use an account with sufficient privileges to access all mailboxes.

To do this:

1. Login to Exchange 2007 server as Administrator

2. Either create a new mailbox called ‘journal’ in Microsoft Exchange or use an existing one that does not already have any Admin privileges assigned.

2. Run the Exchange Management Shell

3. Type the following into the Exchange Management shell:

```
Get-ClientAccessServer | Add-AdPermission -User journal -ExtendedRights ms-Exch-EPI-
Impersonation

Get-MailboxDatabase | Add-AdPermission -User journal -ExtendedRights ms-Exch-EPI-
May-Impersonate
```

Note: If necessary, substitute ‘journal’ to be the chosen username of the impersonation account.

Note: It is also possible to use the computer account created in Active Directory for the purposes of login (as described in . In this case, the mailbox name should resemble the following: “service$@smallbusiness.local”.

10.1.3. Microsoft Exchange 2010 Preparation Steps

MailArchiva similarly integrates with Microsoft Exchange 2010 via the web services API.

1. Login to Exchange 2010 server as Administrator
2. Create a new mailbox called ‘journal’ in Microsoft Exchange or use an existing one that does not already have any Admin privileges assigned.
2. Open the Exchange Management Shell
3. Run the New-ManagementRoleAssignment cmdlet

To enable the ‘journal’ account to impersonate all other users in an organization.

```
New-ManagementRoleAssignment -Name:MailArchivaImpersonation -
Role:ApplicationImpersonation -User:journal
```

11. EMAIL MIGRATION

MailArchiva archives emails from the point at which the server is deployed. If there is a need to archive older (pre-deployment) emails, they need to be imported separately into the system. The MailArchiva server has inbuilt facilities for importing messages from .EML (RFC2822) format, MBOX format, PST format, and direct import from MS Exchange.

11.1. EML / MBOX / PST Import

To import messages from .EML (RFC2822) , MBOX or PST format, ensure a volume is available for archiving and then click the Import Messages button in the Volumes section of the MailArchiva configuration.

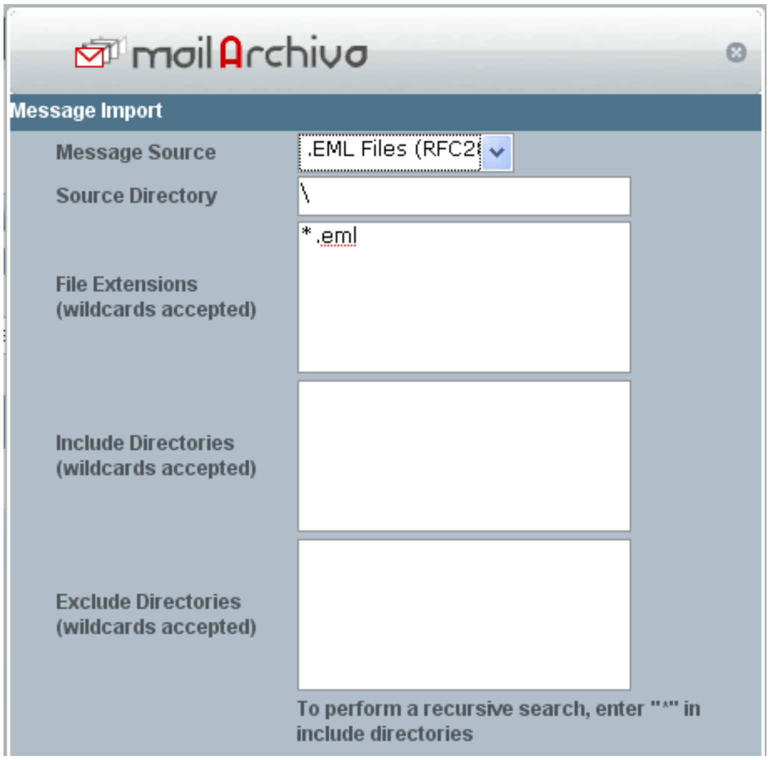

Message Import

Message Source: .EML Files (RFC2

Source Directory: \

File Extensions (wildcards accepted): *.eml

Include Directories (wildcards accepted):

Exclude Directories (wildcards accepted):

To perform a recursive search, enter "*" in include directories

Enter a local source directory and press the Import button.

! If you wish MailArchiva to import messages recursively, enter “*” in the include directories field.

11.2. Exchange Direct Import

MailArchiva is capable of performing a direct import of messages from Exchange 2003, 2007 and 2010. Before an import is attempted, it is necessary to define an impersonation account that has sufficient rights to access all mailboxes.

Follow the steps outlined in Section 10 to setup an impersonation account and prepare your Exchange server for the import process. Thereafter, be sure that a volume is configured in the MailArchiva console configuration.

Click the Import Messages button in the Volumes section. Enter the fully qualified domain name of the Exchange server. Enter the account details for ‘journal’ in the impersonate username and password fields. Thereafter, press the import button.

! If the direct Exchange import methods described in this section do not work, visit the MailArchiva Knowledge Base for further information on how to use alternative methods such as the M-Drive approach.

11.1. Google Apps Import

The Google Import process consists of:

1. Initiating mailbox dump requests for all users in an organization

2. Waiting a day or more for those requests to be completed
3. Importing all completed requests

Due to the batch nature of Google's dump process, the import process can be cumbersome. Follow the below steps to complete the import:

1. In Configuration->Volumes click the Import Email button
2. Select Google as the Message Source
3. In the Import Dialog enter the Google Apps admin email address and password
4. Click the Import Button
5. In the Progress Window ensure that login to Google is successful
6. Close the Window and wait a day or more for requests to be completed
7. Monitor the process of the Import, in System Status->Processes

12. OUTLOOK INTEGRATION (OPTIONAL)

For convenience reasons and to preserve the familiar Outlook experience, MailArchiva's search capabilities can be made accessible from within Outlook. MailArchiva's search interface has an inbuilt Outlook theme that is automatically activated when the console is accessed from within Outlook. Follow the below steps to integrate MailArchiva with Outlook:

1. Ensure that both MailArchiva and the user's Internet Explorer browser is configured for NTLM authentication as described in section 9.4.3.1.
2. Login to the user's Outlook client.
3. Right click the user's mailbox in the tree view on the left
4. Create a new folder called "MailArchiva" and click OK
5. Right click the MailArchiva folder
6. Select the Home Page tab
7. Check show home page by default for this folder

8. Enter the address: http://mailarchiva.smallbusiness.local:8090/outlook.do (replace the FQDN with the FQDN of your MailArchiva server. Do not use the IP Address of the server as NTLM authentication will not work!
9. Click on the MailArchiva folder and the MailArchiva Outlook search interface should appear. If it does not, check the NTLM authentication is working correctly when connecting to the MailArchiva console outside of Outlook. If not, revisit the steps in section 9.4.3.1.
10.

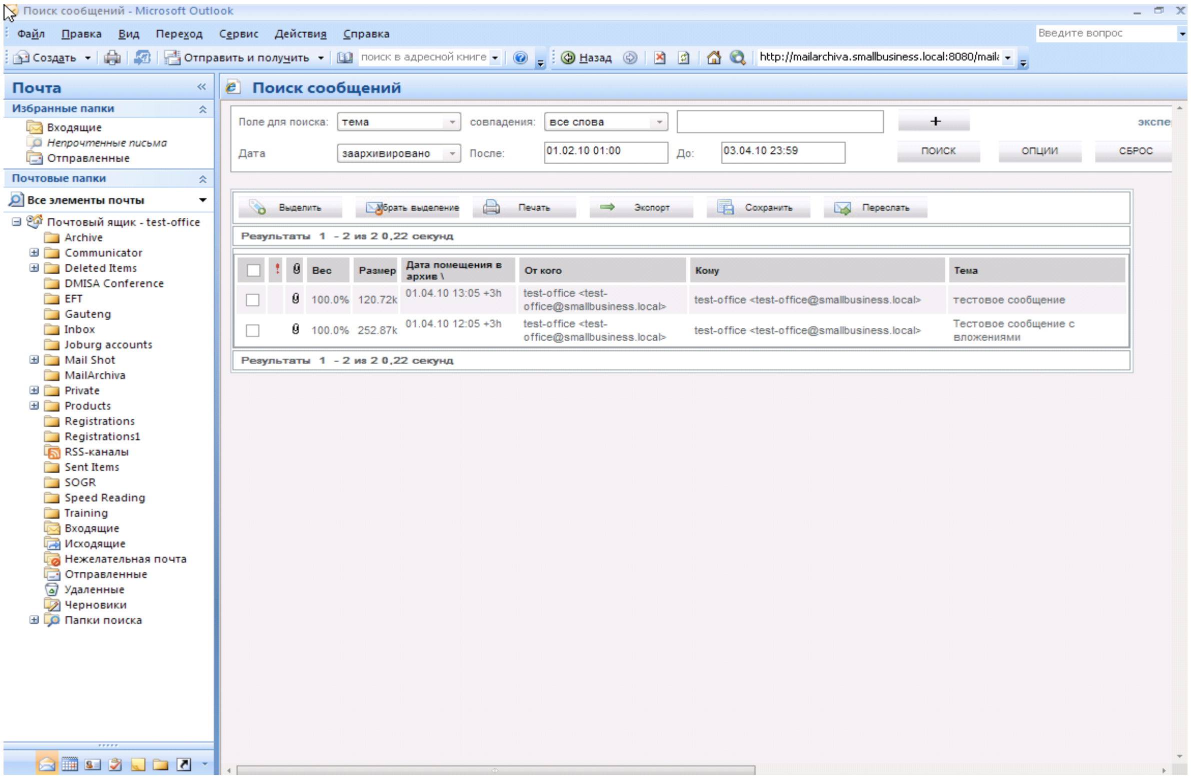

Поле для поиска: тема совпадения: все слова

Дата заархивировано После: 01.02.10 01:00 До: 03.04.10 23:59

Результаты 1 - 2 из 2 0,22 секунд

□	!	🖇	Вес	Размер	Дата помещения в архив \	От кого	Кому	Тема
□		🖇	100.0%	120.72k	01.04.10 13:05 +3h	test-office <test-office@smallbusiness.local>	test-office <test-office@smallbusiness.local>	тестовое сообщение
□		🖇	100.0%	252.87k	01.04.10 12:05 +3h	test-office <test-office@smallbusiness.local>	test-office <test-office@smallbusiness.local>	Тестовое сообщение с вложениями

Результаты 1 - 2 из 2 0,22 секунд

Figure 16 MailArchiva's Outlook Integration

13. MICROSOFT EXCHANGE MESSAGE STUBBING (OPTIONAL)

MailArchiva’s stubbing feature will strip the attachments of older e-mails in Microsoft Exchange and provide a link to the original attachments in the archive. The purpose of this feature is twofold (1) preserve the familiar Outlook experience by enabling users to access archived messages and attachments from within Outlook and in accordance with the user's existing folder structure (2) relieve pressure from the mail server by removing attachments from older messages.

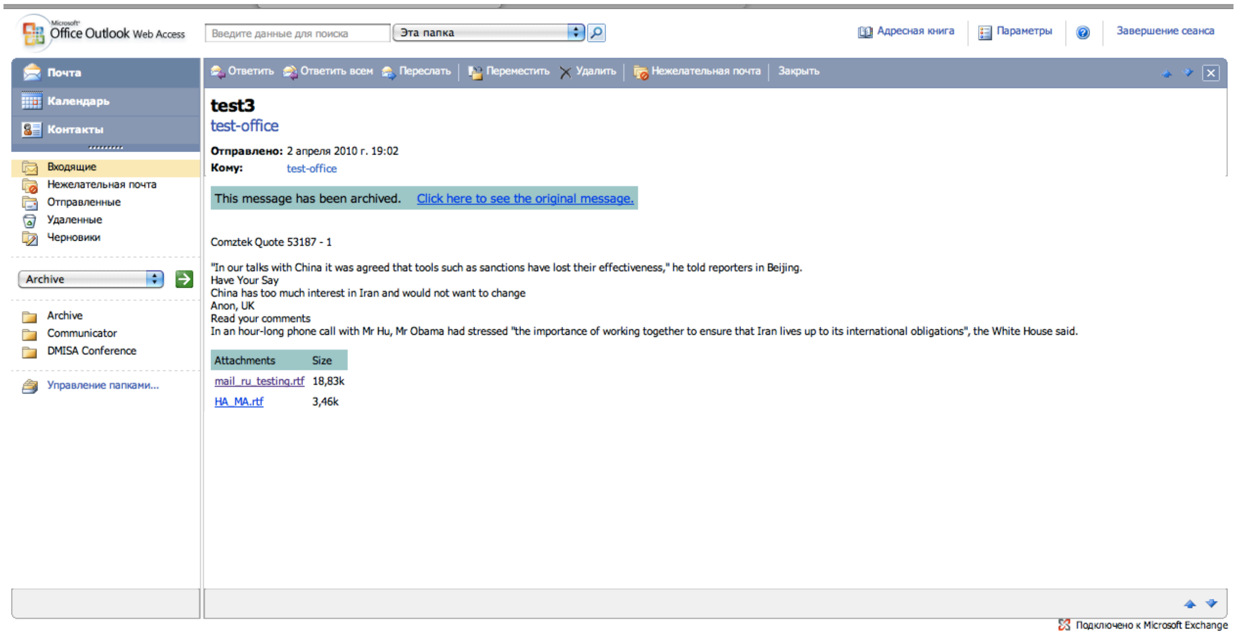

test3
test-office

Отправлено: 2 апреля 2010 г. 19:02
Кому: test-office

This message has been archived. Click here to see the original message.

Comztek Quote 53187 - 1

"In our talks with China it was agreed that tools such as sanctions have lost their effectiveness," he told reporters in Beijing.
Have Your Say
China has too much interest in Iran and would not want to change
Anon, UK
Read your comments
In an hour-long phone call with Mr Hu, Mr Obama had stressed "the importance of working together to ensure that Iran lives up to its international obligations", the White House said.

Attachments	Size
mail_ru_testing.rtf	18,83k
HA_MA.rtf	3,46k

Figure 17 Stub Message Example

As illustrated in Figure 15, a stubbed message has its attachment stripped and contains a link to the original message in the archive. Links to individual attachments in the archive are also available. When a user clicks on one of these attachment links, the user is automatically logged in to MailArchiva and the attachment opens immediately.

MailArchiva is prevented from stubbing messages that are not available in the archive. This restriction is designed to ensure that no message is modified unless there is a backup copy available. As a further safety measure, the stubbing engine is prevented from stubbing messages younger than sixty days.

A stubbed message contains a link to an original message in the archive. When clicking on the link, the Outlook user is taken directly to the message in the archive without having to supply their login information. In reality, the user is authenticated by way of single-sign-on authentication with Windows.

Since the stubbing process modifies information in Microsoft Exchange, care should be taken to ensure that the MailArchiva stubbing rules are defined correctly.

13.1.1. Active Directory Authentication

MailArchiva's stubbing engine authenticates with Microsoft Exchange using Active Directory authentication. As such as, it is necessary to ensure that Active Directory authentication is selected and functioning correctly as described in section 9.4.3. Since stubbing relies on NTLM authentication, care should be taken to ensure that NTLM authentication is checked in the AD configuration settings.

❗ Stubbing can be enabled at any stage after the MailArchiva server is deployed. However, stubbing is not backwards compatible with earlier versions of MailArchiva. For instance, emails archived by MailArchiva EE v1.9.12 or lower will not be stubbed by the server.

13.1.2. Microsoft Exchange 2003 Preparation Steps

The stubbing procedure requires the use of an impersonation account with sufficient privileges to access all mailboxes in Microsoft Exchange. Complete the steps outlined in Section 10 to create an impersonation account and assign the appropriate rights to it.

13.1.3. Web Console Stubbing Configuration

Figure 18 illustrates the configuration settings for MailArchiva's stubbing engine. Care should be taken to ensure that the link prefix correctly points to your instance of MailArchiva since once a message has been stubbed the link prefix cannot be changed. The link prefix should point to the MailArchiva instance as follows:

```
http://archive.company.com:8090
```

❗ **In the link prefix, be sure to use the fully qualified domain name of the MailArchiva server rather than its IP address. In doing so, you will have the flexibility to change the IP address of the MailArchiva server.**

Before enabling stubbing, it is strongly recommended to test your stubbing rules individually. To do this:

1. Ensure that stubbing enabled global option is unchecked
2. Create a new stubbing rule and choose the desired options
3. Save your configuration settings
4. Create a test mailbox in MS Exchange
5. Send a few test messages to it with attachments, varied text and so on
6. Ensure that the message is archived successfully in the archive
7. Click the "Test" button next to the newly created stubbing rule, enter the user principal name (UPN) of the test account (i.e. test@stimulus.local)
8. Verify that the messages are stubbed correctly and according to your preferences

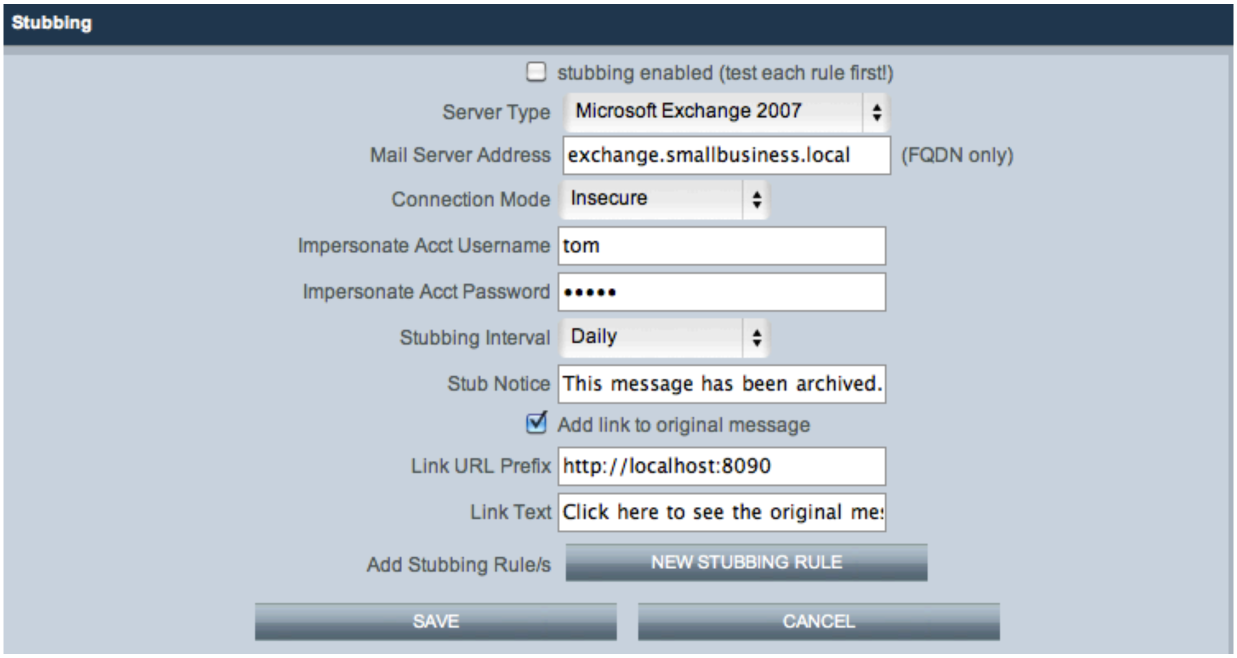

Stubbing

☐ stubbing enabled (test each rule first!)

Server Type: Microsoft Exchange 2007

Mail Server Address: exchange.smallbusiness.local (FQDN only)

Connection Mode: Insecure

Impersonate Acct Username: tom

Impersonate Acct Password: •••••

Stubbing Interval: Daily

Stub Notice: This message has been archived.

☑ Add link to original message

Link URL Prefix: http://localhost:8090

Link Text: Click here to see the original me:

Add Stubbing Rule/s

Figure 18 Stubbing Configuration

Parameter	Description
Server Type	Mail server type
Mail Server Address	The fully qualified domain name of your mail server address (e.g. exchange.mailarchiva.com)
Connection Mode	Whether or not to use a secure connection (i.e. TLS)
Impersonate Account Username	The ‘journal’ account username. The account should have sufficient privileges to read from all mailboxes in Exchange
Impersonate Account Password	The password to the ‘journal’ account .
Stubbing Interval	Specifies how frequently stubbing should occur
Stub Notice	Message appended to each message stub
Link URL Prefix	Link prefix to your MailArchiva instance
Link Text	Link text
Days Since Sent	No. Days since the message was sent

9.

The following stubbing operations are available:

Delete Message	Remove message entirely
Strip Attachments	Strip attachments from the message

Due to the potential for unwanted data loss, it is very important to test each and every one of your stubbing rules against a test account containing a sample of messages. After everything is tested thoroughly, you may enable message stubbing.

❗ Though every effort has been made to ensure that stubbing is as safe as possible, Stimulus Software will not be held liable for data loss due to misconfiguration, program faults, or for any other reason.

14. DIGITAL SIGNING AND VERIFICATION (OPTIONAL)

MailArchiva provides optional support for the digital signing of archived data using digital certificates. The use of digital signatures is an advanced feature that makes it possible to perform powerful integrity checks on a volume.

❗ In most jurisdictions, it is not necessary to enable the use of the digital signatures. The feature is provided as an advanced function for those companies that need it.

Assuming all verification checks on a volume succeed, the following is reasonably assumed:

- That all received emails were not modified since archival
- That no messages were surreptitiously added to the archive
- That no messages were deleted from the archive

MailArchiva's digital signature capability is designed to assist customers in complying with US and EU archiving legislation by ensuring that all archived information has not been tampered with.

14.1.1. Generate and Install Digital Certificates

The digital signing procedure requires the use of a signing certificate and CA certificates obtained from a Certificate Authority (CA) such as Verisign. Follow the instructions outlined in Section 9.7 to generate a server certificate and install trusted CA certificates.

14.1.2. Enabling Digital Signing

Step 1. Enable General Signing

1. Click the Digital Signing Enabled check box in the Signing tab
2. Choose the appropriate signing and verification intervals
3. Enter your current location in the signature production place fields.

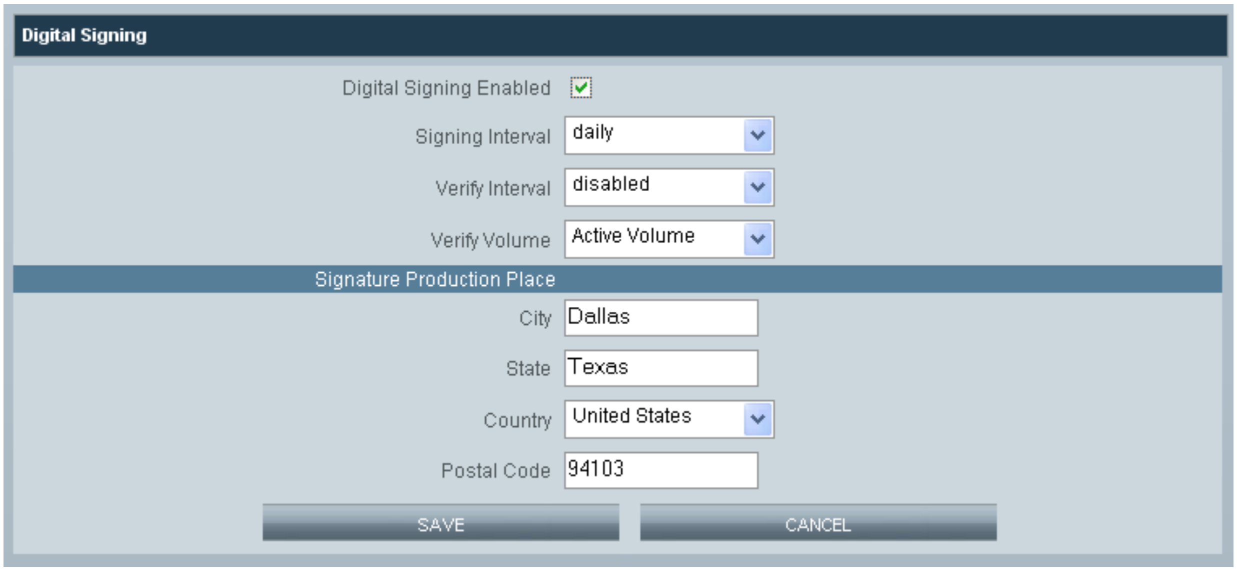

Digital Signing

Digital Signing Enabled ☑

Signing Interval: daily

Verify Interval: disabled

Verify Volume: Active Volume

Signature Production Place

City: Dallas

State: Texas

Country: United States

Postal Code: 94103

Figure 19 Enable Digital Signing

❗ Please note: Verification is very time consuming and resource intensive procedure. Your server performance may be degraded during the verification process. If you do decide to enable automatic verifications, it advisable to verify the active volume only and to set the interval on a monthly basis.

Once digital signing is enabled, the email delete function in the MailArchiva GUI is automatically disabled. Furthermore, the use of retention rules is not permitted. The reason for this is that if an email is deleted from the store for any reason, the signature checks on the volume will fail. If required, the ability to delete emails from the store, MailArchiva's digital signature capability should not be used.

Step 2. Add a Volume

Next, you need to add a new volume while carefully selecting the alias of the signing certificate you imported. Once the volume is ACTIVE, MailArchiva will begin to flag all archived messages for signing.

When the signing interval has lapsed, the server will initiate the signing process as described earlier in the chapter.

14.1.3. Verifying Signatures

Once signing has taken place on the active volume, the Verify button will appear next to the volume in the Volumes tab. The verification of all volumes will take place automatically according the schedule defined in the Signing tab.

To manually verify a volume, click the Verify button. After a while, you will receive a notification that verification has completed. Your verification report will be listed in the Signing tab.

```
**** Volume Signature Verification (ETSI TS 101 903 XAdES) Thu Mar 19 19:46:11 CAT 2009 ****

volume id:5fc01f8e-6e53-4a31-98e8-07f864bc6c17
volume store path:\store\store1557
verified by: Mail Archiva Server 1.9.0-beta2

--- begin metafest \store\store1557\manifest\metafest verification ---
  signed on Thu Mar 19 19:45:48 CAT 2009
  verified OK
  \store\store1557\manifest does not contain any orphaned manifests
--- end metafest \store\store1557\manifest\metafest verification ---

--- begin manifest \store\store1557\manifest\20090317194228.manifest verification ---
  signed on Tue Mar 17 19:42:37 CAT 2009
  verified OK
--- end manifest \store\store1557\manifest\20090317194228.manifest verification ---

--- begin manifest \store\store1557\manifest\20090318194305.manifest verification ---
  signed on Wed Mar 18 19:43:06 CAT 2009
  verified OK
--- end manifest \store\store1557\manifest\20090318194305.manifest verification ---

--- begin manifest \store\store1557\manifest\20090319194316.manifest verification ---
  signed on Thu Mar 19 19:43:18 CAT 2009
  verified OK
--- end manifest \store\store1557\manifest\20090319194316.manifest verification ---

--- begin manifest \store\store1557\manifest\20090319194548.manifest verification ---
  signed on Thu Mar 19 19:45:48 CAT 2009
  verified OK
--- end manifest \store\store1557\manifest\20090319194548.manifest verification ---
--- begin check orphaned messages ---
volume \store\store1557 does not contain any orphaned messages
--- end check orphaned messages ---

volume \store\store1557 verified OK

**** Volume Signature Verification Ended ****
```

Figure 20 Verification Report

The verification report indicates whether the integrity of the volume’s store is intact. If an email is modified, deleted or surreptitiously added (an orphaned message) to the store, it will be noted in the report. Due to the fact that on active volumes, the server is still heavily modifying the store, the check for orphaned message is not performed. Thus, while a volume is active, an email could be added to the store and it would not become known until it was closed.

❗ **The system will only check for orphaned messages on closed volumes (not active ones)**

14.1.4. Technical Background

❗ **Please skip this explanation, if you do not have a technical background in software security technology.**

Once digital signing is enabled, MailArchiva periodically creates a signed manifest in the volume store containing hashes of all emails archived for a specified time period (e.g. one day). A manifest is a digitally signed file containing references and hashes of all emails archived for a specific time period. Once a manifest file is created, it is verified either by manual procedure or automatically.

During the verification process, the system checks the hashes of every email in all manifests. It also, checks the signatures on every manifest. If an email in a volume is modified, the integrity checks will fail and an alert will be sent to the administrator identifying the exact email which was modified.

To ensure a manifest file cannot be deleted from the system without being detected, the system adds every manifest to a signed metafest file. The metafest file contains references to every manifest file. If a manifest file is deleted from the system, the metafest signature check will fail and an alert will be sent to the administrator.

During the verification procedure, the system checks for orphaned messages and manifest files. By orphaned, it is meant, files that may have been surreptitiously added to the archive. Thus, if an intruder was able to gain access to the file system and add a message to the store, the system would identify which email was added.

The digital signatures outputted by MailArchiva are compliant with the Advanced XML Digital Signature Standard (XAdES). This standard is ratified by the European Telecommunications Standards Institute (ETSI).

To support the digital signature functionality, MailArchiva creates a manifest directory in the root of the volume store directory. The contents of the manifest directory are described in Table 9.

File Name	Description
Current	Encrypted file containing the filename and hash of every email that is due to be processed for signing
current.bak	Backup of the above file
*.manifest	XAdES digital signature containing references to every email

	archived for a specified time period.
Metafest	XAdES digital signature containing references to every manifest file in the volume

Table 9 Manifest Directory Contents

15. ENVELOPE JOURNALLING

These days many mail servers support envelope journalling in some form or fashion. The task of a mail server's journalling function is to ensure that an archive server has a copy of every message that passes through the server. It achieves this by forwarding journal messages to a temporary mailbox where they are later picked up by an archive server such as MailArchiva, or forwarded directly to the archive server via the SMTP protocol. The messages outputted by the journalling function are typically encapsulated in a message envelope.

An envelope is in most cases a normal RFC2822 email in of itself. The body of the message contains meta data about the original message that is being journalled. The original message is typically attached to the envelope message. The purpose of the envelope is to provide further information about the message such as the origin, destination, timestamp, recipients, sender, bcc recipients, newsgroup recipients and so forth. All this info can be put together to help auditors provide hard evidence.

Generally, when MailArchiva archives an envelope journal message, it stores both the envelope and the attached original message. When viewing the message, MailArchiva EE unpacks the message envelope and displays just the original message. Furthermore, at the time of archiving, MailArchiva examines the envelope and correctly indexes the fields contained within. The exact indexing behaviour varies according to the specific type of envelope MailArchiva encounters. More detail on this is provided below.

If at any stage you need access to the original envelope, in Search settings, you can elect to export the Journal Message. With this option set, when exporting a message from MailArchiva, the entire envelope along with the attached original message will be exported. You should be able to open the message up in a text editor and see the envelope data, along with the original message. If you don't, it can mean that envelope journalling is not functioning correctly or it is not enabled.

15.1. Exchange 2003

Exchange 2003 supports basic envelope journaling features. When envelope journalling is enabled, an envelope message is created with the original message as an attachment, before it is sent to the journal mailbox. The envelope journal message comprises of the following:

```
Sender: "External E-mail Support" <smtp:Administrator@contoso.com>
Message-ID: <72F2A6CEB90C7F4C8D051364BF4A9FA41A89@lag.contoso.com
<mailto:72F2A6CEB90C7F4C8D051364BF4A9FA41A89@lag.contoso.com>>
Recipients:
"External E-mail Support" <smtp:Administrator@contoso.com>,
"Lene Aalling" <smtp:lenea@contoso.com>,
"Katja Heidemann" <smtp:katjaheidemann@contoso.com>,
"Doug Hite" <smtp:doughite@contoso.com>,
"Chris" <smtp:chris@contoso.com>,
"Katja folder" <smtp:Katjafolder@contoso.com>,
"Wide World Importers Folder" <smtp:WWIFolder@contoso.com>,
"Jeff Low" <smtp:JLow@contoso.com>
```

All BCC and newsgroup recipients are listed as recipients. MailArchiva indexes these recipients in the "recipient" field. When searching for these expanded recipients, users must be assigned a role with the "view recipient" permission. They should select the "recipient" field in the search. When the view recipient permission is assigned, when viewing a message the "recipient" and "sender" fields will be displayed in the message header. This is by design as these fields show who actually received the message.

Note: Granting users the "view recipient" permission will make it possible for users to search both the group recipient and BCC recipient data. With Exchange 2003, it is not possible to separate this information due to a limited envelope journal format. At the very least, it is possible to see who actually received the message and this should be good enough for legal purposes (although, don't take our word for it, we are not legal experts!).

Note: in Exchange 2003, envelope Journalling is not enabled by default. When setting up Exchange 2003 journalling, ensure that you have enable envelope journalling as described in the Administration Guide. You need to have downloaded the exejcfg utility and have executed it!

15.2. Exchange 2007 / 2010

In Exchange 2007/2010, envelope journalling is similar to as described earlier, except the 2007/2010 journal envelopes contain far more information. For one thing, MailArchiva is able to separate BCC and group recipient information, so this information is indexed as one would expect in the TO and BCC fields (they are also available in the recipient fields). For more information on Exchange 2007/2010 journal envelope format, please refer to http://technet.microsoft.com/en-us/library/bb331962.asp in the Microsoft Technet library.

15.3. Google Audit Envelopes

MailArchiva EE v2.6 and higher support the display and indexing of Google Audit envelope messages. These messages contain meta data about an original attached message. The attached message contains a message.txt file, which contains the original message in RFC2822 format. When archiving a Google audit envelope, MailArchiva stores both the envelope and its attached original message.

15.4. IpSwitch iMail Journalling

When journalling is enabled in iMail, bcc recipients are added to a BCC field in the message. Furthermore, groups recipients are automatically expanded in the To, From, CC and BCC fields. As one would expect, MailArchiva simply indexes this information in the correct fields.

15.5. Postfix/Sendmail Journalling

Postfix and Sendmail don't offer envelope journal or BCC data. MailArchiva has a workaround that may help obtain further recipient addressing information from these mails servers. When an email addressed to recipient A and B is handled by Postfix/Sendmail, the mail is sent to MailArchiva twice, once for each recipient. It is possible to configure MailArchiva such that it caches this recipient information between each connection and updates the message with the additional recipient information. The combined rcptto information is indexed in the "recipient" field. Thus, it is possible to search for the actual recipients of a message by using the recipient field. In order to search using the recipient field, the view recipient permission must be added to the user's role. To enable this option, in Configuration->Archive Settings, check "Combine RCPTTO info for identical messages".

15.6. Additional MailArchiva Headers

In additional to parsing and indexing envelope journalled messages, MailArchiva appends additional information to the headers of an email. Firstly, MailArchiva adds an archive time stamp to a message as follows:

```
X-MailArchiva-Archive-Time: Tue, 1 Dec 2010 08:48:15 -0500 (EST)
X-MailArchiva-RcptTo: archive@mailarchiva.com
X-MailArchiva-MailFrom: admin@company.com
```

Note: The time used is taken from the system clock of the archive server. It is imperative to ensure that the time is set correctly on the server at all times. The use of an NTP time server is highly recommended.

Headers containing the rcpt to and mail from informationr received from MailArchiva's SMTP or milter servers is also outputted.

15.7. Timestamps

For all its genius and simplicity, in the modern era, the SMTP protocol is now regarded as flawed (its ubiquity has meant that all efforts to fundamentally change it have failed dismally). There are many reasons why the protocol is not meeting modern needs. One of them is that the information in an email can easily be spoofed. The smtp protocol assumes that the communicating clients should determine who sent a particular email and when. For instance, the sent date of an email is often determined by time on the client machine. It remains a challenge for system administrators to ensure the time is set correctly on thousands of client machines on a network. Once more, it is even more improbable the sent date is correct when the message originates from outside the network.

For this reason, the sent dates on email messsages should be regarded as only a guide. Also of interest are the receive headers that show when the message was received by the various mail servers along the journey to delivery. The timestamps on the receive headers are set by the relaying mail servers, thus lend themselves to a greater degree of accuracy. Finally, a more authoritative indication of when the communication took place comes from the timestamp set by the archive server. In MailArchiva's case, it is set in the field "X-MailArchiva-Archive-Time".

16. AUDIT LOGS

All system activity that is logged is recorded in the audit log. Where possible, when an activity occurs, the system will record when it occurred, the username of the user that caused it, the user's IP address and the name of the specific activity that was performed. This information is indexed and searchable in the Audit section of the console. The audit logs are stored in C:\Program Files\MailArchiva\Server\webapps\ROOT\WEB-INF\logs

(Windows) or /usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/logs (Linux). For convenience reasons, the audit information is stored in CSV format and therefore can be imported into programs like Excel and Crystal Reports for further analysis. In the event that the audit index is corrupted, it can be regenerated by executing a re-index in the Logging section of the console configuration.

17. SYSTEM STATUS

The Status section of the server console screen provides an overall status of the health of the system. It includes:

- System Charts – provides a real time graphical depiction of the no. of messages that are received, ignored, duplicated and archived in the system.
- Summary – provides general information about the system including server up time and license status information.
- System alerts – any message alerts, such as low disk space, will be reported here.
- Volumes – lists all the volumes that are configured in the system and their current status.
- Processes – lists all running processes or ones that ran in the past. An example of a process may be a re-index or bulk message import operation. It is also possible to stop a running process or check its progress by viewing the output associated with it.

18. FAULT-TOLERANCE AND RECOVERY

The MailArchiva server has an inbuilt fault tolerance and recovery mechanism to cope with situations where emails cannot be archived due to faulty network connections with external

NAS's or if a message is malformed to such an extent that it cannot be construed as an RFC2822 message.

If the MailArchiva server cannot access an external storage device during the archiving process, it will immediately begin to archive messages to an internal queue called the "no-archive queue". The location of this queue can be configured in the Archive section of the server console configuration. The default location is as follows:

C:\Program Files\MailArchiva\Server\Webapps\ROOT\WEB-INF \noarchive (Windows)
/usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/noarchive (Linux)

When the connection with external storage device is finally re-established, over time the server will automatically begin to re-archive the messages in the no-archive queue and the queue will be emptied. It is also possible, though not usually necessary, to initiate the manual recovery of messages in the no-archive queue by clicking the Recover button in the Archive settings of the server console configuration.

By default, the no-archive queue resides on the same partition as the MailArchiva server software itself. Thus, it is important to ensure that there is additional hard drive space available on the partition to account for the possibility of email messages building up in the queue. For obvious reasons, it is generally not a good idea to specify the location of the no-archive queue as residing on a remote storage device.

If MailArchiva finds a situation where there is no space left on the local device where the no-archive queue is stored, it will attempt to push the emails back to the mail server. In such circumstances the following behaviour is observed:

- ♦ IMAP archiving: MailArchiva will not delete processed messages from the journal account. Rather, the messages will be marked as READ, meaning MailArchiva's inbuilt IMAP client will not attempt to process them until such time as they are manually marked as unread. If one notices that the server has not been archiving for some time and the no-archive queue is full, it is advisable to login to the journal account using OWA and mark all the messages in the journal account as read so that MailArchiva’s journal client can begin to process all unprocessed messages residing in the journal account.

- SMTP archiving: MailArchiva will inform the SMTP client to re-queue the message and send again after a long period.
- Milter: MailArchiva will inform the milter client to re-queue the message and resend.

In the unlikely event that MailArchiva is unable to process a severely malformed message, meaning that is does not comply with the specification, it will be forwarded to the quarantine location. Like the no-archive queue, the quarantine location is also, by default, located on the same partition at the MailArchiva server software.

19. USER PREFERENCES

Individual users logging to the MailArchiva can change their preferred theme in the User Preferences section. Furthermore, if basic authentication is enabled, they have the ability to change their assigned password.

20. STATUS REPORTS

To save administrators having to manually check up on the health of the system, MailArchiva includes the ability to email a status report at regular intervals to an administrator. The status report includes information such as the status of the volumes, available disk space, last known errors and various statistics.

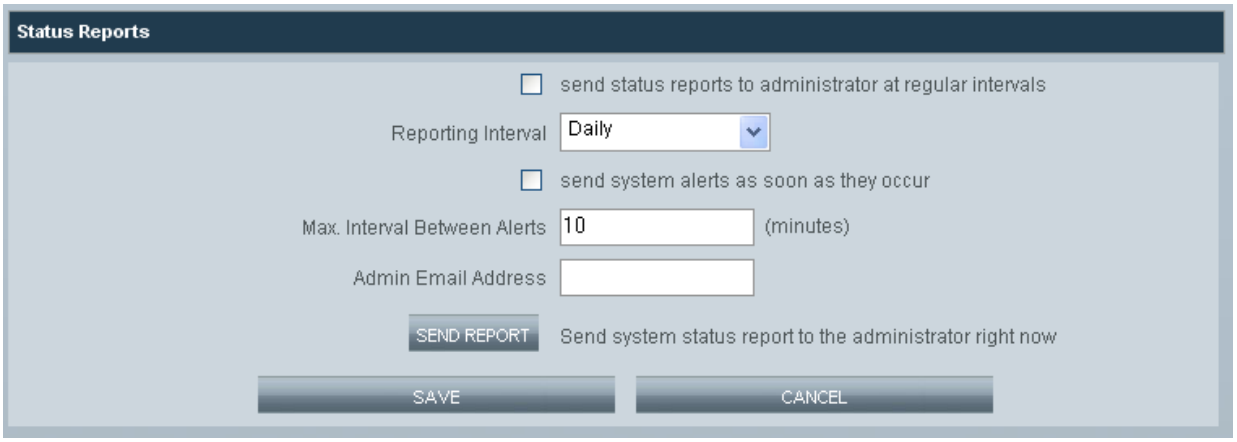

Status Reports

□ send status reports to administrator at regular intervals

Reporting Interval: Daily

□ send system alerts as soon as they occur

Max. Interval Between Alerts: 10 (minutes)

Admin Email Address:

Send system status report to the administrator right now

Figure 21 Status Report

In addition, the system can be configured to send an alert as soon as it occurs. For instance, if alerts are enabled and the server is about to run out of disk space, the administrator will be notified immediately.

The status report and alert features require that the SMTP settings in the General tab are completed. If you do not receive a status report for any reason, please refer to the MailArchiva debug log for an explanation.

21. ADVANCED CONFIGURATION OPTIONS

In addition to the configuration options accessible in the server console, there are a variety of hidden options that one can use to fine tune the server. All configuration outlined in Table 10 below are set in server.conf located in mailarchiva\server\webapps\MailArchiva \WEB-INF\conf.

Key	Values	Description
Volume.diskspace.wait	Seconds	seconds to wait between disk space checks
Volume.diskspace.warn	Megabytes	megabytes remaining on volume before disk space warning is outputted in debug log

Volume.diskspace.threshold	Megabytes	megabytes remaining on volume before disk space is considered used
Volume.diskspace.check	yes/no	Whether to perform diskspace checks
security.pbealgorithm	algorithm name	Java password-based encryption (PBE) algorithm used for encrypting messages. Default is “PBEWithMD5AndTripleDES”. See JCE API for more details
search.maxresults	Number	Default maximum search results
search.analyzer.language	two letter language identifier	Used in conjunction with search.analyzer.class to specify custom Lucene analyzer for specific a language. e.g. “en”
search.analyzer.class	java class name	Specifies the Java class name of a custom Lucene analyzer. Binds index/search languages to bespoke analyzers.
e-mailaddress.map.attribute	LDAP attribute	The LDAP attribute containing the smtp email address in Active Directory. This attribute is used to extract the user’s email address for the purposes of limiting search results for those users who are assigned the “user role”.
emailaddress.map.pattern	Regex Pattern	Used in conjunction with emailaddress.map.attribute to extract smtp addresses from users in Active Directory.
smart.attachment.minimum.size	Bytes	Minimum size of an attachment before it is separated from the body of an email.
ldap.binddn	String	The domain part of the DN used to bind to LDAP.

Table 10 Advanced Configuration Options

To change the port that you use to connect to the Web Console edit the file server\conf\server.xml and change all references from “8080” to the desired port.

22. MAILARCHIVA SHELL ACCESS

MailArchiva provides a shell environment for those who prefer the convenience of the commandline. The shell offer an easy way to check on the health of the system, view the last few entries in a log or perform more advanced configuration operations.

To access the MailArchiva Shell, simply SSH to the shell. The shell is accessible via ssh prompt. To access the MailArchiva Shell environment, ssh using Putty to localhost@2022 or type “ssh admin@localhost –p 2022” on a Linux box.

```
Ssh admin@localhost –p 2022
Password authentication
admin's password:
Welcome to the MailArchiva Shell (v2.6) on Thu Feb 03 22:13:32 SAST 2011
Unauthorized access is strictly prohibited.
Type help to see a list of commands.
Knowledge is realizing that the street is one-way, wisdom is looking both directions anyway
mailarchiva $help
configdomain        - domain settings
configvolume        - volume settings
configjournal       - journal accounts
configlisteners     - smtp/milter listeners
logs                - log settings
configstatus        - status report settings
configarchive       - archive settings
configsearch        - search settings
configindex         - index settings
configproxy         - proxy settings
configbackup        - backup settings
```

```
configgeneral        - general settings
configupdate         - software update settings
about                - view about
import               - import emails
status               - view system status

mailarchiva $status summary
server version:  2.5.10
host name:  58-b0-35-86-3f-e6.dummy.porta.siemens.net
application name:  mailarchiva_enterprise_server
ip addresses:  172.16.13.1 / 192.168.252.1 / 10.0.0.249 / fe80:0:0:0:5ab0:35ff:fe86:3fe6%5
/ fe80:0:0:0:0:0:0:1%1
operating system:  Mac OS X x86_64
server time:  2/3/11 10:14 PM;  uptime:0 days 0 hours 28 mins   start time:2/3/11 9:45 PM
last received time:  2/3/11 9:45 PM
duration since last received:  0 days 0 hours 28 mins
no. since start received:  0;  archived: 0;  duplicated: 0
last archival time:  2/3/11 9:45 PM; duration since last archival:  0 days 0 hours 28 mins
no. messages in no archive queue:  0
license status:  valid license
licensed to:  30 Day Trial
max. mailboxes:  10000
days left:  29
```

For a list of available commands, type “help”.

23. APPLIANCE MODE

When MailArchiva is deployed under the Ubuntu Operating System, the server is automatically configured to run in ‘Appliance Mode’. In this mode, MailArchiva is setup to run as a headless appliance whose basic operating system settings can be controlled via the web console. Once MailArchiva recognizes the host OS to be Ubuntu, a new ‘Server’ tab appears in the Configuration section of the web console. The server tab provides administrators with the ability to:

- Change the time and time zone settings of the machine
- Shutdown and reboot the machine
- Change the hostname and domain
- Define and configure Ethernet and DNS settings

Furthermore, in appliance mode, one has the ability to define external mount points in the Volumes section of the web console configuration.

24. PERFORMANCE TUNING

To ensure that MailArchiva is able to start up on most servers, the server is, out-of-the-box, tuned for smaller organizations (1-100). If you're planning to run MailArchiva in a larger company, it will need to be tuned to support high loads.

Tune	Description
Operating System	The MailArchiva install includes both 32 bit and 64 bit builds. If you're planning to use more than 2 GB of memory, it is necessary to install MailArchiva on a 64 bit OS.
Archiving Threads	To cope with high load, the MailArchiva server can be instructed to perform more simultaneous archiving operations. To do this, simply increase the number of archiving threads in the General section of the MailArchiva Configuration. For example, in a 1000 mailbox site, it is common to use at least 6 archiving threads.
Memory Tuning	As the number of archiving threads is increased, so too does MailArchiva's consumption of memory. Each archiving thread can potentially consume an additional extra 100 MB of memory. Since more archiving threads are needed in high load environments, the server will need more memory. To increase the amount of memory

	allocated to the server, On Windows: right click task tray applet, click Configure..., select Java tab, enter the maximum memory size you wish to allocate (e.g. 1512m). On Linux, run the /usr/local/mailarchiva/server/configure.sh script. In both cases, the server must be restarted for the memory allocation changes to take effect.
Message Retrieval	There are a number of additional settings that affect archiving performance when retrieving messages from a mail server via either IMAP or POP. The "maximum messages to process" field refers to the total number of simultaneous messages retrieved from the mail server in a single moment. Increasing this value will generally increase the speed of archiving until the Ethernet pipe is saturated. In high volume environments, it is also recommended to disable IMAP Idling and rather poll for new messages continuously. This is due the fact that the IMAP Idle function relies on the mail server to signal to MailArchiva when a new message is available. If the mail server does not signal often or fast enough, the archiving of messages may occur slower than the speed at which the journal account is filled up.

The above options are a small sample of the tuning possibilities associated with the MailArchiva system. After experimenting with them, if you still find that MailArchiva is unable to cope with your mail flow, consider whether your hardware configuration is causing any bottlenecks. For example, slow access times to your NAS / SAN may be causing the archiving process to slow. Or perhaps, it is possible you do not have a big enough Ethernet pipe connecting MailArchiva to your mail server? Taking all of the above into account, if you find that, for any reason, your MailArchiva server is unable to cope with your mail flow, please contact us or consult the Knowledge Base for further guidance.

25. SERVER MONITORING

The MailArchiva server is designed to run in a hands-free manner, although as with all enterprise software, it is necessary for administrators to keep an eye on its operation for an unusual activity. Here are some suggestions on how to monitor the server:

- Always keep the logging level at troubleshooting (debug)
- Ensure that system alerts are setup such that you will be notified when a problem occurs.
- Occasionally login to the server and click the Status window to check if the server is functioning correctly.
- If a problem is found, immediately refer to the debug log for a detailed explanation

26. BACKUP

26.1. System Configuration

```
C:\Program Files\MailArchiva\server\webapps\ROOT\WEB-INF\conf (Windows)
/usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/conf (Linux)
```

26.2. Log Files

```
C:\Program Files\MailArchiva\server\webapps\ROOT\WEB-INF\logs (Windows)
/usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/logs (Linux)
```

26.3. License File

```
/usr/local/mailarchiva/server/webapps/ROOT/WEB-INF/classes/license.lic
C:\Program Files\MailArchiva\server\webapps\ROOT\WEB-INF\classes\license.lic (windows)
```

26.4. Volume Data

To backup the store and data, perform a recursive copy of all directories and their contents under the store and index paths, respectively. Note: It is typically not necessary to backup the index data as the index can easily be regenerated using MailArchiva's volume index function.

For safety reasons, it is strongly advised to use a dedicated backup product to backup the volume store data. For additional redundancy, one can also use MailArchiva's internal backup facility accessible in Configuration->Backup. Two types of backup procedures are available: live backups and scheduled backups. The Live Backup option will archive the emails to a specified location. The Scheduled Backup option, will accordingly to a defined schedule, backup the contents of all CLOSED and ACTIVE volumes in the store.

You are strongly advised to backup your volume store data to an external off site location. In accordance with the software license agreement, Stimulus Software will not be held liable for any loss of data.

27. SERVER TROUBLESHOOTING

27.1.1. Audit & Debug Logging

The MailArchiva server has comprehensive logging facilities. There are two logs: the audit log and debug log:

- Audit Log - used for audit and forensic analysis purposes. It records all archiving and user activities in a concise manner.

- Debug Log – used for troubleshooting and debugging purposes. All errors and exceptions are reported in the debug log.

A shortened summary of each log file is accessible from the server console configuration screen. Table 11 outlines where the full log files can be found on disk.

Log	Location
Audit Log	MailArchiva\Server\logs\audit.vsv (Windows) /usr/local/mailarchiva/server/webapps/ROOT/logs/audit.csv (Linux)
Debug Log	MailArchiva\Server\logs\debug.log (Windows) /usr/local/mailarchiva/server/webapps/ROOT/logs/debug.log (Linux)

Table 11 Log File Locations

If you are experiencing problems with the server, the debug log is an invaluable tool that will assist you in getting to the root of the problem. By default, the server will output all warnings, exceptions and errors to the debug log. To enable detailed logging (i.e. to include troubleshooting messages) set the log level to troubleshoot in the Logging tab of the server console configuration. Alternatively, edit the file log4j.properties in server\webapps\WEB-INF\classes and replace all references to “info” with “debug”. You will need to restart the server before the settings will take effect.

```
Log4j.logger.com.stimulus.MailArchiva.audit=debug, MailArchivaaudit
Log4j.logger.com.stimulus.MailArchiva=debug, MailArchivadebug
```

Figure 22 log4j.properties

During normal operations, for performance reasons, it is not recommended to run the server with detailed debug logging enabled.

27.1.2. Common Problems

Please refer to the MailArchiva Knowledge Base for troubleshooting tips. A few common Server problems are described Table 12 below.

Problem	Resolution
The server wont start	(1) The server is not pointing to a valid JRE (2) There is not enough memory allocated to the JVM (3) There is not enough physical memory on the machine Please check all log files in mailarchiva/ server/logs. Run the file MailArchivaServer .exe in mailarchiva/server /bin manually.
The server archives zero byte messages	The Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files are not installed correctly.
Archived emails are not showing up in the console.	This could be any of the following: (1) Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files are not installed correctly (2) No UNUSED volumes are available (3) You've run out of disk space (4) An encryption password is not yet set (5) You have an older version of the JRE installed. You need V1.6. (6) The server is running out of memory. You need to increase the heap space allocated to the server JVM.
Server has run out of memory	The server may require large amount of memory to process extremely large emails. If you see out of memory exceptions in the server logs, you need to increase the heap space allocated to the server JVM. You do this by following the below: (1) Right click the server task tray icon at bottom right corner of the screen (2) Click Configure (3) Click Java tab (4) Increase the Maximum memory pool size (5) Restart the server (6) Test by sending a very large message (7) Examine debug.log to determine if successful

Table 12 Common Server Problems

28. SEARCH QUERIES

The search function in the server console is sufficiently intuitive that it does not warrant detailed discussion. However, it's worth mentioning that MailArchiva supports multiple and single character wildcard searches. The "?" symbol is used to indicate a single character wildcard, while the "*" symbol indicates a multiple character wildcard. For example, to search for "text" or "test" you can use the search term "te?t". To search for "test", "tests" or "tester", the search term "test*" can be used. Wildcards may be used anywhere in a search term, except at the beginning of the term. Thus, "?est" and "*est" are both invalid.

By default, when performing a search, up to 50,000 result items will be retrieved at a time. You can change this setting if you so desire, by clicking "Options" and changing the Max Results setting. It is also possible to sort the search results according to size, sent date, from, to and subject. Simply click on their respective column labels in the search results page to search in ascending and descending order. As an added benefit, you can also search for emails in multiple languages. Refer to Section 30 for an explanation of MailArchiva's internationalization capabilities.

MailArchiva embeds the highly regarded Lucene search engine. As such it can perform a range of other search functions such as fuzzy and proximity based searches. For a comprehensive description of these capabilities, please refer to Lucene's documentation available at http://lucene.apache.org.

29. EMAIL OPERATIONS

The bulk email-related operations described in Table 13 are available in MailArchiva Enterprise Edition only. To perform an operation such as export a set of emails:

(1) perform a search
(2) select the concerned emails

(3) click the icon appropriate icon in the toolbar

In (2), you may select emails individually, in the currently displayed page, or across the entire search results.

Icon	Description
🏷	Select every email in the entire search results (across all search pages)
🏷	Deselect every email
🖨	Print selected emails
✉	Delete selected emails
💾	Save the search results to a CSV file
➡	Export the selected emails to a compressed ZIP file
✉	Restore the selected to emails to a given email address
✉	View the selected emails

Table 13 Bulk Email Operations

30. INTERNATIONALIZATION

MailArchiva is an internationalized email archiving system. By default, MailArchiva supports the indexing, search and retrieval of emails written in English, Portuguese, Chinese, Czech, German, Greek, French, Dutch, Russian, Japanese, Korean and Thai.

As part of the email archiving process, MailArchiva will automatically attempt to determine the language of the email using N-GRAM analysis. The algorithm requires that there is sufficient text available to determine the language that was used. If there is not sufficient text, MailArchiva will assume that the email is written in the default language. To change the default language, refer to the Section 21.

The MailArchiva administration console user interface is currently available in English, French, German, Dutch, Chinese and Spanish. MailArchiva will automatically determine the appropriate language to display based on the user's browser settings. Furthermore, all entered and displayed dates are formatted according to the locale of the user's computer.

If you would like MailArchiva to support any other language, simply edit the file application.properties in webapps\MailArchiva\WEB-INF\classes\properties. If you do this, it would be most appreciated if you could send us a copy of your translation file for inclusion in future releases.

31. LICENSE

MailArchiva Enterprise Edition is licensed under a proprietary license agreement. Please refer to the license agreement that is bundled with the software.

32. APPENDIX

Construct	Matches
Characters	
X	The character *x*
\\	The backslash character
\0***n***	The character with octal value 0*n* (0 <= *n* <= 7)
\0***nn***	The character with octal value 0*nn* (0 <= *n* <= 7)
\0**mnn**	The character with octal value 0*mnn* (0 <= *m* <= 3, 0 <= *n* <= 7)
\x***hh***	The character with hexadecimal value 0x*hh*
\u**hhhh**	The character with hexadecimal value 0x*hhhh*
\t	The tab character ('\u0009')
\n	The newline (line feed) character ('\u000A')
\r	The carriage-return character ('\u000D')
\f	The form-feed character ('\u000C')
\a	The alert (bell) character ('\u0007')
\e	The escape character ('\u001B')

\c***x***	The control character corresponding to *x*
Character classes	
[abc]	a, b, or c (simple class)
[^abc]	Any character except a, b, or c (negation)
[a-zA-Z]	a through z or A through Z, inclusive (range)
[a-d[m-p]]	a through d, or m through p: [a-dm-p] (union)
[a-z&&[def]]	d, e, or f (intersection)
[a-z&&[^bc]]	a through z, except for b and c: [ad-z] (subtraction)
[a-z&&[^m-p]]	a through z, and not m through p: [a-lq-z](subtraction)
Predefined character classes	
.	Any character (may or may not match line terminators)
\d	A digit: [0-9]
\D	A non-digit: [^0-9]
\s	A whitespace character: [\t\n\x0B\f\r]
\S	A non-whitespace character: [^\s]
\w	A word character: [a-zA-Z_0-9]
\W	A non-word character: [^\w]
POSIX character classes (US-ASCII only)	
\p{Lower}	A lower-case alphabetic character: [a-z]
\p{Upper}	An upper-case alphabetic character:[A-Z]
\p{ASCII}	All ASCII:[\x00-\x7F]
\p{Alpha}	An alphabetic character:[\p{Lower}\p{Upper}]
\p{Digit}	A decimal digit: [0-9]
\p{Alnum}	An alphanumeric character:[\p{Alpha}\p{Digit}]
\p{Punct}	Punctuation: One of !"#$%&'()*+,-./:;<=>?@[\]^_`{|}~
\p{Graph}	A visible character: [\p{Alnum}\p{Punct}]
\p{Print}	A printable character: [\p{Graph}]
\p{Blank}	A space or a tab: [\t]
\p{Cntrl}	A control character: [\x00-\x1F\x7F]
\p{XDigit}	A hexadecimal digit: [0-9a-fA-F]
\p{Space}	A whitespace character: [\t\n\x0B\f\r]
Classes for Unicode blocks and categories	
\p{InGreek}	A character in the Greek block (simple block)
\p{Lu}	An uppercase letter (simple category)
\p{Sc}	A currency symbol
\P{InGreek}	Any character except one in the Greek block (negation)
[\p{L}&&[^\p{Lu}]]	Any letter except an uppercase letter (subtraction)
Boundary matchers	
^	The beginning of a line

$	The end of a line
\b	A word boundary
\B	A non-word boundary
\A	The beginning of the input
\G	The end of the previous match
\Z	The end of the input but for the final terminator, if any
\z	The end of the input
Greedy quantifiers	
X?	X, once or not at all
X*	X, zero or more times
X+	X, one or more times
X{n}	X, exactly n times
X{n,}	X, at least n times
X{n,m}	X, at least n but not more than m times
Reluctant quantifiers	
X??	X, once or not at all
X*?	X, zero or more times
X+?	X, one or more times
X{n}?	X, exactly n times
X{n,}?	X, at least n times
X{n,m}?	X, at least n but not more than m times
Possessive quantifiers	
X?+	X, once or not at all
X*+	X, zero or more times
X++	X, one or more times
X{n}+	X, exactly n times
X{n,}+	X, at least n times
X{n,m}+	X, at least n but not more than m times
Logical operators	
XY	X followed by Y
X|Y	Either X or Y
(X)	X, as a capturing group
Back references	
\n	Whatever the n^{th} capturing group matched
Quotation	
\	Nothing, but quotes the following character
\Q	Nothing, but quotes all characters until \E
\E	Nothing, but ends quoting started by \Q

Special constructs (non-capturing)	
`(?:`***X***`)`	*X*, as a non-capturing group
`(?idmsux-idmsux)`	Nothing, but turns match flags on - off
`(?idmsux-idmsux:`***X***`)`	*X*, as a non-capturing group with the given flags on - off
`(?=`***X***`)`	*X*, via zero-width positive lookahead
`(?!`***X***`)`	*X*, via zero-width negative lookahead
`(?<=`***X***`)`	*X*, via zero-width positive lookbehind
`(?<!`***X***`)`	*X*, via zero-width negative lookbehind
`(?>`***X***`)`	*X*, as an independent, non-capturing group

Backslashes, escapes, and quoting

The backslash character (`'\'`) serves to introduce escaped constructs, as defined in the table above, as well as to quote characters that otherwise would be interpreted as unescaped constructs. Thus the expression `\\` matches a single backslash and `\{` matches a left brace.

Character Classes

Character classes may appear within other character classes, and may be composed by the union operator (implicit) and the intersection operator (`&&`). The union operator denotes a class that contains every character that is in at least one of its operand classes. The intersection operator denotes a class that contains every character that is in both of its operand classes.

The precedence of character-class operators is as follows, from highest to lowest:

1	Literal escape	`\x`
2	Grouping	`[...]`
3	Range	`a-z`
4	Union	`[a-e][i-u]`
5	Intersection	`[a-z&&[aeiou]]`

Note that a different set of metacharacters are in effect inside a character class than outside a character class. For instance, the regular expression `.` loses its special meaning inside a character class, while the expression `-` becomes a range forming metacharacter.

Line terminators

A *line terminator* is a one- or two-character sequence that marks the end of a line of the input character sequence. The following are recognized as line terminators:

A newline (line feed) character (`'\n'`),

A carriage-return character followed immediately by a newline character (`"\r\n"`),

A standalone carriage-return character (`'\r'`),

A next-line character (`'\u0085'`),

A line-separator character (`'\u2028'`), or

A paragraph-separator character (`'\u2029`).

Groups and capturing

Capturing groups are numbered by counting their opening parentheses from left to right. In the expression `((A)(B(C)))`, for example, there are four such groups:

1 `((A)(B(C)))`

2 `(A)`

3	`(B(C))`
4	`(C)`

Group zero always stands for the entire expression.

Capturing groups are so named because, during a match, each subsequence of the input sequence that matches such a group is saved. The captured subsequence may be used later in the expression, via a back reference, and may also be retrieved from the matcher once the match operation is complete.

The captured input associated with a group is always the subsequence that the group most recently matched. If a group is evaluated a second time because of quantification then its previously-captured value, if any, will be retained if the second evaluation fails. Matching the string `"aba"` against the expression `(a(b)?)+`, for example, leaves group two set to `"b"`. All captured input is discarded at the beginning of each match.

Groups beginning with `(?` are pure, *non-capturing* groups that do not capture text and do not count towards the group total.

Comparison to Perl 5

Perl constructs not supported by this class:

The conditional constructs `(?{`*X*`})` and `(?(`*condition*`)`*X*`|`*Y*`)`,

The embedded code constructs `(?{`*code*`})` and `(??{`*code*`})`,

The embedded comment syntax `(?#comment)`, and

The preprocessing operations `\l` `\u`, `\L`, and `\U`.

Constructs supported by MailArchiva but not by Perl:

Possessive quantifiers, which greedily match as much as they can and do not back off, even when doing so would allow the overall match to succeed.

Character-class union and intersection as described above.