

Willkommen
Bienvenidos
Welkom
Bienvenue
Welcome



MailArchiva
Enterprise Edition
v1.7

INSTALLATION AND ADMINISTRATION GUIDE

For Windows / Linux




1 INDEX

1	INDEX	2
2	IMPORTANT NOTICE	4
3	CONTACT INFORMATION	4
	TECHNICAL REQUIREMENTS	5
4	OVERVIEW	6
5	HIGH-LEVEL FEATURES	7
6	ARCHITECTURE	9
7	INSTALLATION	10
7.1	EXCHANGE SERVER CONFIGURATION.....	11
7.2	SERVER INSTALLATION (ON WINDOWS)	14
7.3	SERVER INSTALLATION (ON LINUX)	16
7.4	MICROSOFT EXCHANGE	18
7.5	SENDMAIL.....	19
7.6	POSTFIX	19
7.7	QMAIL / EXIM	19
7.8	OTHER MAIL SERVERS	20
8	CONFIGURATION	20
8.1	SERVER CONFIGURATION	20
8.1.1	<i>Local Domains.....</i>	20
8.1.2	<i>E-mail Encryption Password</i>	20
8.1.3	<i>Volumes</i>	21
8.1.4	<i>Console Access.....</i>	22
8.1.4.1	Basic Authentication	22
8.1.4.2	Active Directory Authentication	23
8.1.5	<i>LDAP Authentication.....</i>	25
8.1.6	<i>Roles</i>	25
8.1.7	<i>Archive Rules</i>	26
8.1.8	<i>Retention Policy.....</i>	27
8.1.9	<i>Advanced Configuration Options</i>	27
8.2	SERVER TROUBLESHOOTING.....	28
8.2.1	<i>Audit & Debug Logging.....</i>	28
8.2.2	<i>Common Problems.....</i>	29
9	SEARCH QUERIES.....	30
10	EMAIL OPERATIONS	31
11	EMAIL MIGRATION	31
12	DISTRIBUTED SEARCH.....	31
13	INTERNATIONALIZATION	32
14	BACKWARDS COMPATIBILITY.....	32
14.1	VERSION 1.7	32
14.2	VERSION 1.5	33

14.3	VERSION 1.3	33
14.4	VERSION 1.2	33
15	DECRYPTION SOURCE CODE	34
16	LICENSE.....	34
17	APPENDIX.....	34
17.1	REGULAR EXPRESSION SYNTAX	34

2 IMPORTANT NOTICE

This Administration Guide covers installation and configuration 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)-366-8072 (EU) +44-20-80991035
FAX	(USA) +1-(713)-366-8072 (EU) +44-20-80991035
E-mail Queries	info@mailarchiva.com
Enterprise Support	support@mailarchiva.com
Knowledge Base	http://Knowledge.stimulussoft.com

TECHNICAL REQUIREMENTS

Operating Systems	Windows XP Professional, Advanced Server 2003 Ubuntu or Redhat Linux
Disk Storage	Compatible with most Storage Area Networks (SANs) and Network Attached Storage (NAS) devices
Hardware	CPU: 1x2 GHZ CPU core per 500 mailboxes RAM: 1 GB (additional 0.3 GB per 500 mailboxes)
Mail Servers	Microsoft Exchange 5.5 / 2000 / 2003 / 2007 IPSwitch IMail Lotus Domino AXIGEN Mail Server Postfix, Sendmail, Qmail, Exim, Zimbra



If you are using an older Windows NT Domain Controller, you will need to use MailArchiva's in built authentication mechanism (i.e. Basic Authentication).

4 OVERVIEW

MailArchiva is an easy-to-use, yet feature-rich, e-mail archiving system. It integrates with Microsoft Exchange and other popular mail systems. It enables you to enforce strict e-mail retention, monitoring and compliance policies throughout your organisation.

The main business benefits are:

- ◆ Preserve and access vital company knowledge
- ◆ Monitor and audit employee e-mail communications
- ◆ Ensure strict compliance with US and EU legislation (e.g. Sarbanes Oxley Act)
- ◆ Protect against law suits and legal actions
- ◆ Enhance the performance of Exchange by storing e-mails off the server
- ◆ Lower the cost of storing e-mails (compression)
- ◆ Avoid vendor lock in with proprietary formats
- ◆ End PST Hell (and not replace it with SQL hell)
- ◆ Give users long-term access to their emails

In contrast to many other e-mail archiving systems, MailArchiva stores e-mails directly on the file system. This design allows you 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.

MailArchiva stores your e-mail in standard Internet mail format (RFC822). RFC822 is the standard format for storing and transporting e-mail messages on the Internet. Thus, MailArchiva ensures that your information will remain accessible over the long-run.

5 HIGH-LEVEL FEATURES

Feature	EE
Admin console available in English, French, Dutch, Chinese, German and Spanish	✓
Archive all incoming, outgoing and internal e-mails	✓
Search, view, sort archived e-mails	✓
Find e-mails using complex search criteria	✓
Search inside Word, Powerpoint, Excel, PDF, RTF, ZIP, tar, gz attachments	✓
Define granular archiving rules according to policy	✓
Save on storage costs by storing e-mail in compressed format	✓
Easy-to-use web-based user interface	✓
Login to web console using Windows authentication	✓
Login to web console using in built authentication	✓
Permit/restrict employee access to their own e-mails	✓
Restrict access to web console based on Windows user groups and attributes	✓
Comprehensive audit trail and system logging	✓
No database required - messages are stored directly on the file system	✓
Supports multiple volumes (disk drives)	✓
Messages stored in standard RFC822 format	✓
Tight security – all e-mails encrypted	✓
Interfaces with multiple mail systems via MAPI, IMAP, POP	✓
IPSwitch IMail support	✓
Sendmail support	✓
Postfix support	✓
QMail support	✓
Exim support	✓
Fully internationalized – multiple language support	✓
Do-it-yourself, easy setup	✓
Export e-mails en bulk	✓
Print e-mails en bulk	✓
Delete e-mails en bulk	✓
Microsoft Exchange 2007 Support	✓
Microsoft Exchange Envelope Journaling – stores all e-mail header / addressing information (including BCC fields and Exchange groups)	✓
Multiple Microsoft Exchange Stores – archive e-mails in multiple exchange stores	✓
Multiple Exchange Servers – archive e-mails from multiple Exchange Servers	✓
Smart attachment storage - save space by storing only one copy of an attachment across several e-mails	✓
Restoration of e-mails en bulk to multiple mail systems	✓
Email retention policy	✓

Distributed search – search across multiple machines simultaneously	✓
Authenticate to the console using LDAP authentication	✓
Authenticate to the console using IPSwitch IMail authentication	✓
Single instance message storage	✓
Message failover and recovery	✓
Convert messages in PST files to MailArchiva	✓
Convert messages from Exchange to MailArchiva	✓
Flexible Role management system	✓
Built-in IMAP and POP client	✓
Convert messages from MBOX to MailArchiva	✓

Table 1 MailArchiva EE Features

6 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, then there must be TCP/IP connectivity between the MailArchiva server and the server hosting Active Directory. 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 or a user role. Each of these roles implies a different set of entitlements, which are discussed later in this guide.

For simplicity 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.

If there is a firewall running between any of the components in the architecture, you will need to the communications ports as described in Table 2.

Source	Destination	Protocol	Ports
MailArchiva Server	Active Directory	Kereberos, LDAP	88, 389
MailArchiva Server	Microsoft Exchange	IMAP	143, 993
Sendmail/Postfix	MailArchiva Server	Sendmail milter	8092*
Mail Server	MailArchiva Server	SMTP	8091

* by default, you can change this port

Table 2 Communication Ports

7 INSTALLATION

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

Mail Server	Description	Sections
Exchange	You are using Microsoft Exchange	7.1, 7.2, 7.3, 7.4
IMail	You are using IPSwitch IMail	Refer to IMail Docs
Sendmail	You are using the Sendmail mail server	7.2, 7.3, 7.5
Postfix	You are using the Postfix mail server	7.2, 7.3, 7.6
Qmail/Exim	You are using the Qmail mail server	7.2, 7.3, 7.7
Alternate	All other mail servers	7.2, 7.3, 7.8

Table 3 Installation Sections to Complete

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.

7.1 Exchange Server Configuration

The Microsoft Exchange product includes a message journaling feature that saves a copy of every e-mail message that is sent from or received on a specific mail store. To archive all messages processed by Exchange, the MailArchiva server requires that this 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.

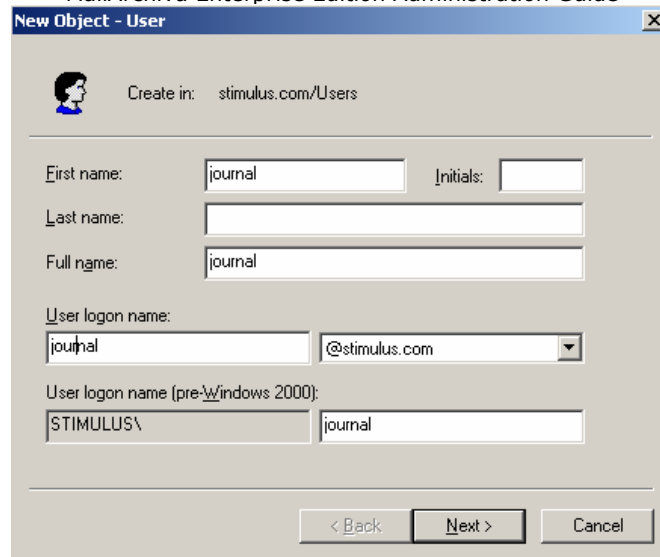
Please refer to Table 4 for an overview of Microsoft Exchange related features supported by MailArchiva OSE and EE.

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

Table 4 MailArchiva Exchange Features

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).



New Object - User

Create in: stimulus.com/Users

First name: journal Initials:

Last name:

Full name: journal

User login name: journal @stimulus.com

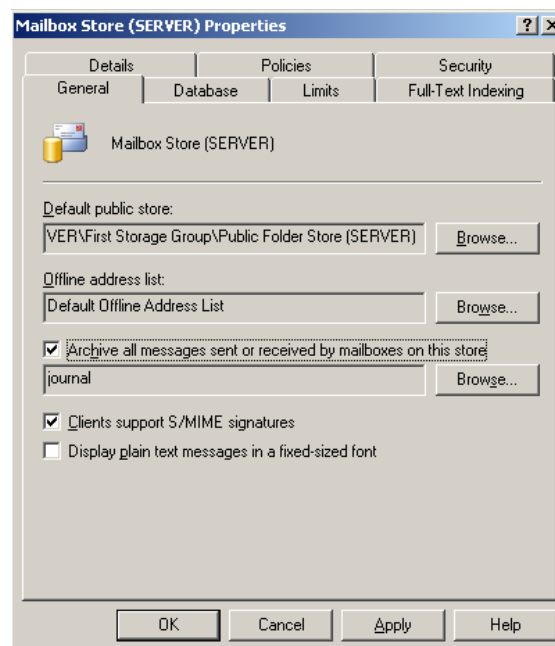
User login name (pre-Windows 2000): STIMULUS\ journal

< Back Next > Cancel

Figure 1 Journaling Account Creation

Step 3. 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 2. Click Browse and enter "journal" for the object name. Click OK. Journaling is now enabled for the Mailbox Store.



Mailbox Store (SERVER) Properties

Details Policies Security

General Database Limits Full-Text Indexing

Mailbox Store (SERVER)

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

Offline address list: Default Offline Address List Browse...

☒ Archive all messages sent or received by mailboxes on this store
journal Browse...

☒ Clients support S/MIME signatures

☐ Display plain text messages in a fixed-sized font

OK Cancel Apply Help

Figure 2 Enable Journaling

Step 3. Enable Envelope Journaling



IGNORE this step if you running Microsoft Exchange 2007.

(envelope journaling is enabled by default in Microsoft Exchange 2007)

If you are running Exchange 2003:

Install the latest Service Pack

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>

7.2 Server Installation (on Windows)

-  **You are no longer required to install the Java Runtime Environment (JRE) as this package is now embedded with the server.**
-  **If you are upgrading from MailArchiva v1.2, you MUST immediately backup your server.conf file before uninstalling the old version.**
-  **If you are upgrading, please read Section 14 on backwards compatibility before continuing.**

Step 1. Install MailArchiva

Run the MailArchiva Server Setup and follow the instructions on screen. It is strongly recommended that you install both the MailArchiva Server and Application Server components. In the event that you wish to install the server application on an existing instance of Apache Tomcat, you may install the .WAR file on its own.

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 4. Modify Hosts File (optional)

If you plan on authenticating users residing in Active Directory, then you must add the IP address, fully qualified domain name (FQDN) and name of the server running Active Directory to your c:\windows\system32\drivers\etc\hosts file. Figure 3 illustrates the entry that must be added to the hosts file.

```
192.168.0.100 ACTIVEDIRECTORY.COMPANY.LOCAL ACTIVEDIRECTORY
```

(NB: Please replace ACTIVEDIRECTORY.COMPANY.LOCAL with the ACTUAL fully qualified domain name of your Active Directory server!)

Figure 3 c:\windows\system32\drivers\etc\hosts

-  **The above hosts file entry must be present if you intend to authenticate users in Active Directory. Furthermore, all letters of the FQDN and server name must be capitalized as illustrated above.**

Please also ensure that the MailArchiva Server time is equivalent to the time reported by the Active Directory Server.

Step 5. Install the License File

If you are running MailArchiva Enterprise Edition, it is necessary to install the license.lic file sent to you via email. Copy your license file to the following location:

MailArchiva\Server\webapps\mailarchiva\WEB-INF\classes

Step 6. 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.

Step 7. Configure Server Settings

This final step involves configuring the server. There are at minimum four configuration tasks that need to be performed before the server is ready to start archiving e-mails: (a) set an encryption password (b) create one volume (c) add your local domains. To complete these tasks, follow the instructions Chapter 8.

7.3 Server Installation (on Linux)

 **If you are upgrading, please read Section 14 on backwards compatibility before continuing.**

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_v1_7_5_linux.tar.gz
cd mailarchiva_dist
sudo install.sh
```

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. Modify Hosts File

If you plan on authenticating users residing in Active Directory, then you **must** add the IP address, fully qualified domain name (FQDN) [uppercase] and name [uppercase] of the server running Active Directory to your `/etc/hosts` file. Figure 4 illustrates the entry that must be appended to the hosts file.

```
192.168.0.100 ACTIVEDIRECTORY.COMPANY.LOCAL ACTIVEDIRECTORY
```

(NB: Please replace ACTIVEDIRECTORY.COMPANY.LOCAL with the ACTUAL fully qualified domain name of your Active Directory server!)

Figure 4 /etc/hosts

 **The above etc/hosts entry must be present if you intend to authenticate users in Active Directory. Furthermore, the FQDN and server name must be uppercase as illustrated above.**

Please also ensure that the MailArchiva Server time is equivalent to the time reported by the Active Directory Server.

Step 4. Install the License

Copy your license file to the following location (you may overwrite the demo license):

mailarchiva/server/webapps/mailarchiva/WEB-INF/classes

Step 5. Start MailArchiva Server

To start the MailArchiva Server from the commandline type:

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

Type: "<http://localhost:8090/mailarchiva>" 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 sh /etc/init.d/mailarchiva stop
```

Step 6. Configure Server Settings

This final step involves configuring the server. There are at minimum three configuration tasks that need to be performed before the server is ready to start archiving e-mails: (a) set an encryption password (b) create one volume (c) add your local domains. To complete these tasks, follow the instructions Chapter 8.

7.4 Microsoft Exchange

MailArchiva can interface with Microsoft Exchange in a variety of ways. The easiest way is to configure MailArchiva to fetch emails from Exchange's IMAP connector. On a fresh install of the Exchange product, the IMAP connector is switched on and ready for action. Before continuing, ensure that Microsoft Exchange's IMAP service is switched on. 1

Step 1. Add a Mailbox Connection

In the Mailboxes tab of the MailArchiva server console configuration screen, click Add Mailbox Connection and do the following:

Select IMAP as the preferred protocol

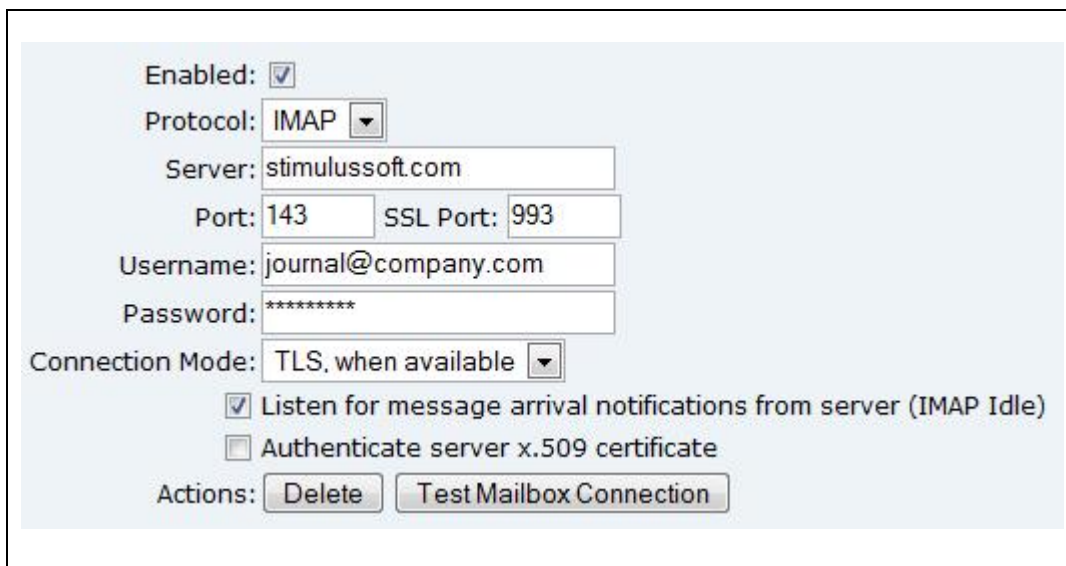
Enter the server address of your Exchange server

Enter the Microsoft Exchange journal account username and password

For "Connection Mode", select TLS when available

Ensure "Auth Certs" is unchecked.

Ensure "Listen for message arrival notifications from server" is checked.



The screenshot shows the 'Add Mailbox Connection' configuration window. It includes the following fields and options:

- Enabled:** ☒
- Protocol:** IMAP (dropdown menu)
- Server:** stimulussoft.com
- Port:** 143
- SSL Port:** 993
- Username:** journal@company.com
- Password:** *****
- Connection Mode:** TLS, when available (dropdown menu)
- ☒ Listen for message arrival notifications from server (IMAP Idle)
- ☐ Authenticate server x.509 certificate
- Actions:** Delete, Test Mailbox Connection

Figure 5 Mailbox Connection To Microsoft Exchange

Step 2. Test Mailbox Connection

Click the Test 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 SMTP 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).

 **If you do not wish to have Exchange's IMAP service enabled, you can also interface with MailArchiva by configuring Exchange to forward SMTP traffic to MailArchiva. Alternatively, you can also use the MailArchiva Exchange agent. Refer to <http://knowledge.mailarchiva.com> for more information on these approaches.**

7.5 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
```

7.6 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 send mail

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

7.7 Qmail / Exim

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

7.8 Other Mail Servers

There are four main ways to connect your mail server to MailArchiva, namely:

- By forwarding SMTP traffic to MailArchiva
- By forwarding Sendmail filter traffic to MailArchiva
- By configuring MailArchiva to fetch mail from your mail server using POP
- By configuring MailArchiva to fetch mail from your mail server using IMAP

The MailArchiva knowledge base (<http://knowledge.mailarchiva.com>) includes further instructions on how to interface MailArchiva with a variety of mail systems.

8 CONFIGURATION

8.1 Server Configuration

The server configuration settings are accessible from the Configuration screen in the web console. Only users with 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/mailarchiva>" 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, you may simply hit enter to login. By default, authentication to the web console is disabled.

8.1.1 Local Domains

When configuring MailArchiva for the first time, you need to add one or more of your organization's domains. To do this, click "Add Domain" in the domain section of the configuration screen. An example domain is "company.com" or "company.local". The entered domains are used by the server to assess whether the origin and destination of e-mails are internal or external to your organization. When applying archive rules, the server will match the domain of a given e-mail address with all of the domains entered here. If your organization has an internal domain called "company.local" and an external one called "company.com", you need to include both these domains.

8.1.2 E-mail Encryption Password

All e-mails are stored encrypted using triple DES password-based encryption. Before using the server to archive e-mails, you need to choose and enter an E-mail 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 e-mails, 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 e-mails.

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 e-mail encryption purposes. If you lose either of these, you will be unable to access the e-mails 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.

❗ Java source code to decrypt archived messages to RFC822 is available on sourceforge.net.

8.1.3 Volumes

Archived e-mails are organised into one or more 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 one or more hard disks. 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.

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

Once you've created a volume, you'll notice that it is assigned the "NEW" status, as described in Table 5 below. Volumes have a lifecycle of their own. Once the archiving process begins, the server will automatically switch over to the first unused volume on the list. This volume will become the active volume until such time as its maximum size is exceeded, the disk is full, or you explicitly close the volume. Once a volume is closed, no further data can be written to it and it cannot be reopened.

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.

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.
REMOTE	The volume's index resides on a remote machine. The volumes store must still be held locally.

Table 5 Volume Status

When using removable disks, it is not recommended to remove the disk containing the active volume data without closing the volume first. You may remove any physical disk containing a closed volume. When doing so, is it usually a good idea to explicitly unmount the volume, although this is not absolutely necessary.

When users search for e-mails, the search is conducted across all 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 recommended in the event of data loss. To re-index a volume, you need to close it first, and click on the "Re-Index" button.

8.1.4 Console Access

The MailArchiva server can be configured such that only authorised users are able to access the web console and its features. There are two types of authentication mechanisms currently supported: Basic, Active Directory and LDAP authentication. When the server is set to use Basic Authentication, users are authenticated using credentials stored in a configuration file. When Active Directory Authentication is enabled, users login to the web console using their normal Windows login credentials. In addition, MailArchiva can be configured to authenticate against users in standard directory 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. Refer to Section X to learn more about MailArchiva's role mechanism.

8.1.4.1 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 6, contains a list of users, each of which has an e-mail address, role and a password. The users listed

in users.conf will login using an e-mail address and password that corresponds with an entry in the file. Once a user is authenticated, the user will be assigned the specified role.



```
<Users version="1.0">
  <User e-mail="admin@company.local" role="administrator" password="123"/>
  <User e-mail="user@company.local" role="user" password="abc"/>
  <User e-mail="auditor@company.local" role="auditor" password="xyz"/>
</Users>
```

Figure 6 Users.conf

-  **If basic authentication is enabled and no users are defined in the users.conf file, the system will assume that any user can login using any credentials.**
-  **If you update users.conf using a text editor, the server will pick the changes up automatically without having to restart.**

8.1.4.2 Active Directory Authentication

In Active Directory (AD) Authentication mode, the server uses the Kerberos and LDAP protocols to authenticate users residing in Active Directory. Thus, when enabling AD Authentication, in the Kerberos Server and LDAP Server fields enter the fully qualified domain name of your Active Directory server (e.g. exchange.company.com). By default, ports 88 and 389 are used for the Kerberos and LDAP server, respectively.

-  **You must enter fully qualified domain names. An IP address or short name will not work.**
-  **Make sure you have added the FQDN, server name and IP address of the Windows Server to the hosts file. See Section 7 for more details on how to do this.**

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 6 Role Assignment Fields

To complete the attribute and match criterion fields, you need an understanding of 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 your attribute selection is retrieved from Active Directory. This value is compared against the value you enter in the match criterion field. If there is a match, the role you select is assigned to the user.

To assign a role to a Windows user, simply 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 7 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. You do this 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 you've configured your role assignments, execute a Test Login to ensure that your Kerberos settings, LDAP settings and user roles have been configured correctly. If you encounter problems, enable server debugging as described in Section 8.2.1 to determine the source of the problem.

 **Note: If you get locked out from the web console during the below configuration procedure, do the following:**

- 1. Edit the file server.conf located in
mailarchiva\server\webapps\MailArchiva\WEB-INF\conf**
- 2. Set "security.loginmethod=basic" and edit the file users.conf**
- 3. Restart the server**

8.1.5 LDAP Authentication

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

The screenshot shows a web-based configuration interface for LDAP authentication. It contains several input fields and two buttons. The fields are labeled as follows: 'LDAP Server Address' with the value 'openldap.company.com:389' and a hint '(FQDN:port)'; 'Base DN'; 'Service DN'; 'Service Password'; 'Bind Attribute' with the value 'uid'; and 'Email Attribute' with the value 'email'. At the bottom, there is a label 'Assign Roles to User/s:' followed by two buttons: 'New Role Assignment' and 'Test Login'.

Figure 7 LDAP Login Attributes

The following process occurs during LDAP console login:

- MailArchiva authenticates with the directory using a Bind 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 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).

8.1.6 Roles

Once a user has logged into the console, 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 X and X, 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
-------	-----	-----	-----	-----	-----	-----	-----

Table 8 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 view any email

Table 9 Built-In Role Email Filters

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 Custom Role tab 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.

⚠ 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 (see Figure X below).

Figure 8 Custom Role Definition

8.1.7 Archive Rules

In some circumstances, it may not be desirable to archive all e-mails. Archive rules are used to determine whether or not an e-mail should be archived. As an administrator, you can choose to archive incoming, outgoing and/or internal e-mails. If these basic rules are not granular enough, advanced rules may be defined that will determine whether or not to archive an e-mail based on specific criteria.

The sequence in which the archiving rules are processed is significant. By design, advanced rules are always processed before basic rules. Furthermore, an advanced rule that appears before another will always be processed first. If during processing, an

advanced rule determines that an e-mail should not be archived then the action will be applied, irrespective of whether a subsequent rule contradicts the decision.

An advanced rule consists of one or more clauses. By selecting "any of the following" or "all of the following", any or all of the clauses in the rule must match for it to apply. 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 e-mail 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 e-mails addressed to john@company.com are archived, you would simply select the field "to", select the "contains" operator and enter "john@company.com".

8.1.8 Retention Policy

Using this configuration option, MailArchiva can be configured to automatically purge emails older than a specific time period.

8.1.9 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 disk space checks
security.pbealgorithm	algorithm name	Java password-based encryption (PBE) algorithm used for encrypting messages. Default is "PBESWithMD5AndTripleDES". 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	The LDAP attribute containing the smtp e-

Key	Values	Description
	attribute	mail address in Active Directory. This attribute is used to extract the user's e-mail 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.

8.2 Server Troubleshooting

8.2.1 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 X outlines where the log files can be found on disk.

Log	Location
Audit Log	MailArchiva\Server\logs\audit.log (Windows) /usr/local/mailarchiva/server/logs/audit.log (Linux)
Debug Log	MailArchiva\Server\logs\debug.log (Windows) /usr/local/mailarchiva/server/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 9 log4j.properties

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

8.2.2 Common Problems

The most commonly encountered Server problems are described Table 12.

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 <code>mailarchiva/server/logs</code> . Run the file <code>MailArchivaServer.exe</code> in <code>mailarchiva/server/bin</code> 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 (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

9 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 multiple languages. Refer to Section 13 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>.

10 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

11 EMAIL MIGRATION

MailArchiva archives emails as they are processed by the email server. If you'd like to access older (pre-deployment) emails, you will need to import them using a variety of utilities available in the MailArchiva utilities package.

EML/MSG Files	Ex2MailArchiva Utility
Direct MS Exchange Import	Ex2MailArchiva Utility
MBOX Files	Ex2MailArchiva Utility
Maildir	Ex2MailArchiva Utility
PST Files	PST2MailArchiva Utility
Direct IPSwitch IMail Import	IMail2MailArchiva Utility

For all other import requirements, please contact us.

12 DISTRIBUTED SEARCH

The MailArchiva server includes a distributed search capability for large deployments. Using this feature, you can execute simultaneous searches on multiple indexes distributed on several machines. By doing this, you can leverage the resources of multiple server simultaneously and thereby increase search performance.

Here is how to implement distributed search:

- (1) Copy an old index to a second machine
- (2) Mount an NFS share, accessible to both machines, on the volume store directory
- (3) Install MailArchiva on the second machine
- (4) Copy the MailArchiva server.conf file from the first machine
- (5) Create a volume on the second machine
 - a) Set archive store to be the mount point that refers to the shared NFS drive
 - b) Set index to be the local index that was copied from the first machine
- (6) On the second machine, click the publish volume for distributed search in the Volumes screen
- (7) On the first machine, in the volume's index location, type "rmi//ipaddress" (where ipaddress is the IP address of the second machine)

You should see that the volume status changes to "REMOTE". This means that the volume is pointing to the index on the external machine.

When you do a search on the first machine, it will use the processing resources of the second machine. The first search operation is slow because it has to initialize the connection. Once this is established, the process works very smoothly.

13 INTERNATIONALIZATION

MailArchiva is an internationalized e-mail 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 e-mail archiving process, MailArchiva will automatically attempt to determine the language of the e-mail 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 e-mail is written in the default language. To change the default language, refer to the Section 8.1.9.

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.

14 BACKWARDS COMPATIBILITY

14.1 Version 1.7

To upgrade from earlier versions, run the Repair Volume Utility in the MailArchiva Utilities package.

14.2 Version 1.5


Be sure to re-index your older volumes to take advantage of the improved search capabilities of MailArchiva Enterprise Edition Version 1.5.


14.3 Version 1.3


If you are upgrading from v1.2 or 1.3 and you intend to use the user role functionality, you may need to reindex your older volumes.

To upgrade from v1.3, run the MailArchiva Server v1.4 setup directly and install the product over the existing v1.3 installation on your hard disk.

14.4 Version 1.2

 **WARNING:** Ensure that you have backed up your server.conf file before uninstalling version 1.2, as it contains a salt and password that is needed to access your data. *Knowing just your password is not sufficient.*

 It is recommended you close older volumes before upgrading to version 1.4.

 **It is highly recommended that you run the MailArchiva v1.4 setup without uninstalling MailArchiva v1.2. In this way, your settings will be automatically carried over to the new version.**

When upgrading to MailArchiva v1.4 from v1.2, you may initially experience problems accessing your v1.2 volumes. The reason for this behaviour is that the v1.2 server did not explicitly state that you needed to install the Java Security Policy files. The consequence of not having installed the policy files is that your archived emails may be encrypted using DES encryption (as opposed to triple-DES encryption). For security reasons, the v1.3 server, by default, uses triple-DES encryption.

To ensure v1.3 compatibility with older volumes, you can either decide to continue using the DES encryption for all volumes (including the new ones that you create), or you can use a commandline utility to upgrade your volumes to triple-DES encryption.

To configure the server to use DES encryption edit the server.conf file as described in below and restart your server.

```
security.pbealgorithm = PBESWithMD5AndDES
```

Table 14 Security Algorithm Change For v1.2 Backward Compatibility

If you decide, you'd like to upgrade your existing v1.2 volumes, there is a ConvertVolume utility available in mailarchiva/server/utilities that will recursively convert all messages in a given volume store to triple-DES encryption. Before you run the utility, you must backup

all the affected volumes for safety sake. ConvertVolume accepts a salt, passPhrase and storeDirectory as commandline parameters. The salt and passPhrase can be copied directly from your v1.2 server.conf file. The storeDirectory refers to the location of the target volume's store directory.

```
ConvertVolume e7150baa58927558 password C:\volume1\store
```

Table 15 Volume Conversion to Triple-DES Encryption

In addition, the various bug fixes and enhancements in v1.3 have necessitated a change in how volumes are indexed. Notably, on volumes created with v1.2, the following side effects may be observed:

- ◆ The search results cannot be sorted correctly
- ◆ Users cannot search for attachments, priority and flags (as these fields have been added to the index)
- ◆ The user role capability will not work on v1.2 volumes

If you desire any of these capabilities on your v1.2 volumes, you will need to re-index them. For safety sake, only re-index if you absolutely require these capabilities. Ensure that your volume indexes are backed up before doing this.

15 DECRYPTION SOURCE CODE

To ensure that your e-mails are attainable over the long-term, a DecryptMessage utility, along with source code, is available in the mailarchiva/server/utilities directory. If you study the code, you'll notice that it is very straight forward to decrypt and decompress messages in the store. You are free to modify this source code as you see fit.

16 LICENSE

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

17 APPENDIX

17.1 Regular Expression Syntax

Construct	Matches
Characters	

X	The character <i>x</i>
\	The backslash character
\0<i>n</i>	The character with octal value <i>0n</i> ($0 \leq n \leq 7$)
\0<i>nn</i>	The character with octal value <i>0nn</i> ($0 \leq n \leq 7$)
\0<i>mnn</i>	The character with octal value <i>0mnn</i> ($0 \leq m \leq 3, 0 \leq n \leq 7$)
\x<i>hh</i>	The character with hexadecimal value <i>0xhh</i>
\u<i>hhhh</i>	The character with hexadecimal value <i>0xhhhh</i>
\t	The tab character (' <i>\u0009</i> ')
\n	The newline (line feed) character (' <i>\u000A</i> ')
\r	The carriage-return character (' <i>\u000D</i> ')
\f	The form-feed character (' <i>\u000C</i> ')
\a	The alert (bell) character (' <i>\u0007</i> ')
\e	The escape character (' <i>\u001B</i> ')
\c<i>X</i>	The control character corresponding to <i>x</i>

Character classes

[<i>abc</i>]	<i>a</i> , <i>b</i> , or <i>c</i> (simple class)
[^<i>abc</i>]	Any character except <i>a</i> , <i>b</i> , or <i>c</i> (negation)
[<i>a-zA-Z</i>]	<i>a</i> through <i>z</i> or <i>A</i> through <i>Z</i> , inclusive (range)
[<i>a-d[m-p]</i>]	<i>a</i> through <i>d</i> , or <i>m</i> through <i>p</i> : [<i>a-dm-p</i>] (union)
[<i>a-z&&[def]</i>]	<i>d</i> , <i>e</i> , or <i>f</i> (intersection)
[<i>a-z&&[^bc]</i>]	<i>a</i> through <i>z</i> , except for <i>b</i> and <i>c</i> : [<i>ad-z</i>] (subtraction)
[<i>a-z&&[^m-p]</i>]	<i>a</i> through <i>z</i> , and not <i>m</i> through <i>p</i> : [<i>a-lq-z</i>] (subtraction)

Predefined character classes

.	Any character (may or may not match line terminators)
\d	A digit: [<i>0-9</i>]
\D	A non-digit: [<i>^0-9</i>]
\s	A whitespace character: [<i>\t\n\x0B\f\r</i>]
\S	A non-whitespace character: [<i>^\s</i>]
\w	A word character: [<i>a-zA-Z_0-9</i>]
\W	A non-word character: [<i>^\w</i>]

POSIX character classes (US-ASCII only)

\p{Lower}	A lower-case alphabetic character: [<i>a-z</i>]
\p{Upper}	An upper-case alphabetic character: [<i>A-Z</i>]
\p{ASCII}	All ASCII: [<i>\x00-\x7F</i>]
\p{Alpha}	An alphabetic character: [<i>\p{Lower}\p{Upper}</i>]
\p{Digit}	A decimal digit: [<i>0-9</i>]
\p{Alnum}	An alphanumeric character: [<i>\p{Alpha}\p{Digit}</i>]
\p{Punct}	Punctuation: One of <i>!"#\$%&'()*+,-./:;<=>?@[\\]^_`{ }~</i>
\p{Graph}	A visible character: [<i>\p{Alnum}\p{Punct}</i>]
\p{Print}	A printable character: [<i>\p{Graph}</i>]

<code>\p{Blank}</code>	A space or a tab: [\t]
<code>\p{Cntrl}</code>	A control character: [\x00-\x1F\x7F]
<code>\p{XDigit}</code>	A hexadecimal digit: [0-9a-fA-F]
<code>\p{Space}</code>	A whitespace character: [\t\n\x0B\f\r]

Classes for Unicode blocks and categories

<code>\p{InGreek}</code>	A character in the Greek block (simple block)
<code>\p{Lu}</code>	An uppercase letter (simple category)
<code>\p{Sc}</code>	A currency symbol
<code>\P{InGreek}</code>	Any character except one in the Greek block (negation)
<code>[\p{L}&&[^\p{Lu}]]</code>	Any letter except an uppercase letter (subtraction)

Boundary matchers

<code>^</code>	The beginning of a line
<code>\$</code>	The end of a line
<code>\b</code>	A word boundary
<code>\B</code>	A non-word boundary
<code>\A</code>	The beginning of the input
<code>\G</code>	The end of the previous match
<code>\Z</code>	The end of the input but for the final terminator, if any
<code>\z</code>	The end of the input

Greedy quantifiers

<code>X?</code>	<i>X</i> , once or not at all
<code>X*</code>	<i>X</i> , zero or more times
<code>X+</code>	<i>X</i> , one or more times
<code>X{n}</code>	<i>X</i> , exactly <i>n</i> times
<code>X{n, }</code>	<i>X</i> , at least <i>n</i> times
<code>X{n,m}</code>	<i>X</i> , at least <i>n</i> but not more than <i>m</i> times

Reluctant quantifiers

<code>X??</code>	<i>X</i> , once or not at all
<code>X*?</code>	<i>X</i> , zero or more times
<code>X+?</code>	<i>X</i> , one or more times
<code>X{n}?</code>	<i>X</i> , exactly <i>n</i> times
<code>X{n, }?</code>	<i>X</i> , at least <i>n</i> times
<code>X{n,m}?</code>	<i>X</i> , at least <i>n</i> but not more than <i>m</i> times

Possessive quantifiers

<code>X?+</code>	<i>X</i> , once or not at all
<code>X*+</code>	<i>X</i> , zero or more times
<code>X++</code>	<i>X</i> , one or more times
<code>X{n}+</code>	<i>X</i> , exactly <i>n</i> 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	
$\backslash n$	Whatever the n^{th} capturing group matched
Quotation	
\backslash	Nothing, but quotes the following character
$\backslash Q$	Nothing, but quotes all characters until $\backslash E$
$\backslash E$	Nothing, but ends quoting started by $\backslash Q$
Special constructs (non-capturing)	
$(?:X)$	X , as a non-capturing group
$(?idsux-idsux)$	Nothing, but turns match flags on - off
$(?idsux-idsux:X)$	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.