INTERNATIONAL MESSAGING
ASSOCIATES

Internet Exchange Messaging Server 7 Site Planning Guide

White Paper

Document ID: IEMS7SITPLN001
Publication Date: June 2003
Website: www.ima.com

Table of Contents

Internet Exchange Messaging Server 7 - Site Planning Guide

INTRODUCTION

The purpose of this document is to provide email administrators, technology managers and information technology professionals an overview on how to design and implement messaging systems based upon Internet Exchange Messaging Server 7 (IEMS). This document is not intended to be a complete design or implementation guide but instead a general rollout guide based the current and projected needs of a site.

This document outlines the following areas:

- Internet Exchange Messaging Server Overview
- Capacity Planning
- Message Load Profiling
- Designing and implementing IEMS on a single machine
- Designing and implementing IEMS in a distributed environment
- Designing and implementing IEMS in an Intranet
- Frequently Asked Questions

As this document only provides an overview and does not detail the features and functions of IEMS components. It is, recommended, therefore, to read both the Internet Exchange Messaging Server 7 Principles of Operation, and the Internet Exchange Messaging Server Administrators' Manual. These documents are part of the IEMS 7 document set, which is made up of the following volumes:

- Internet Exchange Messaging Server 7 Principles of Operation
- Internet Exchange Messaging Server 7 Site Planning Guide
- Internet Exchange Messaging Server 7 User's Guide
- Internet Exchange Messaging Server 7 Installation Guide
- Internet Exchange Messaging Server 7 Administrator's Manual
- Internet Exchange Messaging Server 7 cc:Mail Connector
- Internet Exchange Messaging Server 7 Lotus Notes Connector
- Internet Exchange Messaging Server 7 Programmer's Manual

These documents and others can be found by visiting http://www.ima.com/documents or http://www.ima.com/support/faq.

OVERVIEW

The Internet Exchange Messaging Server (IEMS) is a highly modular and scalable open architecture system. It can be used from small single machine installations to fully distributed systems linking geographically distributed sites into a common set of logical domains (see Figure 1). Its various components can be run on a single machine or in a distributed environment.

IEMS 7 introduces a new integrated Anti-Spam approach to message reception and delivery. The MTA Pass-Through technology employed by IEMS 7 allows end users (message store accounts), individual distribution list maintainers, and connector modules to define their own security profiles independent of the rest of the system. At the same time the messaging system administrator can still define an overall global security policy, where some anti-spam measures will be handled directly by the MTA (such as reliable DNS-BL identified traffic). Other measures which may be desired by part of the user community, such as DNS-BL's with known high false positive rates (at the time of this writing, SpamCop and a few others have received a lot of industry coverage for their perceived indiscriminate listing practices) can then be passed through to the users for consultation on a case by case basis.

In most conventional messaging systems, security measures are employed on a system wide basis, making the choice of tools, such as DNS-BL's, critical. IEMS MTA Pass-Through technology changes this by allowing the administrator to be able to employ many more countermeasures, enabling only those that have been proven to be universally effective at the MTA level, and letting users pick and choose what additional measures they may or may not wish to apply to their individual message traffic.

Other IEMS modules include:

- **MTA / Preprocessor**
  The MTA is a message switch responsible for routing mail messages received by the Preprocessor to the intended channels. Upon receiving messages, the MTA temporarily stores the messages locally in a shared message queue while analyzing the recipient's address. It will either route the message to the recipient's local address or forward the mail to another MTA.

  The Preprocessor Unit is an integrated subsystem of the MTA. It is equipped with anti-spam and anti-virus plug-in modules to protect the system against viruses and spam mail. It incorporates an auto text insertion engine, providing the capability to insert disclaimers into messages passing through the MTA. Channel Action Matrices provide the system administrator with a flexible tool in configuring which plug-in modules should be run for a particular message based upon message flow through the system.

- **Directory Services**
  IEMS Directory Services are designed to effectively manage information about users, groups, mailing lists, alias processing and mail routing. It has a rich set of searching capabilities that makes directory lookup fast and efficient.

- **Distribution List Manager**
  The Distribution List Manager allows messages to be sent to list subscribers by simply submitting messages to a single address. It enables the system administrator to create electronic mailing lists that support the following features: mail blocking, automatic mailing list subscription and un-subscription, and setting the preferred delivery options. It also provides the system administrator with an option to accept or reject subscribers to the mailing list.

- **SMTPC** (**S**imple **M**ail **T**ransfer **P**rotocol **C**lient)
  SMTPC delivers messages to the Internet. It provides fast mail delivery by processing messages based on their priority weight and by assigning different processors for deferred and pending messages.

- **SMTPD** (**S**imple **M**ail **T**ransfer **P**rotocol **D**aemon)
  SMTPD listens for incoming messages on the Internet. It is capable of sustaining simultaneous SMTP connections by creating multiple threads, thereby minimizing delay in message delivery.

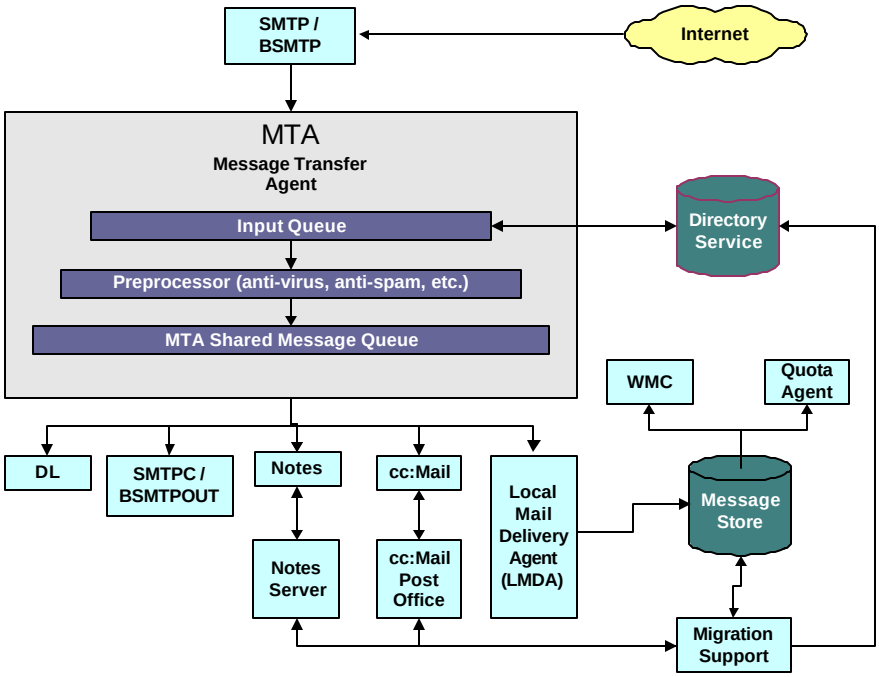

Figure 1: IEMS Architecture

- **BSMTP**
  Batch SMTP (BSMTP) tunnels messages so that they can pass through non-SMTP transports, such as POP3 (Post Office Protocol version 3). The original envelope and delivery information of each message is maintained across the tunnel.

- **Message Store**
  The Message Store acts as a dedicated mail repository for storing, retrieving and manipulating messages, while also enabling users to access their mailboxes via any POP3- and/or IMAP4-capable client. Users may also access their mail from the Message Store using the IEMS Web Mail Client, or any third-party application written around the Open Client API.

- **Bayesian Filtering / MailSort (LMDA)**
  Messages destined for a local user's Message Store account pass through the Local Mail Delivery Agent (LMDA) prior to delivery. The LMDA consists of the User Spam Controls, Bayesian Filter Engine, and the MailSort Engine (see Figure 2). Each of these three modules perform certain filtering and/or message filing operations on behalf of the user. Unlike similar operations that some mail clients use, these actions are performed by the

messaging system, and at the time of message delivery. Once these modules are optionally configured by the user, their actions are transparent, as their mail client is not involved, and the actions happen as soon as messages arrive.

The User Spam Control module looks for messages that have been tagged by the MTA as potential spam for one or more reasons. This can be due to DNS-BL tagging (with the offending BL or BL's identified), and/or as a result of content filtering. Users can choose for each control if they want to act upon the tagging or not, and an appropriate action to take (ignore, discard, or file in the user specified spam folder).

The Bayesian Filtering module utilizes a statistical technique for spam detection based upon the users database of offending spam messages. Each user will have a different database based upon message they have individually categorized as spam according to their wishes. Messages caught by the Bayesian Filtering module can be either discarded, filed in a spam or suspicious folder, or passed to the MailSort engine for further processing.

The MailSort engine performs simple header pattern matching for the purpose of automatic filing of incoming messages as well as a vacation utility that is capable of sending vacation notifications during periods when the recipient is away.

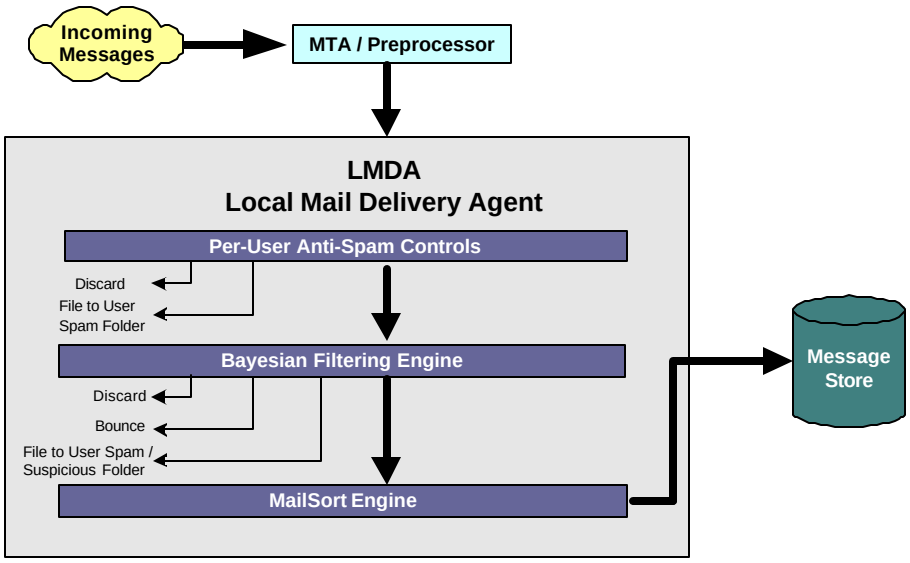

Figure 2: LMDA Architecture

Note: *The different IEMS components are also discussed in detail of the* ***Internet Exchange Messaging Server 7 Principles of Operation*** *and the* ***Internet Exchange Messaging Server 7 Administrator's Guide****.*

Aside from the components mentioned above, the following queues also contribute to the message transfer process.

- **Input Queue**
  Contains messages received from any of the configured input channels, including SMTP, BSMTP, and others. Messages in the input queue are retrieved by the Preprocessor Unit

for later processing and routing.

- **MTA Shared Message Queue**
  Temporarily stores messages inserted by the Preprocessor Unit after preprocessing the messages for virus scanning, spam control, among others. Later they will be retrieved by the respective output channel processors for delivery to the intended recipients or downstream MTA's.

Due to the IEMS modular design combined with a rich set of features, the resulting system can be configured to serve many different types of messaging requirements. Some of these include:

- A high performance **Message Switch**, acting as a hub for connecting disparate e-mail systems deployed within an organization.
- A full-featured, standalone **Internet Mail Server**, providing e-mail services directly to clients using Internet standards, such as SMTP (Simple Mail Transfer Protocol), ESMTP (Extended SMTP), BSTMP (Batch SMTP), POP3, IMAP4, and LDAP (Lightweight Directory Access Protocol).
- A **Gateway Server** providing seamless integration with either Lotus cc:Mail and/or Lotus Notes environments.
- A **Firewall Messaging Server**, providing front-end anti-virus, anti-spam, and attachment filtering for the entire organization.
- A **Distribution List Processor**, capable of handling both internal as well as external lists of different types.
- A **Public Web Mail Platform** – create your own web mail service using the ISP/ASP toolkit.

IEMS EDITIONS

The Internet Exchange Messaging Server 7 (IEMS) is packaged as a single distribution, however can be licensed depending upon the needs of the organization. Three Editions can be enabled depending on the type of license purchased. This flexibility allows sites to obtain just the features needed, while at the same time providing for a simple upgrade path. Editions can be upgraded without need for re-installation by simply applying a new license.

A summary of the features provided in the different IEMS 7 Editions can be found in the table below.

Feature	Free 3-User	Standard Enterprise	Professional Enterprise
MTA / Preprocessor	✓	✓	✓
Directory Services	✓	✓	✓
Message Store	✓	✓	✓
Web Mail Client	✓	✓	✓
MailSort	✓	✓	✓
IMAP4 Server	✓	✓	✓
POP3 Server	✓	✓	✓
Anti-Spam: Content Filtering	✓	✓	✓
Anti-Spam: Bayesian Filtering	✓	✓	✓
Anti-Spam: Multiple DNS-BL	✓	✓	✓
Anti-Spam: Header Filtering	✓	✓	✓
Anti-Spam: Connection Control	✓	✓	✓
Anti-Spam: Sender Site Verification	✓	✓	✓
Anti-Virus		✓	✓
Web Folders (Online Storage)		✓	✓
Web Online Bookmarks		✓	✓
SMTPD Message Flow Control		✓	✓
SMTP Authorization		✓	✓
SMTPD SSL Support		✓	✓
Automatic Attachment Removal		✓	✓
Automatic Disclaimer Insertion		✓	✓
BSMTP Client		✓	✓
MTA Pass-Through Capabilities			✓
Distribution Lists			✓
Distributed Operations			✓
Multi-Domain Administration			✓
cc:Mail Connector			✓
Lotus Notes Connector			✓
MS Exchange Migration Tools			✓
Calendaring / Scheduling			✓
BSMTP Server			✓
ISP / ASP Toolkit			✓
Open MTA API	✓	✓	✓
Open Web Mail / Submission API	✓	✓	✓
Users	3	50 - 250	250+

IEMS 7 Edition Summary

LICENSING AND ENABLING OF IEMS

After installing the Internet Exchange Messaging Server, a license certificate containing license information is required in order to fully activate the software. License certificates can be requested from any authorized license manager (either IMA or who you purchased Internet Exchange from). After registration, a certificate containing information on the licensed modules is issued. This certificate identifies and validates the user when installing the license key.

There are three types of licenses for the fully functional version of Internet Exchange: Evaluation, Interim, and Permanent. Evaluation licenses are time-limited licenses (normally 30-days) and are used with the freely available evaluation copies of Internet Exchange. Once a registration form is received, the authorized license manager generates this license and provides it to the user. Evaluation licenses may be also be obtained over the Internet by visiting *http://www.ima.com/iems/reg.html*, as well as through the IEMS Installation Registration procedures (assuming a live Internet connection).

Interim licenses are also time limited, except that an interim license can be updated to a permanent license at a later date. These licenses are used for serialized or purchased copies of Internet Exchange.

Permanent licenses are permanent and never have to be renewed. Permanent licenses are usually provided to holders of Interim licenses after the Internet Exchange software has been purchased. They are applied only to serialized copies of the software. Only an IMA authorized license manager generates the permanent license.

IEMS Licenses are tied with the FQDN, or fully qualified domain name of the machine IEMS is running on. Licenses can be upgraded at any time without requiring re-installation of the software. This makes both the migration from evaluation to full production use simple and painless, as well as any potential license upgrades in the future. Site planners can safely license for today's messaging needs knowing that future upgrades can be made with minimal disruption to the users.

Licenses are not required for the Free 3-User mode of IEMS.

How IEMS Components Work Together

IEMS is responsible for sending and receiving messages over the Internet using either SMTP or BSMTP protocols. Once messages are received, the MTA then routes messages to an appropriate output channel (SMTP, Message Store, Distribution List Manager, cc:Mail/Notes Connector, among others) and performs preset preprocessing on each message. Message routing is facilitated by data stored in the IEMS Directory.

Messages received from the Internet arrive via either SMTPD or BSMTP (BSMTPIN) input channels (see Figure 3 below). These messages arrive in the Input Queue. From the Input Queue, messages are fetched by the Preprocessor, performing directory look up from the Directory Server.

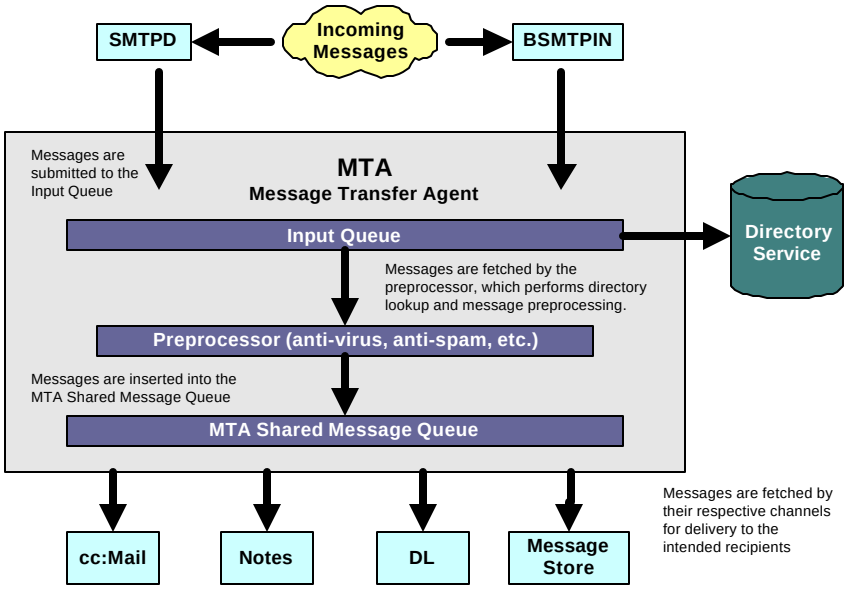

Figure 3: Message flow for incoming messages

If the domain name (e.g. *@ima.com*) of the recipient address is defined as local in the Directory, the Preprocessor performs further look up's, this time to obtain the channels / connectors needed to route the message within the local MTA. After this information is acquired, a new message envelope with proper routing information is created for the message. When proper routing information for all messages is determined, messages undergo Preprocessor operations, including virus scanning, spam checks and disclaimer insertion depending on the configuration. After preprocessing is complete, messages are inserted into the MTA Shared Message Queue where they are later fetched by their respective channel processors for delivery to their intended recipients.

If the domain name of the recipient address is not defined as local (e.g. *@ima.com*), the Preprocessor will pass the message to the default output channel (usually SMTPC), which will relay the message to the next downstream MTA for further routing.

Messages created within the IEMS domain, bound for the Internet (see Figure 4 below), are directly submitted by the MTA to the Preprocessor, which will then perform virus scanning, spam checks and disclaimer insertion on messages depending on the system configuration. After preprocessing, messages are inserted into the MTA Shared Message Queue to be picked up

later by the appropriate output channel, such as SMTPC or BSMTP. SMTPC routed messages are then sent to the next MTA on the Internet for eventual delivery to the intended recipients.

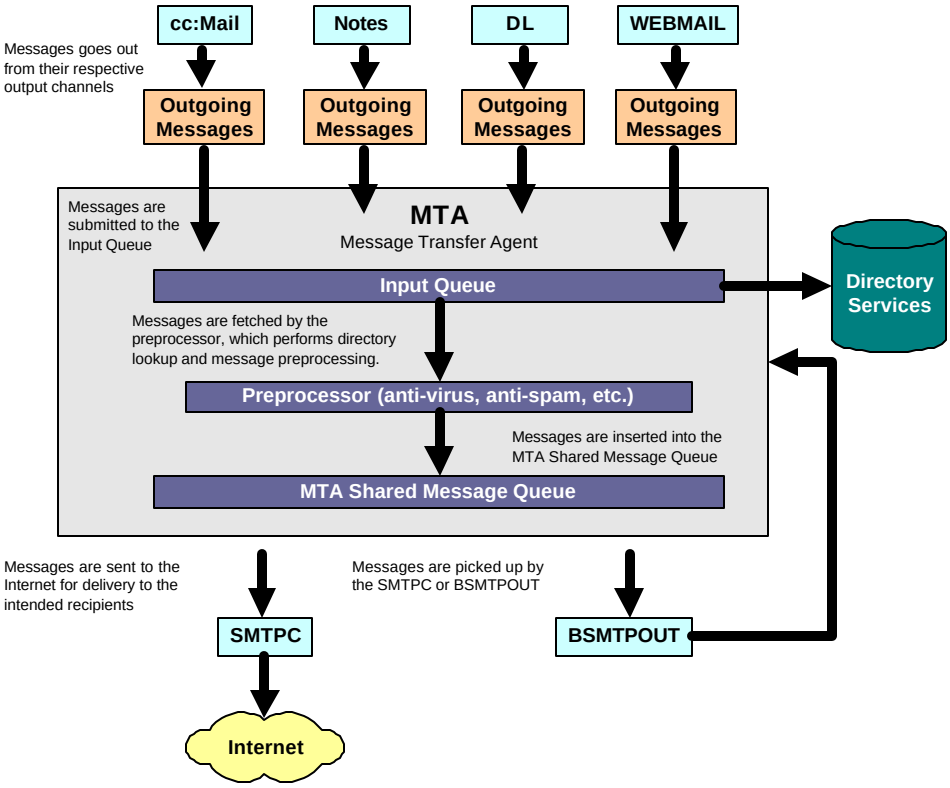

Figure 4: Message flow for outgoing messages

CAPACITY PLANNING

When either designing a new network, or analyzing a current setup for expansion, traditional computer system and network architectural procedures should be employed. Issues particular to IEMS planning include storage management, and performance related issues on both an individual computer as well as network basis. In general, the focus should be on the identification of potential bottlenecks in the system, and how to minimize or eliminate them. In some cases, understanding where these potential bottlenecks are can save in expenses elsewhere, such as the utilization of cheaper disks where network bandwidth is the limiting factor.

Message Store

IEMS use of disk storage, with the exception of the Message Store is quite modest. Basic Message Store (where the user mailboxes are stored) usage should be fairly simple to compute – simply estimate the average amount of space required per user, and multiply by the total number of users, allocating an appropriate amount of space for expansion, as usage increases. For large sites, this can be better controlled through Message Store quotas. Regardless if quotas are used or not, the Message Store reports can pinpoint how much storage each user is utilizing at any point in time. Figure 5 below shows the system architecture of IEMS Message Store.

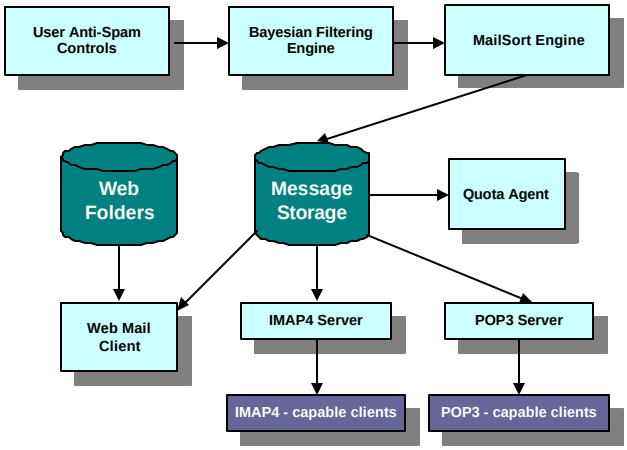

Figure 5: Message Store system architecture

For sites that employ Bayesian Filtering, additional storage is required in order to store the Bayesian Filter databases. IEMS stores these in a special directory called ".Bayesian" under each Message Store user home directory. These can over time can consume considerable disk space, so planning for this additional storage is essential. When using quotas, the space allocation specified applies only to message files, and not the filtering databases. There are two databases that are used by the filter – one that stores characteristics of good email, and another that stores characteristics of mail identified by the user as spam. In general, over time the Bayesian databases will tend to grow in size, the extent of which depends upon message volume. In addition, these databases retain their size even when users remove email, so for users that tend to remove a high percentage of mail rather than store it, the filter databases can actually grow to consume more space than retained messages.

Initially, Bayesian databases tend to grow linearly with message volume as they start to learn. Over time the increase in database usage will level out after the system has a better representation of good and bad messages. Initially, in order to approximate the amount of total disk usage for a given user, a good rule of thumb is to add 50% for the good message database,

and 25% for the spam message database, or a total of 75% of expected message volume per user. Operational data regarding growth over time of more mature databases is still being obtained at the time of this writing, but it clear that database size tends to slow in growth over time.

Aside from the capacity requirements for a Message Store, the performance of this storage can make a big difference in perceived end-user performance. Heavily used Message Store machines are usually I/O bound. Machines that are hosting many simultaneous users should attempt to optimize the disk subsystem for speed. This typically means using fast SCSI devices, preferably spread across multiple devices (to take advantage of the detached command queuing nature of SCSI devices, and even RAID systems). In these environments, separating the total storage across multiple machines can also help. Assuming basic physical memory requirements are met, as usage increases, the ability to spread accesses across multiple disk spindles, either via separate filesystems or RAID, will dictate the number of simultaneous users that can be accommodated.

The last issue affecting the Message Store (other than reasonably powered CPU's) is memory. Each POP3 and IMAP4 connection incrementally consumes a fixed 120K of memory. In addition, to this fixed amount, approximately 100 bytes is allocated for each message present in the selected mailbox (IMAP4). So for a sample installation that needs to support 50 simultaneous connections, with an average mailbox size of 200 messages, you would need 50 * 140K or about 7MB extra physical memory in addition to the needs of the rest of the system. It is important to scale physical memory to match the normal operating requirements. If the machine runs out of physical memory, and has to page or swap, this will considerably slow the system, while at the same time adding additional requirements to the already heavily utilized disk I/O system. A good rule of thumb here is to calculate the memory requirements for normal operation, and then factor in a reasonable (20+%) margin on top of this when determining physical memory requirements, to cover most peak demand times smoothly.

Directory Server

The IEMS Directory also utilizes storage. For all but the largest and most dynamic sites, this storage is small. As the Directory Server caches information, physical memory should be plentiful on this machine to result in quick replies to queries. Disk access speeds and capacities are not critical, however network responsiveness is. Fast and reliable network boards are crucial to fast response times in congested networks.

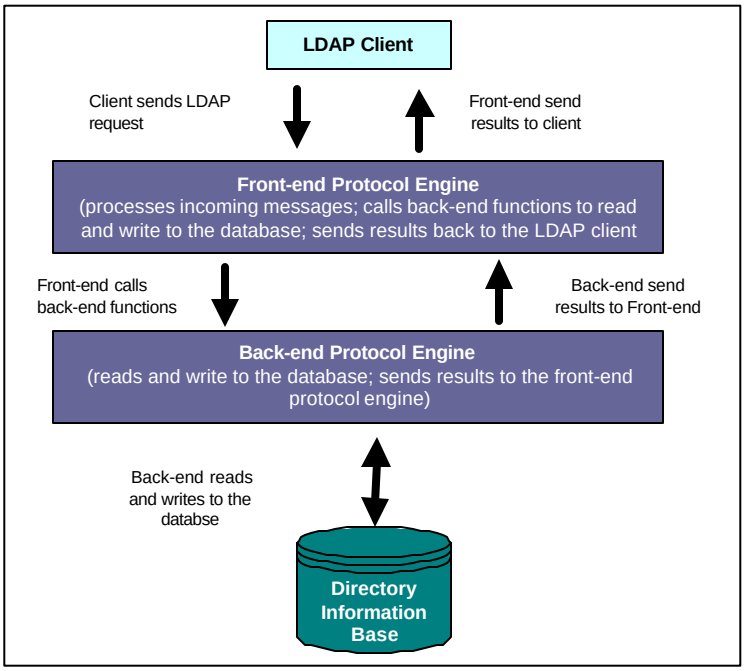

Figure 6: Directory Service system architecture

IEMS Queues

The remaining storage that IEMS uses is for its various queues. These include the SMTPD / SMTPC inbound / outbound queues for Internet bound messages, as well as the Shared Message Queue. The amount of storage needed for these queues can vary considerably depending upon the traffic patterns at a given site. For outbound SMTP, simply estimate the maximum number of queued messages that are expected, and them multiply by the expected average message size. The Shared Message Queue usage should be minimal, as local channel processors under normal operating conditions will consume these messages almost as fast as they are created.

As disk storage is relatively inexpensive these days, a good rule of thumb is to allocate at least 2 to 4x the maximum size needed for outbound SMTP queue. If the MTA log files are located on the same volume, you'll need to factor this in as well. If at all possible, either automatically remove log files after sending to postmaster, or place them on a separate volume from the queues. This way, if they go unattended for some period of time, and consume the remaining disk space on the volume, it will not affect message delivery.

In general, disk access speed is not critical for the IEMS queues, as network bandwidth will be the primary bottleneck. If your network connection is faster than T1/E1, this may not be the case however.

Internet Connection Speed

Internet connection speed is usually not a factor in the exercise of capacity planning, other than making sure that the available bandwidth can adequately handle the expected traffic for the site. This includes not only mail related protocols, but any other services that are required, including http, ftp, and others.

One exception to this however is for sites that are connected by slow links. If your site for instance is connected by say a 56K dialup connection, it is safe to assume that bandwidth utilization at any given point in time will be low (even it if is 100% of the available circuit). With

this in mind, there is no need for disks or systems handling SMTP processing to be extremely fast. SCSI disk subsystems can be replaced with significantly cheaper IDE systems, with no effect on overall system performance. Unless you are planning to replace that 56K dialup with multiple T1/E1 circuits, and have your messaging traffic increase by a similar amount, it is safe to use smaller systems for these purposes.

Local Mail Delivery Agent

The Local Mail Delivery Agent (LMDA) is responsible for taking messages from the Shared Message Queue and delivering them to local Message Store accounts (see Figure 2). An integral component of the LDMA is the MailSort engine and Bayesian Filter engine, which handles the automatic spam scanning and sorting of individual email at message delivery time. Depending upon the extent to which the local user community makes use of the Bayesian and MailSort filters, this can represent a significant CPU load on the system. This is one of the few IEMS modules where CPU speed can be a factor. Basically, the faster the machine, the faster MailSort will be able to process incoming messages.

SECURITY FRAMEWORK

System administrators are often caught in the middle of conflicting sets of requirements. On one hand, it is their responsibility to protect their organization and systems from outside (and sometimes inside) attacks from virus infected messages as well as spam. At the same time, they serve the users of these systems.

Traditional spam fighting techniques are performed by the MTA based upon policies set by the administrator. These global policies normally are set to ensure the maximum protection for the organization with minimal impact on the end user. In the case of spam detection and handling, the definition of what constitutes spam can vary widely from community to community, as well as from user to user within a single organization. Sales and marketing related messages may be very welcome in a sales group, while not being tolerated in a nearby engineering group. Advertisements pitching lower mortgage rates may be undesirable by most but a small group of people looking to purchase a new home. Viagra advertisements and other personal enhancement types of advertisements may not be at home for any users, especially if the site caters to the young or corporate users.

To assist the IEMS administrator in providing for both system security as well as keeping the collateral damage associated with spam detection and handling to an absolute minimum, several new tools can be applied. These can be applied on a system wide basis (global) and/or on an individual basis. Some tools such as virus scanning, some SMTP connection controls, site-wide blacklists, and SMTP Authentication affect an entire site and are global in scope. Others such as Bayesian Filtering and mail sorting based upon pattern matching are tools end users can apply.

Other tools such as DNS Blacklists (DNS-BL), header analysis, and message content analysis occur within the MTA, however can be acted upon either as directed by a system security policy, or end user security policy. The ability for end users to be able to set security policies on actions normally only associated with system activities is made possible by the IEMS MTA Pass-Through features, which allow for the optional tagging of suspect messages by the MTA. The local mail delivery agent (working on behalf of the user) can then act upon these tagged messages later. This allows for both much more aggressive checking at the MTA level, as well as far more control of what messages are rejected at the user level (see Figure 7).

System Wide Security Settings

It is usually desirable to apply some security measures to all messages that pass through your systems. Some of the tools that by their nature are applicable to all message traffic include the following:

- Anti-Virus Scanning
- Reverse DNS Lookup of SMTP Data
- SMTP Authentication
- Loop Detection
- Mail Relay Control
- SMTP Sender Checks

The application of the tools listed above apply to all messages that pass through an IEMS system, and cannot be overridden by user preferences. The only exception is anti-virus scanning, whose behavior can be modified through the anti-virus channel action matrix settings.

MTA Pass-Through

IEMS 7 Pass-Through technology allows the system administrator to be able to perform MTA level checks on messages, and then to optionally defer any action until being handled by an agent controlled by the end user. These agents are typically output channel processors, such as the Local Mail Delivery Agent, the Distribution List Processor, and others. As not all output channels are capable of handling deferred actions (such as the cc:Mail and Notes connector

modules), the administrator can define default actions to be performed on a channel by channel basis, which will then be carried out by the preprocessor.

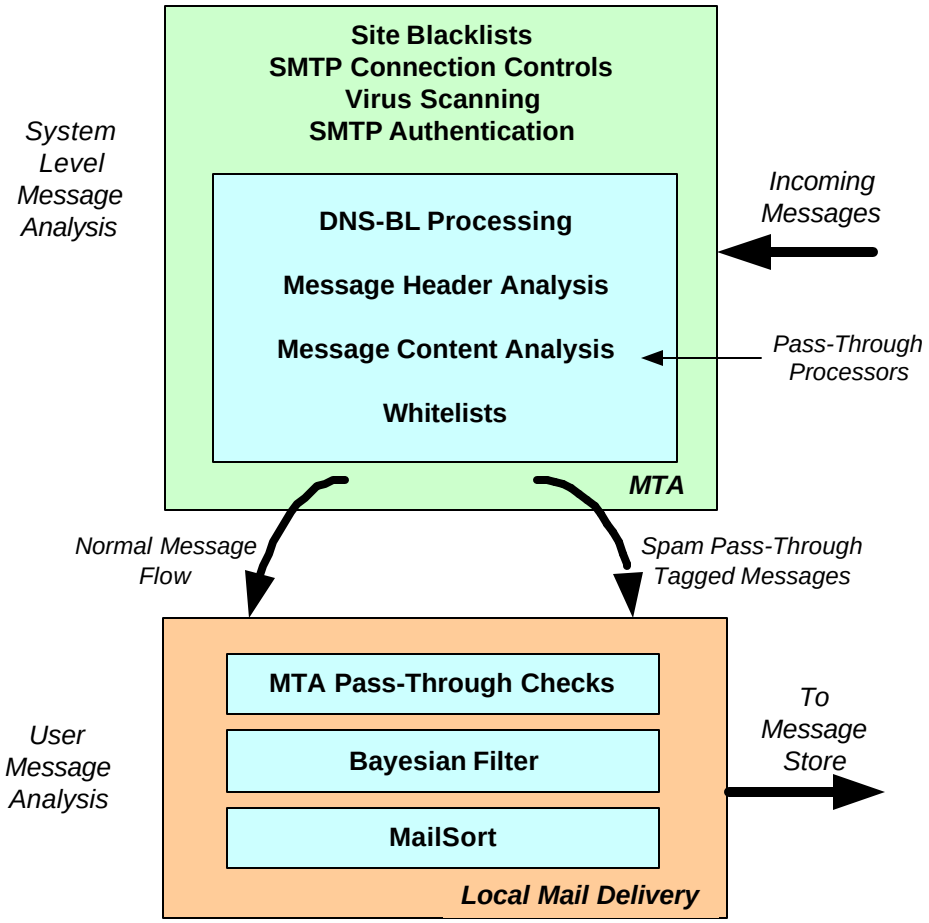

Figure 7: IEMS Pass-Through Architecture

The primary Pass-Through processors in the MTA consist of the DNS-BL lists and the Content Filtering system. The DNS-BL's are blacklist databases that are consulted at SMTP connection time by SMTPD. Each configured list can be setup to either immediately reject mail from the identified remote MTA, or accept the mail, and tag it as having been positively identified by the corresponding DNS blacklist. Each list in the system is configured separately, as the nature of the different lists as well as reliability of the various blacklists vary widely. Normally, an administrator will choose to act directly upon messages that have been identified only by very reliable lists. Examples of some of these include a few of the open relay lists. If the chances of false positive identification of spam by these lists is sufficiently low, then it is reasonable to remove the message before it can even get into the system, reducing overall MTA loading.

For other blacklists whose false positive identification rates are significantly high (SpamCop is one recent example), the administrator can choose to accept messages from hosts identified by these lists, but tag them so that the user can at their choosing perform special processing on these messages. In the case of the local mail delivery agent, this involves message removal, bouncing, automatic filing in a user defined spam folder, or ignoring the tagging altogether.

For distributed configurations, IEMS can be setup so that Pass-Through information is retained across MTA's, while at the same time stripped out before messages are sent outside the system. This capability can be very useful in situations where a Firewall MTA is used in a DMZ network to handle frontend tasks such as Anti-Spam and Anti-Virus scanning, while delivery is handled by an interior network MTA.

Direct End User Controls

The Local Mail Delivery Agent (LMDA) and the Distribution List Engine perform actions on behalf of their respective users (Message Store, and Distribution Lists). Both of these channel processors can be configured on a per DNS-BL basis as to what actions to perform. The LMDA components are shown above in Figure 7. In addition to MTA Pass-Through processing, the LMDA can configured to perform Bayesian messaging filtering on behalf of the user. This filtering technique utilizes per-user message databases made up of user identified spam as the basis for its message blocking. Users, using either the Web Mail Client, or any IMAP client can place received SPAM into a special folder where the system can later process and update the individual Bayesian Filter databases. After an initial learning phase, accuracy rates for Bayesian filters can exceed 98%.

The combination of SMTP controls, Content Filters, Bayesian Filters, DNS-BL's, and the extension of these controls to the end users allows for an extremely flexible protection system, designed to block the maximum number of problem messages.

MESSAGE LOAD PROFILING

The previous section on *Capacity Planning* provided an overview to the resource requirements of the various IEMS modules. Basic factors that can affect performance, as well as the primary limiting factors for the various modules were discussed. When planning for the computer and network hardware requirements, the following should be matched for the load requirements place upon them:

- Physical Memory (DRAM)
- Disk Subsystems
- CPU Type / Speed
- Network Bandwidth
- Operating System Limits

Network bandwidth and disk subsystems were covered in the previous section. In general, it is always best to go with high reliability components. For network interfaces, given the current state of the market, a minimum configuration should be at least 100 MB speed. Be careful however that only reliable components are used. For example IMA tests uncovered many intermittent problems when RealTek 8139 based 100BastT network interfaces were used under various Linux systems (regardless of the NIC manufacturer). Cost savings of a few dollars at this level simply are not effective, as reliability and subsequent outages and debugging costs dwarf any short-term savings. Internal tests have shown that 3Com and Intel NIC's provide good reliability, however other solutions may also provide similar or better results.

In order to best plan and design an IEMS based messaging environment, current and anticipated message loads must be matched against the network and hardware resources used to implement a given solution. While it is desirable to be able to establish general guidelines that indicate how many users a Message Store will accommodate, the truth is that different usage profiles result in significantly different numbers. For example, an ISP providing web based email access where the users only occasionally connect to download a few small messages once every few days will have significantly different system requirements from an organization of heavy mail users, connecting via IMAP4 and each dealing with 100MB of mail on a daily basis. As a result, it is necessary to be able to establish a simple model of usage, and match this against the IEMS architecture and operating system limits to come up with network and hardware recommendations.

Critical Paths

As the performance requirements of different parts of the IEMS system are not always the same, it is necessary to not only identify total traffic through a site over a given period of time, but also to determine the maximum expected loads at key points in the system. The most important loading figures to estimate are the following:

- Total traffic into the system over a 24 hour period
- Average load (bytes transferred) by email clients (web, IMAP4, POP3)
- Anti-Virus processing
- Content Filtering
- Bayesian Filtering
- Average Number of concurrent SMTPD and SMTPC Connections
- Average Number of concurrent POP3 and IMPA Connections
- Average Number of concurrent Web Mail Client Users
- Additional Message traffic attributable to MTA Pass-Through deferred blocking

For Internet connected IEMS sites, the total traffic into an IEMS system is approximately the same as the load placed on SMTPD. Unless the site is being used primarily for distribution list expansion for non-local recipients, the SMTPD load will approximate the load placed on the Shared Message Queue, and later the LDMA. If extensive use of distribution lists is expected, then the LMDA load may be significantly different – and should be taken into consideration when estimating LMDA loading.

In test laboratory environments, a single IEMS machine running on modest hardware (AMD K6-2/500 processor, 128M memory, 20G IDE disk) is able to sustain data input to SMTPD data at the rate of 95 MB / minute (not taking into consideration any other tasks that may be required on this machine).

LDMA throughput on similar hardware, where no other processing was being done has shown sustained throughput rates of 110 MB / minute. In both cases, the test environment was sending messages of 250K in size through the system. No MailSort filters were used for this particular LMDA testing, so if extensive MailSort usage is expected, this LMDA throughput estimate should be revised downward.

IMAP4 throughput in this environment has been shown to be able to sustain throughput rates of 90 MB / minute. In a test laboratory environment, 18 simultaneous connections were maintained where each connection continuously downloaded new messages. No other IEMS processes were active in this test environment. For POP3 usage – the throughput figures should be the same or better.

Module	Limiting Factors	Expected Performance
LMDA	CPU	110 MB / min
SMTPD	Network	95 MB / min
Directory	CPU	
IMAP4 / POP3	Disk I/O	90 MB / min
Anti-Virus	CPU	Vendor specific

For sites that choose to run anti-virus on messages that pass through their systems (highly recommended), this can place a significant additional workload on the system. IEMS allows the administrator to configure one or more anti-virus packages into the system. For more than one configured anti-virus solution, the processing is chained. As would be expected, moving from one anti-virus plug-in to two would roughly double the processing load per message.

Heavily loaded sites should consider offloading anti-virus processing from the MTA machine to a separate machine. This way, the hardware / software requirements of virus scanning can be separated from that of the MTA, while at the same time removing this processing overhead from the main MTA.

As with anti-virus scanning, message content filtering takes place within the MTA Preprocessor module. The SpamAssassin plug-in, like the anti-virus plug-in supports a client – server mode of operation, and it is recommended in heavily loaded system to offload this processing as well.

In addition to content filtering done within the MTA Preprocessor, Bayesian Filtering can take place within the LMDA prior to final delivery to a users mailbox. Bayesian filtering works by consulting a user's database of known spam patterns, and then applying this data against each message that is subject to delivery to the message store.

Sites that run with MTA Pass-Through options enabled are likely to encounter additional loads within the MTA and any configured Preprocessor modules (anti-virus, content filtering, etc). The MTA Pass-Through technology works by instead of rejecting email at the point where a potential problem is detected, it instead tags the message and lets it continue through the system. For sites configured in this manner, the appropriate connector modules are then responsible for how they handle such tagged message.

While this approach to message tagging is essential in order to allow connector modules such as the LMDA (acting on behalf of individual users) to make their own decisions, the deferment of handling ensures that the messages must now flow through the MTA Preprocessor and undergo normal message processing. In the case of DNS-BL tagged traffic, this can also result in additional network usage, as this traffic would otherwise be dropped before the sending site has the chance to transfer it over the network.

Memory Requirements

Regardless of configuration, each IEMS machine must have enough physical memory in order to properly operate. Without adequate memory, the machine will swap or page, resulting in the machine spending more time moving processes to temporary disk storage than in handling the messaging tasks which it is setup to perform.

The memory requirements of a given machine is made up of the following components:

- Base Operating System
- Base IEMS Components
- Incremental IEMS Component Requirements
- Non-IEMS Applications

Due to the difficulty in determining non-IEMS application processor and memory needs, it is strongly recommended that non-IEMS applications not be run together with IEMS. If you do decide however to run them in parallel, their machine utilization should be added to what is anticipated for IEMS to require.

Operating System

Every operating system needs a certain resources in order to function properly. This minimum memory requirement for Windows 98 and Linux machines is 64MB. For Windows XP, NT and 2000, this should be set to 96MB. Please note that this is enough to get basic operating system services up and running, but not for any additional applications. For Linux, it is possible to run in less than 64MB, but this value is highly recommended as a base requirement as Linux uses the extra memory for network buffers, and other buffers useful for IEMS. This does not include memory needed to run the X-Windows graphical front-end under Linux.

For installations that need to support a significant number of concurrent users, it is important to take into consideration limits that may be imposed by the underlying operating system. In many cases these values can be adjusted and the system tuned to support higher loads, however the default installation configuration may not be adequate for some situations. In particular, the operating system will impose limits on several critical resources, including:

- Number of Open File Descriptors
- Process Table Size
- Number of Allowed Threads
- Number of TCP Connections (sockets)
- Network Timeout Values

Other issues can also impact performance, including:

- File System Tuning (including disk layout)
- Service Maintenance (removing unnecessary services)
- Memory
- Logging Levels

At the time of this writing, IMA engineers are working on profiling several of the more popular operating systems so that a more precise method can be established for determining hardware and software requirements for large installations. This will be documented in the next version of this document.

ISP / ASP CONSIDERATIONS

Sites looking to implement a messaging solution for an ISP or ASP environment have a unique set of requirements to meet. In particular, the messaging solution must be able to handle several domains within a single system. While IEMS has always been a multi-domain solution, Version 7 adds additional administrator controls and capabilities designed to make the life of the ISP / ASP much easier. These include command line tools for local account creation, domain based administration, automated user account creation, user password recovery, and SMTP Authentication for remote users.

Account Creation Tools

New user accounts are normally created using the IEMS Administrator Interface to the Message Store. This is a convenient and simple way to add a single or a few accounts at a time. In situations where many accounts need to be created at the same time, or using other applications as a front end, the primary IEMS Administrator screens are not the optimal solution.

IEMS 7 includes the ***iemsuser*** user creation utility, which can be used inside of administrator scripts or called by third party applications. The *iemsuser* utility is available for both Linux as well as Windows versions of IEMS. This utility can be used to create, delete, or list Message Store users.

Domain Based Administration

For large sites, or those who are managing domains on behalf of other organizations, it can be desirable to delegate portions of the administrative responsibilities. This can take the form of distributing the administrative load across a single organization, or allowing customers whose domains a site is handling to manage portions of their own messaging environment.

IEMS Version 7 introduced new domain based administration tools which allow the main system administrator (what IEMS refers to as the ***super administrator*** to be able to setup domains whose message store account administration can be delegated to a domain administrator. System wide settings, such as the MTA Anti-Spam and Anti-Virus, routing, etc, are still controlled by the Super Administrator. The Domain Administrator can setup domain accounts, set quotas, and manage domain based distribution lists.

Automated User Account Creation

For installations where users are allowed to establish their own accounts, IEMS 7 provides new tools to facilitate this. The Account Registration system can be used to ease the administrative overhead in account creation for organizations on one hand, or as the basis for providing a free email service, with web based automatic account creation.

The system provides a form where the new user can enter their name, desired email address, and password. If enabled by the administrator, a secret question and answer section is also included which works with the new Password Recovery feature.

Password Recovery

When enabled by the system administrator, the Password Recovery system, introduced in IEMS Version 7, provides a mechanism where users can automatically recover forgotten passwords. When enabled, the appropriate question and answer fields in the Account Registration form are displayed. Questions can be determined by the administrator, or supplied by the user. The question and answer are stored in the user's Message Store account. Should the user forget his account password at a later date, they can retrieve it after supplying the correct answer to the

secret question. Options in the Web Mail login are provided to easily retrieve this information on Password Recover enabled systems.

SMTP Authentication

In most corporate environments where employees are either using IMAP4 connected clients, or the Web Mail client, it is simple to configure systems to be able to relay outbound mail on behalf of local users. For users who are remote, and are connecting through foreign networks, or whose identity cannot easily be established by their network connection, allowing the relay of mail for these users can be extremely dangerous. The reason is that it is impossible to determine the difference between friendly local staff and spammers looking for open mail relays.

The solution to this problem is the use of authentication within the SMTP protocol. STMP is the language spoken by mail clients such as Outlook Express, Eudora, and many others when they need to send mail. After a user composes a mail message, the mail program tries to talk to a friendly mail relay that will take responsibility for the transmission of the message over the Internet.

If the client is communicating over a trusted network (say a local area network), then most likely the mail relay can assume the client is trusted based upon it's network id. If not, then the client can first authenticate itself to the mail server using SMTP Authentication. Once the IEMS mail server has established the identity of the remote user, it can allow relay of messages. This is particularly useful for people on the road that need to dial in from random places, and still get their mail through.

IEMS IN A SINGLE MACHINE ENVIRONMENT

IEMS may be installed on either a Single Machine or in a Distributed Environment. Hardware requirements can vary depending on the number of users and bulk of incoming messages. For small sites with few simultaneous Message Store users, or those that do not run anti-virus software, this is the preferred configuration, as it the most cost effective, and simple to setup. For small sites that need to run Anti-Virus scanning, it is strongly recommended that the anti-virus scanning be moved to a separate machine, so as to not interfere with Message Store response time. More on this configuration can be found in the following pages.

On a single machine, all IEMS components such as the MTA, Preprocessor, Directory Server, Message Store, among others, will reside and run on a single machine. To know more about these components, please read Chapter 1 of the Internet Exchange Messaging Server 7 Administrators' Manual.

Figure 8 shows the network architecture of IEMS on a single machine. This is the most typical type of installation for small sites.

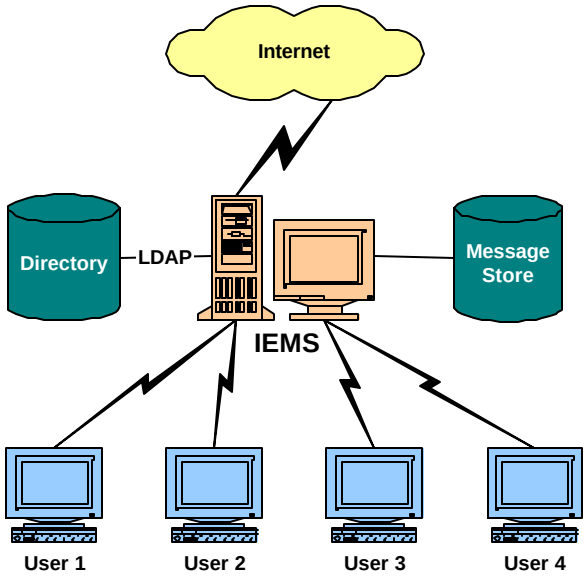

Figure 8: IEMS on a Single Machine

The "Single Machine" installation is simple to install and works seamlessly, especially for small organizations. As the messaging load increases, and performance starts to be affected, the IEMS server can either be upgraded to a faster machine, or the functionality of IEMS can be broken apart and spread across several cooperating machines on the local area network.

IEMS IN A DISTRIBUTED ENVIRONMENT

In a distributed environment, IEMS may be installed on multiple Linux and/or Windows machines. Unlike the "Single Machine" installation, different components in a distributed environment are installed on separate machines. See figure 9 below for a sample configuration.

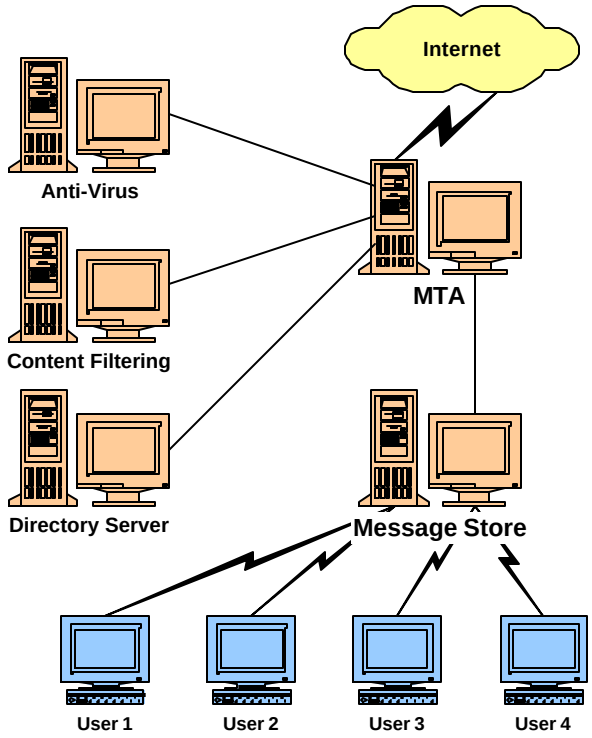

Figure 9: IEMS in a Distributed Environment

The following components can be run on their own machine, or together with other modules running in a distributed mode:

- smtpc / smtpd
- Anti-Virus
- Directory Server
- Message Store
- cc:Mail Connector
- LMDA
- Preprocessor
- Lotus Notes Connector
- Content Filtering

With the exception of the SpamAssassin Content Filtering plug-in, a *Professional Enterprise* Edition license is required to enable distributed operations.

When first expanding from a single machine configuration, the most effective functions to run separately are usually anti-virus and the Message Store. Anti-virus processing tends to be very CPU intensive. Separating this out makes sense if you want to isolate the processing load from other functions, such as the Directory Server or Message Store – both of which are expected to provide services to the rest of the system and end users. For Linux environments, breaking out the Anti-virus also is necessary if Windows based solutions are to be integrated with an otherwise Linux-only environment. For example, the Anti-virus can be run on a Windows 98 machine, while

everything else is Linux, taking advantage of more extensive anti-virus solutions that may be available only for Windows environments.

In managing a messaging environment, back-end functions, such as the MTA, SMTP subsystems, and anti-virus have to run reliably with minimum delays. As long as the total system can keep up with the overall message flow, users will tend to be happy. Real time loading factors on backend machines are not usually a factor, as long as overall message flow can be adequately maintained.

The same however cannot be said for front-end functions, such as the Web Mail Client, IMAP4, and POP3 services. These services work directly with client applications on the desktop, such as email clients (Outlook, Eudora, Netscape Communicator, etc), or a web browser. A delay of a minute or two while a backend machine is doing anti-virus scanning may be acceptable, but it generally is not considered an acceptable delay for end user applications, where the user is in front of the machine literally waiting for a response.

As your installation grows in size or in message load, the response time of the client applications can start to be affected. When this starts to become evident to the users, you may want to consider moving the Message Store away from the rest of the system. This will tend to isolate this part of the system from other modules that can consume significant amounts of system resources from time to time.

If this is the first time the Message Store has been separated, both the LMDA and other Message Store modules such as the IMAP4 / POP3 servers can be moved together. In time as the system continues to grow, you may want to consider separating the LMDA, as it can consume considerable CPU resources in heavy message load environments.

IEMS In An Intranet Environment

For sites that outgrow the basic distributed setup discussed above, it is possible to link together different IEMS systems (single machine and/or distributed) into a much larger messaging solution. For solutions of this type, message routers are usually used to isolate the Internet from the rest of the organization's messaging environment. Internet related MTA functions, such as anti-virus scanning, and anti-spam processing are done at this level. Departmental servers are then used to provide services to local user communities. As these servers are linked to the IEMS message routers, the need for them to perform anti-spam and anti-virus scanning is reduced, or potentially eliminated completely. This results in much simpler administration for the growing organization, as anti-virus and anti-spam administration is centralized.

An sample configuration can be seen in the figure below:

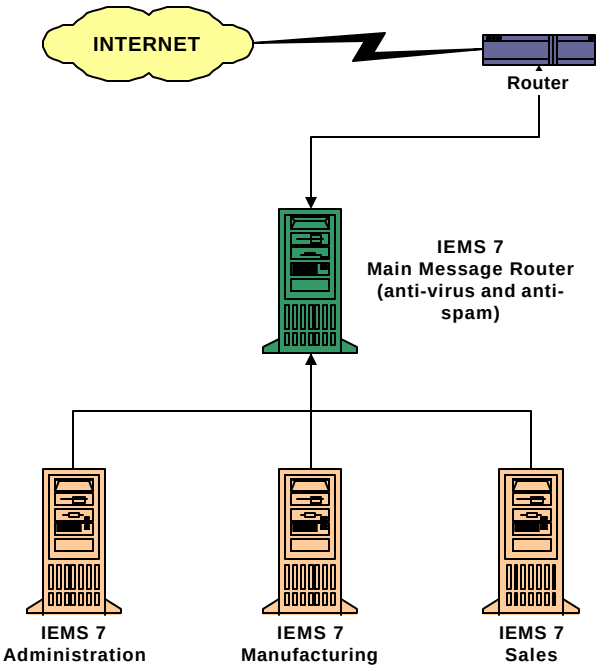

Figure 10: Multiple IEMS Systems on a Corporate Network

In this configuration, messages from the Internet arrive at the main message router. Here, common messaging processing such as anti-virus, anti-spam, and auto-text insertion takes place. Once the basic preprocessing is complete, messages are then routed via SMTP to the appropriate department server.

The message router system can also be used as a centralized distribution list processor. This has the effect of centralizing the administration of maintenance of organization wide distribution lists, while isolating distribution list processing from the department servers.

The above setup is also highly recommended in situations where department servers need to be isolated via a firewall from the rest of the Internet. Services provided on these machines can then be selectively made available through the firewall.

Once received by the departmental servers, messages can be stored in the local Message Store for later retrieval. Messages sent by local email clients can be sent either to the local SMTPD process, or directly to the main message router, depending upon local policy. If the departments are separated geographically, or through firewalls from the main message route, connection to the local SMTP service is suggested. This greatly simplifies the firewall configuration between the main message router and the department servers.

Note: If a single domain / namespace is to be maintained across all department servers, separate, but coordinated directories need to be maintained. As the number of servers and users grow, the manual administration of changes across multiple directories becomes problematic. This is a known limitation, which will be resolved in IEMS Version 8. IEMS 8 will support replicated directories, as well as the storage of Intranet routing information, so that a single directory can be maintained across multiple servers.

If message loads become high enough that a single SMTP server is unable to handle all the traffic for an organization, processing can be split across multiple routing MTA's. Please see the figure below for an example of such a configuration:

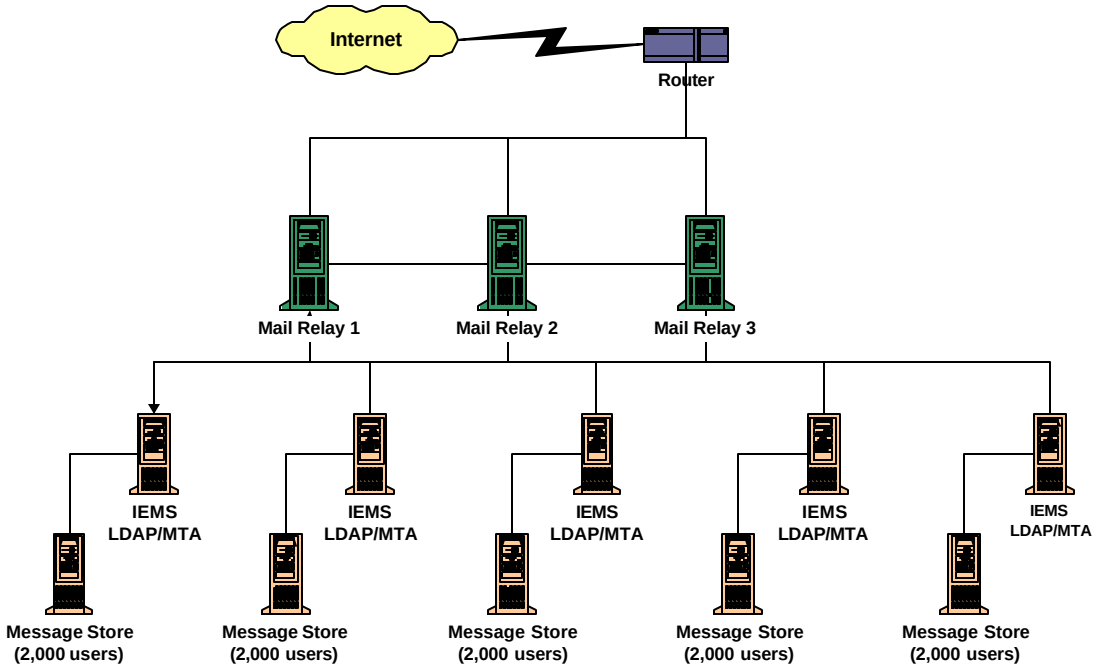

Figure 11: Multiple IEMS Systems with Redundant Message Routers / Relays

In this configuration, all three message routing MTA's (relays) share a common Directory, and hence messages can arrive from the Internet to any of the three and be treated identically. In this case MX records need to be setup with equal priorities pointing to each of the three relays.

Once received by any of the three relay MTA's, messages are routed to the appropriate local IEMS MTA via SMTPC. Once there, the situation is identical to the single message routing MTA scenario described earlier.

Note: The indication of 2,000 users per Message Store in the figure above is an example only. Please see the previous sections on *Capacity Planning* and *Message Load Profiling* to get a better indication of what is appropriate for your environment. The number of Message Store users supported may be larger or smaller than this number depending upon the anticipated message load for your site.

Like the single message router scenario, outbound SMTP traffic can be routed first through the department server, and then upstream to any of the available relay MTA's. This configuration in addition to being friendly to firewall situations, has the added advantage of redundant MTA's – ensuring message delivery even in the event of a single MTA hardware failure.

TYPICAL INTERNET EXCHANGE INSTALLATIONS

Internet Exchange is designed to cater to the messaging requirements of a wide range of organizations and individuals, including:

- Cc:Mail and/or Lotus Notes Gateway customers who are currently using earlier versions of Internet Exchange or other gateway products.
- Sites migrating from other messaging systems, such as Lotus cc:Mail, Lotus Notes, Microsoft Exchange. A period of co-existence between these environments and IEMS should be expected while a phased migration is taking place.
- Sites migrating from proprietary email systems to an open Internet standards-based messaging system.
- Organizations looking to implement a messaging solution for the first time who are interested in using open Internet standards-based products.
- Organizations looking to add messaging capabilities to existing networks, such as anti-spam, anti-virus, attachment removal, disclaimer insertion, distribution list processing, etc.
- Sites with non-dedicated Internet access, needing a total messaging solution not dependent upon fixed-IP addresses.

Several sample configurations are described in the following sections describing how IEMS can be configured and used to solve different messaging requirements. The following examples are typical of what can be expected in a test laboratory environment. All network environments differ in one way or another, and as a result it is normal to expect variations in expected performance. When designing your messaging infrastructure, be sure to take this into consideration as well as expected growth and the corresponding changing needs of the users and systems.

Existing IMA Gateway Installation

This sample scenario is useful for either Lotus cc:Mail and/or Lotus Notes sites planning to maintain their current use of Lotus technologies, while at the same time looking for a more robust and capable Internet gateway solution. This solution œvers sites currently running earlier Internet Exchange Gateway versions, or other gateway products.

IEMS running in a gateway only configuration, while not taking full advantage of all IEMS features (such as the Message Store and related functions) still offers considerable value to cc:Mail and/or Notes users. IEMS capabilities such as Anti-Spam, Anti-Virus, Distribution Lists, Disclaimer Insertion, TNEF Attachment decoding, and Batch SMTP integration are all seamlessly made available to Lotus users.

This sample site is running a Lotus environment (either cc:Mail and/or Notes) with 500 users. Message volume is what could be considered a medium load – 40 messages inbound per day per user, each averaging 50K each. This message load translates to 1GB inbound traffic for the site daily. For most sites, inbound traffic is considerably higher than outbound, however for estimating purposes we will assume they are the same.

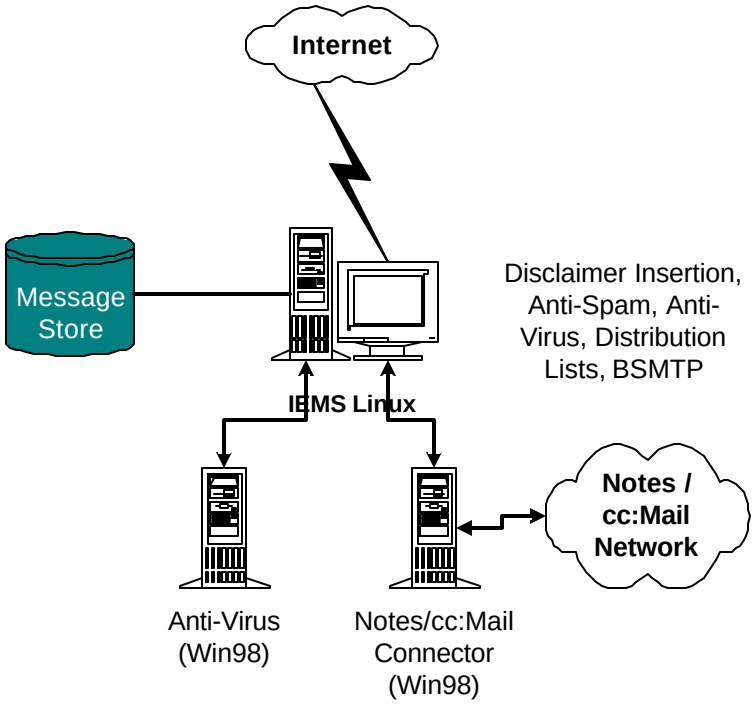

Figure 12: IEMS Gateway Configuration

The site requires Anti-Virus scanning for all messages that enter via the Internet, as well as all Internet bound messages. IEMS is used for Distribution List processing (20 lists of 300 members each). Normal Anti-Spam controls are in place, as well as Disclaimer Insertion for all Internet bound messages.

Due to the number of messages that have to be scanned during the Anti-Virus phase, it is recommended that a separate machine be used for this purpose. Virus scanning can be very CPU and Disk I/O intensive. As long as the machine selected is powerful enough to handle the required message traffic, real time machine load should not be an issue. Even if the machine runs at capacity for periods of time, the impact on the rest of the messaging environment should be minimal. A Windows-98 or more recent machine should do fine for this purpose.

The cc:Mail and Notes connectors only run in Microsoft Windows environments. As message load is not particularly high, the site administrator has the choice of running the rest of the messaging environment on a single machine (Windows NT or 2000 Server recommended), or splitting the processing between a cc:Mail/Notes connector machine (Windows 98 or better), and the rest of the MTA/Preprocessor (Linux or Windows NT/2000 Server). If the site has expertise in both Windows as well as Linux, then a cheaper Windows 98 machine running the connector(s), with a Linux MTA/Preprocessor/Directory implementation will work nicely. This configuration is shown above.

Microsoft Exchange to IEMS Migration

Current Microsoft Exchange sites wishing to migrate away from Exchange Server and/or Windows operating systems can move the backend server to IEMS while still retaining full Microsoft Outlook functionality on the desktop. IEMS includes support for Outlook client features such as Calendaring and Scheduling. New migration tools are provided to assist Microsoft Exchange sites in their conversion to IEMS.

In our sample scenario, we have a Microsoft Exchange site running with 500 users, 150 active at any given point in time. Message volume is average a 40 messages per day, each averaging

50K each for total system inbound traffic of 1GB per day. The existing mail server is Exchange 2000 running on a Windows 2000 Server.

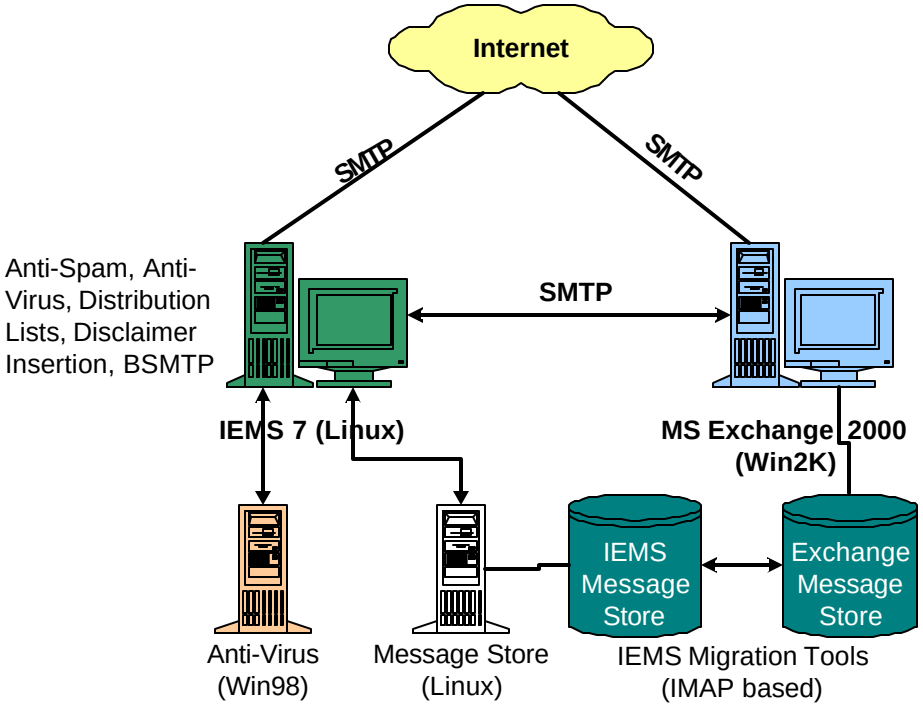

Figure 13: Migration From Microsoft Exchange

The recommended setup for this site is made up of a central IEMS MTA/Preprocessor/Directory Server connected to a high performance server running the Message Store. If Anti-Virus capabilities are needed, these should be run on a separate machine in order not to burden the main MTA/Preprocessor (whose clients will be directly connecting via SMTP for message submission).

MTA/Preprocessor

This machine is responsible for the reception and propagation of approximately 2GB of mail daily. While this is a respectable amount of email to handle on a daily basis, as long as the MTA is not slowed down by functions such as virus scanning, it should be able to accomplish this on a modest machine. For this machine, a recent model Celeron 800 MHz or faster machine, with a DMA-66 EIDE drive (or faster), and 128MB (minimum) memory running Linux should work fine. The Ethernet NIC should be of good quality, preferably one of the newer 100 Mbit types (3Com 3c905 or equivalent).

In addition to SMTP reception from client workstations, this machine will receive SMTP based mail from the Internet, handle all anti-spam, distribution lists, disclaimer insertion, and potential TNEF attachment handling. The shared Message Queue will also reside here. For this machine it is important to keep the machine loading relatively light, but at the same time a reasonable sustained load is acceptable. The main issue is that enough resources are available to handle SMTP requests reasonably fast, while at the same time making items in the Shared Message Queue available in a timely manner to the other IEMS machines in the cluster.

Anti-Virus

As most of the commercially available anti-virus software is still primarily available in Microsoft Windows environments, this machine can be a simple Windows 98 machine. The main issue here is that the machine is fast enough, with enough memory to handle the handling of

attachment virus scanning. The determination of exact hardware here is highly dependent on the efficiency of the virus scanning package(s) selected for the site. The loading factor for this machine is not important as long as over a reasonable period of time, the machine can keep up with the overall message volume. As neither users nor remote SMTP sessions are impacted, 100% utilization of this machine for sustained periods of time is acceptable.

To best determine hardware requirements here, check with the anti-virus vendors of choice for their recommendations. Alternatively, a hardware configuration similar to the MTA/Preprocessor machine can be used for testing purposes. Chances are good that if only a single anti-virus profile is configured that this configuration should be workable.

Message Store
The Message Store machine is the most critical machine in term of hardware requirements. User client workstations will be directly connecting with this machine for the retrieval and maintenance of their messages and folders. Good response time is critical here if end users are to be kept happy.

As disk I/O is apt to be the primary system bottleneck here, a good rule of thumb is the faster the disks the better. Ultra-160 SCSI or faster drives storing the message store accounts and the system files spread across volumes is highly recommended. In addition, in order to avoid swapping, which further places additional load on the disk I/O subsystem, make sure that the machine has adequate memory.

For the basic system requirements, a Pentium III, 1GHz machine or faster, with a minimum of 256MB of memory running a recent release of Linux is recommended. A fast SCSI controller and matching SCSI drives are necessary. If at all possible try to avoid mixed SCSI / IDE configurations for this machine, as the IDE drives tend to slow the machine down. As with the other machines, try to use as high quality network interface as you can – this is after all where all data has to pass. It should be quick and responsive, and above all highly reliable.

Migration Tools
IEMS Migration tools are included with IEMS 7 and work with the Exchange Server IMAP4 service (Exchange 2000 and later). The migration tools can be run together with the Exchange server to migrate users over to the IEMS Message Store (Linux or Windows based).

Messaging Firewall Applications
For sites that have existing Messaging infrastructure in place that they need to retain, but still need the added features of IEMS, IEMS can be installed in a *Messaging Firewall* configuration. This front-end setup is responsible for the reception of all Internet based email into the organization, and all outbound messages from the organization are sent through IEMS for final processing. This setup allows a site to transparently add capabilities such as IEMS Anti-Spam, Anti-Virus, Content Filtering, Disclaimer Insertion, Distribution List, and TNEF attachment handling to the underlying (existing) messaging environment.

This type of setup is also very useful for IEMS (or other messaging environments) sites that are operating on the internal network side of a firewall. If SMTP communication through the firewall is via port forwarding, many times the originating IP address is lost, making IP Address based Anti-Spam prevention difficult at best. For situations such as this, having a firewall IEMS server in the DMZ filtering connections and content is quite effective.

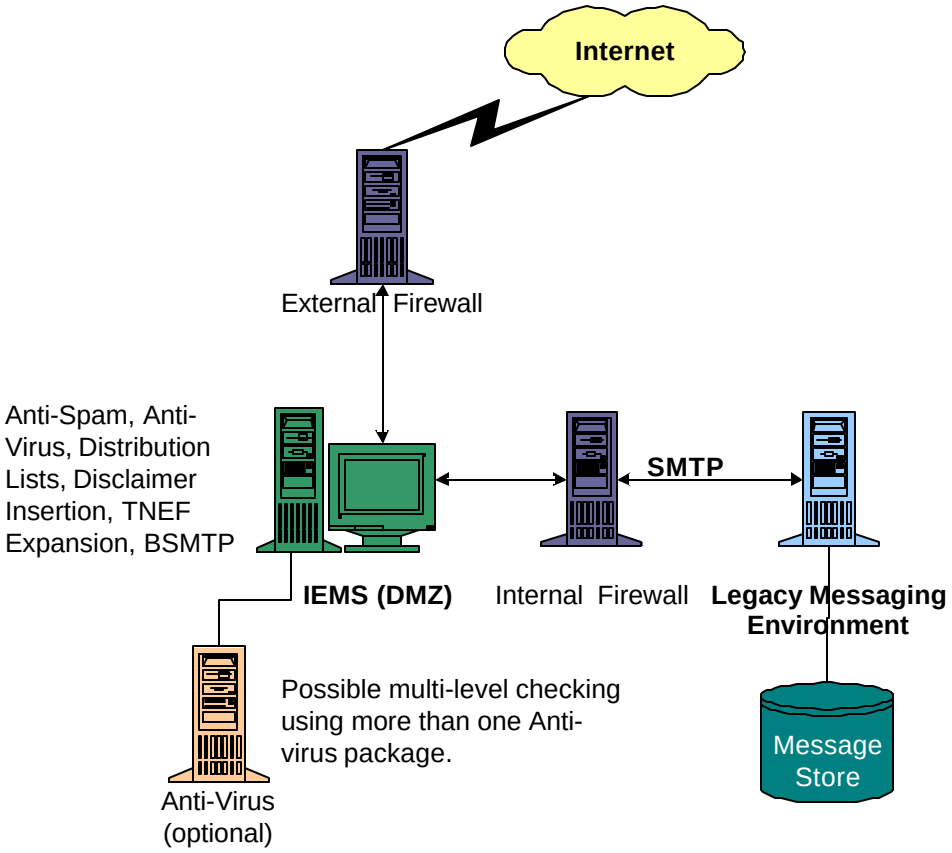

Figure 14: IEMS As A Messaging Firewall

The setup of a firewall IEMS server is simple. No connector modules, or Message Store accounts need to be configured, as they are not used. With the possible exception of optional virus scanning, the primary purpose of this machine will be for selective relaying of mail into the internal network, and the final relaying of outbound messages out over the Internet. Note that using IEMS as a mail relay for internal machines will significantly lessen the load on the internal network SMTP servers, as SMTPC queues on the internal network should almost always be empty, assuming high availability of the IEMS relay in the DMZ.

If Anti-Virus scanning and/or Content Filtering is necessary, and message load is respectable, it is highly recommended that they be run on a separate machine or machines. Please see the section above on Microsoft Exchange Migration for appropriate guidelines on the Anti-Virus component here.

Other Preprocessor functions, such as Anti-Spam (other than content filtering), Distribution Lists, and Disclaimer Insertion all place minimal load on the MTA/Preprocessor. The main concern when choosing hardware for this configuration is reliability. Most recent Celeron / Pentium II or faster processors should do fine here in most situations. The Celeron 800 MHz Linux configuration with 128MB of memory described above should be more than sufficient for most applications of this type.

For sites that are running IEMS 7 or later inside the firewall, MTA Pass-Through capabilities can be enabled on the DMZ and propagated across the internal firewall to the internal systems. If message loading in the DMZ becomes significant due to Anti-Virus and/or Content Filtering, redundant IEMS configurations can be established in the DMZ. In this configuration multiple MX records for the domain need to be established – one for each MTA. Each of these MTAs can then independently forward mail across the internal firewall.

Configuration

While the network and hardware configuration for this scenario is rather simple, it is important that the message routing settings of IEMS and the internal network messaging servers be setup properly. For the internal network machines, they should be configured in Mail Relay Only mode, to send all outbound mail directly through the internal firewall to the DMZ IEMS server.

Internet originated mail is received by the IEMS server in the DMZ. If this machine is handling Distribution List expansion, DL's can be setup as with any other IEMS machine with no special consideration for this application. All other mail however for the mail domain or domains being serviced however need to be passed off to an internal network server for further routing (after passing all preprocessing tests).

To accomplish this, the IEMS Firewall machine routing options need to be set for "HOST TABLE THEN DNS." In addition, the names and IP addresses (relative to the DMZ IEMS machine) of all internal machines that mail will be forwarded to need to be configured in the local host table. With this configuration, mail can be manually routed (via specific host table values) for internal machines, while retaining proper routing for externally bound mail.

Finally, in order to enforce the above routing policies, make sure that both the external as well as internal firewalls are configured to allow SMTP message flow as outlined above, while at the same time, blocking all other SMTP requests.

Small Office – Low Volume – 25 Users

Installing IEMS in a small office environment is very simple and straightforward. If message volume is very small, or no virus scanning is required, a single machine configuration will work quite well. If the number of messages are large, and anti-virus scanning is required, this can potentially place intermittent loads on the server, slowing down remote clients trying to access their message store accounts. If such intermittent loads are infrequent, or if the temporary slowdown of message store access is acceptable, then the single machine approach is recommended. Regardless, if a single machine configuration is first put into operation, and later it is found that virus scanning interferes too much with remote users, this functionality can be moved to a remote virus scanning machine at a later date.

In our sample setup, if we assume 25 users exchanging an average of 40 messages per day, each averaging 50K each, this comes to a total inbound daily volume of only 2MB. This is a trivial amount and will not represent a significant load on the MTA or Anti-Virus scanner. Even if the average number of messages is increased to 300 with an average size of 250K, the total inbound message volume is still only 1.9GB, an amount easily handled by a single IEMS server on most modern hardware.

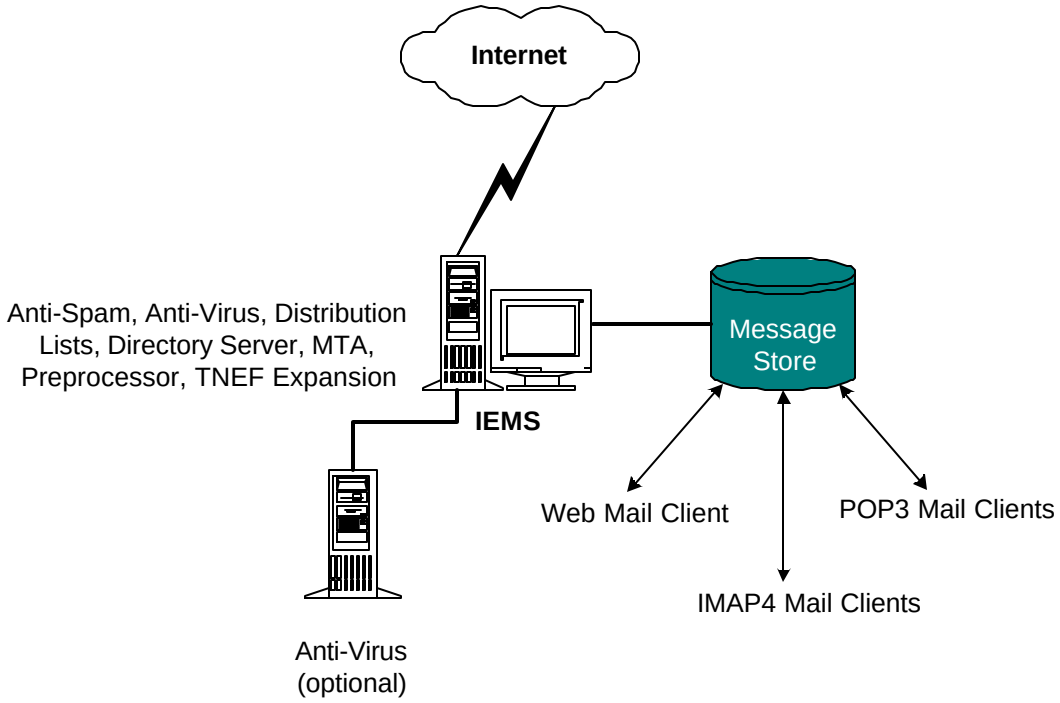

Figure 15: IEMS Small Office Configuration

Medium / Large Office – High Volume – 500 Users

For larger sites, with high message loads, it is important to make sure that the different IEMS components are balanced across a local network. In this scenario, the site has to support 500 local users, 200 assumed to be active at any point in time. Message volume averages 300 messages per day with an average message size of 250K. This translates to 38GB total inbound traffic daily.

The main challenge with a site such as this is making sure that the high message volume can be adequately handled, both in terms of exchanging messages with other systems (SMTP) as well as conveying these messages to the users (IMAP4 and/or POP3). With adequate hardware, a single cluster IEMS configuration should be able to work in this environment, however most systems will be approaching maximum suggested loading for reasonable end user response times. If and when expansion of the system is necessary, a move to a more distributed model with separate Message Servers servicing portions of the user community may become necessary.

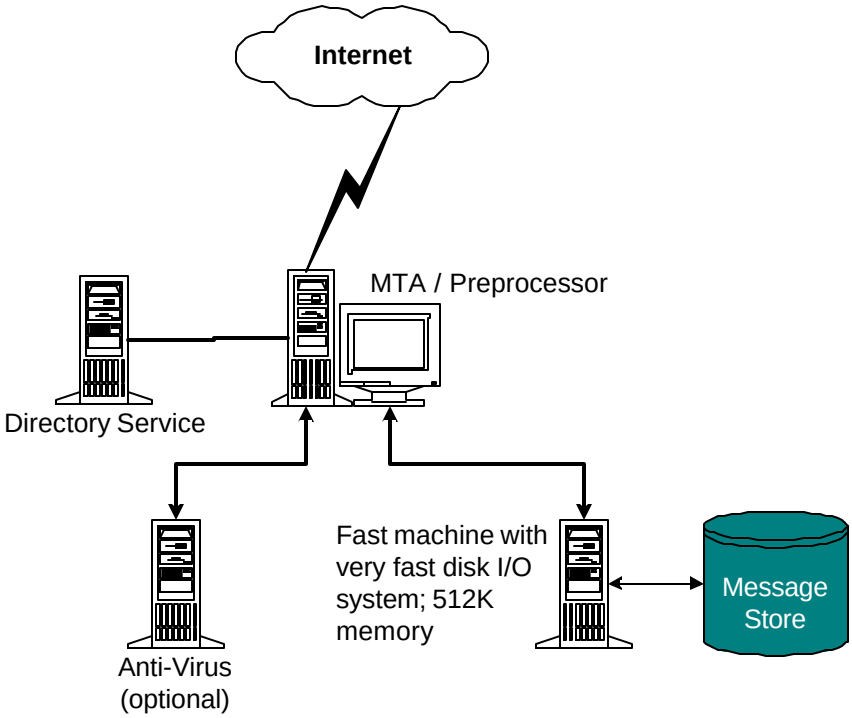

Figure 16: IEMS Large Office Configuration

MTA/Preprocessor
While the speed of the processor is not critical here, the disk I/O channel should be fast. This is so that information received via SMTP from client workstations and the Internet can be handled quickly, as well as fast access to the IEMS Shared Message Queue by other IEMS components on the network. DMA-66 EIDE or faster drives are recommended here. It you are running in a Linux environment, make sure that the DMA drivers are enabled in the kernel – most stock Linux installations have this turned off by default. Also make sure that adequate storage is available for the various IEMS queues. Most modern drives should be able to handle this requirement easily.

As message volume is high, and the number of concurrent network connections can also be high, a minimum recommended configuration here is Celeron/Pentium II, with 256MB memory, with fast drives as described above, and a fast and reliable Ethernet NIC.

Directory Server
IEMS modules obtain configuration and message routing information from the Directory. Because of this, in a high load environment, calls to the Directory will be frequent. The faster the response time, the faster messages will flow through the system.

While providing directory access is not really CPU intensive, the more information that can be cached in memory the better. Like the MTA/Preprocessor machine above, a reasonably fast drive, and adequate memory is necessary. The same requirements for the MTA/Preprocessor above apply here.

Anti-Virus
Please see the description of Anti-Virus requirements for the Microsoft Exchange migration scenario. These basic requirements also apply to this situation. However, as message load is high, make sure that the machine is reasonably fast. A 1 GHz Pentium III with fast disks (as

described in the Message Store requirement sections) is recommended. Due to the high message load, make sure that this machine is capable of handling the overall processing requirements placed upon it over an extended period of time. Like the other scenarios, machine loading is not important as neither users nor remote SMTP servers directly interact with it. The main issue here is to make sure you've got enough processing ability to handle the high message traffic.

Message Store

The same design issues described in the Microsoft Exchange migration scenario apply here. In fact as message load is higher, they are even more critical in this environment. Get the fastest drives you can, while maintaining quality (reliability). If running in a Linux environment, a minimum of 512MB of memory is required, with 768MB or more recommended. Like the previous examples, don't try to save on the network board – this needs to be fast and reliable.

WHERE TO GO FOR MORE INFORMATION

The information described in this document is intended to provide an overview to designing messaging environments based upon IEMS. Detailed information on the operation, configuration, and maintenance can be found in the various IEMS manuals, and on the IMA Web Site (see the Introduction of this document for how to obtain these over the Internet).

IMA engineers are also always available to help answer questions on how to install, configure, and run IEMS systems and are available at all times to assist. For a complete list of ways in which IMA can be contacted, please see http://www.ima.com/contacts/.

CONCLUSIONS

The Internet Exchange Messaging Server is a highly functional messaging environment capable of running in many diverse messaging environments. Its ability to defer what would typically be decisions made by the MTA, and put that decision making capability into the hands of the end users allows the administrator to be able to provide and environment where users are far more in control of what they receive. While IEMS is primarily designed to be a complete backend messaging environment, it also can act as high-end gateway solutions, messaging firewalls, and other configurations.

This document describes the architectural issues involved in deploying IEMS in various situations. As with all such descriptions, as a given site configuration will most likely differ from the examples given here, these differences need to be factored in by the network engineer designing the messaging environment. The information provided here is intended to provide the network engineer with enough information to make reasonable decisions in the design of their networks.

It is acknowledged that not all situations can be documented in publications such as this. If issues or questions arise that are not covered here, please feel free to call upon the IMA Technical Support group at any time to lend a helping hand.