

Architecture Guide

March 2006

Contents

MailMarshal 6.0 SMTP Architecture	2
MailMarshal 6.0 SMTP Components	3
Installation Scenarios	9
Customer Deployment Scenarios	13

This paper introduces the concepts behind the architecture model for MailMarshal 6.0 SMTP. The second section of the paper details the individual components and explains how each interacts with the other components. It also discusses the benefits that this architecture provides and how you can take advantage of them.

The third section covers the different installation scenarios for MailMarshal. It presents and discusses five different environments where the product could be deployed.

The final section looks at different company sizing scenarios, and discusses several different ways that MailMarshal could be implemented for them. These sizing scenarios are based on small, medium, large and enterprise companies.

Author: Bradley Anstis

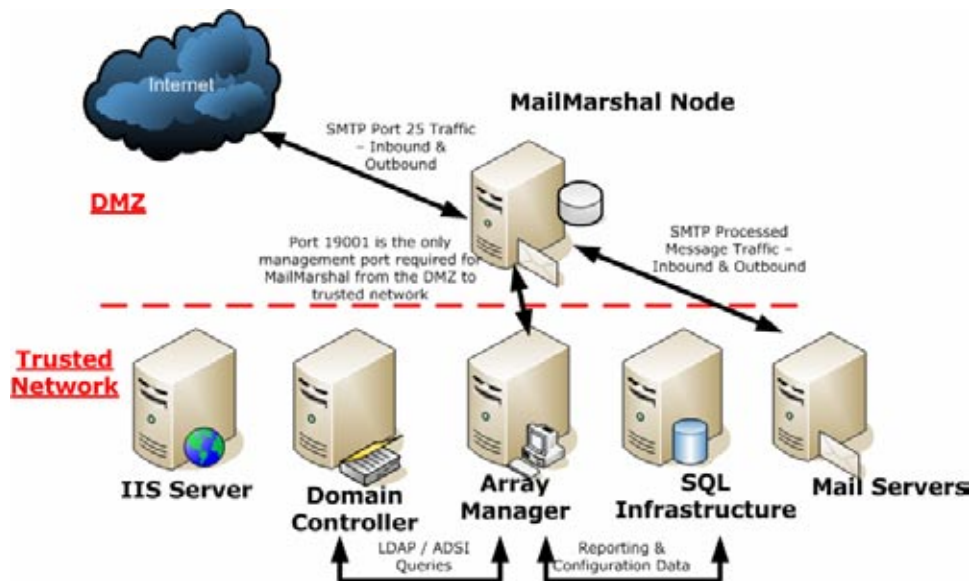
WHITEPAPER - MailMarshal 6.0 SMTP

MailMarshal v6.0 SMTP Architecture

The main thinking behind the architecture of this version was to improve scalability and manageability as compared to the MailMarshal 5.5 product, especially to meet the requirements and needs of Enterprise size installations.

The main goals were to:

- Provide a single configuration console enterprise wide
- Provide a simplified, more secure security model for DMZ deployment
- Allow the MailMarshal processing node to update their configuration using minimal bandwidth
- Provide a unified management console enterprise wide with the ability to manage high volumes of quarantined messages no matter what MailMarshal processing node Server they are quarantined on
- Provide rolled up enterprise wide reports, with remote data delivered over low bandwidth links



The resulting design introduces a split security model where Configuration and policy is managed using a new component called an 'Array Manager' and subsequently stored into the SQL Database (now a requirement not an option). Configuration and policy is sent from the Array Manager to the MailMarshal processing node servers and is applied on each server.

Note: Each of the components can be seen in the expanded diagram below. Note that the components do not require separate servers. A decision about the number of servers and which components to install on which servers depends on the existing company infrastructure and size.

WHITEPAPER - MailMarshal 6.0 SMTP

MailMarshal SMTP v6.0 Components

Array Manager

Recommendations

The location of the Array Manager can affect the performance of the administration and configuration tools used in MailMarshal, but will not necessarily affect the mail processing performance.

The recommended locations for the Array Manager are listed below, from most to least preferred:-

1. On the same server as the SQL Server holding the MailMarshal database. This location is preferred because the Array Manager is the only component that actually communicates directly to the MailMarshal database. All the other components communicate to the database through the Array Manager.
2. On a server that is as close as possible to the SQL Server holding the MailMarshal database, and has a LAN speed network connection to it.
3. On the same server as an Active Directory Global Catalog. The Array Manager communicates to the Global Catalog on a regular basis if the installation uses Active directory user groups. The Array Manager can also communicate through LDAP to another existing directory server.
4. On a smaller site, the Array Manager, the MailMarshal processing node and the SQL Database (SQL Server or MSDE installation) could all be on the same server.

Configuration Changes

As mentioned above, in the new architecture the Array Manager is the point where you define Configuration and Policy.

The MailMarshal Configurator tool connects to the Array Manager to retrieve the current configuration, the configuration that is displayed when the Configurator starts is retrieved from the registry on the Array Manager. As policy or configuration information is changed in the configurator the changes are written to the registry on the Array Manager.

When the administrator 'Commits Configurator Changes' a snapshot of the configuration stored in the registry configuration is stored in the SQL database and also exported to a XML Document. The array Manager then notifies the MailMarshal processing node servers that an update is available. The node servers also check for updates periodically.

The MailMarshal Node servers are then notified that an update is available through the single TCP 19001 port. The MailMarshal Nodes also poll for updates periodically.

When the MailMarshal Node server recognizes a new update, the XML document is requested from the Array Manager. The Array Manager compresses the XML document and delivers it through the TCP 19001 port to the MailMarshal Node.

WHITEPAPER - MailMarshal 6.0 SMTP

When this new XML configuration document arrives at the MailMarshal Node and is uncompressed, the new configuration settings are dynamically applied to any new threads that start after the new configuration is available, except for some server values like local domain lists which will require the associated service to be restarted.

SpamCensor Updates

The Array Manager will also require access to DNS for it to resolve the location from which to pull updated SpamCensor scripts. The Array manager will check for these updates daily and if an update is received then the standard process for configuration changes as explained above will take place to get the updates out the MailMarshal nodes.

When an update is requested, the following process happens...

1. Https connection made to marshal.com to request a list of files, versions and checksums.
2. If the files are deemed to be of a later version than the current rules then an http connection is made to download the compressed files.
3. When the files are downloaded, they are unpacked and the checksums are compared to ensure that

the downloaded files have not been tampered with
4. Finally the updated rules are merged into the SpamCensor rule set, and the configuration changes are committed.

Database Communications

The Array Manager is the only component that communicates directly to the MailMarshal Database.

This communication uses TCP and defaults to the standard SQL port, port 1433.

Some examples of data passed between the Array Manager and the database are:

- Logging data passed on from the MailMarshal Node Servers
- Lists of quarantined messages retrieved to populate the folder listing views in the Console
- Message details such as the quarantine location, used when viewing the details of a quarantined message in the Console
- Transactional information from the MailMarshal processing nodes relating to administrative tasks like message release, forward etc

WHITEPAPER - MailMarshal 6.0 SMTP

Directory Connections

The Array Manager is also the component that communicates with any organizational directories, as defined by the 'Connectors' set up through the MailMarshal Configurator tool. Depending on the type of connector, the Array Manager communicates with Active Directory natively using ADSI, or with a LDAP directory using TCP port 389 or a custom port (including SSL).

The Array Manager uses the directory to maintain a list of groups and members referenced in the MailMarshal policy.

Native Active Directory support provides greater flexibility in queries. MailMarshal can retrieve all the email aliases assigned to each user for use in policy or in the Spam Quarantine Management tool.

The administrator can specify how often each directory is checked for changes to group membership. If any changes have been made the Array Manager refreshes the affected group to ensure it is up to date. The changes to group membership take effect immediately. They do not require you to commit configuration.

MailMarshal MMC Console Actions

The MailMarshal Console now provides an enterprise wide view of the MailMarshal implementation.

It connects to the Array Manager and uses data from the MailMarshal database.

For example if an administrator running the MailMarshal Console application views the contents of a quarantine folder, the information comes from the MailMarshal database and is a list of all the messages quarantined in that particular folder on all MailMarshal nodes enterprise wide.

When the administrator selects a message from the folder list and views the details of that message, the Array Manager communicates with the node where the message is quarantined. The node invokes the unpacking engine which returns the message data. The Console provides a safe view of the message components in a protected 'sandbox.'

When an administrator decides to forward, release or delete a particular message, the Array Manager communicates this request to the MailMarshal node where the message has been quarantined, and the node takes the appropriate action. The Array Manager also updates the information about the message in the database.

All the other views in the MailMarshal Console, such as the MailMarshal Today page and the domain views, are populated by data the Array Manager has retrieved from the database. The MailMarshal Console does not communicate directly with the database.

WHITEPAPER - MailMarshal 6.0 SMTP

MailMarshal WebConsole Actions

The Web based version of the MailMarshal console provides exactly the same functionality as the MMC based version described above. The pages are created by an ASP.NET application running on a Microsoft Internet Information Services server, using content provided by the Array Manager.

Spam Quarantine Management Actions

The new end-user focused Spam Quarantine Management tool is also an ASP.NET application running on an IIS server, using content provided by the Array Manager.

The process of viewing quarantined messages in the Spam Quarantine Management tool is similar to that used in the Console. A user will only be able to view the messages that have been quarantined for their email address, including any aliases they may have.

MailMarshal Node Server

The function of the MailMarshal node servers is solely to process incoming and outgoing SMTP messages according to the policy defined by the organization. A MailMarshal node generates SMTP traffic and management communication.

The SMTP traffic generated by the standard SMTP gateway function consists of the incoming and outgoing mail streams. You can also configure MailMarshal as a POP3 mail server, which adds POP3 traffic. (This option is only available with there is a single node.)

All management communication traffic is between the node and the Array Manager using a single TCP port. By default MailMarshal uses port 19001 for this purpose.

Nodes retrieve configuration from a local XML settings file that is local on the node itself. The configuration is loaded when the MailMarshal services are started, and reloaded after the node downloads an updated XML file from the Array.

The only other incoming management communication happens when the Array Manager requests the node to take some action on a quarantined message. For example the Array Manager can request the node to send the data required to view the message in the Console, or to release, forward or delete a particular message.

The only outgoing management communication initiated by the node is delivery of the batched logging information to the Array Manager to be written to the MailMarshal database.

A MailMarshal node creates one or more log entries for each message it processes. These log entries are written to a local file queue on each MailMarshal node. A 'Log Follower' process batches this queue content to the Array Manager on a regular basis (typically every 15 seconds, or 100 records per batch). If the Log Follower cannot communicate with the Array Manager, the node continues to store records locally until it can contact the Array Manager.

This design provides a highly efficient and fault tolerant method of getting the logging information into the MailMarshal database.



WHITEPAPER - MailMarshal 6.0 SMTP

SQL Database Server

The SQL database server (SQL 2000 or MSDE 2000) hosts the MailMarshal Database. The database contains logging information, configuration and policy information, and quarantine folder information. As previously described the only component that talks directly to the SQL server is the Array Manager.

It is good practice to place the Array Manager on the same server as the SQL Server, or as close as possible. For sizing recommendations, please see the MailMarshal v6.0 SMTP sizing whitepaper. MailMarshal includes a tool to migrate the contents of an older (4.x or 5.x) database. MailMarshal also provides a tool to checks that all the messages quarantined on a particular MailMarshal node are detailed in the MailMarshal database, so that they will be shown in the folder listings within the Consoles and the Spam Quarantine Management tool.

IIS Web Server

The MailMarshal Web components install on an Internet Information Services (IIS) server. These components are only required if an organization wants to make use of the Spam Quarantine Management tool or the Web Console.

The Web components require IIS 5.0 or higher with ASP.NET 1.1. ASP.NET is part of the .NET Framework 1.1, available on the MailMarshal CD-ROM.

You can use a secure Web site (HTTPS) to protect user data and authentication information.

The web application on the IIS server communicates with the Array Manager to get the content for these pages.

Internal Mail Server

An organization's internal email server communicates through SMTP to the MailMarshal processing node servers, typically using TCP port 25 in both directions. The processing nodes do not require any other access to the internal email server.

If the internal email server is the organization's LDAP directory server, and you are using LDAP to provide user and group information in support of the MailMarshal policy, the Array Manager uses LDAP to retrieve this information.

WHITEPAPER - MailMarshal 6.0 SMTP

Network Ports Required

Typically MailMarshal uses the following network protocols and ports:

- SMTP email (port 25) from MailMarshal email processing servers to the Internet and to the internal email servers and/or clients.
- DNS (port 53, both TCP and UDP) from MailMarshal email processing servers for resolution of external email server names.
- TCP port 19001 for communication between the email processing servers, Array Manager, and Console. This connection can be changed to another port of your choice.
- HTTP and HTTPS from the Array Manager to the Internet for access to SpamCensor updates (ports 80 and 443). You can use a proxy server for Web access if your environment requires it.
- SQL server connection (port 1433 by default) between the Array Manager and the SQL database, and between the database and any Reports Console.
- The Array Manager and the Configurator would typically be deployed within the trusted LAN, not through a firewall for both performance and security reasons.
- POP3 (port 110) if MailMarshal is functioning as a POP3 server.


WHITEPAPER - MailMarshal 6.0 SMTP

Installation Scenario's

This section discusses some typical options for the deployment of the MailMarshal components. Each option provides all the required functions of an email gateway. Five different scenarios are explained below:

1. Install MailMarshal on its own physical server, as an email relay within an organization
2. Install MailMarshal as a POP3/SMTP server providing all email server functions for a small organization
3. Install MailMarshal on a Proxy Server
4. Install MailMarshal on the same physical server as the organizations email server software
5. Install MailMarshal nodes into the DMZ, with the Array Manager in the trusted network.

1. Install MailMarshal on its own physical server, as an email relay within an organization



In this case the MailMarshal installation is a "standalone server" including all management and email processing components. This option is suitable for small to medium sized organizations with a single Internet gateway and email server.

All workstations within the organization send email through the email server. The email server forwards all external messages to the MailMarshal server for processing and delivery.


The DNS MX record (or the firewall's relay setting) is set so that the MailMarshal server receives all email inbound to the organization.

Install the MailMarshal database on an available SQL server in the local network. In smaller organizations, it will be possible to install SQL Server or MSDE on the MailMarshal server.

Install the MailMarshal Spam Quarantine Management Web site on an intranet Web server. Install MailMarshal Reports and optionally management consoles on one or more workstations in the local network.

WHITEPAPER - MailMarshal 6.0 SMTP


2. Install MailMarshal as a POP3/SMTP server providing all email server functions for a small organization



In this example, workstations within the organization send email to the MailMarshal server on port 25 for processing. MailMarshal delivers email for internal addresses to MailMarshal POP3 mailboxes for collection by email clients. Retrieve and send email to and from external addresses over a dialup or other link to an ISP.

In this case the MailMarshal installation is a "standalone server" including all management and processing components. In most organizations that choose this scenario, it will be possible to install SQL Server or MSDE on the MailMarshal server. Install the MailMarshal Spam Quarantine Management Web site on an intranet Web server. Install MailMarshal Reports and optionally management consoles on one or more workstations in the local network.

3. Install MailMarshal on a Proxy Server




The above environment is simple to implement and fits 'logically' with the concept of a proxy server and no environmental or configuration changes need be made to the proxy server.

MailMarshal effectively proxy's SMTP between the interfaces of the proxy server and passes them to the existing email server.

The existing email server simply forwards all outbound email to the internal interface of the proxy. If no firewall is used in the environment the MX records for the organisation will point to the external interface of the proxy.

4. Install MailMarshal on the same physical server as the organizations email server software



In this case, all email sent from outside the organization arrives at the email server computer on the default SMTP port, port 25.

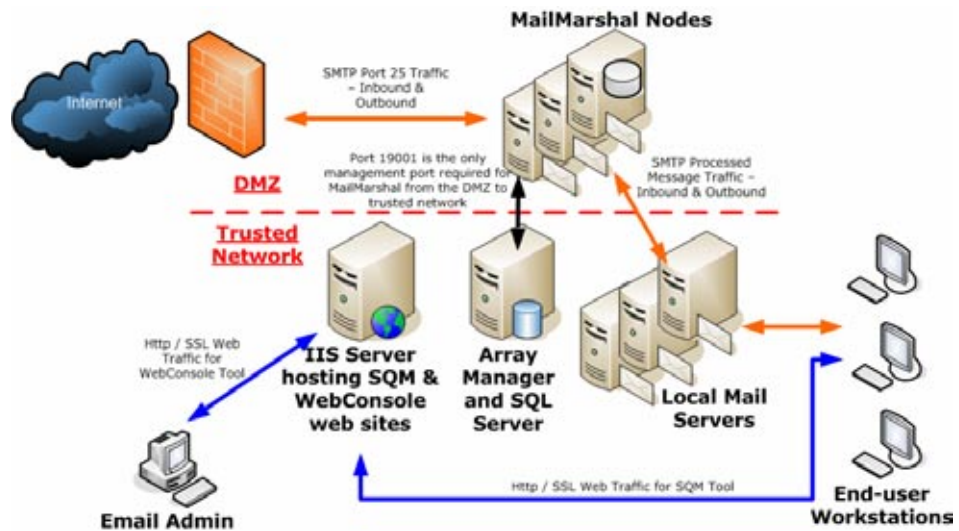
MailMarshal forwards processed inbound email to the other server software using the "localhost" IP address and port 97. The other server sends email for outside delivery to MailMarshal using the "localhost" IP address and port 25.

Install the MailMarshal database on an available SQL server in the local network. MailMarshal is installed as a "standalone server" including all management and processing components.

Install the MailMarshal Spam Quarantine Management Web site on an intranet Web server. Install MailMarshal Reports on one or more workstations in the local network

Note: This installation option depends on the server having sufficient resources to support both MailMarshal and another email server application.

5. Install MailMarshal nodes in the DMZ, with the ArrayManager in the trusted network



The above environment is straightforward to implement as well, and would be the environment of choice for larger organizations.

In this example, the MailMarshal installation includes a load balanced array of MailMarshal email processing servers in the DMZ.

All email sent from within the organization passes through the local email server, which delivers outbound messages to the MailMarshal servers on port 25. MailMarshal delivers incoming email to the local email server.

Install the MailMarshal Array Manager on a SQL server or a dedicated server within the LAN to perform configuration and connect to the MailMarshal database.

Install the MailMarshal Spam Quarantine Management Web site and the MailMarshal Web console on an intranet Web server. Open TCP port 19001 (or a single port of your choice) in the firewall between the DMZ and the Array Manager, to allow MailMarshal configuration and logging traffic.

A distributed enterprise with more than one email gateway can install one or more MailMarshal email processing servers at each gateway.

If the enterprise uses the same policies at all locations, it can use a single MailMarshal Array Manager to control configuration and perform logging for all locations.

All the email processing servers must be able to communicate with the Array Manager on port 19001. Install MailMarshal Reports and optionally management consoles on one or more workstations in the local network.

WHITEPAPER - MailMarshal 6.0 SMTP

Customer Deployment Scenarios

This final section looks at some typical actual customer sites and shows how you might integrate not only MailMarshal, but also the WebMarshal product as well.

The section is divided up into four different size companies:

- Small Company – 200 Users, single site
- Medium Company – 750 users, single site
- Large Company – 3000 users, multiple sites, 1 Internet Gateway
- Enterprise Company – 50,000 users, multiple sites, multiple Internet Gateways

The information and guidance provided here is only a guide and is not guaranteed to meet every customer's specific needs. Please contact a Marshal representative for detailed guidance about your specific situation.

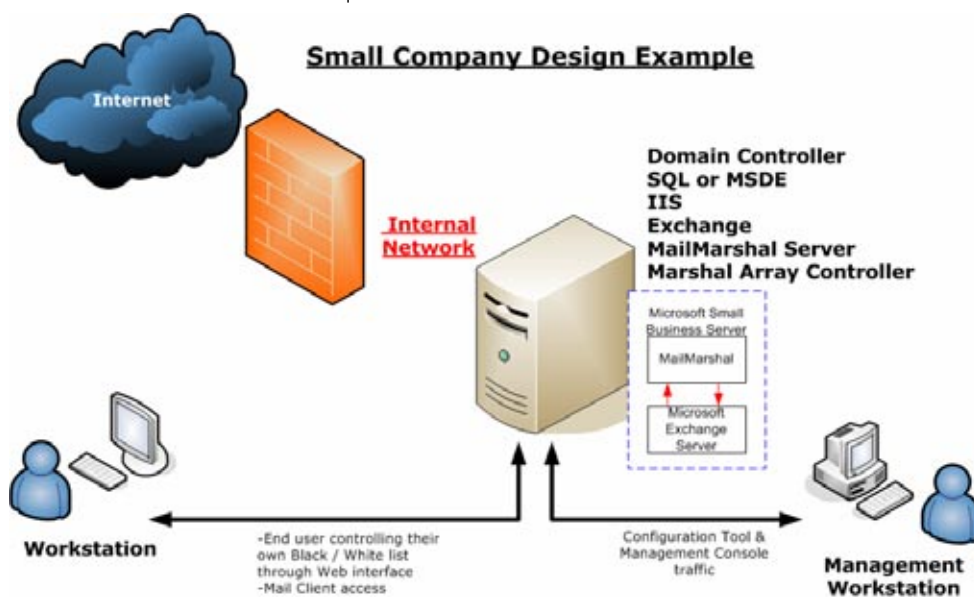
This section does not provide hardware sizing guidelines. For more information on hardware sizing, please refer to the MailMarshal 6.0 SMTP Sizing whitepaper.

WHITEPAPER - MailMarshal 6.0 SMTP

Small Company

The first example is a small company with 200 users at a single site. The company has a firewall device which connects through a router to the internet. The firewall has only two security domains, Public (Un-trusted) and the Internal Network (Trusted).

The company uses a Microsoft Small business server environment. The integrated Exchange server is the internal email server, and the integrated Microsoft ISA server is their web access server. MailMarshal is installed on this same server along with the other required resources such as IIS and the Domain Controller.



This solution is based on a single server providing all of these functions which is realistic for a site with only 200 users. In this scenario, remember to configure MailMarshal to send the inbound SMTP email on a different port to the internal mail server, perhaps port 97. The internal mail server will need to receive on port 97 as well.

This configuration scenario is the same as option 4 of the Installation scenarios.

WHITEPAPER - MailMarshal 6.0 SMTP

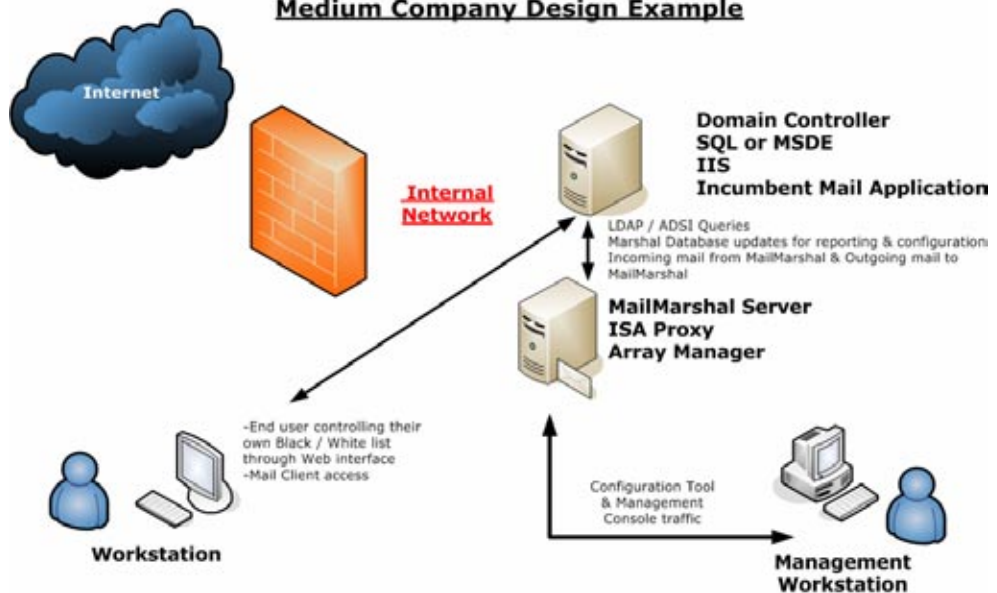
Medium Company

For this example, a medium company is defined as having 750 users in a single site. They have a firewall device which connects through a router to the internet. The firewall has only two security domains, Public (Un-trusted) and the Internal Network (Trusted).

The organization uses two main servers. One is used as a Domain Controller and also has SQL or MSDE installed, as well as IIS and the incumbent mail application such as Lotus Notes, GroupWise, or Exchange.

The second server runs all the Content Security products with MailMarshal email processing, and the Array Manager together.

Medium Company Design Example

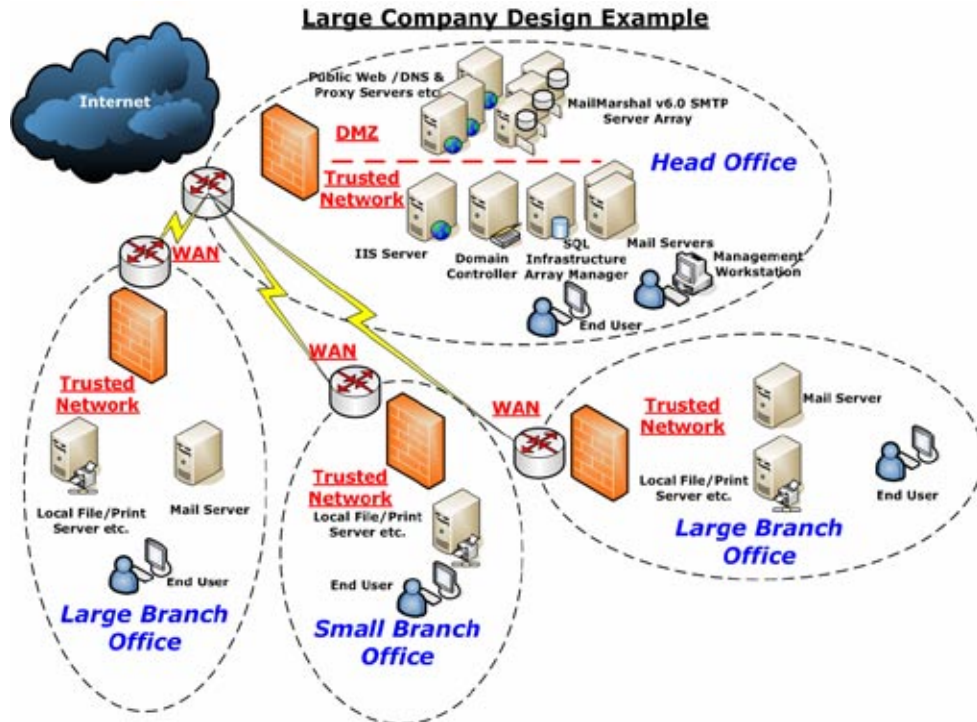


This medium company has two servers, with the load split between them. The configuration is simply the same as option 1 of the Installation scenarios.

WHITEPAPER - MailMarshal 6.0 SMTP

Large Company

The Large Company is defined as 3000 users spread across multiple sites. They have one central point for Internet access where all the SMTP mail and other Internet traffic is received and sent. The remote sites utilize Frame Relay based Wide Area Networks (WAN) for connectivity with the main office. The larger remote sites have their own exchange environment while the smaller ones rely on the main office Exchange infrastructure.



This large company uses the same configuration scenario as option 5 above, and since they have the MailMarshal nodes in the DMZ they are able to take advantage of the new security features.

Since there is only a single Internet Gateway, there is only one group of MailMarshal processing nodes.

This scenario is the same as option 5 of the Installation scenarios.

WHITEPAPER - MailMarshal 6.0 SMTP

Enterprise Company

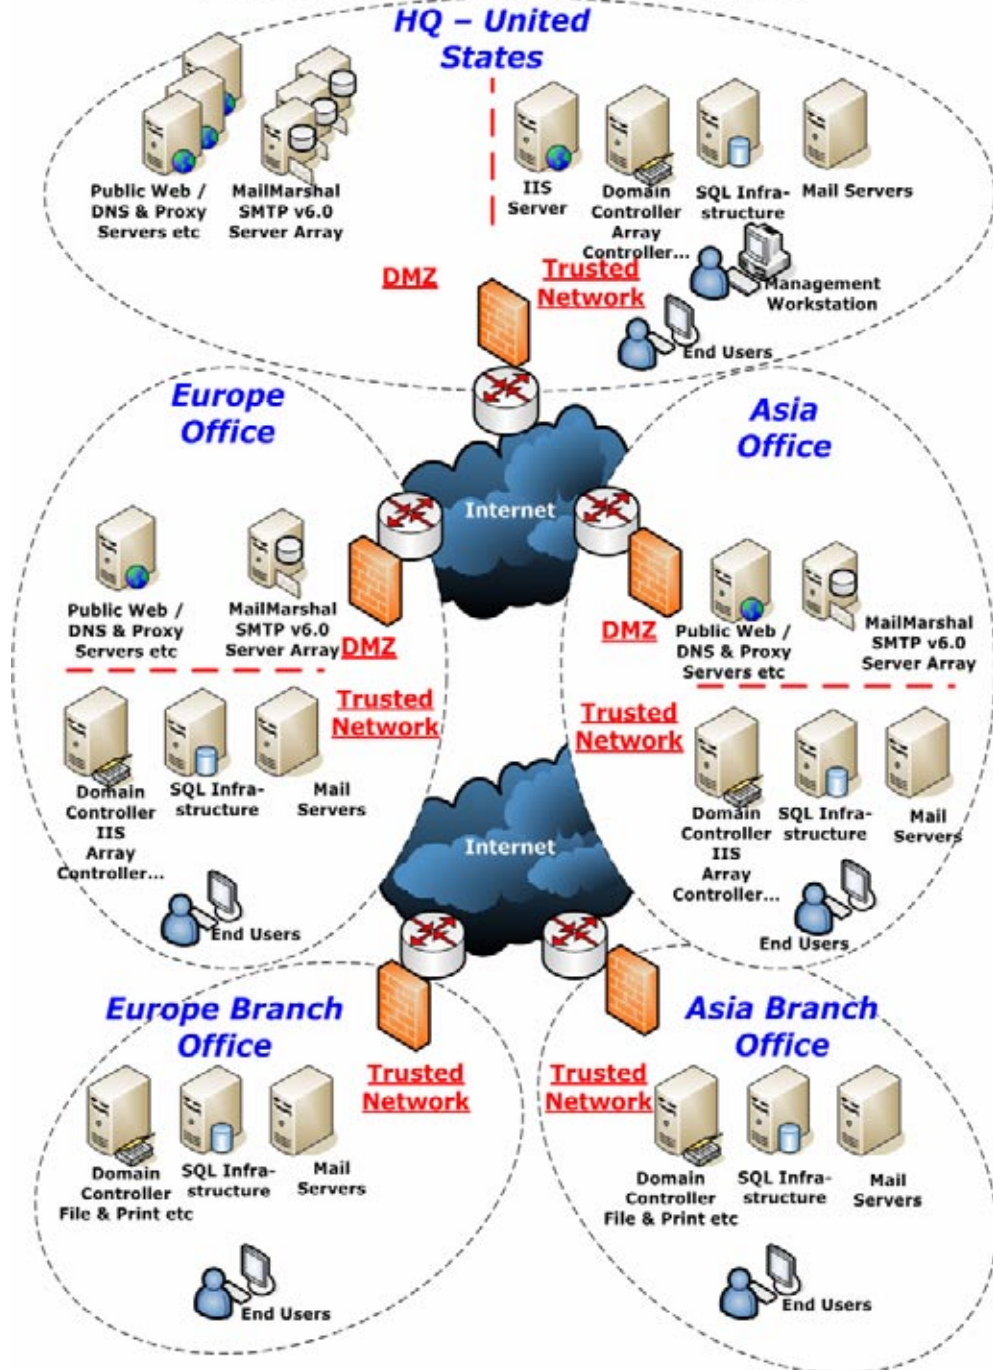
For this last example an Enterprise company is defined as having 50,000 employees spread over several countries, and multiple office locations in each country. As such they have multiple points for Internet access where SMTP mail and other internet traffic is routed. These multiple points are set up to load balance between each other based on load and location.

All the sites are joined to each other by an Internet based VPN solution. There is one main office for each country that handles the core services for that country such as Internet access but all the core services in each country are managed from the world headquarters.

The Configuration of this solution is also the same as option 5 of the Installation scenarios.

Diagram for this configuration is on the following page.

Enterprise Company Design Example





Marshal's Worldwide and EMEA HQ
Marshal Limited,
Renaissance 2200,
Basing View,
Basingstoke,
Hampshire RG21 4EQ
United Kingdom

Phone: +44 (0) 1256 848080
Fax: +44 (0) 1256 848060

Email: emea.sales@marshal.com

Americas
Marshal, Inc.
5909 Peachtree-Dunwoody Rd
Suite 770
Atlanta
GA 30328
USA

Phone: +1 404-564-5800
Fax: +1 404-564-5801

Email: americas.sales@marshal.com
info@marshal.com | www.marshal.com

Asia-Pacific
Marshal Software (NZ) Ltd
Suite 1, Level 1, Building C
Millennium Centre
600 Great South Road
Greenlane, Auckland
New Zealand

Phone: +64 9 984 5700
Fax: +64 9 984 5720

Email: apac.sales@marshal.com