Symantec™ Data Loss Prevention MTA Integration Guide for Network Prevent (Email)

Version 11.0

Symantec

Symantec Data Loss Prevention MTA Integration Guide for Network Prevent (Email)

Documentation version: 11.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrades protection
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our web site at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Managed Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Education Services	Education Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about enterprise services, please visit our web site at the following URL:

www.symantec.com/business/services/

Select your country or language from the site index.

Contents

Chapter 1

Introduction

This chapter includes the following topics:

About the Network Prevent (Email) Server

Network Prevent (Email) Server is the detection server that analyzes email messages and blocks or modifies them as required by your policies. It can receive email messages from one or more Mail Transfer Agents (MTAs) in your network.

The Network Prevent (Email) Server supports SMTP error response relay, the SMTP command verb `EHLO`, and the following extensions to SMTP:

- `8BITMIME`
- `VRFY`
- `DSN`
- `HELP`
- `PIPELINING`
- `SIZE`
- `ENHANCEDSTATUSCODES`
- `STARTTLS`

The Network Prevent (Email) Server does not store messages locally, and it is therefore not an MTA. The Network Prevent (Email) Server is never the only message handler holding the message, because it maintains each inbound SMTP message transaction only until the outbound transaction has been closed.

Network Prevent (Email) Server can receive TLS-encrypted email from an upstream MTA, and can initiate a TLS session to an outbound MTA, a hosted email service, or the reflecting-mode MTA as necessary.

See "About integration architectures" on page 17.

See "About TLS authentication" on page 30.

Note: You must implement the Network Prevent (Email) Server only into your outbound SMTP message stream.

Operating modes for Network Prevent (Email) Server

You can configure the Network Prevent (Email) Server to operate in one of the following modes:

- Reflecting mode
  In reflecting mode, the Network Prevent (Email) Server receives messages from an MTA, analyzes them, and then reflects them back to the same MTA.
- Forwarding Mode
  In forwarding mode, the Network Prevent (Email) Server receives messages from an upstream MTA, analyzes them, and sends them on to a downstream MTA or to a hosted email service such as the MessageLabs Email Content Control Service.

See "Configuring the Network Prevent Server (Email)" in the Symantec Data Loss Prevention Administration Guide for information about configuring either mode.

See "Configuring Network Prevent (Email) Server for reflecting or forwarding mode" on page 38.

About hosted Network Prevent deployments

Symantec Data Loss Prevention supports deploying one or more Network Prevent detection servers in a hosted service provider network, or in a network location that requires communication across a Wide Area Network (WAN). You may want to deploy a Network Prevent server in a hosted environment if you use a service provider's mail server or Web proxy. In this way, the Network Prevent server can be easily integrated with the remote proxy to prevent confidential data loss via email or HTTP posts.

The Enforce Server and all other detection servers must reside in the corporate network and communicate over a LAN. Only Network Prevent (Email) and Network Prevent (Web) can be deployed to a hosted environment.

When you choose to install a detection server, the Symantec Data Loss Prevention installation program asks if you want to install Network Prevent in a hosted environment.

If you choose to install a Network Prevent detection server in a hosted environment, you must use the `sslkeytool` utility to create multiple, user-generated certificates to use with both internal (corporate) and hosted detection servers. This ensures secure communication from the Enforce server to the hosted Network Prevent server, and to all other detection servers that you install. You cannot use the built-in Symantec Data Loss Prevention certificate when you deploy a hosted Network Prevent detection server.

The *Symantec Data Loss Prevention Installation Guide* describes how to install and configure the Network Prevent server in either a LAN environment or a hosted environment.

Environment compatibility and requirements for Network Prevent (Email)

The Network Prevent (Email) Server is compatible with a wide range of enterprise-grade third-party SMTP-compliant MTAs and hosted email services. Consult your MTA vendor or hosted email service for specific support questions.

Network Prevent (Email) Server can integrate with an MTA or hosted email service that meets the following requirements:

- The MTA or hosted email service must be capable of strict SMTP compliance. It must be able to send and receive mail using only the following command verbs: HELO (or EHLO), RCPT TO, MAIL FROM, QUIT, NOOP, and DATA.
- When running the Network Prevent (Email) Server in reflecting mode, the upstream MTA must be able to route messages to the Network Prevent (Email) Server once and only once for each message.

In practice, these requirements mean that you can use an SMTP-compliant MTA that can route outbound messages from your internal mail infrastructure to the Network Prevent (Email) Server. For reflecting mode compatibility, the MTA must also be able to route messages that are returned from the Network Prevent (Email) Server out to their intended recipients.

Network Prevent (Email) Server attempts to initiate a TLS connection with a downstream MTA only when the upstream MTA issues the STARTTLS command. The TLS connection succeeds only if the downstream MTA or hosted email service supports TLS and can authenticate itself to the Network Prevent (Email) Server. Successful authentication requires that the appropriate keys and X509 certificates are available for each mail server in the proxied message chain.

See "About TLS authentication" on page 30.

About selecting an integration architecture

This manual describes several suggested integration architectures for Network Prevent (Email) Server.

See "About integration architectures" on page 17.

The architecture you implement depends on your existing messaging architecture, the capabilities of your MTA, and your organization's messaging needs. Work closely with your messaging team to identify the best solution for your environment. You may decide that the best solution for your environment requires an integration architecture other than one of those that Symantec suggests.

Chapter 2

Network Prevent (Email) Server Response Rules

This chapter includes the following topics:

About Network Prevent (Email) response rules

This chapter describes the behavior and functionality of Network Prevent (Email) Server. It discusses blocking and redirecting messages, as well as tagging message headers.

Network Prevent (Email) Server monitors and analyzes outbound email traffic in-line and (optionally) blocks, redirects, or modifies email messages as specified in your policies. You create polices for the Network Prevent (Email) Server in the Enforce Server administration console. Policy authors can configure a policy for prevention (in-line management) or for monitoring only on a per-policy basis.

See the *Symantec Data Loss Prevention Administration Guide* for details on creating response rules and policies.

About message blocking

You can configure Network Prevent (Email) Server to block delivery of those messages that violate a policy. Network Prevent (Email) Server blocks messages by returning an SMTP 5xx failure response code.

You can also specify that a customized non-delivery report be sent back to the message sender when a message is blocked. To use a non-delivery report, create a Block SMTP Message response rule in the Enforce Server administration console. The non-delivery report (or bounced message) contains whatever text you specify in the response rule. The report is generated by the MTA at the moment the message is blocked.

MTA-generated non-delivery reports are different from sender notifications, which can be configured as another type of response rule. The Enforce Server generates sends sender notifications. When connectivity between the Network Prevent (Email) Server and the Enforce Server is normal, only a few seconds should elapse before the sender notification message is generated. However, if connectivity is interrupted between the Network Prevent (Email) Server and Enforce Server, the sender notification message is not generated until connectivity is restored.

You can configure email message blocking and Enforce server-generated sender notification actions in the Enforce Server administration console on the **Add/Edit Response Rule** screen. Then you can include response rules in the appropriate policies.

For details on response rules and policies, see the *Symantec Data Loss Prevention Administration Guide*.

About messages redirecting

You can redirect messages violating a policy to an address that is configured in a Block SMTP Message response rule. This address is typically a mailbox or list used by administrators or managers to review and release the messages. These mailboxes are outside of the Symantec Data Loss Prevention system. For this feature to work correctly, you must configure all such redirect addresses as individual sender exceptions on each Prevent-integrated MTA or hosted email service.

Keep redirect addresses in policies synchronized with sender exception addresses configured on the Prevent-integrated MTAs or hosted email service.

To enable and configure message redirection in a Block SMTP Message response rule, enter an address in the **Redirect Message to this Address** field of the Enforce Server administration console's **Add/Edit Response Rule** page.

About downstream message tagging

Gateway-based message encryption systems can be configured to take specified actions based on keywords in the message subject. Certain RFC-2822 message

headers can be used for the same purpose. The typical practice is to specify actions based on extension headers that start with "X-".

You can configure a policy to modify a message in one or all of the following ways:

- Replace, append, or change the beginning of the subject line.
- Generate a new header that can trigger further processing in the Prevent-integrated MTA or hosted email service. The processing may include message encryption, message quarantine, message archiving, or some other action.

If your MTA or hosted email service is capable of interpreting headers to process message routing rules, you can configure further actions to perform when violations are detected. Create a Modify SMTP Message response rule on the Enforce Server **Add/Edit Response Rule** screen. You can add up to three RFC 2822 header lines. Symantec recommends using the `-Cfilter` header with different values depending upon the scan verdict. You may also change or replace the Subject header.

Table 2-1 shows some common applications of these headers.

Table 2-1 Examples of Network Prevent (Email) Server-added headers

Example Header	Description
X-CFilter: Encrypt	Requests end-to-end encryption for the message.
X-CFilter: Quarantine	Requests quarantining for the message.
X-CFilter: Archive	Requests archiving for the message.

Be sure to keep the configuration of your message encryption system synchronized with the relevant details of Symantec Data Loss Prevention policies.

You can enable and configure message tagging for downstream encryption by creating Modify SMTP Message response rules on the **Add/Edit Response Rule** screen. Then you can use the rules to set up an appropriate incident remediation workflow on a per-policy basis.

Chapter 3

MTA Integration Architecture

This chapter includes the following topics:

About integration architectures

This chapter explains how the Network Prevent (Email) Server integrates into the message chain, and it describes several architectures for achieving this integration.

You can configure the Network Prevent (Email) Server to operate in either of the following modes:

- Reflecting mode
  In reflecting mode, the Network Prevent (Email) Server acts as an RFC-2821–compliant SMTP proxy. It receives messages from an MTA, analyzes them, and then sends them back to the same MTA. The Network Prevent (Email) Server blocks or modifies messages when your policies require it.
- Forwarding Mode
  In forwarding mode, the Network Prevent (Email) Server acts as an RFC-2821-compliant SMTP proxy that receives messages from an upstream

MTA. It analyzes messages and then sends them on to a downstream MTA or to a hosted email service such as the MessageLabs Email Content Control Service (instead of reflecting them back to the original MTA). Because the server supports SMTP error response relay and the DNS SMTP extension, it can relay message status as a proxy between two MTAs.

You can configure the Network Prevent Server (Email) to proxy messages to specific IP addresses or hostnames you specify in your server configuration. Or, you can configure Network Prevent (Email) Server to perform MX record lookups for the hostnames you specify in the configuration. By performing MX record lookups, Network Prevent (Email) Server can utilize DNS load balancing and failover capabilities when selecting the next hop MTA or hosted mail server.

About the Network Prevent (Email) Server message chain

The Network Prevent (Email) Server works by integrating into your organization's message chain.

Figure 3-1 shows an example implementation in which the Network Prevent (Email) Server operates in reflecting mode in the message chain.

Figure 3-1 The Network Prevent (Email) Server operating in reflecting mode

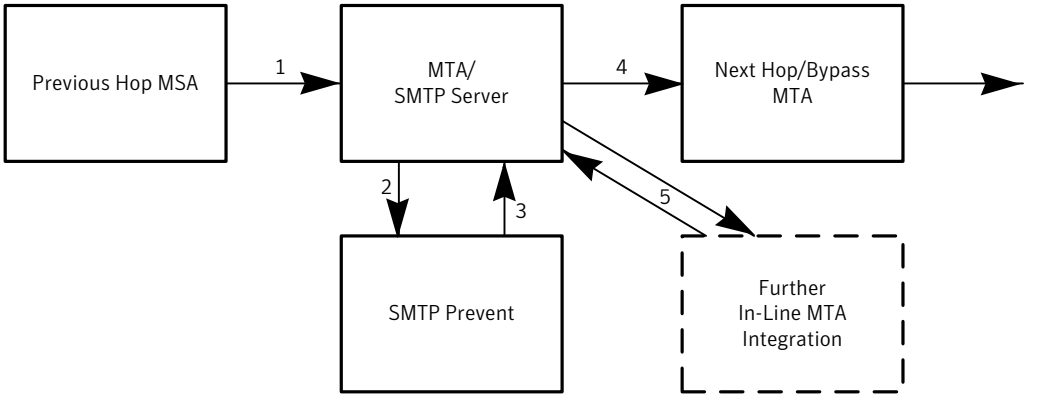

The following figure shows an implementation in which the detection server is operating in forwarding mode in the message chain.

Figure 3-2 The Network Prevent (Email) Server operating in forwarding mode

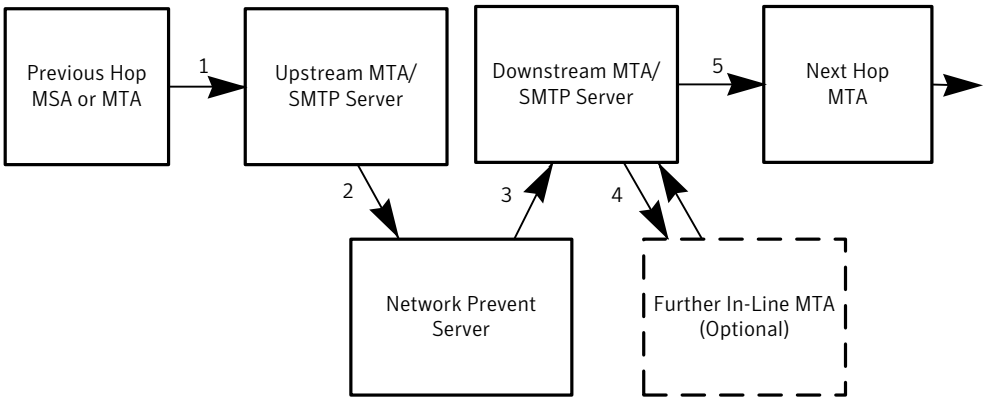

The following list describes the message chain for Figure 3-1 and Figure 3-2:

- A message submission agent (MSA) or an MTA sends an SMTP message to the Prevent-integrated MTA or the upstream MTA.
- Depending on your setup, one of the following occurs. The Prevent-integrated MTA determines that the message has not come from the Network Prevent (Email) Server, and the MTA routes it to that server. Or, the upstream MTA routes the message to the Network Prevent (Email) Server.
  When the Network Prevent (Email) Server receives the SMTP message, it analyzes the message against your Symantec Data Loss Prevention policies. The Network Prevent (Email) Server does not end the SMTP session until it forwards the message and ends the forwarding session.
- The Network Prevent (Email) Server handles the message in one of several ways, based on your policies and response rules as shown in Table 3-1.
- Optionally, the Prevent-integrated MTA or a downstream MTA can send the SMTP message to other in-line MTAs for further processing (such as encryption). The SMTP message returns to the MTA.
  For the Network Prevent (Email) Server to trigger message encryption, the Prevent-integrated MTA or the downstream MTA must be able to encrypt the message itself or route the message to an in-line MTA that encrypts. The Prevent-integrated MTA or downstream MTA must be able to use header information to determine the appropriate action.
- The Prevent-integrated MTA, or the downstream MTA or hosted email service provider, sends the SMTP message to the next-hop MTA or out to the Internet to a selected mail server.

Table 3-1 outlines how a response is triggered and how the message is handled.

Table 3-1 Message handling

Trigger	Configured Response	Message Handling
Message does not violate a policy	None	Symantec Data Loss Prevention sends the message (unchanged) back to the Prevent-integrated MTA or to the downstream MTA or hosted email service.
Message violates a policy	Block SMTP Message	This rule blocks the message by returning a 550 SMTP message to the Prevent-integrated MTA or to the upstream MTA. You can configure the 550 response text to contain a reason for the failure or a contact address. You can also configure Symantec Data Loss Prevention to replace the envelope recipient.
Message violates a policy	Modify SMTP Message	This rule lets you automatically modify or replace the message subject line and add as many as three SMTP headers to the message. Modified subject lines and extra SMTP headers can trigger downstream processing. Modified message subject lines can also make a message more user-friendly. See "About downstream message tagging" on page 14.
Message violates a policy	Send Email Notification	This rule lets you automatically send an incident email notification to a list of recipients and the original sender of the message.

Integration architectures for reflecting mode

Four integration architecture options are compatible with the Network Prevent (Email) Server operating in reflecting mode:

- Second SMTP Listener-Based Routing
  See "About second SMTP listener-based routing" on page 21.
- SMTP Client IP Address-Based Routing
  See "About SMTP client IP address-based routing" on page 23.

- HELO Identification String-Based Routing
  See "About HELO identification string-based routing" on page 24.
- Message Header-Based Routing
  See "About message header-based routing" on page 26.

The options are listed in order from most secure to least secure. You should choose the first integration architecture on the list that matches the capabilities of your MTA. If none of these integration architectures fit your message stream, that there may be other ways to successfully integrate the Network Prevent (Email) Server into your messaging system. Contact your messaging group and your MTA vendor to discuss alternative possibilities for integration.

Note: Symantec recommends you configure outbound messages automatically generated by the Prevent-integrated MTA to bypass the Network Prevent (Email) Server.

The following sections describe each of the integration options in detail.

About second SMTP listener-based routing

To implement Second SMTP Listener-Based Routing, you configure the Prevent-integrated MTA to determine how to process and route an SMTP message based on which of two TCP ports received the message.

You configure one SMTP listener (TCP port 25) to route SMTP messages received from the Prevent-integrated MTA to the Network Prevent (Email) Server—unless the message meets some criteria that requires it to bypass the Network Prevent (Email) Server. You configure the second SMTP listener (TCP port 10026) to listen for SMTP sessions only from Network Prevent (Email) Servers. Messages received on the SMTP Prevent listener are forwarded further on the message chain. The Second SMTP Listener-based routing is the most secure integration architecture.

Figure 3-3 shows SMTP listener port-based routing.

Figure 3-3 Second SMTP listener port-based routing

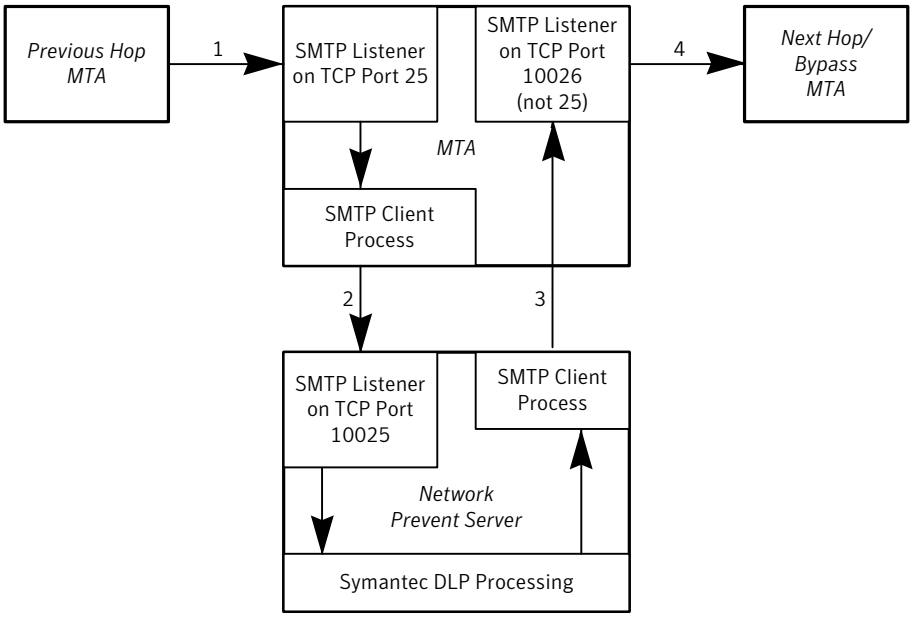

Details about SMTP listener port-based routing are as follows:

- The Prevent-integrated MTA SMTP listener receives a message on port 25.
- The Prevent-integrated MTA SMTP sender routes the message for inspection to the Network Prevent (Email) Server SMTP listener on TCP port 10025. (TCP port 10025 is the default port number. You can change this number.)
- The Network Prevent (Email) Server inspects the message and, if the server does not block the message (based on the relevant policy), it reflects the message back to the Prevent-integrated MTA on TCP port 10026. (The message is reflected back to TCP port 10026 by default, but you can set any port, other than 25, to receive the message.)
- The Prevent-integrated MTA SMTP listener receives the message on port 10026 and determines the message comes from a valid Network Prevent (Email) Server because of the TCP port number. Your message stream architecture and any headers the Network Prevent (Email) Server modifies determine the next hop in the message delivery.

About SMTP client IP address-based routing

To implement SMTP Client IP Address-Based Routing, you configure the Prevent-integrated MTA to determine how to process and route an SMTP message based on the IP address of the previous hop.

SMTP Client IP Address-based routing is secure as long as the IP address is obtained from a reliable source. The most reliable source of the true IP address is directly from the TCP connection information, rather than extracting the IP address from header information, which can be forged. If you use an alternative method (for example, reading the received header that the SMTP Listener placed into the message), then the SMTP Client IP Address-based routing integration architecture is only as secure as the SMTP Listener's ability to ascertain the IP address.

Figure 3-4 shows SMTP client IP-based routing.

Figure 3-4 SMTP client IP address-based routing

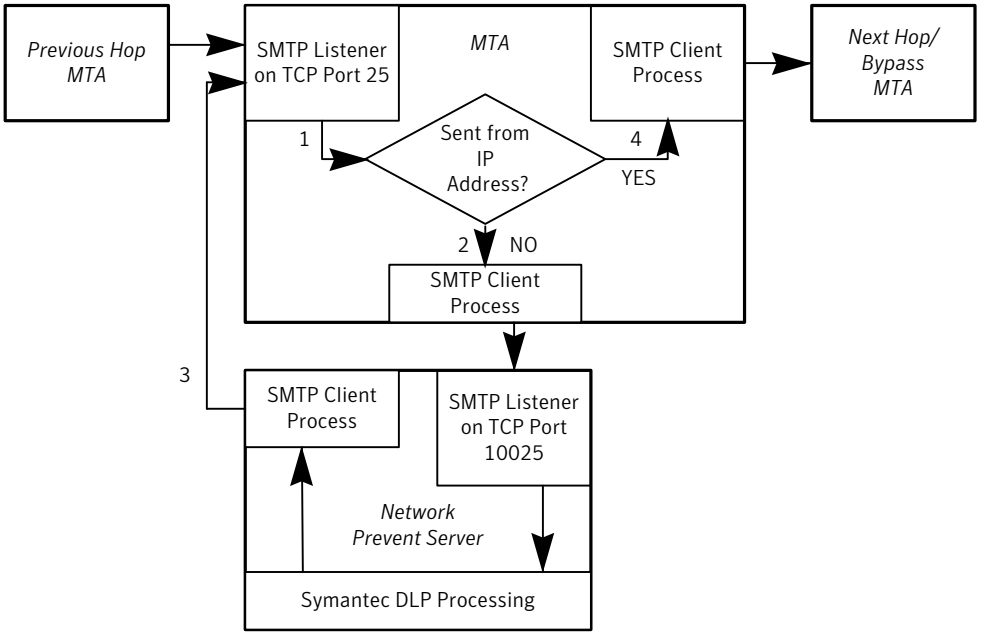

Details about SMTP client IP-based routing are as follows:

- The Prevent-integrated MTA SMTP listener receives a message on port 25.
- The Prevent-integrated MTA examines the message to determine the sender IP address. If the sender IP address does not match an IP on the Prevent-integrated MTA delivery list of Network Prevent (Email) Server IP

addresses, then the Prevent-integrated MTA routes the message for inspection to the Network Prevent (Email) Server SMTP listener on TCP port 10025. (TCP port 10025 is the default port number. You can change this number.)

- The Network Prevent (Email) Server inspects the message and, if the server does not block the message (based on the relevant policy), then the message is reflected back to the Prevent-integrated MTA on TCP port 25.
- The Prevent-integrated MTA determines, based on the IP address, whether or not the message comes from a valid Network Prevent (Email) Server. It matches the Network Prevent (Email) Server SMTP client IP address against the Prevent-integrated MTA list of SMTP client IP addresses. If the sender IP address matches an IP on the Prevent-integrated MTA list of Network Prevent (Email) Server IP addresses, then the Prevent-integrated MTA processes the message based upon the appropriate next hop. Your message stream architecture and any headers the Network Prevent (Email) Server modifies determine the next hop in the message delivery.

About HELO identification string-based routing

To implement the HELO Identification String-Based Routing integration architecture, you configure the Prevent-integrated MTA to determine how to process SMTP messages based on the HELO identification string.

Determining the next hop based on the HELO identification string is a relatively secure way to integrate the Network Prevent (Email) Server into your message stream. It is difficult to alter the HELO response of an email client; however, forcing the use of an IP address is more secure.

Figure 3-5 shows HELO Identification String-Based Routing.

Figure 3-5 HELO identification string-based routing

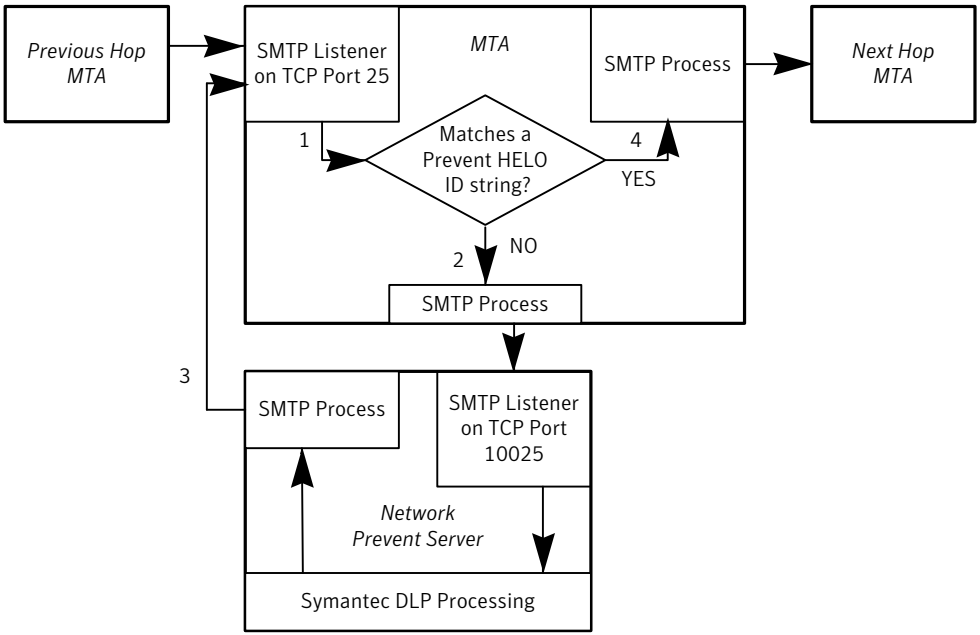

Details about HELO identification String-Based Routing are as follows:

- The Prevent-integrated MTA SMTP listener receives a message on port 25, and the Prevent-integrated MTA captures the HELO ID string.
- The Prevent-integrated MTA examines the message to determine the sender's HELO ID string. If the sender HELO ID string does not match a HELO ID string on the Prevent-integrated MTA list of Network Prevent (Email) Server HELO ID strings, then the Prevent-integrated MTA routes the message for inspection to the Network Prevent (Email) Server SMTP listener on TCP port 10025. (TCP port 10025 is the default port number. You can change this number.)
- The Network Prevent (Email) Server inspects the message and, if the server does not block the message (based on the relevant policy), it reflects the message back to the Prevent-integrated MTA on TCP port 25.
- The Prevent-integrated MTA examines the message to determine the sender's HELO ID string. If the sender HELO ID string matches a HELO ID string on the Prevent-integrated MTA list of Network Prevent (Email) Server HELO ID strings, then the Prevent-integrated MTA processes the message based upon the appropriate next hop. Your message stream architecture and any headers the Network Prevent (Email) Server modifies determine the next hop in the message delivery.

About message header-based routing

To implement the Message Header-Based Routing integration architecture, you configure the Prevent-integrated MTA to determine how to process SMTP messages based on the existence of an informational message header.

If your MTA uses message headers to determine the next hop for a message, a user can easily circumvent detection using common Mail User Agents (MUA). Symantec recommends choosing another integration method if any is available; however, this method does provide a fully functional integration.

Figure 3-6 shows Message Header-Based Routing.

Figure 3-6 Message header-based routing

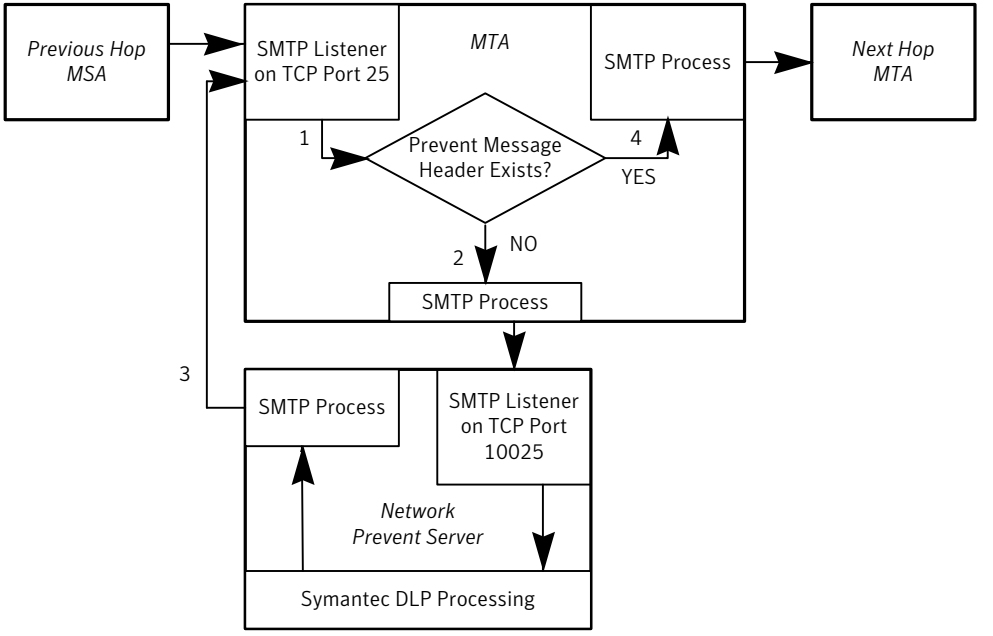

Details about message header-based routing are as follows:

- The Prevent-integrated MTA SMTP listener receives a message on port 25.
- The Prevent-integrated MTA examines the message's headers. If the MTA finds no header inserted by the Network Prevent (Email) Server, then the MTA routes the message for inspection to the Network Prevent (Email) Server SMTP listener on TCP port 10025. (TCP port 10025 is the default port number. You can change this number.)

- The Network Prevent (Email) Server inspects the message and, if the server does not block the message (based on the relevant policy), then it inserts a header and reflects the message back to the Prevent-integrated MTA on TCP port 25.
- The Prevent-integrated MTA examines the message to determine whether the Network Prevent (Email) Server header exists. If the header exists, then the Prevent-integrated MTA processes the message based upon the appropriate next hop. If the Network Prevent (Email) Server added additional headers to the message, then the Prevent-integrated MTA might make a different routing decision. Your message stream architecture and any headers the Network Prevent (Email) Server modifies determine the next hop in the message delivery.

About the integration architecture for forwarding mode

In forwarding mode, the Network Prevent (Email) Server operates as an SMTP proxy between an upstream MTA and a downstream MTA or hosted email service provider. The Network Prevent (Email) Server relays responses from the downstream host back to the upstream MTA. To configure the Network Prevent (Email) Server to operate in forwarding mode, you must select the **Forward** option in the Inline SMTP section of the Network Prevent (Email) Server configuration page and configure the **Next MTA** field.

See "Configuring Network Prevent (Email) Server for reflecting or forwarding mode" on page 38.

Figure 3-7 shows the Network Prevent (Email) Server in forwarding mode.

Figure 3-7 Architecture of Network Prevent (Email) Server operating in forwarding mode

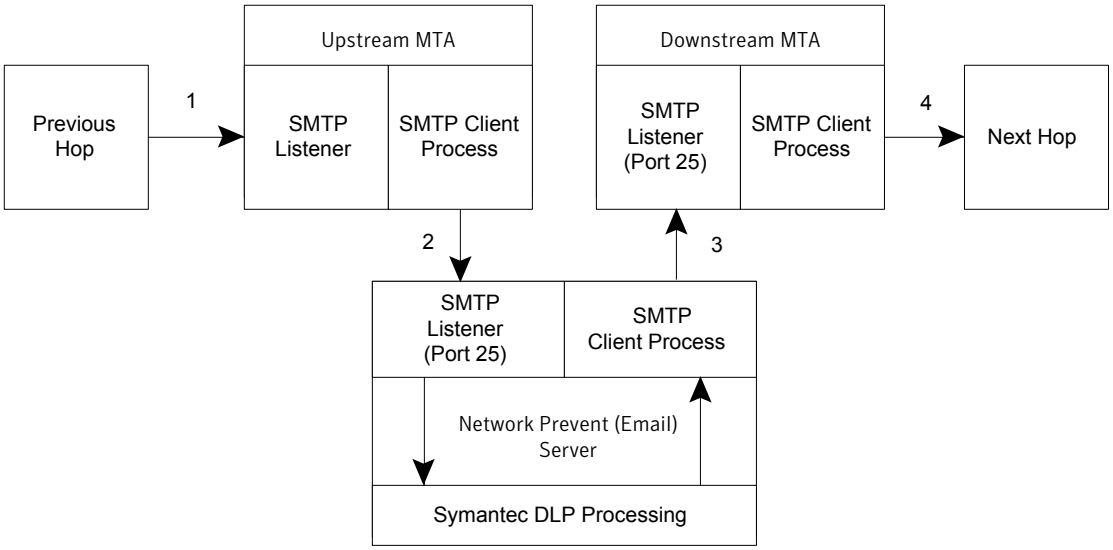

Details about Network Prevent (Email) Server operating in forwarding mode are as follows:

- The upstream MTA SMTP listener receives a message.
- The upstream MTA routes the message to the Network Prevent (Email) Server SMTP listener on TCP port 25. (By default, the configured TCP port number is 10025, but it is recommended to change this to 25.)
- The Network Prevent (Email) Server inspects the message and, if the server does not block the message (based on a configured policy), it proxies the message to a downstream MTA or hosted email server. The IP address address for this next hop in the message chain can either be specified in the Network Prevent (Email) Server configuration or obtained through an MX record lookup of a configured domain name.
- The downstream MTA or hosted email server sends the SMTP message to the next-hop MTA hosted mail server. A hosted email service may perform additional tasks, such as detecting viruses in the message or encrypting the message contents, before proxying to the receiving MTA.

About next-hop MTA selection

When Network Prevent (Email) Server proxies a message after performing detection, it uses a configured list of mail server addresses to determine the next-hop server. Network Prevent (Email) Server can use the mail server addresses (DNS names or IP addresses) as they are configured, or can perform MX-record lookups for a configured domain.

If MX-record lookups are not enabled, Network Prevent (Email) Server simply attempts to forward messages to the first configured mail server address in the list. If it cannot establish a connection to the mail server, it tries the subsequent server addresses in the listed order.

If you enable MX-record lookups, Network Prevent (Email) Server performs a DNS query to obtain the mail exchange (MX) records for a configured domain. Network Prevent (Email) Server uses the returned MX records to select the next-hop mail server in the proxy chain.

An MX record specifies the address of a mail server for a particular domain, as well as a MX preference number. The MX preference assigns priority to multiple MX records that are returned for the same domain. All SMTP clients, including Network Prevent (Email) Server, observe the following rules associated with the MX preference:

- Mail servers that have lower-numbered MX preferences are used before servers that have higher MX preferences.
- If multiple mail servers have the same MX preference, clients must randomize access to those mail servers.

For example, assume that you have configured Network Prevent (Email) Server to perform MX lookups, and that you have entered a single address, `emailcompanyname.com`, in the list of next-hop MTAs. Network Prevent (Email) Server performs a DNS lookup and receives the following MX records:

```
emailcompanyname.com   10  smtp1.emailcompanyname.com
emailcompanyname.com   10  smtp2.emailcompanyname.com
emailcompanyname.com   10  smtp3.emailcompanyname.com
emailcompanyname.com   40  smtp4.emailcompanyname.com
```

In the above case, Network Prevent (Email) Server chooses the first of the three servers with MX preference 10 (smtp1, smtp2, and smtp3) as the next hop MTA. The DNS `rrset-order` determines the order of MX records based on the configured load-balancing algorithm, which is generally cyclic (round-robin). The final server, smtp4, would be chosen only if none of the servers with MX preference 10 are available.

You use the Enforce console to configure the valid list of next-hop mail server addresses.

About TLS authentication

Network Prevent (Email) Server uses TLS with a downstream MTA only when:

- The upstream MTA requests a TLS connection using the STARTTLS command, and
- The downstream MTA or hosted email service supports TLS and can authenticate itself.

These conditions also apply when Network Prevent (Email) operates in reflecting mode, where a single MTA acts as both the upstream and downstream MTA.

When TLS is requested, each successive proxy in the email chain must authenticate itself to the previous server in order to establish an end-to-end TLS connection. Successful authentication requires that each mail server stores a valid certificate for the next-hop mail server in its trust store. For example, Network Prevent (Email) Server must authenticate itself to the sending MTA, and the downstream MTA or hosted email service must authenticate itself to Network Prevent (Email) Server.

Note: Each MTA performs its own authentication setup, and an MTA in the email chain could potentially choose to ignore certificate validation. This practice is not recommended for production email configurations.

If you configure an upstream MTA to bypass Network Prevent (Email) Server if the server is unavailable, then the upstream MTA must store a certificate for the downstream mail server as well as the certificate for Network Prevent (Email) Server.

Configuring keys and certificates for TLS

In a typical forwarding-mode MTA integration, the following keys and certificates are required to support TLS:

- The keystore of the upstream MTA must contain the public key certificate for Network Prevent (Email) Server. This key is required if the upstream MTA decides to authenticate Network Prevent (Email) as part of the TLS session.
- The Network Prevent (Email) Server keystore must contain its own private key as well as a public key certificate for the downstream MTA or hosted email server.

- If the upstream MTA is configured to bypass Network Prevent (Email) Server when the server is unavailable, the upstream MTA trust store must also contain a valid certificate for the downstream MTA or hosted email service.

In a reflecting-mode MTA integration, a single MTA acts as both the upstream MTA and downstream MTA. The reflecting-mode MTA must contain the public key certificate for the Network Prevent (Email) Server. The Network Prevent (Email) Server keystore must contain its own private key as well as the public key certificate for the integrated reflecting-mode MTA. If the reflecting-mode MTA is configured to bypass Network Prevent (Email) Server when the server is unavailable, the MTA trust store must also contain a valid certificate for the downstream MTA or hosted email service.

Hosted email servers generally use a public key certificate that is digitally signed by a root certificate authority (CA). You must obtain the CA-signed public key certificate from your hosted email service provider and add it to the Network Prevent (Email) Server keystore for forwarding-mode configurations. Add the key to the reflecting-mode MTA keystore in reflecting-mode configurations.

Any certificate that you add to the Network Prevent (Email) keystore must be an X.509 certificate in Private Enhanced Mail (`.pem`) Base64-encoded Distinguished Encoding Rules (DER) certificate format, enclosed within `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` strings in the certificate file.

Table 3-2 outlines the process of setting up the required keys and certificates.

Table 3-2 Configuring keys and certificates for TLS

Step	Action	Description
Step 1	Change the default keystore password for Network Prevent (Email) Server.	Use the Java `keytool` utility to change the default Network Prevent (Email) Server keystore password to a secure password. Then use the Enforce Server administration console to configure Network Prevent (Email) Server to use the updated password. See "Changing the Network Prevent (Email) Server keystore password" on page 33.

Table 3-2 Configuring keys and certificates for TLS *(continued)*

Step	Action	Description
Step 2	Generate the key pair for Network Prevent (Email) Server.	Use the Java `keytool` utility to generate a public/private key pair for Network Prevent (Email) Server. See "Generating Network Prevent (Email) Server keys" on page 34.
Step 3	Export the public key certificate from the Network Prevent (Email) Server keystore.	Use the `keytool` utility to export the self-signed certificate for the public key you generated in Step 2. See "Exporting the Network Prevent (Email) Server public key certificate" on page 36.
Step 4	Import the Network Prevent (Email) Server public key certificate into the upstream MTA keystore or reflecting-mode MTA keystore.	Use `keytool` to import the public key certificate file you exported in Step 3 into the upstream MTA keystore. This enables the MTA to authenticate Network Prevent (Email) Server for TLS communication. See your MTA documentation for instructions about how to import public key certificates.

Table 3-2 Configuring keys and certificates for TLS *(continued)*

Step	Action	Description
Step 5	Obtain the public key certificate for the next-hop MTA or hosted email service.	Obtain the public key certificate file for any next-hop MTA that you manage in the network. See your MTA documentation for instructions about how to export the certificate. If you are accessing an external, hosted email server as the next hop in the TLS proxy chain, obtain the public key certificate from your provider. See your email hosting service provider documentation for instructions.
Step 6	For forwarding-mode integrations, add the next-hop public key certificate to the Network Prevent (Email) Server keystore.	Use the Java `keytool` utility to import the downstream MTA's or hosted email server's public key certificate into the Network Prevent (Email) Server keystore. See "Importing public key certificates to the Network Prevent (Email) Server keystore" on page 37.
Step 7	For reflecting-mode integrations, add the next-hop public key certificate to the reflecting-mode MTA keystore.	See your MTA documentation for instructions about how to import public key certificates.

Changing the Network Prevent (Email) Server keystore password

When you install Network Prevent (Email) Server, the installer creates an empty keystore file in *installdir*`\Vontu\Protect\keystore\prevent.ks`. This keystore file has an initial password, **dummypassword**. Use the following procedure to change the keystore password.

Changing the Network Prevent (Email) Server keystore password

1. Change to the `c:\Vontu\jre\bin` directory on the Network Prevent (Email) Server computer.
2. Execute the `keytool` utility with the `-storepasswd` option to change the default password. For example:

   ```
   keytool -storepasswd -new prevent_keystore_password -keystore
   c:\Vontu\Protect\keystore\prevent.ks -storepass dummypassword
   ```

   Replace *prevent_keystore_password* with a secure password for the keystore. On Linux systems, the default keystore location is `/opt/Vontu/Protect/keystore/prevent.ks`.

   Note: The Network Prevent (Email) Server keystore password and key password values must match. Remember to use the same *prevent_keystore_password* when you generate the key for Network Prevent (Email) Server.

   See "Generating Network Prevent (Email) Server keys" on page 34.

3. Log onto the Enforce console that manages Network Prevent (Email) Server.
4. Select **System > Servers > Overview** from the main menu bar.
5. Click the name of the Network Prevent (Email) Server you want to configure.
6. Click **Configure**.
7. In the **Security Configuration** section, fill in the fields as follows:

Field	Description
Keystore Password	Enter the correct password for the keystore file. (Use the *new_password* you specified in Step 2.)
Confirm keystore Password	Re-enter the keystore file password.

8. Click **Save**.

Generating Network Prevent (Email) Server keys

Each mail server that you manage must have a keystore that contains the keys and X.509 certificates required to authenticate TLS communication. Use the

following procedure to create a new public and private key pair in the keystore file.

Creating a public and private key pair for Network Prevent (Email) Server

1 Change to the `c:\Vontu\jre\bin` directory on the Network Prevent (Email) Server computer.

2 Execute the `keytool` utility with the `-genkeypair` and `-keystore` options to add a new public and private key to the keystore :

```
keytool -genkeypair -dname "dname_string" -alias smtp_prevent
-keypass key_password -keystore c:\Vontu\Protect\keystore\prevent.ks
-storepass store_password -validity expiration_days
```

Table 3-3 describes the tokens that are used in the command.

For example, the following command generates a new key pair that expires in 90 days:

```
keytool -genkeypair -dname "CN=John Doe, OU=DLP_Development, O=Symantec
L=SanFrancisco, S=California, C=USA" -alias smtp_prevent
-keypass prevent_keystore_password
-keystore c:\Vontu\Protect\keystore\prevent.ks
-storepass prevent_keystore_password -validity 90
```

3 Export the public key certificate that you just created. You must import the certificate to any upstream MTAs that need to authenticate Network Prevent (Email) Server in the TLS session.

See "Exporting the Network Prevent (Email) Server public key certificate" on page 36.

Table 3-3 Keytool token reference

Token	Description
dname_string	The X.500 distinguished name to bind with the public key. The distinguished name generally contains a series of codes for the common name of the person, the organization and organizational unit, and the location associated with the key. For example: `-dname "CN=John Doe, OU=DLP_Development, O=Symantec, L=SanFrancisco, S=California, C=USA"` See the `keytool` help or Sun `keytool` documentation for more information about the format of a distinguished name string.

Table 3-3 Keytool token reference *(continued)*

Token	Description
smtp_prevent	The alias for the new key.
key_password	The password for the new key you are creating. **Note:** The *key_password* and *store_password* values must be identical for Network Prevent (Email) Server. See "Changing the Network Prevent (Email) Server keystore password" on page 33.
store_password	The password to modify the keystore file. **Note:** The *key_password* and *store_password* values must be identical for Network Prevent (Email) Server. See "Changing the Network Prevent (Email) Server keystore password" on page 33.
expiration_days	The number of days before the new key pair becomes invalid.

Exporting the Network Prevent (Email) Server public key certificate

In order to authenticate Network Prevent (Email) Server, the upstream MTA (or reflecting-mode MTA) must store a public key certificate for Network Prevent (Email) Server in its local keystore. Follow the procedure below to export the public key certificate for Network Prevent (Email) Server to a file. You can then import the certificate from the file into the keystore for your upstream MTA.

Exporting the public key certificate

1. Change to the `c:\Vontu\jre\bin` directory on the Network Prevent (Email) Server computer.
2. Execute the `keytool` utility with the `-exportcert` option to export the public key certificate to a new file:

```
keytool -exportcert -alias smtp_prevent -file smtp_prevent.cer
-keystore c:\Vontu\Protect\keystore\prevent.ks -storepass prevent_key_password
```

   In the above command, *smtp_prevent.cer* is the filename in which you store the public key certificate and *prevent_key_password* is the password to the keystore and the Network Prevent (Email) Server key.

3. Import the public key certificate into the keystore of each upstream MTA that must authenticate Network Prevent (Email) Server, or to the reflecting-mode MTA. See your MTA documentation for instructions.

Importing public key certificates to the Network Prevent (Email) Server keystore

Each mail server in the TLS proxy chain must authenticate the next-hop mail server. Authentication requires that you add the next-hop mail server certificate to the upstream mail server trust store. Follow the procedure below to import a next-hop MTA server or hosted email server public key certificate into the Network Prevent (Email) Server keystore.

Note: Any certificate that you add to the Network Prevent (Email) keystore must be an X.509 certificate in Private Enhanced Mail (`.pem`) Base64-encoded Distinguished Encoding Rules (DER) certificate format, enclosed within `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` strings in the certificate file.

Importing public key certificates

1. Begin by copying the certificate file you want to import onto the Network Prevent (Email) Server computer.

   If you manage the next-hop MTA in your network, refer to the MTA documentation for information about exporting the public key certificate.

   If you use a hosted email service as the next-hop server, consult your service provider for information about obtaining the certificate.

2. Change to the `c:\Vontu\jre\bin` directory on the Network Prevent (Email) Server computer.

3 Execute the `keytool` utility with the `-importcert` option to import the public key certificate into the Network Prevent (Email) Server keystore:

```
keytool -importcert -alias new_mta_alias
-file certificate_file
-keystore c:\Vontu\Protect\keystore\prevent.ks -storepass prevent_key_password
```

In the above command, *certificate_file* is the full path to the public key certificate file you want to import and *store_password* is the keystore password. *new_mta_alias* is a new alias to assign to the imported certificate.

If you are importing a public key chain that includes a root CA certificate, include the `-trustcacerts` option to verify the full chain, as in:

```
keytool -importcert -alias prevent_alias -file .\smtp_prevent.cer
-keystore c:\Vontu\Protect\keystore\prevent.ks -trustcacerts
```

4 Repeat the above commands for each MTA or hosted email server that Network Prevent (Email) Server might need to authenticate.

Configuring Network Prevent (Email) Server for reflecting or forwarding mode

Use the following instructions to configure Network Prevent (Email) Server to operate either in reflecting or forwarding mode.

To configure the Network Prevent (Email) Server

1 Log on to the Enforce Server administration console for the Symantec Data Loss Prevention system you want to configure.

2 Select **System > Servers > Overview** to display the list of configured servers.

3 Click the name of the Network Prevent (Email) Server that you want to configure.

4 Click **Configure**.

5 Deselect **Trial Mode** to enable blocking of email messages that are found to violate Symantec Data Loss Prevention policies.

6 Configure reflecting mode or forwarding mode by modifying the following fields:

Field	Description
Next Hop Configuration	Select **Reflect** to operate Network Prevent (Email) Server in reflecting mode. Select **Forward** to operate in forwarding mode. **Note:** If you select **Forward** you must also select **Enable MX Lookup** or **Disable MX Lookup** to configure the method used to determine the next-hop MTA.
Enable MX Lookup	This option applies only to forwarding mode configurations. Select **Enable MX Lookup** to perform a DNS query on a domain name to obtain the mail exchange (MX) records for the server. Network Prevent (Email) Server uses the returned MX records to select the address of the next hop mail server. If you select **Enable MX Lookup**, also add one or more domain names in the **Enter Domains** text box. For example: `companyname.com` Network Prevent (Email) Server performs MX record queries for the domain names that you specify. **Note:** You must include at least one valid entry in the **Enter Domains** text box to successfully configure forwarding mode behavior. See "About next-hop MTA selection" on page 29.

Field	Description

Disable MX Lookup	This field applies only to forwarding mode configurations.
	
	Select **Disable MX Lookup** if you want to specify the exact hostname or IP address of one or more next-hop MTAs. Network Prevent (Email) Server uses the hostnames or addresses that you specify and does not perform an MX record lookup.
	
	If you select **Disable MX Lookup**, also add one or more hostnames or IP addresses for next-hop MTAs in the **Enter Hostnames** text box. You can specify multiple entries by placing each entry on a separate line. For example:
	
	`smtp1.companyname.com`
	`smtp2.companyname.com`
	`smtp3.companyname.com`
	
	Network Prevent (Email) Server always tries to proxy to the first MTA that you specify in the list. If that MTA is not available, Network Prevent (Email) Server tries the next available entry in the list.
	
	Note: You must include at least one valid entry in the **Enter Hostnames** text box to successfully configure forwarding mode behavior.
	
	See "About next-hop MTA selection" on page 29.

7 Click **Save**.

8 Click **Server Settings** to verify or configure these advanced settings:

Field	Description
RequestProcessor.ServerSocketPort	Ensure that this value matches the number of the SMTP Listener port to which the upstream MTA sends email messages. The default is 10025. **Note:** Many Linux systems restrict ports below 1024 to root access. Network Prevent (Email) cannot bind to these restricted ports. If the computer receives mail for inspection on a restricted port (for example, port 25), reconfigure the computer to route traffic from the restricted port to the non-restricted Network Prevent (Email) port (port 10025 by default). See "Configuring Linux IP tables to reroute traffic from a restricted port" on page 42.
RequestProcessor.MTAResubmitPort	Ensure that this value matches the number of the SMTP Listener port on the upstream MTA to which the Network Prevent (Email) Server returns mail. The default is 10026.
RequestProcessor.AddDefaultHeader	By default, Network Prevent (Email) Server uses a header to identify all email messages that it has processed. The header and value are specified in the **RequestProcessor.DefaultPassHeader** field. Change the value of this field to false if you do not want to add a header to each message.

Field	Description
RequestProcessor.AddDefaultPassHeader	This field specifies the header and value that Network Prevent (Email) Server adds to each email message that it processes. The default header and value is `X-CFilter-Loop: Reflected`. Change the value of this field if you want to add a different header to each processed message. If you do not want to add a header to each email message, set the **AddDefaultPassHeader** field to False.

Note: Always configure both **RequestProcessor.ServerSocketPort** and **RequestProcessor.MTAResubmitPort**, whether you implement reflecting or forwarding mode. With forwarding mode, **RequestProcessor.ServerSocketPort** specifies the SMTP Listener port on the detection server to which the upstream MTA sends email messages. **RequestProcessor.MTAResubmitPort** is the SMTP Listener port on the downstream MTA to which the detection server sends email messages.

9 Click **Save**.

10 Click **Done**.

11 If your email delivery system uses TLS communication in forwarding mode, each next-hop mail server in the proxy chain must support TLS and must authenticate itself to the previous hop. This means that Network Prevent (Email) Server must authenticate itself to the upstream MTA, and the next-hop MTA must authenticate itself to Network Prevent (Email) Server. Proper authentication requires that each mail server stores the public key certificate for the next hop mail server in its local keystore file.

See "About TLS authentication" on page 30.

Configuring Linux IP tables to reroute traffic from a restricted port

Many Linux systems restrict ports below 1024 to root access. Network Prevent (Email) cannot bind to these restricted ports.

If the computer receives mail for inspection on a restricted port (for example, port 25), use the `iptables` command to route that traffic to a non-restricted port,

such as the Network Prevent (Email) default port 10025. Then ensure that Network Prevent (Email) listens on the non-restricted port to inspect email.

Use the following instructions to configure a Linux system to route from port 25 to port 10025. If you use a different restricted port or Network Prevent (Email) port, enter the correct values in the `iptables` commands.

To configure route traffic from port 25 to port 10025

1. Configure Network Prevent (Email) to use the default port 10025 if necessary.

   See "Configuring Network Prevent (Email) Server for reflecting or forwarding mode" on page 38.

2. In a terminal window on the Network Prevent (Email) computer, enter the following commands to reroute traffic from port 25 to port 10025:

```
iptables -N Vontu-INPUT
iptables -A Vontu-INPUT -s 0/0 -p tcp --dport 25 -j ACCEPT
iptables -I INPUT 1 -s 0/0 -p tcp -j Vontu-INPUT
iptables -t nat -I PREROUTING -p tcp --destination-port 25 -j REDIRECT --to-ports=10025
iptables-save > /etc/sysconfig/iptables
```

Note: If you only want to test local IP routing between the ports with Telnet, use the command: `iptables -t nat -I OUTPUT -o lo -p tcp --destination-port 25 -j REDIRECT --to-ports=10025`

If later you decide to delete the IP tables entry, use the command:

```
iptables -t nat -D OUTPUT -o lo -p tcp --destination-port 25 -j REDIRECT --to-ports=10025
```

Chapter 4

Capacity and Fault Tolerance

This chapter includes the following topics:

About capacity and fault tolerance

Your message architecture should be designed to accommodate a maximum message load even if one of your MTAs or Network Prevent Servers (Email) is temporarily unavailable (whether because of maintenance or for other reasons). This chapter provides tips and suggestions for integrating one or more servers into your message architecture while managing capacity and maintaining fault tolerance. You can tailor the exact capacity and fault tolerance specifications to your own requirements. In this chapter, the term message handler refers to any MTA, Network Prevent Server (Email), or other in-line SMTP processor in your message architecture.

About capacity management and fault tolerance implementation

This section introduces common methods for managing capacity and implementing fault tolerance. Subsequent sections provide more details as well as example implementations of these methods in your message architecture.

A common way to add capacity and fault tolerance to your architecture is to create or expand clusters, which are sets of load sharing systems that perform the same step in your message architecture. You can cluster multiple MTAs and multiple Network Prevent Servers (Email). You can increase the number of Network Prevent Servers (Email) independently of the number of MTAs.

Some ways to add capacity and fault tolerance to your architecture are as follows:

- MX-based clusters are useful for managing capacity and implementing fault tolerance. You can use mail exchange records (MX records) to distribute email services over a cluster of message handlers. This method of clustering works both for MTAs and Network Prevent Servers (Email). To create this kind of cluster, each message handler must have SMTP listeners on the same TCP ports and each must process inbound messages in exactly the same way. The MX record for each message handler in the cluster must have equal MX preference values.
  See "About the DNS system" on page 57.
- IP load balancer-based clusters are useful for managing capacity. IP load balancers are devices that provide a virtual IP address that distributes traffic among several back-end servers over an internal IP network.
- MTA-based queue management is useful for implementing fault tolerance. One or more MTAs in a cluster may be able to check the age of messages queued for a Network Prevent Server (Email). If any messages have been in the queue longer than the configured limit, such MTAs move them to a queue for a different message handler.

About capacity management

Determine how many MTAs and Network Prevent Servers (Email) are needed to ensure your message architecture can accommodate a peak message load. A conservative estimate of the amount of traffic a Network Prevent (Email) Server can handle is approximately 20 messages per second or up to 30 megabytes per second (Mbps) of throughput. Differences in the characteristics of the message stream (such as the distribution of message sizes and message content types) affect the performance of an individual Network Prevent (Email) Server. As you become more familiar with the characteristics of your traffic, you can adjust capacity plans appropriately.

When Network Prevent Servers (Email) operate in reflecting mode, Prevent-integrated MTAs handle each outbound message twice, which adds to their processing and CPU load.

See "Integration architectures for reflecting mode" on page 20.

About MX-Based clusters

To use an MX-based cluster to inspect your outbound mail, you can define an equivalent MX preference for each Network Prevent Server (Email) in the cluster. If load balancing is required, each record should have the same priority value and point to the fully-qualified domain name (FQDN) of one of the Network Prevent Servers (Email). Any MTA that sends a message to the Network Prevent Servers (Email) will go to one of the servers in the cluster.

If the Network Prevent Servers (Email) are operating in reflecting mode, they reflect messages back to the IP address of the MTA from which the messages arrived. If the Network Prevent Servers (Email) are operating in forwarding mode, they forward messages to the host or IP addresses that you configure. MX record lookups can also be performed for valid DNS names configured in the next-hop MTA list.

See "About next-hop MTA selection" on page 29.

For details on configuring forwarding mode and on configuring the Network Prevent (Email) Server in general, see the *Symantec Data Loss Prevention Administration Guide*.

About IP load balancer-based clusters

When you use an IP load balancer to implement clusters of MTAs and Network Prevent Servers (Email), make sure that every Network Prevent (Email) Server can connect back to the MTA cluster. The particular architecture you implement depends on the capabilities of your load balancer and the available routes in your network.

If the load balancer is bi-directional, you can operate the Network Prevent Servers (Email) in either reflecting mode or forwarding mode. If the load balancer is uni-directional, you must operate the server in forwarding mode.

See "Example of bi-directional load balancing" on page 47.

See "Example of uni-directional load balancing" on page 49.

Example of bi-directional load balancing

Figure 4-1 shows an architecture that includes a bi-directional load balancer and a cluster of Network Prevent Servers (Email) operating in forwarding mode. After receiving and analyzing messages, the Network Prevent Servers (Email) forward them to a virtual IP address (VIP) specified in the advanced settings of each Network Prevent (Email) Server (in the RequestProcesser.NextMTA field).

For details on configuring forwarding mode and on configuring the Network Prevent (Email) Server in general, see the *Symantec Data Loss Prevention Administration Guide*.

See "Example of uni-directional load balancing" on page 49.

Figure 4-1 Load balancing with Network Prevent Servers (Email) in forwarding mode

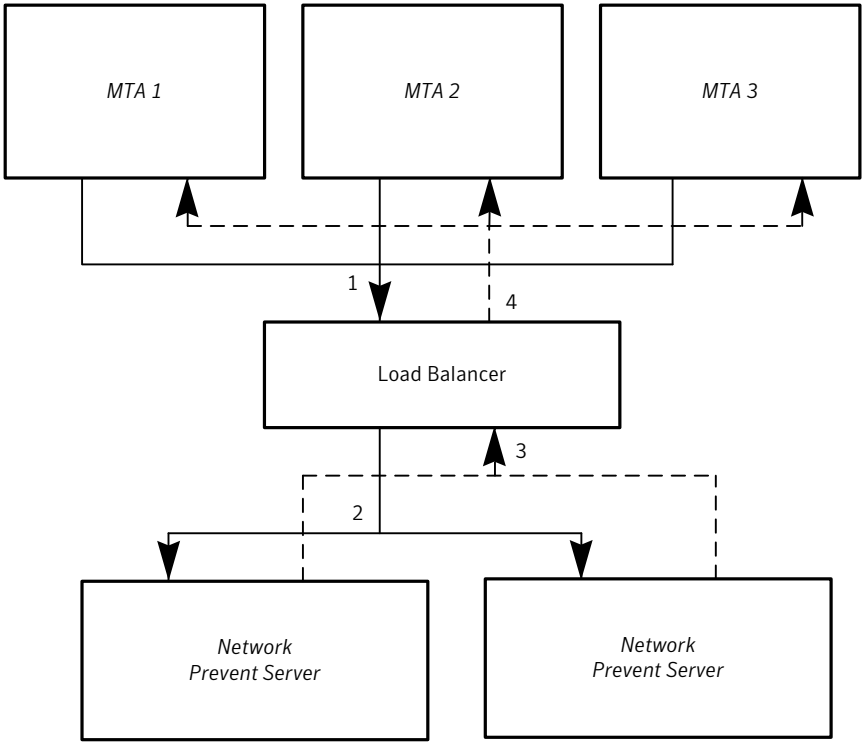

Details about load balancing with Network Prevent Servers (Email) in forwarding mode are as follows:

- An upstream MTA creates an SMTP connection to the Network Prevent Server (Email) VIP.
- The load balancer rewrites the destination address of each packet in each SMTP session as the IP address of one of the Network Prevent Servers (Email).
- The Network Prevent (Email) Server creates a connection to the MTA VIP (specified in the RequestProcesser.NextMTA field in the advanced settings of the Network Prevent (Email) Server).

- The load balancer rewrites the destination address of each packet in the SMTP session as the IP address of one of the MTAs in the cluster.

Note that you can also operate the Network Prevent Servers (Email) in reflecting mode with a bi-directional load balancer. One way to do this is to configure the load balancer to rewrite the source IP address to point to the second (return) VIP on the load balancer. The Network Prevent Servers (Email) would then reflect messages back to the return VIP.

Example of uni-directional load balancing

Figure 4-2 shows an architecture that includes uni-directional load balancing and a cluster of Network Prevent Servers (Email) operating in reflecting mode. The servers reflect messages back to the source address, which has not been virtualized by the load balancer.

Figure 4-2 Load balancing with Network Prevent Servers (Email) in reflecting mode

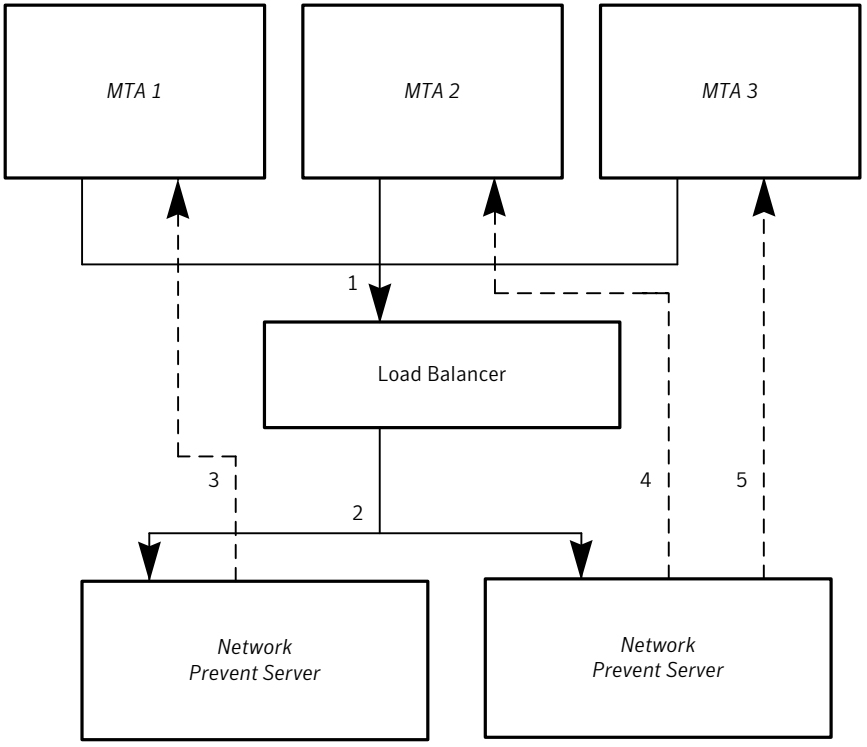

Details about load balancing with Network Prevent Servers (Email) in reflecting mode are as follows:

- An MTA creates an SMTP connection to the Network Prevent Server (Email) VIP.
- The load balancer rewrites the destination address of each packet in each SMTP session as the IP address of one of the Network Prevent Servers (Email). The load balancer does not rewrite the source address.
- The Network Prevent (Email) Server creates a connection directly to the originating MTA based on the unaltered source address of the incoming packet. The same holds for 4 and 5 in the diagram.

About fault tolerance planning

You should plan how to handle the message stream in the event that an MTA or a Network Prevent (Email) Server is not available. If you want to bypass an unavailable Network Prevent (Email) Server you can use MX records to define a second cluster for handling mail. Alternately, you can use MTA queue management to to move waiting messages to another outbound queue.

See "About MX-based bypass" on page 50.

See "About MTA-based queue management" on page 51.

As an alternative to these methods, you may simply want to prevent email messages from exiting your network until a Network Prevent (Email) Server is available to analyze them. If you plan to hold messages until a Network Prevent (Email) Server becomes available, you do not need to implement either either MX-based bypass or MTA queue management. However, you should make sure your that clusters include enough capacity to handle any service disruptions.

About MX-based bypass

In the unlikely event that a Network Prevent (Email) Server cluster is completely unavailable, you can use MX records to define a second cluster for handling mail. You must define the MX records for the second cluster with a higher precedence (lower priority) than the main Network Prevent Server (Email) cluster. Each member of the bypass cluster must have an SMTP listener on the same port on which the Network Prevent Server (Email) listens (as defined in the RequestProcessor.ServerSocketPort field in the advanced settings of the server). A common configuration creates an alternative SMTP listener on a virtual interface of the sending MTA.

The advantages of this method are:

- It works with appliances
- It is standards-based

The disadvantage is that email is not forwarded when a Network Prevent (Email) Server is still running but is overloaded.

See "About fault tolerance planning" on page 50.

See "About MTA-based queue management" on page 51.

About MTA-based queue management

If your MTAs let you manipulate their message queues, you can write a program that examines messages bound for the Network Prevent (Email) Server and then moves old messages (as defined by the configured limit) to an outbound queue.

The advantage of this method is that it works for both failure and system overload. The disadvantage is that you must write custom code to deal with your MTA mail queues, which may not be possible with some appliances.

See "About fault tolerance planning" on page 50.

See "About MX-based bypass" on page 50.

Chapter 5

Integration Testing

This chapter includes the following topics:

About Network Prevent (Email) Server integration testing

This chapter outlines the various functional and basic failover tests that you need to perform to ensure you have successfully integrated with the Network Prevent (Email) Server.

These tests assume an email generator, one or more Prevent-integrated MTAs, two Network Prevent (Email) Servers with MX-record failover and load balancing, and a downstream MTA/destination mail host. (These tests involve Symantec Data Loss Prevention policies and response rules.)

For details on response rules and policies, see the *Symantec Data Loss Prevention Administration Guide*.

About functional tests

Table 5-1 describes common functional tests that you can perform.

Table 5-1 MTA functional tests

Functional Test	Description
No Policy Violated	Confirm that the message stream is functioning, and that test email messages are handled once and only once by a Network Prevent (Email) Server.
Block SMTP Message	Create a response rule that redirects a message to a different address by changing the address. Confirm that the alternative address receives the message. Create a response rule that returns a non-delivery message to the sender. Confirm that the bounced message is received by the sender and that it contains the configured text from the Symantec Data Loss Prevention policy.
Send Email Notification	Create a response rule that causes an incident message to be sent to a list of addresses (including the sender).
Modify SMTP Message	Create a Modify SMTP Message rule, confirm the Network Prevent (Email) Server can modify the subject header, and that it can add additional headers to the message. If you want the Prevent-integrated MTA or the downstream MTA to act on the responses, confirm it has taken the appropriate action in response to the modified header. For instance, in response to an X-CFilter: Encrypt header, confirm the MTA routed the message to your in-line SMTP encryption server.

About basic failover tests

You can perform any of the following basic failover tests:

- Disconnect any subset or the entire set of Network Prevent Servers (Email) from the message stream, and confirm that the message stream continues to operate.
- Disconnect any subset or the entire set of Network Prevent Servers (Email) from the message stream, and confirm that the message stream continues to operate. Reconnect the Network Prevent (Email) Server, and ensure the email

message stream resumes through the Network Prevent (Email) Server. Repeat this test multiple times.

See "About fault tolerance planning" on page 50.

Appendix A

Email Message Systems

This appendix includes the following topics:

About store and forward email systems

Email systems are different from other types of network communications. While other communications systems are often end-to-end applications, email message systems are always store-and-forward systems. Email is never delivered directly from a sender application to an email reader application. The sender application sends the mail through a series of Message Transmission Agents (MTAs) that read the message, store it, and then forward it when the next hop is available. The last hop is a message store that holds the message until the email reader client views or downloads the mail. No MTA will remove a message from its message queue until it has been successfully delivered to the next hop. The message queue on an MTA is capable of storing a message for anywhere from a few hours to a few days before sending a non-delivery message back to the sender. The final message store is often capable of storing the message indefinitely.

The Network Prevent (Email) Server is different from an MTA because it does not store any messages. The Network Prevent (Email) Server does not end the SMTP session through which it receives a message until it forwards the message and ends the forwarding session.

See "About the Network Prevent (Email) Server message chain" on page 18.

About the DNS system

The route the email takes from sender to recipient is determined by information stored in the global DNS systems. The DNS system is a hierarchical system that

manages the mapping between named elements on the Internet and the underlying systems supporting those names. For instance, an address record (A record) relates a host's fully-qualified domain name (FQDN) to its IP address or the address that routers can use to send a packet to that host. The reverse of that mapping is stored in a pointer record (PTR record). A mail exchange record (MX record) is a special record that is used only by mail systems. An MX record identifies a mail domain name (the part of an email address that comes after the @) to a system that provides mail service for that mail domain.

All hosts that run SMTP-compliant MTAs are required to have valid A and PTR records. SMTP-compliant MTAs might not function correctly if, for example, a PTR record does not exist for that host or the PTR record for that host is not synchronized with the A record. An MX record relates the email domain to a valid A record for a host that knows how to deliver email for that domain, and it associates a delivery priority value for that particular host when sending email to that email domain.

When the DNS system is queried for the MX records that match a given email domain, it returns a list of all of the matching MX records with both the address values and their priorities. An SMTP-compliant MTA that receives more than one matching record first attempts to make a connection to the lowest precedence value (or highest priority) server for that mail domain. If that connection fails, a connection to the next-highest priority server for that domain is attempted. This process continues until either a connection is made or all of the hosts on the list are exhausted. When all of the hosts on the list are exhausted, the message is stored for a while, and another attempt to deliver the message is made. An elementary form of load sharing can be implemented by having multiple MX records listed with the same priority. Priority values can have any value that is represented by a 16-bit integer, but most often the values used are powers of 5.

See "About MX-based bypass" on page 50.

Appendix B

MTA Integration Checklist

This appendix includes the following topics:

About the MTA integration checklist

This appendix provides general guidance, in the form of a checklist, on how to integrate Network Prevent (Email) Server into your SMTP messaging architecture.

Completing the Network Prevent (Email) Server integration prerequisites

This section outlines tasks to complete before you decide which MTA integration architecture to implement. Research your company's existing SMTP-message routing architecture and MTA implementations by performing the following tasks.

To complete the prerequisites

1. Identify the key personnel in charge of the creating DNS entries, managing the network for the servers on the message chain, and information security.
2. Gather computer and function maps for all message processing servers, including the MTA to be integrated.

3 For each MTA host to be integrated, gather the IP address, subnet mask, default gateway IP address, hostname (short), fully-qualified domain name, and the DNS server IP addresses for each NIC.

4 For each MTA host to be integrated, obtain the administrator account with user name and password.

5 If you use TLS for communication between MTAs, obtain the location and password of each MTA keystore. In addition, you must obtain the public certificate for each MTA that must authenticate itself during TLS communication. Network Prevent (Email) Server requires an X.509 certificate in `.pem` format.

See "About TLS authentication" on page 30.

6 Read the section on MTA compatibility and requirements, and confirm that your existing MTA meets these requirements.

See "Environment compatibility and requirements for Network Prevent (Email)" on page 11.

If your existing MTA does not meet the requirements or cannot perform the suggested integration architectures, consider using an MTA that does meet the requirements.

7 Read the section about the Network Prevent (Email) Server in the message chain to understand how the server fits into the message chain.

See "About the Network Prevent (Email) Server message chain" on page 18.

8 Review the sample integration architectures.

See "Integration architectures for reflecting mode" on page 20.

See "About the integration architecture for forwarding mode" on page 27.

Selecting an integration architecture

This section describes how to select an integration architecture that best suits your messaging environment.

To select an integration architecture

1 Select one of the following Network Prevent (Email) Server failure modes: Block on Prevent failure or Pass through on Prevent failure.

2 If you chose the Pass through on Prevent failure mode, select a failover implementation: MX-record-based pass through on Prevent failure or MTA code-based pass through on Prevent failure.

3 Choose an MTA integration architecture.

See “Integration architectures for reflecting mode” on page 20.

See “About the integration architecture for forwarding mode” on page 27.

See “About fault tolerance planning” on page 50.

Evaluating message stream component capacity

This section outlines the tasks you need to accomplish to determine your messaging environment’s message stream component capacity.

To determine your messaging environment’s message stream component capacity

1. Determine number of MTAs to be integrated.
2. Determine the number of Network Prevent (Email) Servers.
3. Determine number of MX-record-based pass-through MTAs (real or virtual).

See “About integration architectures” on page 17.

See “About fault tolerance planning” on page 50.

Implementing Network Prevent (Email)

This section provides a high-level overview of major steps in implementing Network Prevent (Email).

See “About integration architectures” on page 17.

To implement Network Prevent (Email)

1. Install and configure your MTA software depending on your chosen integration architecture.

   If you are running the Network Prevent (Email) Server in reflecting mode, configure each Prevent-integrated MTA to receive reflected messages on the port you specified in the RequestProcessor.MTAResubmitPort advanced setting of the Network Prevent (Email) Server. Configure each MTA to send any mail not already inspected by a Network Prevent Server to the Network Prevent Server cluster, unless the cluster is unavailable and you want to forward mail to a bypass MTA.

   If you are running the Network Prevent (Email) Server in forwarding mode, configure each upstream MTA to forward all mail to the Network Prevent Server cluster, unless the cluster is unavailable and you want to forward mail to a downstream MTA. Then configure each downstream MTA to forward all mail received from any upstream MTA or Network Prevent Server to the appropriate next hop. If you use a hosted email service provider instead of a next-hop MTA, configure it to perform any additional processing of messages before delivering mail. If the hosted email service uses TLS communication, obtain the root CA-signed public key certificate for the email service.

   See "About downstream message tagging" on page 14.

   See "About TLS authentication" on page 30.

   Contact your MTA vendor for support, if required.

2. Install and configure your Network Prevent Servers (Email).

   See the appropriate *Symantec Data Loss Prevention Installation Guide* for installation instructions.

   See the *Symantec Data Loss Prevention Administration Guide* for configuration information.

3. Create the appropriate DNS records for all message handlers (MTAs, Network Prevent Servers, and other in-line SMTP processors), including the following:

   - A (address) records (for all MTA interfaces, real and virtual) associated with any MX records required to implement your MTA and Network Prevent Server clusters.
   - PTR records (for all interfaces, real and virtual) associated with any MX records required to implement your MTA and Network Prevent Server clusters.
   - Load balancing MX records for the MTAs to be integrated.

- Load balancing MX records for your Network Prevent Server cluster and for the bypass MTA cluster, if any.
  See "About the DNS system" on page 57.

4 Test your integration before deploying to your production environment.

See "About Network Prevent (Email) Server integration testing" on page 53.

Index

X