Symantec Brightmail™ Gateway 9.0 Administration Guide

Symantec Brightmail™ Gateway 9.0 Administration Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 9.0

Legal Notice

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party (“Third Party Programs”). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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

Appendix C Content Filtering templates 819

Chapter 1

Introducing Symantec Brightmail Gateway

This chapter includes the following topics:

About Symantec Brightmail Gateway

Symantec Brightmail™ Gateway (formerly branded as Symantec Mail Security) offers enterprises a comprehensive gateway-based message-security solution. Symantec Brightmail Gateway provides a solution that integrates email security, IM security, and premium data loss prevention capabilities in one appliance.

Symantec Brightmail Gateway does the following to protect your environment:

- Detects spam, denial-of-service attacks, and other inbound email threats
- Leverages a global sender reputation and local sender reputation analysis to reduce email infrastructure costs by restricting unwanted connections
- Filters email to remove unwanted content, demonstrate regulatory compliance, and protect against intellectual property and data loss over email
- Secures and protects public instant messaging communications with the same management console as it uses to secure and protect email

- Obtains visibility into messaging trends and events with minimal administrative burden

Note: To access Symantec Brightmail Gateway documentation, including translated versions, go to this location:

http://www.symantec.com/business/support/brightmail-gateway/documentation

See "What's new in Symantec Brightmail Gateway" on page 26.

See "What you can do with Symantec Brightmail Gateway" on page 38.

What's new in Symantec Brightmail Gateway

Table 1-1 lists the new features and enhanced features for this release of Symantec Brightmail Gateway.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features

New feature or enhancement	Description
Directory data service	You can configure the new directory data service to use the information that is stored in your Lightweight Directory Access Protocol (LDAP) directories for features in the Symantec Brightmail Gateway. For example, based on your directory data service configuration, the system can use your directory data to do the following: ■ Authenticate end-user login to the Control Center. ■ Authenticate the users that want to send email via SMTP. ■ Manage Spam Quarantine through improved handling of aliases, distribution lists, and invalid addresses. ■ Validate email addresses against directory data and drop messages or reject connections for invalid recipients. ■ Route messages to alternate addresses or mail hosts on a per-user basis using directory data. ■ Apply policies to users and groups consistently. Improvements over Symantec Brightmail Gateway version 8.0.3 and previous releases include the optional use of SSL to encrypt the LDAP server connection, the ability to use more than one authentication source, and the ability to follow LDAP referrals. See "About using the directory data service" on page 479.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features *(continued)*

New feature or enhancement	Description
Real time rule updates and secured rulesets	Symantec Brightmail Gateway improves antispam effectiveness by enabling faster rule distribution. Incremental rules can arrive as frequently as once a minute to improve responsiveness to new antispam threats and increase overall antispam effectiveness. The incremental updates reduce the frequency of full ruleset downloads. If the ruleset download fails for any reason, an error is logged that indicates the nature of the failure. Ensure that your configuration has access to the aztec.brightmail.com and liveupdate.symantec.com. See the *Symantec Brightmail Gateway Installation Guide* for details on connectivity requirements.
Probe Participation	Symantec Brightmail Gateway offers you the option to participate in Symantec's Probe Network, part of its Global Intelligence Network. The feature lets you forward to the probe network the email that is destined to unused email addresses or invalid email addresses. The probe accounts help Symantec detect spam attacks and create the spam filters that can better detect localized spam attacks. See "About probe accounts" on page 249. See "About the Symantec Probe Network" on page 248.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features *(continued)*

New feature or enhancement	Description
New features and enhancements to content filtering	Content compliance is now known as content filtering. The new content filtering features in this release are as follows: ■ Content incident folders As in previous releases, you can still create your own custom folders to collect incidents of violations. But instead of one type of folder to collect incidents, you can now monitor incidents different types of incident folders. The types of folders that you can use are as follows: ■ Hold for Review (Content Quarantine) Use this type of folder for the incidents that you want to review. This folder lets you retain the messages that trigger content filtering violations so that you can review them and determine how to act on them. Any additional actions for that policy are deferred until the incident is reviewed. ■ Informational Incidents Use this type of folder to track the incidents that are at a lower priority than the ones that you want to hold for review. See "About content incident folders" on page 438. ■ Expunger Symantec Brightmail Gateway provides an Expunger that can manage the size of your content incident folders. The Expunger automatically runs at the frequency that you specify. Symantec Brightmail Gateway can expunge incidents based on either the maximum size of the content incident folder. Or incidents can be expunged based on the maximum number of days that you want to store an incident. See "About managing the size of content incident folders" on page 439. ■ Enhanced actions for deleting attachments You can specify a content filtering policy action to delete the following types of attachments: ■ The attachment that violates the policy ■ All attachments ■ Specific types of attachments (such as archive files) See "Content filtering policy actions" on page 365. ■ Additional file extensions can be detected An attachment list contains the file extensions and the file application types that you want Symantec Brightmail Gateway to detect. Symantec Brightmail Gateway comes with pre-populated attachment lists that you can use. See "About attachment lists" on page 424.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features *(continued)*

New feature or enhancement	Description
DKIM authentication	In addition to Sender ID and SPF (Sender Policy Framework) sender authentication, you can now use DKIM (DomainKeys Identified Mail) authentication. You can configure DKIM signing on a per-domain basis. You can configure DKIM validation on system-wide basis. You can create a content filtering policy to apply actions based on the results of DKIM validation. See "Configuring DKIM authentication" on page 137.
Content encryption	Content encryption is a separately licensed feature that lets you encrypt outbound messages for greater security and track statistics for encrypted messages through the Control Center. Symantec content encryption uses Symantec Hosted Services, which is powered by MessageLabs. See "About encrypting messages with Symantec Content Encryption" on page 458.
SMTP authentication	You can use SMTP authentication to allow remote users to send email through Symantec Brightmail Gateway. SMTP authentication allows an MTA to authenticate an email client before it permits the client to send messages. A typical use of SMTP authentication is to allow authorized users to relay mail. You can use an LDAP authentication source. Or you can forward the credentials that the MUA supplies to another SMTP server for authentication. See "Using SMTP authentication" on page 145. See "Best practices for using SMTP authentication" on page 152.
Logging enhancements	Symantec Brightmail Gateway includes an enhanced logging feature that allows system administrators to enable alerts when logging disk space on Scanners nears or reaches capacity. When the default threshold levels are reached, the system shifts to a reduced or halted logging mode. It sends an email notification to the specified administrator of the change. See "About log disk space alerts" on page 633.
Commands	The commands in the command line interface have been redesigned for better usability. Some old commands have been removed or merged into new or existing commands. The command documentation has been rewritten in UNIX MAN page format. Type `help` *`command`* to view the command documentation.
Ability to use the same IP and port for both inbound and outbound traffic	You can use the same IP address and port for both inbound messages and outbound messages. The Scanner then uses outbound mail acceptance settings to determine if a message is inbound or outbound. See "About Scanner email settings " on page 86. See "Changing Scanner outbound mail acceptance settings" on page 93.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features *(continued)*

New feature or enhancement	Description
Per-domain non-local SMTP delivery bindings	You can set domain-specific delivery bindings for non-local messages. For each of your local domains, this feature enables you to define one or more IP addresses from which messages from that domain are sent. See "SMTP advanced settings for delivery bindings" on page 102.
Per-domain multiple downstream routes	You can route the inbound email that is addressed to different local domains using an unlimited number of default relays. See "About email domains" on page 115. See "Adding or editing domains" on page 117.
Message queue improvements	You can search the delivery queue for messages from a specific route or set of routes. If your search includes messages from only one route, you can reroute all of those messages in one action. See "Viewing queued messages" on page 658. See "Rerouting messages in the delivery queue" on page 658.
Link to online information for error level Scanner logs	The Control Center displays a question mark icon next to the description of error level Scanner logs. Click the icon to display a Web page that contains more information about the error if information is available. Symantec tracks the error information requests and adds new error information continually. You can disable this feature. See "Viewing log files" on page 627.
New rule for advance fee scam messages	Statsig is a new antispam module designed to detect 419 spam (advance fee scams). When your organization receives a 419 spam message and the statsig module identifies the message as spam, Symantec Brightmail Gateway takes the actions that you specify in your spam policy.
Scheduled Tasks and Failed Scheduled Tasks	Symantec Brightmail Gateway lets you track the status of all your scheduled tasks on one page. Scheduled tasks are grouped into tabs by associated task types. Under each tab, a page contains the name and status of the scheduled task. If a task completes successfully, users see the start time, the finish time, and the time that the task is scheduled to run next. If a task fails to complete, the **Finished** column displays a status of **Failed.** See "About scheduled tasks" on page 621.

Table 1-1 Symantec Brightmail Gateway new features and enhanced features *(continued)*

New feature or enhancement	Description
Ability to create a diagnostic package from the Control Center	You can now create a diagnostic package for the host that you specify through the Control Center. Symantec Support uses this package to troubleshoot issues with your product. All packages contain default components. You can create a default package or specify what additional components you want to include in the package. When you generate a diagnostic package, Symantec Brightmail Gateway creates the package on the host for which you are running the report. Symantec Brightmail Gateway then transfers a copy of the package to the location that you specify. You can transfer the copy of the package through a transfer protocol to a remote location. You can also save the copy on the computer on which you run your browser. See "Specifying what to include in diagnostic packages " on page 705.
Configurable date and time format	Administrators can configure date formats and time formats. Configurable date and time formats apply to the following: ■ Detail views ■ List views ■ Mailbox views ■ Report components See "Customizing the date format and time format" on page 703.

Components of Symantec Brightmail Gateway

Symantec Brightmail Gateway is an all-in-one appliance that secures your gateway. Symantec Brightmail Gateway integrates the core hardware and software pieces necessary for a comprehensive, secure, and easy-to-deploy message security solution. Symantec Brightmail Gateway provides multiple layers of protection that result in infrastructure cost savings and ease of management. By managing connections and messages at each point in the SMTP conversation, Symantec Brightmail Gateway ensures that your infrastructure is used in a cost-effective manner.

The following describes the ways that Symantec Brightmail Gateway protects your environment:

Gateway	A surging increase in spam volumes strains customer email infrastructures as email scanners must process more messages. When a product filters mail for spam and viruses, the process slows down email delivery and taxes CPU and mail server resources. The increased demand in resources and mail delivery delay occurs because each message needs to be opened and processed. Symantec Brightmail Gateway can detect undesirable connections and block or defer them. It features a set of automated and configurable connection management features that go into effect as soon as an incoming connection is detected. It serves as a “gatekeeper” in front of the more CPU-intensive pieces of the filtering engine, including the antispam and antivirus layers. It can be configured to block spam attacks, directory harvest attacks, connections from the senders that are identified as spammers by Symantec, and more—automatically. The gateway is the first stage in the inbound protection process. It protects the internal infrastructure by detecting and examining the incoming IP connection before the mail server in Symantec Brightmail Gateway accepts a message. It can then take preventive action such as rejecting the SMTP connection. The gateway can also recognize and block directory harvest attacks and help defend against denial-of-service attacks. The ability to stop potential attacks is another reason to ensure that certain mail does not breach the gateway in the first place. The best email security solutions accurately reject unwanted mail at the gateway that is based on its IP address. Such SMTP connection management features are an increasingly effective method of dealing with the side effects of increased email volume.

Brightmail Engine

Accounting for over 80 percent of all email traffic, spam chokes the messaging infrastructure, saps mail server and storage resources, and clutters user inboxes. Offensive and fraudulent spam can create liability issues for organizations. Given the unbeatable economics and ineffective legislation, spammers continues to flood organizations with unsolicited mail. As always, spammers continue to adjust tactics and increase their volume to get around the defenses that IT organizations deploy.

Multi-layered spam protection is the cornerstone of the Symantec Brightmail Gateway. Driven by technologies and response capabilities originally developed by Brightmail, the Brightmail Engine harnesses a robust arsenal for filtering techniques. These techniques include spam signatures, heuristics, URL filtering, reputation-based filters, and other standard and proprietary approaches.

Viruses can wreak extreme havoc in an organization. The damage ranges from email server crashes to system downtime and the destruction of company data. From an email security perspective, the worlds of spam and viruses are intricately tied. Internet-delivered email accounts for approximately 80 percent of virus incidents. In addition, the actual payload of many viruses and email-borne worms includes the software that turns the target computer into a spam "zombie." Spammers then access these zombies and use them to launch spam and other email-based attacks. Given the damage resulting from viruses, it is essential to employ virus protection at the earliest point of network entry: the email gateway. Symantec Brightmail Gateway scans and detects viruses by integrating award-winning Symantec AntiVirus technology.

Antivirus protection includes the following features:

- Automatic virus definition updates
- Flexible policies to handle messages with viruses
- Specific defenses against mass-mailing worms and the associated spawned email messages

MTA

Symantec Brightmail Gateway includes a Message Transfer Agent (MTA) that processes, routes, and delivers email messages in cooperation with the Brightmail Engine. You can use the facilities of the MTA to configure custom message handling for different domains or mail servers.

Data loss prevention

To conform to IT, regulatory, HR guidelines, organizations increasingly look to email security appliances to assist in managing policies for email. Symantec Brightmail Gateway includes several features to support a company's regulatory and internal governance requirements. It also provides the tools to enable development of robust content filtering policies.

About the Symantec Brightmail Gateway software components

A hardened, preinstalled Linux-based operating system powers Symantec Brightmail Gateway. The filtering and management platform software also resides on the appliance. In addition, there is an IM relay and a mail transfer agent (MTA) that enable email communication. Software updates are easily applied, which helps to ensure minimal disruptions for updates.

Symantec Brightmail Gateway software consists of the following subcomponents:

Scanner

Scanners do the following tasks:

- Process the inbound messages and outbound messages and route messages for delivery.
- Download virus definitions, spam and spim signatures, and other security updates from Symantec Security Response.
- Run filters, render verdicts, and apply actions to messages in accordance with the appropriate policies and settings.

See “About Scanner email settings ” on page 86.

Each Symantec Brightmail Gateway Scanner uses a separate mail transfer agent, or MTA, when it scans email messages.

See “MTA and message queue behavior” on page 613.

Instant messages are handled through the IM Relay proxy.

See “Working with Services” on page 112.

Control Center	The Control Center provides message-management services, such as centralized administration, reporting, and monitoring. The Control Center also houses a Web server and the databases that store system-wide information. The Control Center collects and aggregates statistics from connected and enabled Scanners and provides information on their status and maintains system logs. The Control Center also collects statistics on types and levels of security threats. These statistics can be displayed in a variety of reports and distributed in different formats. The Control Center also hosts Spam Quarantine and Suspect Virus Quarantine. It may also be configured to store Information that is related to messages that trigger content filtering policies. See "About quarantining spam" on page 258. See "About quarantining suspected viruses" on page 229. See "About content incident folders" on page 438.

About the Symantec Brightmail Gateway hardware component

You can deploy Symantec Brightmail Gateway on the Brightmail / Mail Security 8300 Series. The appliance is rack mountable and includes features such as redundant storage with RAID and dual power supplies and fans.

See "Viewing the status of your hardware" on page 607.

See "Viewing information about your hardware" on page 608.

How Symantec Brightmail Gateway works

Figure 1-1 shows how Symantec Brightmail Gateway processes an email message. This diagram assumes that the message passes through the Filtering Engine to the Transformation Engine without being rejected. The diagram also shows the path IM traffic takes through the system.

Figure 1-1 Symantec Brightmail Gateway Architecture

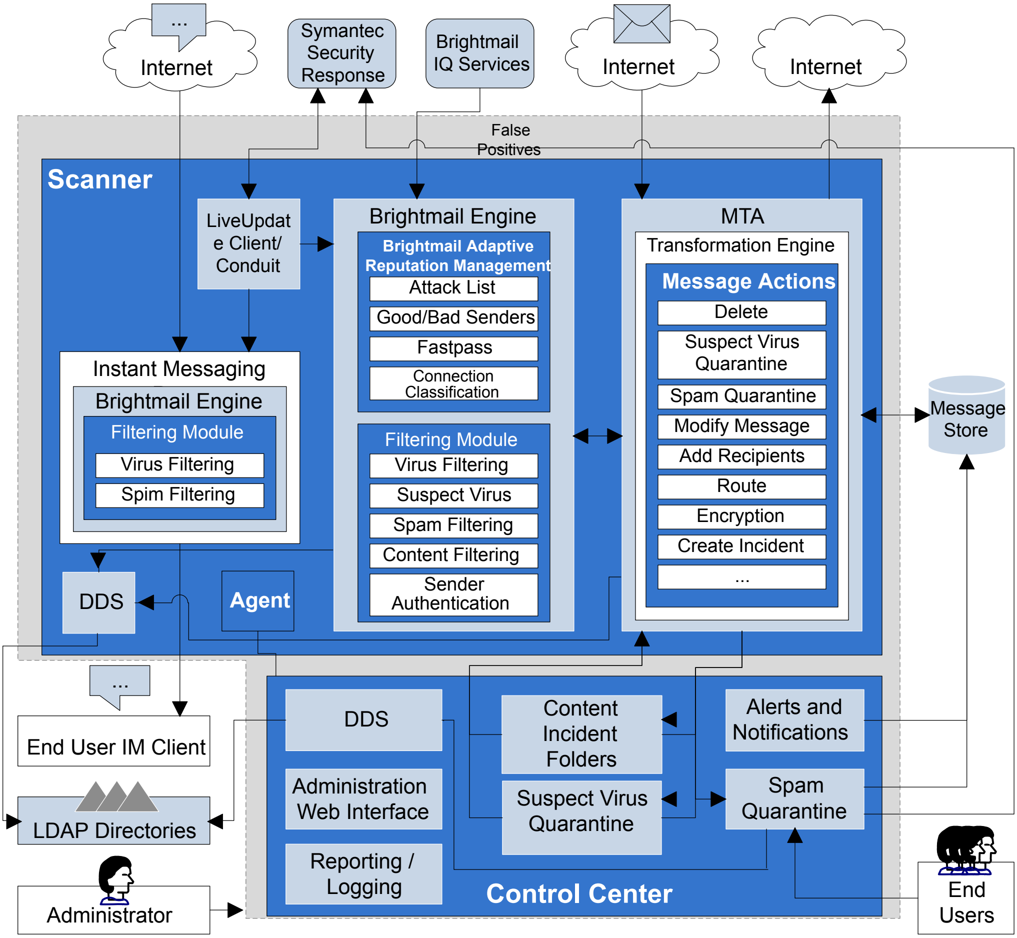

A description of the path that email messages and instant messages take is as follows:

Email messages

The path an email message takes through the system is as follows:

- At the gateway, Connection Classification classifies the sending IP into one of 9 classes based on local reputation. It either accepts or defers the connection based on class membership. New senders are placed in a tenth, default class. Symantec Brightmail Gateway also checks the IP address to determine if it belongs to a good sender group or bad sender group. It then blocks or permits the connection accordingly.
- Before the MTA accepts the message, it checks the domain address and email address. The MTA determines if it belongs to the Local Good Sender Domains or Local Bad Sender Domains groups. If it does, applies the configured action to the message. If appropriate, the MTA moves the message to its inbound queue.
- The Brightmail Engine consults the directory data service to expand the message’s distribution list.
- The Brightmail Engine determines each recipient’s filtering policies.
- Antivirus filters determine whether the message is infected.
- Content filtering policy filters scan the message for restricted attachment types or words, as defined in configurable dictionaries.
- If the sending IP is granted a pass by Fastpass, antispam filtering is bypassed. If not, the antispam filters that use the latest rules from Symantec Security Response determine whether the message is spam. The message may also be checked against user-defined Language settings.
- The Transformation Engine performs actions according to filtering results and configurable policies and applies them to each recipient's message based on policy group membership.

Instant messages

The path an instant message takes through the IM message flow (from an external source) is as follows:

- IM traffic enters your network and is redirected to the IM proxy by your enterprise DNS servers.
- The IM proxy filters IM traffic according to your settings and compares the traffic with current filters Symantec Security Response publishes. These filters determine whether a message is spim or contains a virus. If a message is determined to contain spim or a virus, you can choose to block this traffic.
- The IM traffic reaches the internal user's IM client.
- If you have enabled outbound IM filtering, outbound messages are routed through the IM proxy before they are sent to an external user's IM client.

Note: Symantec Brightmail Gateway does not filter any messages that do not flow through the SMTP gateway. For example, it does not filter the messages that are sent between mailboxes on the same Microsoft Exchange Server. Nor does it filter the messages on different servers within a Microsoft Exchange organization.

See "About email message flow" on page 131.

What you can do with Symantec Brightmail Gateway

Table 1-2 describes what you can do with Symantec Brightmail Gateway.

Table 1-2 What you can do with Symantec Brightmail Gateway

Tasks	Description
Block unwanted email	Symantec Brightmail Gateway provides several features that let you block email from entering your network. When you block unwanted email, you reduce your risk of getting a virus. You also reduce the resources that are needed to scan messages. See "About blocking and allowing messages at connection time" on page 163.
Create policy groups and policies	You can manage users through policy groups. You can specify these groups of users according to email addresses, domain names, or LDAP groups. Then you can apply filtering policies to specific policy groups. Symantec Brightmail Gateway installs with a Default policy group that consists of all of the users. See "About policy groups" on page 315.

Table 1-2 What you can do with Symantec Brightmail Gateway *(continued)*

Tasks	Description
Detect spam	Symantec Brightmail Gateway can detect spam with a high level of accuracy. Depending on your settings you can: ■ Define policies for handling the messages that are identified as spam and set thresholds for suspected spam. ■ Store spam messages in Spam Quarantine until they can be reviewed. If you configure user access to Spam Quarantine, recipients receive notification when they have messages in their quarantine. Users can then review these messages and take appropriate action. ■ Configure Symantec Brightmail Gateway to allow messages from specified domains to bypass antispam scanning altogether. See "About filtering spam" on page 239. See "About quarantining spam" on page 258. ■ Participate in Symantec's Probe Network and further improve spam filtering effectiveness. See "About probe accounts" on page 249.
Detect viruses and other malicious attacks	You can create policies and configure settings to detect viruses and other malicious attacks. See "About detecting viruses and malicious attacks" on page 204.
Filtering messages to enforce content policies	Keyword dictionaries and templates help you create the policies that filter the email messages. Such policies can be used to monitor and enforce compliance with corporate and regulatory requirements and to prevent data loss. See "About content filtering" on page 331.
Detect IM threats	Symantec Brightmail Gateway offers enterprises a gateway-based instant messaging (IM) traffic filter solution. Along with its email security solutions, Symantec Brightmail Gateway provides threat protection solutions to your enterprise for IM through the features and settings. See "About IM" on page 289.

Table 1-2 What you can do with Symantec Brightmail Gateway *(continued)*

Tasks	Description
Monitor performance	The Control Center contains a Dashboard that displays the overall system status. It provides statistics about the types of threats that inbound and outbound messages pose to your system. Statistics include data about the messages that are addressed to invalid recipients or that come from the addresses that have bad reputations. They also include the number of messages that have triggered virus, spam, and content filtering policies. See "About monitoring the status of your product" on page 602. Symantec Brightmail Gateway includes over 50 reports that provide statistics on content filtering, email messages, instant messages, IP connections, spam, and viruses from all Scanners. You can create reports when you need them or configure them to be emailed daily, weekly, or monthly. See "About working with reports" on page 568.
Obtain definition updates	You can use either of the following methods to obtain virus definition updates: ■ LiveUpdate You can use LiveUpdate to automatically update your protection. When LiveUpdate runs, it downloads and installs any available definitions. ■ Rapid Response You can use Rapid Response when you need quick responses to emerging threats. Rapid Response definitions are most useful for a perimeter defense to mitigate quickly spreading threats. Rapid Response is an alternative to LiveUpdate. See "About updating virus definitions" on page 221.

Table 1-2 What you can do with Symantec Brightmail Gateway *(continued)*

Tasks	Description
Obtain notifications about outbreaks, system issues, and policy violations	You can configure Symantec Brightmail Gateway to automatically send alerts and notifications about a wide variety of events. These events include the following: ■ Policy violations ■ Virus outbreaks ■ System and user preferences replication errors and status ■ Spam and virus quarantine information ■ License expiration and update availability ■ Scheduled tasks failures See "Types of alerts" on page 615.
Manage your appliance	Symantec Brightmail Gateway provides the features that help you do the following tasks: ■ Monitor devices. See "Monitoring devices through SNMP" on page 619. ■ Configure and manage Scanners. See "About Scanner email settings " on page 86. ■ Manage system software. See "Updating your software" on page 700.
Use the command line to configure Symantec Brightmail Gateway	Each appliance has a set of commands you can use to configure, optimize, and administer your system. Access these commands by logging into the system either through SSH or by the VGA or serial connections on the appliance.

Where to get more information

The following resources provide more information about your product:

Documentation

The Symantec Brightmail Gateway documentation set consists of the following manuals:

- *Symantec Brightmail Gateway Administration Guide*
- *Symantec Brightmail Gateway Installation Guide*
- *Symantec Brightmail Gateway Getting Started Guide*

http://www.symantec.com/business/support/documentation.jsp?language=english&view=manuals&pid=53991

Product Help system

Symantec Brightmail Gateway includes a comprehensive help system that contains conceptual and procedural information.

Symantec Web site

Visit the Symantec Web site for more information about your product as follows:

- www.symantec.com/enterprise/support
  Provides you access to the technical support Knowledge Base, newsgroups, contact information, downloads, and mailing list subscriptions
- https://licensing.symantec.com/acctmgmt/index.jsp
  Provides you information about registration, frequently asked questions, how to respond to error messages, and how to contact Symantec License Administration
- www.enterprisesecurity.symantec.com
  Provides you product news and updates
- www.symantec.com/business/security_response/index.jsp
  Provides you access to the Virus Encyclopedia, which contains information about all known threats; information about hoaxes; and access to white papers about threats

Chapter 2

Getting Started with Symantec Brightmail Gateway

This chapter includes the following topics:

- Logging on and logging off
- Feature dependencies
- Preinstalled policies
- Performing initial configuration tasks
- Performing optional configuration tasks

Logging on and logging off

End users manage their Spam Quarantine, personal Good Senders list, Bad Senders list, and email language settings through the Control Center. Use the Control Center to configure an LDAP source, enable LDAP authentication, and enable those features.

Note: Do not create an account for an administrator that is identical to a user account name. Conversely, do not create an account for a user that is identical as an administrator account name. If a naming conflict occurs, the administrator logon takes precedence, and the user is denied access to their account. If an administrator user name and password and a user name and password are identical, the user is granted access to the administrator account.

To log on as a user with an iPlanet, SunONE, or Domino directory server account, your Administrator must enable authentication for the Control Center.

To log on as an administrator

1 Access the Control Center from a browser.

The default logon address is as follows:

https://`<hostname>`

where `<hostname>` is the host name designated for the appliance. Or you can use the IP address in place of `<hostname>`.

2 If you see a security alert message, accept the self-signed certificate to continue.

The Control Center **Login** page appears.

3 Choose the language that you want to use to operate the Quarantine views and user views of the Control Center.

4 In the **User name** box, type the user name that your system administrator assigns to you.

If you are the first administrator to access the Control Center, type **admin**.

5 In the **Password** box, type your administrative password.

Contact your system administrator if you do not know the password.

6 If the system administrator has enabled the **Remember me** feature, the **Remember me on this computer** option appears. Check this option to bypass your logon credentials when you subsequently access the Control Center.

Symantec Brightmail Gateway requires you to re-enter your logon credentials after you loglogging out, or based on the duration that the administrator specifies.

Note that if you use this feature, anyone that has access to your computer has access to the Control Center.

7 Click **Login**.

To log on as a user with an iPlanet or SunONE Directory Server account

1. Access your Control Center from a browser.

   The default logon address is as follows:

   https://`<hostname>`

   where `<hostname>` is the host name designated for the appliance. Or you can use the IP address in place of `<hostname>`.

2. If you see a security alert message, accept the self-signed certificate to continue.

   The Control Center Login page appears.

3. Choose the language that you want to use to operate the Quarantine views and user views of the Control Center.

4. In the **User name** box, type your full email address (for example, kris@symantecexample.com).

5. In the **Password** box, type the password that you normally use to log onto the network.

6. If the system administrator has enabled the **Remember me** feature, the **Remember me on this computer** option appears. Check this option to bypass your logon credentials when you subsequently access the Control Center.

   Symantec Brightmail Gateway requires you to re-enter your logon credentials based on the duration that the administrator specifies.

   Note that if you use this feature, anyone that has access to your computer has access to the Control Center.

7. Click **Login**.

To log on as a user with a Domino account

1. Access your Control Center from a browser.

   The default logon address is as follows:

   https://`<hostname>`

   where `<hostname>` is the host name designated for the appliance. Or you can use the IP address in place of `<hostname>`.

2. If you see a security alert message, accept the self-signed certificate to continue.

   The Control Center Login page appears.

3. Choose the language that you want to use to operate the Quarantine views and user views of the Control Center.

4 In the **User name** box, type your full email address (for example, kris@symantecexample.com).

5 In the **Password** box, type the password that you normally use to log onto the network.

6 If the system administrator has enabled the **Remember me** feature, the **Remember me on this computer** option appears. Check this option to bypass your logon credentials when you subsequently access the Control Center.

Symantec Brightmail Gateway requires you to re-enter your logon credentials after logging out, or based on the duration that the administrator specifies.

Note that if you use this feature, anyone that has access to your computer has access to the Control Center.

7 Click **Login**.

To log on as a user with an Active Directory account

1 Access your Control Center from a browser.

The default logon address is as follows:

https://`<hostname>`

where `<hostname>` is the host name designated for the appliance. Or you can use the IP address in place of `<hostname>`.

2 If you see a security alert message, accept the self-signed certificate to continue.

The Control Center Login page appears.

3 Choose the language that you want to use to operate the Quarantine views and user views of the Control Center.

4 In the **User name** box, type your user name (for example, kris).

5 In the **Password** box, type the password that you normally use to log onto the network.

6 Select the LDAP server that you use to verify your credentials.

7 If the system administrator has enabled the **Remember me** feature, the **Remember me on this computer** option appears. Check this option to bypass your logon credentials when you subsequently access the Control Center.

Symantec Brightmail Gateway requires you to re-enter your logon credentials after logging out, or based on the duration that the administrator specifies.

Note that if you use this feature, anyone that has access to your computer has access to the Control Center.

8 Click **Login**.

To log off

1 In the upper right corner of any page, click the **Log Out** icon.

2 For security purposes, close your browser window to clear your browser's memory.

Feature dependencies

Certain Symantec Brightmail Gateway features require additional configuration tasks or prerequisites to function properly.

The following features require that one of the following two statements is true:

- Your Scanner appliance is positioned at the perimeter of your network, in front of your mail servers, and optionally behind a firewall.
- Your Scanner appliance is not at the perimeter, and you have specified as internal mail servers all servers between the Scanner and the perimeter.
  See "Specifying internal mail hosts for non-gateway deployments" on page 104.
- Connection Classification
- Fastpass
- Bounce Attack Prevention
- Directory Harvest Attack (DHA) Prevention
- Email Virus Attack Prevention
- Local Bad Sender IPs
- Local Good Sender IPs
- Third Party Bad Senders
- Third Party Good Senders
- Symantec Global Bad Senders

- Symantec Global Good Senders

If you want Symantec Brightmail Gateway to perform any of the following functions, you must first create and configure an LDAP data source:

- Reject or drop invalid recipients.
- Validate recipient data.
- Expand distribution lists and apply policies to the list members.
- Use an LDAP directory data source to resolve aliasing.
- Specify a unique route for a specific recipient.
- Allow end users access to Spam Quarantine messages.
- Allow end users to create personal good sender and bad sender lists.
- Allow end users to set email language preferences.
- Enable DHA prevention.
- Create groups of users based on LDAP.
- Use an LDAP server as your authentication source for SMTP authentication.

See "Creating a data source" on page 491.

If you want Symantec Brightmail Gateway to take different actions on the same email or IM messages for different internal recipients, create policy groups. If you do not create policy groups and assign policies to them, Symantec Brightmail Gateway applies the same actions to a message for all internal recipients.

See "Creating policy groups and assigning policies" on page 73.

Implementing any of the following features requires that you first specify your local domains. You can also specify non-local domains:

- DKIM signing
- Drop invalid recipients
- Reject invalid recipients
- Directory harvest attack recognition
- Domain-specific message routing and delivery designations

See "About email domains" on page 115.

Preinstalled policies

Symantec Brightmail Gateway comes with default policies to handle some common tasks for viruses and spam in email and IM messages.

About default policies and features

Symantec Brightmail Gateway is ready to use upon installation, with a set of default policies, and a group of features that are enabled by default.

Upon installation Symantec Brightmail Gateway includes only one policy group, the default group. A policy group is a group of users to which you an assign filtering policies. The default group cannot be deleted and always includes all users. However, you can create additional policy groups.

For each additional policy group you create, you can specify the group members, then assign different policies for spam, virus, and instant messages.

Depending on the results of filtering, including filtering at connection time and message scanning, Symantec Brightmail Gateway assigns one or more verdicts to a message. You can choose the policies that apply for each verdict, per policy group.

See "About filtering" on page 717.

See " Verdicts and actions for email messages" on page 718.

See "Verdicts and actions for instant messages" on page 723.

Upon installation, one default policy for each spam or virus verdict is assigned by default to the default group. Other default policies for spam and virus are provided for your use and initially not assigned to any group. You can create additional policies of any type. Symantec Brightmail Gateway provides no default content filtering policies. You can use predefined templates to create content filtering policies.

See "Default email spam policies" on page 50.

See "Default email virus policies" on page 52.

See "Default IM virus policies" on page 56.

See "Creating content filtering policies" on page 334.

The following features are enabled by default upon installation:

- Connection Classification
- Good and bad sender groups
- Filtering email messages for spam, viruses, and other threats
- Suspected spam
- Hold messages in Suspect Virus Quarantine
- Instant message filtering
- Message audit logging

- A selection of alerts

The following features are not enabled by default upon installation:

- Spam Quarantine
- LDAP directories
- Fastpass
- Directory harvest attack protection
- Email virus attack prevention
- Bounce Attack Prevention
- Sender authentication, including SPF, Sender ID, DKIM
- SMTP authentication
- TLS encryption
- Content filtering
- Invalid recipient handling
- Probe accounts
- Data gathering for all report types
- All alert types

Default email spam policies

Symantec Brightmail Gateway installs with three default email spam policies. These policies are automatically enabled and assigned to the default group. However, if you create a custom email spam policy and apply it to the default group, that policy overrides the default policy.

See “Setting policy group precedence” on page 329.

See “Selecting spam and spim policies for a policy group” on page 324.

Table 2-1 Preconfigured spam policies

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Spam: Modify subject line with "[Spam]" (default)	Inbound and outbound	If a message is spam	Prepends the subject line with "[Spam]"	Default group	Enabled by default
Suspected Spam: Modify subject line with "[Suspected Spam]" (default)	Inbound and outbound	If a message is suspected spam	Prepends the subject line with "[Suspected Spam]"	Default group	Enabled by default
Spam or Suspected Spam: Delete message	Inbound and outbound	If a message is spam or suspected spam	Delete message	None	Enabled by default
Spam or Suspected Spam: Quarantine message	Inbound and outbound	If a message is spam or suspected spam	Hold message in Spam Quarantine	None	Enabled by default
Spam or Suspected Spam: Deliver to Spam folder	Inbound and outbound	If a message is spam or suspected spam	Deliver the message to the recipient's spam folder	None	Enabled by default
Spam or Suspected Spam: Deliver normally	Inbound and outbound	If a message is spam or suspected spam	Deliver message normally	None	Enabled by default

Table 2-1 Preconfigured spam policies *(continued)*

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Failed Bounce Attack Validation: Reject message	Inbound only	If a message fails bounce attack validation	Reject messages failing bounce attack validation	Default group	Enabled by default

See "Creating email spam policies" on page 242.

See "Enabling and disabling spam policies" on page 244.

Default email virus policies

Symantec Brightmail Gateway installs with pre-configured virus policies. These policies are enabled by default and can be applied to a policy group. You can disable or modify the policy actions and the policy groups to which the policies apply. The policy name, the type of message that the policy applies to, and the condition that must be met cannot be modified for the pre-configured virus policies that are labeled as default.

See "Creating email virus policies" on page 210.

See "Selecting virus policies for a policy group" on page 322.

Table 2-2 Default virus policies

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Virus: Clean message (default)	Inbound and outbound messages	Message contains a virus	Clean the message	Default group	Enabled by default
Worm: Delete message (default)	Inbound and outbound messages	Message contains a mass-mailing worm	Delete the message	Default group	Enabled by default

Table 2-2 Default virus policies *(continued)*

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Unscannable: Delete message (default)	Inbound and outbound messages	Message is unscannable for viruses	Delete the message	Default group	Enabled by default
Encrypted Attachment: Modify subject line with "[WARNING - ENCRYPTED ATTACHMENT NOT VIRUS SCANNED]"	Inbound and outbound messages	Message contains an encrypted attachment	Prepend the subject line with "[WARNING - ENCRYPTED ATTACHMENT NOT VIRUS SCANNED]"	Default group	Enabled by default
Virus: Delete message	Inbound and outbound messages	Message contains a virus	Delete the message	None	Enabled by default
Virus: Modify subject line with "[VIRUS INFECTED]"	Inbound and outbound messages	Message contains a virus	Prepend the subject line with "[VIRUS INFECTED]"	None	Enabled by default
Virus: Deliver normally	Inbound and outbound messages	Message contains a virus	Deliver the message normally	None	Enabled by default
Worm: Clean message	Inbound and outbound messages	Message contains a mass-mailing worm	Clean the message	None	Enabled by default

Table 2-2 Default virus policies *(continued)*

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Worm: Modify subject line with "[WORM INFECTED]"	Inbound and outbound messages	Message contains a mass-mailing worm	Prepend the subject line with "[WORM INFECTED]"	None	Enabled by default
Worm: Deliver normally	Inbound and outbound messages	Message contains a mass-mailing worm	Deliver the message normally	None	Enabled by default
Unscannable: Modify subject line with "[WARNING -NOT VIRUS SCANNED]"	Inbound and outbound messages	Message is unscannable for viruses	Prepend the subject line with "[WARNING -NOT VIRUS SCANNED]"	None	Enabled by default
Unscannable: Deliver normally	Inbound and outbound messages	Message is unscannable for viruses	Deliver the message normally	None	Enabled by default
Threat: Modify subject line with "[SPYWARE OR ADWARE INFECTED]" (default)	Inbound and outbound messages	Message contains spyware or adware See "Spyware or adware verdict details" on page 207.	Prepend the subject line with "[SPYWARE OR ADWARE INFECTED]"	Default group	Enabled by default

Table 2-2 Default virus policies *(continued)*

Policy name	Applies to the following messages	If the following condition is met	Actions	Applies to the following policy group	Default status
Inbound suspect virus: Strip attachments and hold message in Suspect Virus Quarantine (default)	Inbound messages	Message contains a suspicious attachment	Strip and Delay in Suspect Virus Quarantine with message "Parts of your message have been stripped because they were considered suspect."	Default group	Enabled by default
Outbound suspect virus: Hold message in Suspect Virus Quarantine (default)	Outbound messages	Message contains a suspicious attachment	Hold message in Suspect Virus Quarantine	Default group	Enabled by default

Default IM spim policies

Symantec Brightmail Gateway installs with the following default IM spim policies. These policies are automatically enabled and assigned to the default group. However, if you create a custom spim policy and apply it to a group, that policy overrides the default policy.

See "Setting policy group precedence" on page 329.

Note: If you upgraded from an earlier version of Symantec Brightmail Gateway, the settings for your default IM spim policies are based on your previous settings. As a result, the settings of your default policies may differ from those listed.

Table 2-3 Default IM Spim Policies

Policy Name	Description
Incoming spim	Deletes inbound IM messages that contain known or suspected as spim.
Outgoing spim	Deletes outbound IM messages that contain known or suspected spim, and sends the default spim notification to the senders of those messages.

Default IM virus policies

Symantec Brightmail Gateway installs with the following default IM virus policies. These policies are automatically enabled and assigned to the default group. However, if you create a custom virus policy and assign it to the default group, that policy overrides the default policy.

See “Selecting virus policies for a policy group” on page 322.

See “Creating IM virus policies” on page 309.

See “Setting policy group precedence” on page 329.

Warning: If a group does not have an enabled IM virus policy, Symantec Brightmail Gateway allows the IM users that belong to that group to send infected, encrypted, and unscannable files.

Table 2-4 Default IM virus policies

Policy	Description
IM Inbound Virus Detected	Blocks the infected files that are sent from external IM users
IM Inbound Encrypted File	Blocks the encrypted files that are sent from external IM users
IM Inbound Unscannable File	Blocks the unscannable files that are sent from external IM users
IM Outbound Virus Detected	Blocks the infected files that are sent from internal IM users and sends a notification to those users

Table 2-4 Default IM virus policies *(continued)*

Policy	Description
IM Outbound Encrypted File	Blocks the encrypted files that are sent from internal IM users and sends a notification to those users
IM Outbound Unscannable File	Blocks the unscannable files that are sent from internal IM users and sends a notification to those users

Performing initial configuration tasks

During installation you set the initial configuration parameters that Symantec Brightmail Gateway uses to operate. Symantec Brightmail Gateway will continue to operate using the initial defaults as well as the specific choices you made during installation. However, most customers benefit from reviewing the initial configuration settings, enabling additional features, and modifying settings that were not a part of the installation process.

Follow the four-step process below to ensure that you are ready to take full advantage of the extensive capabilities of Symantec Brightmail Gateway to meet the specific needs of your installation.

Table 2-5 Initial configuration tasks

Step	Action	Description
Step 1	After installing Symantec Brightmail Gateway, test message flow.	Ensure that your appliance is filtering and delivering mail. See the *Symantec Brightmail Gateway Installation Guide*.

Table 2-5 Initial configuration tasks *(continued)*

Step	Action	Description
Step 2	Configure optional communications and monitoring features.	Symantec Brightmail Gateway provides a variety of powerful communications and monitoring features. You can control SMTP and IM communications parameters and security. You ca n control end user access and communications between your Control Center and your Scanners.You can set up alerts, logs, and reports, as well as SNMP monitoring and UPS backup. See Table 2-6 on page 59.
Step 3	Configure optional directory integration features.	You can use LDAP directory data sources to integrate Symantec Brightmail Gateway with your existing directory data infrastructure. See Table 2-7 on page 60.
Step 4	Configure optional email management and filtering features.	Symantec Brightmail Gateway enables you to manage many aspects of email flow and filtering. These features can vastly increase antispam effectiveness, reduce infrastructure needs, and significantly enhance protection of your users and assets. See Table 2-8 on page 60.

Performing optional configuration tasks

Depending on your network environment, your users, and your processing needs, you may need to change some configuration settings in order to make the Symantec Brightmail Gateway product work optimally in your environment.

Symantec recommends enabling reputation filtering for increased antispam effectiveness and processing efficiency. You may want to enable other optional features. Some optional features require the configuration of an LDAP directory data source, or have other requirements.

See "Feature dependencies" on page 47.

Table 2-6 Communications and monitoring

Action	Description
Configure additional Scanner settings	In addition to the MTA and SMTP choices made during installation, you can configure additional settings as needed. You can enable Scanner email settings, SMTP, and IM filtering. See "Setting up Scanners" on page 62.
Set up alerts, log settings, and report settings	Ensure that the appropriate system administrators are alerted of situations requiring their attention. Set log levels and locations. Ensure that the data required for the types of reports you want to run is collected. See "Setting up system monitoring with alerts, logs, and report settings" on page 63.
Set up certificates and domain keys	Enable certificates to provide secure communications via HTTPS and TLS. Set up domain keys to use for DKIM authentication. See " Setting up certificates for authentication" on page 64.
Configure Control Center settings	Configure certificates, system locale, fallback encoding, listening ports, and SMTP settings for the Control Center. Set up end user logins for access to Spam Quarantine, and manage end user preferences data. See "Configuring Control Center settings" on page 65.

Table 2-6 Communications and monitoring *(continued)*

Action	Description
Set up SNMP and UPS	Set up System Network Management Protocol (SNMP) monitoring of your appliances, and set up a Universal Power Supply (UPS) automated backup power facility. See "Setting up SNMP and UPS monitoring" on page 67.

Table 2-7 Directory integration

Action	Description
Configure directory integration	Create and configure LDAP directory data sources. Some Symantec Brightmail Gateway features require you to configure a directory data source. See "Configuring a directory data service" on page 66.

Table 2-8 Email managment and filtering

Action	Description
Configure email settings	Configure additional local and non-local domains, address masquerading, aliasing, invalid recipient handling, bad message handling, SMTP greetings, postmaster address, and container limits. See "Configuring email settings" on page 68.
Enable reputation filtering	Enable preliminary filtering at connection time via Brightmail Adaptive Reputation Management. By enabling this feature you can dramatically reduce message processing volumes and enhance protection. See "Enabling reputation-based filtering features" on page 69.

Table 2-8 Email managment and filtering *(continued)*

Action	Description
Configure spam, virus, and IM settings and policies	You can set up custom policies that determine what actions Symantec Brightmail Gateway takes on spam, suspected spam, viruses, and IM messages. Or, you can skip this step and use default policies. See "Configuring spam, virus, and IM settings policies" on page 71.
Set up email authentication	You can set up four different types of email authentication: SPF, Sender ID, DKIM, and SMTP. See "Setting up email authentication " on page 72.
Create policy groups	You can set up groups of users, so that you can process email and IM messages differently based on group membership. Assign policies to groups. Or, you can skip this step if you want to apply the same actions to email and IM messages for all users. See "Creating policy groups and assigning policies" on page 73.
Set up content filtering	Set up policies that process email messages based on their content. Set up policies that enforce regulatory requirements in your email message flow, including review prior to the final action on a message. See "What you can do with content filtering" on page 75.

Setting up Scanners

Table 2-9 Processes for configuring Scanners

Phase	Action	Description
Phase 1	Verify Scanner configurations.	After installing the Scanner software, review your configurations. See "Specifying DNS server addresses" on page 80. See "Verifying Scanner time settings" on page 81. See "Specifying proxy settings" on page 81. See "Configuring Ethernet settings and routes" on page 83.
Phase 2	Configure Scanner email acceptance and delivery settings.	Configure scanner email settings. See "Configuring mail flow direction" on page 88. See "Configuring Scanner inbound email delivery settings" on page 91. See "Configuring Scanner outbound mail delivery settings" on page 94. See "About Scanner email settings " on page 86.
Phase 3	Configure advanced Scanner SMTP settings.	Configure advanced Scanner SMTP settings. See "SMTP advanced authentication settings" on page 96. See "SMTP advanced inbound settings" on page 98. See "SMTP advanced outbound settings" on page 99. See "SMTP advanced delivery settings" on page 101. See "Configuring SMTP advanced settings" on page 95.
Phase 4	Configure internal mail hosts for non-gateway deployment.	See "Specifying internal mail hosts for non-gateway deployments" on page 104. See "Internal mail servers: non-gateway deployments" on page 106.

Table 2-9 Processes for configuring Scanners *(continued)*

Phase	Action	Description
Phase 5	Enable filtering.	See "Enabling IM filtering" on page 292.

Setting up system monitoring with alerts, logs, and report settings

Table 2-10 Processes for setting up system monitoring

Step	Action	Description
Step 1	Configure alerts.	Specify the criteria by which you want alerts set. See "Configuring alerts" on page 614. See "Types of alerts" on page 615.
Step 2	Configure log levels.	Specify how much data you want the system to store. See "Configuring log levels" on page 643. You must enable Message Audit log before any auditing information is available for viewing or searching. See "Enabling the Message Audit Log" on page 643.
Step 3	Select report data.	From the available options, select the data you want the system to track. Data you select here is visible when you create reports. See "Selecting the data to track for reports" on page 569.

Setting up certificates for authentication

Table 2-11 Process for setting up Certified Authority-signed certificates

Step	Action	Description
Step 1	View Symantec's preinstallation CA certificates.	Viewing the preinstallation CA certificates enables you to determine if you need additional certificates. See "About certificates" on page 191. See "Viewing existing CA certificates" on page 198.
Step 2	Add an additional certificate.	If you determine you need additional certificates, you can add certificates to the list of available certificates. See "Adding a self-signed certificate" on page 193. See "Adding a CA certificate" on page 193. Ensure that you use standard PEM formats. See "PEM format requirements for certificates and domain keys" on page 196.
Step 3	Submit a certificate request.	If the request is successful, the certificate authority returns an identity certificate that has been digitally signed with the private key of the certificate authority. See "Requesting a Certificate Authority-signed certificate" on page 194.
Step 4	Import the Certificate Authority-signed certificate.	You must import the new certificate to make it available. See "Importing a Certificate Authority-signed certificate" on page 197.
Step 5	Assign a certificate to a Scanner and Control Center.	If you want a Scanner to accept TLS encrypted messages you need to assign an MTA TLS certificate. See "Assigning an MTA TLS certificate to a Scanner" on page 200. Enhance Control Center security. See "Assigning a user interface HTTPS certificate to the Control Center" on page 201.

Configuring Control Center settings

Table 2-12 Process for configuring Control Center settings

Step	Action	Description
Step 1	Set up host and client mappings.	Specify host names to enable Reverse DNS lookup. See "About specifying host names for Control Center access" on page 669. Specify which computers or networks can access the Control Center. See "Specifying which hosts can access the Control Center" on page 669.
Step 2	Add enhanced security for the Control Center.	Designate a user interface certificate HTTPS. See "Designating a Control Center certificate" on page 671. Modify default demo certificate to prevent security warning. See "Bypassing the security warning when you access the Control Center" on page 671.
Step 3	Modify default settings.	Change or disable the listener port. See "Configuring the Control Center listener port" on page 672. Specify host settings for product generated alert and report notifications. See "Configuring Control Center SMTP settings for alerts and reports" on page 673. Configure the Control Center for single- and double-byte character sets and for number, date, and time settings. See "Setting the locale encoding and fallback encoding" on page 670.

Table 2-12 Process for configuring Control Center settings *(continued)*

Step	Action	Description
Step 4	Customize user login experience.	Customize the **Login help** page. See "Specifying a custom user Login help page" on page 674. Enable "Remember me" feature for automatic login. See "Enabling users to bypass Control Center login credentials" on page 674. Extend user preferences to attached Scanners. See "Configuring the replication of end user preference data" on page 676.

Configuring a directory data service

Table 2-13 Processes for setting up a directory data service

Step	Action	Description
Step 1	Create a new data source.	See "Creating a data source" on page 491. See "Adding a data source " on page 492.
Step 2	Enable functions for your new data source.	See "Enabling functions on a new data source" on page 498. See "Creating a recipient validation data source" on page 512. See "Creating an authentication data source" on page 500. See "Creating a routing data source" on page 516. See "Creating an address resolution data source" on page 520.

Table 2-13 Processes for setting up a directory data service *(continued)*

Step	Action	Description
Step 3	Customize and test data source queries.	See "Creating and testing a custom recipient validation query" on page 514. See "Creating and testing a custom authentication and quarantine address resolution query" on page 505. See "Creating and testing a custom routing query" on page 517. See "Creating and testing a custom address resolution query" on page 523.
Step 4	Set directory data cache.	See "About preloading your directory data cache" on page 528. See "Preloading your data cache" on page 528.
Step 5	Configure advanced settings.	See "Configuring data source advanced settings" on page 494.

Setting up SNMP and UPS monitoring

Table 2-14 Process for setting up SNMP and UPS monitoring

Step	Action	Description
Step 1	Establish your Management Information Base (MIB) database to manage your network devices.	Required. Download and import an MIB. Symantec provides an MIB selection for both the hardware appliance and the Symantec Brightmail Gateway application. See "Downloading a Management Information Base for SNMP" on page 620.
Step 2	Set up SNMP for monitoring devices and access privileges to the SNMP agent.	Configure SNMP settings. See "Monitoring devices through SNMP" on page 619.
Step 3	Set the device shut down preferences in event of power failure.	Enable UPS monitoring and select conditions under which the appliance turns itself off. See "Configuring UPS settings" on page 621.

Configuring email settings

Table 2-15 Process for configuring Scanner email settings

Step	Action	Description
Phase 1	Set up domains.	Understand how email domains work. Determine which domain confirfguration is right for you. See "About email domains" on page 115. See "Adding or editing domains" on page 117. Import an existing list of domains. See "Importing a domains list" on page 120. Specify domain acceptance settings. See "About email domain acceptance settings" on page 122.
Phase 2	Configure email aliases and address masquerades.	Understand how aliases and address masquerades work. See "About aliases and address masquerades" on page 123. Map a source email address and destination for each alias. See "Adding or editing aliases" on page 124. See "Alias addresses" on page 126. Import an existing list of aliases. See "Importing aliases" on page 125. Add, edit, and import address masquerades. See "Adding or editing address masquerades " on page 127. See "Importing an address masquerade list" on page 128.

Table 2-15 Process for configuring Scanner email settings *(continued)*

Step	Action	Description
Phase 3	Configure invalid recipient email handling.	IR handling options vary depending on features you want to enable. Determine your needs before specifying up IR handling. See "About invalid recipients" on page 129. Specify what you want the system to do with invalid recipient emails. See "Setting up invalid recipient handling" on page 130.
Phase 4	Understand SMTP (Simple Mail Transfer Protocol) message flow phases.	See "About email message flow" on page 131. Use the message flow phases as a guideline to creating and maintaining your Scanner email configurations. See "Email message flow phases" on page 132.

Enabling reputation-based filtering features

Table 2-16 Process for setting up reputation based filtering features

Phase	Action	Description
Phase 1	Understand Symantec Brightmail Gateway's blocking features and technologies.	See "About blocking and allowing messages at connection time" on page 163. See "About managing connection load at the gateway" on page 166.
Phase 2	Customize message flow parameters.	Configure Connection Classification to customize the parameters for your message flow. See "Configuring Connection Classification" on page 168. See "Connection class default settings" on page 169.

Table 2-16 Process for setting up reputation based filtering features *(continued)*

Phase	Action	Description
Phase 3	Set up email virus attack recognition and specify actions to take.	Configure your system to recognize and block offenders. See "Configuring email virus attack recognition" on page 170. See "Configuring directory harvest attack recognition" on page 172. Specify filter rules for Bad Sender email. See "About blocking and allowing messages using sender groups" on page 174. Use additional tools for IP analysis to verify sender history. See "Researching IP address reputation" on page 186.
Phase 4	Protect system resources.	Understand how Fastpass can prevent drainage of valuable system resources. See "About conserving resources using Fastpass" on page 178. Configure system to recognize and fast track legitimate senders and reduce processing power. See "Configuring Fastpass" on page 179.

Table 2-16 Process for setting up reputation based filtering features *(continued)*

Phase	Action	Description
Phase 5	Set up Bounce Attack Prevention .	Understand bounce attacks and non-deliverable receipt (NDR) messages. See "About defending against bounce attacks" on page 182. Provide a Bounce Attack Prevention seed value in your Control Center. See "Configuring the Control Center for bounce attack prevention" on page 183. Determine and configure the groups to which you want the system to apply Bounce Attack Prevention See "Configuring policy groups for bounce attack prevention" on page 184. Assign a policy (a default policy is provided) to the group that determines the actions to be taken for NDRs that do not pass Bounce Attack Prevention validation. See "Creating an email spam policy for bounce attack prevention" on page 185.

Configuring spam, virus, and IM settings policies

Table 2-17 Processes for configuring spam, virus, and IM settings policies

Step	Action	Description
Step 1	Complete the setup process and configure email spam policies.	See "Configuring spam detection" on page 240. Create spam policies that determine the actions the system must take for specific email spam conditions. See "Creating email spam policies" on page 242.
Step 2	Complete the setup process and configure email virus policies.	See "How to detect virus and malicious threat detection" on page 208. Create virus policies that determine the actions the system must take for specific email virus conditions. See "Creating email virus policies" on page 210.

Table 2-17 Processes for configuring spam, virus, and IM settings policies *(continued)*

Step	Action	Description
Step 3	Set up threat protection policies for your instant messaging network.	Create IM virus and spim policies that determine the actions the system must take for specific IM virus conditions. See "Creating IM virus policies" on page 309. See "Creating IM spim policies" on page 311.

Setting up email authentication

Table 2-18 Process for setting up email authentication

Step	Action	Description
Step 1	Enable SPF and Sender ID authentication.	Set up sender authentication and specify how the system processes inbound email. See "Enabling SPF and Sender ID authentication" on page 135.
Step 2	Enable DKIM authentication.	Enable system to perform DKIM validation on inbound email. Complete the process steps in the following section to set up DKIM authentication. See "Configuring DKIM authentication" on page 137.
Step 3	Enable SMTP authentication.	Set up authentication and specify how the system handles outbound email. Choose an authentication source and method, and configure MUAs to connect to the authentication listener port. See "Using SMTP authentication" on page 145. Complete the process steps in the following section to set up SMTP authentication. See "Configuring SMTP authentication mail settings" on page 149.

Creating policy groups and assigning policies

Table 2-19 Process for creating policy groups and assigning policies

Step	Action	Description
Step 1	Create a policy group and add members	Create a policy group See "Creating a policy group" on page 316. Add members to the policy group See "Adding members to a policy group" on page 317.
Step 2	Assign email message and instant message (IM) **Virus policies** to a policy group.	Symantec pre-installs email and IM virus policies which you can assign to your policy group. See "Selecting virus policies for a policy group" on page 322. See "Default email virus policies" on page 52. See "Default IM virus policies" on page 56. If you want to assign custom email and IM virus policies to a policy group, you must first create these; then, return to this step and apply them. See "Creating email virus policies" on page 210. See "Creating IM virus policies" on page 309.
Step 3	Assign email message and instant message (IM) **Spam/Spim policies** to a policy group.	Symantec pre-installs email spam and IM spim policies which you can assign to your policy group. See "Selecting spam and spim policies for a policy group" on page 324. See "Default email spam policies" on page 50. See "Default IM spim policies" on page 55. If you want to assign custom email spam and IM spim policies to a policy group, you must first create these; then, return to this step and apply them. See "Creating email spam policies" on page 242. See "Creating IM spim policies" on page 311.

Table 2-19 Process for creating policy groups and assigning policies *(continued)*

Step	Action	Description
Step 4	Assign content filtering policies to a policy group	There are no default content filtering policies. You must set these up using predefined templates. See the process topic for assigning content filtering policies. See "Creating content filtering policies" on page 334.
Step 5	Create notifications or annotations you want to include in that policy.	When a policy is violated you can customize a notification to send to the email or IM sender and recipient. See "Creating policy violation notifications" on page 401. See "Annotating messages that contain violations" on page 421.

What you can do with content filtering

Table 2-20 What you can do with content filtering

Task	Description
Create a basic content filtering policy.	Symantec Brightmail Gateway does not provide default content filtering policies, so you must create the policies that your organization needs. You can create basic policies that do not require the use of any resources (such as dictionaries or lists). To create a basic policy that does not use resources, ensure that you select the blank template. The Structured Data template and Described Content template require the use of resources. See "About content filtering policy templates" on page 336. Examples of the basic policies that do not require resources are as follows: ■ Search email for a specific string. ■ Perform an action on email that is sent to a specific person. ■ Search outbound email for words that contain proprietary information. See "Creating content filtering policies" on page 334. See "About content filtering" on page 331.
Create a complex content filtering policy.	To create more complex policies that use resources, configure the resources that you want to use in your policy before you create the policy. See "Annotating messages that contain violations" on page 421. See "About attachment lists" on page 424. See "About content filtering dictionaries" on page 404. See "About policy violation notifications" on page 400. See "About patterns" on page 417. See "About preventing data loss with structured data" on page 375. See "Creating content filtering policies" on page 334.

Table 2-20 What you can do with content filtering *(continued)*

Task	Description
Set up content incident folders.	You can specify actions in your content filtering policies to create incidents in content incident folders. Content incident folders help you organize, monitor, and manage the incidents of content filtering policy violations. See "About content incident folders" on page 438.
Monitor incidents.	Monitor the incidents in your content incidents folders. You can take actions on the incidents (such as to forward them) or do nothing at all. See "About monitoring and acting on incidents" on page 444.
Configure the content incidents folder Expunger.	To manage the size of content incidents folders, by default an Expunger runs every day at midnight. It deletes the oldest incidents in the folders that exceed the default size, 1 MB. However, you can modify the frequency with which the Expunger runs and specify the thresholds that you want the Expunger to use. See "About managing the size of content incident folders" on page 439.
Configure Symantec Network Prevent integration.	Symantec Brightmail Gateway integrates with Symantec Network Prevent to deliver, route, hold, or block email traffic. Symantec Network Prevent is a component of Symantec Data Loss Prevention, which discovers, monitors, and protects confidential data wherever it is stored or used. With Symantec Data Loss Prevention, you can create policies that extend across endpoint, network, and storage systems. See "About Symantec Network Prevent" on page 462.
Configure Symantec Brightmail Gateway for content encryption.	Symantec content encryption leverages Symantec Hosted Services to provide you the ability to encrypt outbound messages. To encrypt messages, you must purchase the Symantec Content Encryption license, configure your system for encryption, and provision an encryption account. You then create and assign the policies that encrypt outbound messages. See "Preparing your system for content encryption" on page 459.

Table 2-20 What you can do with content filtering *(continued)*

Task	Description
Configure actions for DKIM authentication	Domain Keys Identified Mail (DKIM) is an email authentication protocol that uses public key cryptography. If you implement DKIM validation on inbound messages, you can create a content filtering policy to apply actions to messages that fail DKIM validation. See "Configuring DKIM authentication" on page 137. See "Creating a content filtering policy for DKIM validation" on page 141.

Configuring Scanners

This chapter includes the following topics:

Adding Scanners

You must have full administration rights or manage settings modify rights to add a Scanner. After you add a Scanner, you can check its status to ensure that it functions properly.

See "Viewing the status of software and services" on page 609.

Note: If you have provisioned content encryption for your messaging system and then add or change the IP address of a Scanner, you must inform your Symantec provisioning representative. For more information, see the Symantec Content Encryption Provisioning page by clicking **Content > Settings > Content Encryption** and then clicking **Click here**.

See "About encrypting messages with Symantec Content Encryption" on page 458.

To add a Scanner

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Click **Add**.
3. Complete the Add Scanner Wizard.

   Refer to the *Symantec Brightmail Gateway Installation Guide* for information about completing the Add Scanner Wizard.

Verifying Scanner installation settings

Specifying DNS server addresses

Domain Name System (DNS) servers translate domain names into IP addresses. You specify the IP addresses of the DNS servers that a Scanner uses during installation, and you can change the IP addresses after installation. If your Scanner hosts email and instant message scanning, the DNS servers that you specify are used for email delivery and external IM communication.

Note: For DNS servers to function properly, firewall port 53 must be open.

You must have Full Administration rights or Manage Settings modify rights to add or modify DNS server settings.

To specify DNS server addresses

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Click the name of the host whose DNS definitions you want to modify.
3. On the **Edit Host Configuration** page, click the **DNS/Time** tab.

4 Click **Use internal DNS server**, OR

Click **Use the following external DNS servers** and type the IP address of each DNS server.

If you chose **Use internal DNS server** and you need to override the default DNS query port of 53, you can type the new port number in the **Query source port** field.

When necessary, you can also flush buffers for DNS servers or routers.

5 Check **Apply above settings to all hosts** to apply your changes to all hosts.

6 Click **Save**.

Verifying Scanner time settings

You can specify primary, secondary, and tertiary Network Time Protocol (NTP) time servers or you can manually specify the time.

You must have Full Administration rights or Manage Settings modify rights to configure time settings.

To configure Scanner time settings

1 In the Control Center, click **Administration > Configuration**.

2 Check the name of the host whose NTP server definitions you want to modify and click **Edit**.

3 On the **Edit Host Configuration** page, click the **DNS/Time** tab.

4 Select the time zone of the host to which you want to synchronize the time.

5 Type the NTP hostnames or IP addresses.

You can choose to specify up to three NTP servers, or you can set the time manually.

6 Click **Apply above settings to all hosts** to apply your changes to all hosts.

7 Click **Save**.

8 Restart the computer for the time changes to take effect.

Specifying proxy settings

The Conduit service runs on each Scanner. Through this conduit, you can register your licenses, update antispam filters, download new virus definitions, and perform software updates. If you use a proxy host, you must add the proxy server information to your Scanner definition.

Proxy errors are logged to BrightmailLog.log at the warning log level. Before you configure proxy access, you may want to configure BrightmailLog.log to log warning-level errors with the command `cc-config cclog --level warnings`. If the proxy host does not function after you configure it, check BrightmailLog.log for the errors described in Table 3-1.

See cc-config on page 742.

Table 3-1 Potential proxy host errors

Problem	Error in BrightmailLog.log
The proxy host requires a user name and password, but none was specified.	`IOException: Server returned HTTP response code: 407 for URL:`
The wrong user name or password was specified for the proxy host.	`ProtocolException: Server redirected too many times`
The wrong address and port were specified for the proxy host.	Various timeout errors can occur

Note: LiveUpdate uses the proxy that you define for the Scanner to download virus definitions from Symantec. If you download virus definitions from a LAN host, LiveUpdate uses a proxy only if you have defined one.

See "Specifying from where to obtain virus definitions" on page 225.

To specify proxy settings

1. In the Control Center, click **Administration > Hosts > Configuration.**
2. Under **Hostname,** check the name of the Scanner for which you want to specify a proxy and click **Edit**.
3. On the **Edit Host Configuration** page, click the **Proxy** tab.
4. Check **Use proxy server**.
5. In the **Proxy host name** field, type the proxy host name.
6. In the **Proxy host port** field, type the proxy port number.
7. Specify a user name and password if they are required to log on to the proxy host.
8. Click **Save**.

Configuring Ethernet settings and routes

A Scanner's Ethernet settings can be customized to accommodate your site's mail-flow requirements as follows:

- Configure a Scanner's Ethernet settings and, optionally, enable an Ethernet interface to use up to 50 static routes.
  Static routes direct data from one subnet to a different subnet faster than dynamic routes. Static routes must be updated if addresses change.
- Add or delete virtual IP addresses to an Ethernet interface as needed.
- Change the IP address of a Scanner.

You must have Full Administration rights or Manage Settings modify rights to modify Ethernet settings.

Table 3-2 lists the Scanner Ethernet options that you can configure.

Table 3-2 Ethernet connection parameter descriptions

Item	Description
Description	Provides a free-form description for the Ethernet interface.
IP address	Specifies the IP address of the Ethernet interface. You can change the IP address of a Scanner or the Control Center. If you change the IP address of the Control Center, you must subsequently log into each Scanner's command-line interface and use the `agent-config` command to re-allow the secured connection to the Scanner from the new Control Center IP. See agent-config on page 739.
Netmask	Specifies which part of an IP address is the network address and which is the host address.
Broadcast	Specifies a common address that is used to direct (broadcast) a message to all terminals on a network. The broadcast address is based on the IP address and the subnet mask.
Network	Specifies an address identifying the node.
Negotiation	Check **Auto Negotiation** if you want Symantec Brightmail Gateway to negotiate the rate for this connection or choose the speed (10, 100, or 1000 GB per second) and duplex setting (half or full) for your high-speed Ethernet connection. **Note:** If the port you want to use is set to a specific speed and duplex setting, choose that speed and duplex setting and do not choose **Auto Negotiation**.

Table 3-2 Ethernet connection parameter descriptions *(continued)*

Item	Description
Segmentation	Check **Enable TCP Segmentation Offload** if you want the network interface card {NIC) to segment data for outbound transmission. This feature can reduce the load on the CPU resources that are needed to process large volumes of mail for high-bandwidth networks.
Maximum transmission unit	Works in tandem with the TCP segmentation offload. You can enter a value from 512 bytes to 1500 bytes. When TSO is enabled and the default maximum transmission unit is blank, the value is whatever the NIC (network interface card) specifies for its maximum value. Typically, this value is 1500 bytes maximum for an Ethernet link. **Note:** When you change this value, you must then restart the Control Center using the `service controlcenter restart` command. See service on page 805.
Virtual	Allows for a virtual IP address to be defined. Depending on the interfaces you specify, up to 256 virtual networks can appear.
Enable this interface	Activate the second Ethernet interface.
Default gateway	Address of a router to handle IP traffic between networks (required). Specify an IP address.
Destination address	Computer or network destination for a static route. Specify an IP address, IP address with subnet mask, or CIDR address. Static routes direct data from one subnet to a different subnet faster than dynamic routes. Static routes must be updated if addresses change.
Gateway address	Address of a router to use for a static route. Specify an IP address.
Interface	The opening in the back of the appliance that allows for the insertion of an Ethernet cable. You can use this interface for a static route, by indicating the hostname or IP address and port that identifies the interface.

To configure Ethernet settings

1. In the Control Center, click **Administration > Hosts > Configuration.**
2. Under **Hostname**, check the Scanner whose Ethernet settings you want to modify and click **Edit**.
3. On the **Edit Host Configuration** page, click the **Ethernet** tab.
4. Edit the Ethernet interfaces as needed by entering an IP, netmask, broadcast, and network address for each Ethernet interface.

   Change the Scanner or Control Center IP address if needed. You can change the IP address of a Scanner or the Control Center. If you change the IP address of the Control Center, subsequently log into each Scanner's command-line interface. Use the `agent-config` command to re-allow the secured connection to the Scanner from the new Control Center IP.
5. For each Ethernet NIC, check **Auto Negotiation**. Otherwise, select a speed for the connection, and specify half or full duplex operation of the connection.
6. Check **Segmentation** to offload TCP segmentation from the gateway CPU to the Ethernet card.
7. Check **Enable this interface** to activate a second Ethernet interface.

   You can dedicate a second Ethernet interface to handling inbound or outbound traffic only.
8. Add or delete virtual IP addresses to an Ethernet interface as needed.

   See To add or delete virtual IP addresses.

   If you use this Scanner to filter IM traffic, you must configure a virtual IP address. This IP address must be routable on your organization's internal network so that IM clients inside your network can connect to it.
9. Under **Routing**, in the **Default gateway** field, type the IP address of a default gateway.

   A default gateway is required. You must indicate it with an IP address.
10. To optionally define a static route, under **Static Routes**, specify the following:
    - Destination address—IP address, IP address with subnet mask (for example, 128.113.1.0/255.255.255.0), or CIDR notation (for example, 192.30.250.00/18)
    - Gateway address—IP address
    - Interface—None, Ethernet 1, or Ethernet 2

11 Click **Add** to add a static route.

You can add up to 50 static routes per Ethernet interface.

12 Click **Save**.

To add or delete virtual IP addresses

1 Click **Administration > Hosts > Configuration**.

2 Under the **Hostname** check the name of the Scanner whose virtual IP addresses you want to modify and click **Edit**.

3 On the **Edit Host Configuration** page, click the **Ethernet** tab.

4 To add or delete a virtual IP address to either Ethernet interface, perform one of the following tasks:

To add a virtual address	In the **Virtual** field, type the IP address and details for that virtual IP address and then click **Add**. You must have checked **Enable this interface** for additions to the second Ethernet interface to take effect when settings are saved. You can add up to 256 virtual IP addresses per Ethernet interface.
To delete a virtual IP address	Check the box beside the name that is associated with the virtual address and click **Delete**.

5 Click **Save**.

Configuring Scanner email settings

About Scanner email settings

Scanner email settings let you control various aspects of inbound and outbound message flow. You can limit the IP addresses from which Symantec Brightmail Gateway accepts email messages. You can also control where Symantec Brightmail Gateway delivers filtered email.

You can reduce the volume of email that any one Scanner filters by configuring separate inbound and outbound Scanners. Even if you use only one Scanner to filter both inbound and outbound email, you can configure various SMTP settings to control message flow.

Note: Individual Scanner email settings should not be confused with spam or virus scan settings, which control spam and virus scanning activity for all Scanners.

During site setup, you configure at least one combined inbound and outbound Scanner. This Scanner can be on the same computer that runs the Control Center, or on another computer. You can instead configure separate inbound and outbound Scanners to run on separate hosts or add dedicated Scanners later.

Note: If you use the same IP address and port for inbound and outbound email, the Scanner uses outbound mail acceptance settings to determine if a message is inbound or outbound. The Scanner first checks outbound mail acceptance settings. If the mail is not accepted, the Scanner then checks inbound mail acceptance settings.

See "Changing Scanner outbound mail acceptance settings" on page 93.

The Add Scanner Wizard guides you through the process of configuring a Scanner's email settings. These initial email settings include:

- Scanner role – inbound email, outbound email, or both inbound and outbound email.
- Mail filtering – Scanner IP address and port number.
- Mail acceptance – IP addresses from which the Scanner accepts email to be filtered.
- Mail delivery – mail server to which filtered email is relayed and whether to enable MX Lookup.

See "Adding Scanners" on page 79.

Note: A Scanner that is configured to filter email can also be configured to filter instant messages. However, IM connections must be made over a separate, virtual IP address.

After adding a Scanner to your deployment and testing it, you can modify its settings further by:

- Changing the Scanner role or limiting it to IM only
  See "Configuring mail flow direction" on page 88.
- Modifying the Scanner's inbound email settings.
  See "Changing Scanner inbound mail settings" on page 89.
  See "Changing Scanner inbound mail acceptance settings" on page 90.
  See "Configuring Scanner inbound email delivery settings" on page 91.

- Modifying the Scanner's outbound email settings.
  See “Changing Scanner outbound mail settings” on page 92.
  See “Changing Scanner outbound mail acceptance settings” on page 93.
  See “Configuring Scanner outbound mail delivery settings” on page 94.
- Configuring advanced SMTP settings
  See “Configuring SMTP advanced settings” on page 95.

In addition, Connection Classification defers some connections based on local reputation data that is collected and implemented on a per-Scanner basis.

See “About managing connection load at the gateway” on page 166.

Configuring mail flow direction

Scanners can be configured to filter inbound, outbound, or both inbound and outbound email.

To configure mail flow direction for a Scanner

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Under **Hostname** check the Scanner whose role you want to define and click **Edit**.
3. On the **Edit Host Configuration** page, click the **SMTP** tab.
4. In the **Host name** field, modify the name for the Scanner if necessary to identify it by its role.
5. In the **Host definition** field, modify the definition of the Scanner to reflect its new role.
6. Select the option that describes the role of this Scanner.
   - No mail filtering – Select this option if you plan to use this Scanner for IM only and click **Save**.
   - Inbound mail filtering only – Configure inbound SMTP Scanner settings.
   - Outbound mail filtering only – Configure outbound SMTP Scanner settings.
   - Inbound and Outbound mail filtering – Configure both inbound and outbound SMTP Scanner settings.

   See “About Scanner email settings ” on page 86.
7. At the bottom of the page, check **Apply above settings to all hosts** if you want settings to apply to all Scanners.
8. Click **Save** to store your changes.

Changing Scanner inbound mail settings

You can change the IP address or port number through which a Scanner accepts inbound mail connections. You can also designate whether or not the Scanner accepts inbound TLS-encrypted connections.

To change Scanner inbound mail settings

1 In the Control Center, click **Administration > Hosts > Configuration**.

2 Under **Hostname**, check the inbound Scanner whose settings you want to modify and click **Edit**.

3 On the **Edit Host Configuration** page, click the **SMTP** tab.

4 Select the address where you want to receive inbound messages in the **Inbound mail IP address** drop-down menu.

Only those IP addresses, including virtual IP addresses, that have been configured for this Scanner's network interface card appear in the drop-down menu.

5 In **Inbound mail SMTP port** text box, enter the port number where you want inbound mail to be received.

Typically the port number is 25.

6 Check **Accept TLS encryption** if you want the host to accept connections using TLS encryption.

If you leave this option unchecked, Symantec Brightmail Gateway will not advertise support for TLS encryption during the SMTP session.

Note: You must configure an MTA TLS certificate and assign it to this Scanner before it can accept TLS encrypted email from a connection.

7 Select the name of a certificate from the drop-down menu to authenticate the Scanner as a trusted source to clients sending over TLS-encrypted connections.

See "About certificates" on page 191.

8 Check **Request client certificate** if you want the Scanner to request a TLS encryption certificate from a sender before accepting a TLS-encrypted connection.

9 Click **Save** to save settings for this host only.

Changing Scanner inbound mail acceptance settings

You can conserve Scanner resources by limiting inbound connections to only IP addresses from which you want the Scanner to filter email. By allowing connections from only certain IP addresses and domains, you exclude all other inbound clients from sending messages at connection time.

Warning: If you accept mail only from selected hosts, and your Scanner is not at the gateway, you must add all upstream mail servers. Add the IP addresses, CIDR blocks, or netmasks of upstream mail servers on both the **SMTP** and **Internal Mail Hosts** tabs of the **Edit Host Configuration** page. Symantec Brightmail Gateway rejects email from unspecified upstream servers.

See "Specifying internal mail hosts for non-gateway deployments" on page 104.

To change Scanner inbound mail acceptance settings

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Under **Hostname**, check the name of the inbound Scanner whose settings you want to modify and click **Edit**.
3. On the **Edit Host Configuration** page, click the **SMTP** tab.
4. Either:
   - Check **Accept inbound mail connections from all IP addresses** if you want the Scanner to accept connections from senders of all inbound messages, and click **Save**. Skip the rest of this procedure.
   - Check **Accept inbound mail connections from only the following IP addresses** if you want the Scanner to accept only connections from the addresses that you check in the Inbound Mail Acceptance IP Addresses checklist. Proceed to the next step.
5. Check the IP addresses from which you want this Scanner to accept inbound mail.
6. If necessary, click **Add** to add an IP address to the checklist of addresses.
7. Check any other IP address and click **Edit** to alter the address.
8. To delete an address from the checklist, check the IP addresses you want to remove and click **Delete**.
9. Click **Save**.

Configuring Scanner inbound email delivery settings

After a Scanner filters inbound email, the MTA relays filtered email to a mail server for delivery to recipients. When you configure email delivery for an inbound-only or combined inbound and outbound Scanner, you designate the default local and non-local hosts where filtered inbound mail should be routed.

You can designate an unlimited number of local or non-local mail servers for delivery of inbound mail. Multiple downstream mail servers can improve load distribution and fault tolerance for filtered inbound mail if you accept mail that is addressed to different domains. You must specify the order in which a default relay (static route) delivers inbound email by assigning a preference number to each mail host. The Scanner attempts to deliver email to lower- numbered mail hosts first. If only one mail server is configured, its preference defaults to 1. If you want to load balance, specify equal preference numbers for each server.

Additionally, you can assign a mail server to each domain for which the Scanner accepts inbound email. Each mail server can host multiple local domains.

See “About email domains” on page 115.

You can enable MX Lookup for any mail host that you specify using a hostname. MX Lookup then determines which server to use to deliver email that is addressed to a recipient at a local domain.

To configure Scanner inbound email delivery settings

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Under **Hostname**, check the Scanner to which you want to add or edit a local mail server and click **Edit**.
3. On the **Edit Host Configuration** page, click the **SMTP** tab.
4. Under **Inbound Local Mail Delivery**, check the local host whose delivery information you want to change and click **Edit**.

   Alternatively, click **Add** to add delivery information about a local mail host.

   You can configure inbound mail delivery for an unlimited number of local mail hosts.
5. Enter or modify the IP address and port number or hostname in the text boxes that appear.
6. Check **MX Lookup** if you want the local mail host to use MX Lookup to determine which IP address to use for delivery to email recipients.

   You can only check MX Lookup if you specified a hostname.
7. If you have more than one host specified, type a number in the **Preference** field. Connections for lower numbered servers are attempted first.

8 If the Scanner role is **Inbound and outbound mail filtering**, click **Save** to save settings or check **Apply above settings to all hosts** and then click **Save**. Skip the rest of this procedure.

9 If the Scanner role is **Inbound mail filtering only**, click one of the radio buttons under **Inbound Non-Local Delivery**:

- **Use MX Lookup for non-local mail** to have Symantec Brightmail Gateway route mail for non-local recipients by MX query on recipient domains.
- **Relay non-local mail to** to specify mail hosts for delivery of non-local mail.

10 Click **Add** to add the hostname or IP address and port number for a non-local mail host.

You can configure inbound mail delivery for an unlimited number of non-local mail hosts.

You can check the non-local host you want and click **Edit**, then modify the IP address and port number in the text boxes that appear.

11 Check **MX Lookup** if you want the local mail host to use MX Lookup to determine which IP address to use for delivery to email recipients.

You can only check MX Lookup if you specified a hostname.

12 If you have more than one host specified, type a number in the **Preference** field. Connections for lower numbered servers are attempted first.

13 Click **Save** to save the delivery settings for an Inbound-only scanner or check **Apply above settings to all hosts** and then click **Save**.

Changing Scanner outbound mail settings

You can change the IP address or port number where a Scanner accepts outbound mail connections. You can also designate whether or not the Scanner accepts outbound TLS-encrypted connections.

To change Scanner outbound mail settings

1 In the Control Center, click **Administration > Hosts > Configuration**.

2 Under **Hostname**, click the name of the Outbound or Inbound and Outbound Scanner whose settings you want to modify and click **Edit**.

3 On the **Edit Host Configuration** page, click the **SMTP** tab.

4 In **Outbound mail IP address**, enter the IP address where you want outbound messages to be received.

5 In **Outbound mail SMTP port**, enter the port number where outbound mail is received.

Typically the port is 25.

6 Check **Accept TLS encryption** if you want the host to accept TLS-encrypted outbound connections.

Note: You must configure an MTA TLS certificate and assign it to this Scanner before you can accept TLS encrypted outbound mail for filtering.

7 From the **Certificate** drop-down menu, select a certificate to authenticate the Scanner as a trusted source to clients sending over TLS-encrypted connections.

See "About certificates" on page 191.

8 Click **Save** or check **Apply above settings to all hosts** and then click **Save**.

Changing Scanner outbound mail acceptance settings

You can configure a Scanner to accept outbound connections from up to three mail servers. By allowing connections from only certain IP addresses and domains, you exclude all other hosts from sending messages at connection time.

Note: If you use the same IP address and port for inbound and outbound email, the Scanner uses outbound mail acceptance settings to determine if a message is inbound or outbound. The Scanner first checks outbound mail acceptance settings. If the mail is not accepted, the Scanner then checks inbound mail acceptance settings.

To change the outbound mail acceptance settings for a Scanner

1 In the Control Center, click **Administration > Hosts > Configuration**.

2 Under **Hostname**, check the Outbound or Inbound and Outbound Scanner whose settings you want to modify and click **Edit**.

3 On the **Edit Host Configuration** page, click the **SMTP** tab.

4 Under **Outbound Mail Acceptance**, click **Add** and enter the hostname or IP address of a local domain from which you want to filter outbound email. You can enter a range of IP addresses using the CIDR format.

5 Check the hostname or IP address of the email client whose settings you want to change and click **Edit**.

6 Edit the hostname or IP address or CIDR addresses of a range.

7 Select an existing address or domain from which you no longer want to accept outbound connections and click **Delete** to delete it from the list.

8 Click **Save** or check **Apply above settings to all hosts** and then click **Save**.

Configuring Scanner outbound mail delivery settings

An outbound Scanner relays filtered email to a local or non-local mail server for delivery to recipients. When you configure email delivery for an outbound-only or a combined inbound and outbound Scanner, designate the IP addresses and port numbers for both local and non-local mail hosts where outbound mail that has been filtered should be relayed.

You can designate a limited number of mail servers to relay outbound email to the Internet or the next hop in your network. Multiple downstream servers can be used for load balancing or for failover support of the primary mail server for local delivery.

To configure Scanner outbound mail delivery settings for an inbound and outbound Scanner

1 In the Control Center, click **Administration > Hosts > Configuration**.

2 Under **Hostname**, check the name of the Scanner whose settings you want to modify and click **Edit**.

3 On the **Edit Host Configuration page**, click the **SMTP** tab.

4 Scroll to the **Outbound mail settings** section.

5 Under **Outbound Non-Local Mail Delivery**, click one of the following options:

- **Use MX Lookup for non-local mail** to have Symantec Brightmail Gateway route mail for non-local recipients through MX query on recipient domains. Skip the rest of this procedure.
- **Relay non-local mail to** to specify up to three mail hosts for non-local delivery of outbound email.

6 In **Host** list, check the non-local host whose delivery information you want to change and click **Edit**.

7 In the host table, edit the **Host** (IP address) or **Port** (number) of the mail server to which you want the filtered outbound mail delivered.

Alternatively, click **Add** to add delivery information about a local email host or **Delete** to delete a host.

8 Check **MX Lookup** if you want to enable MX lookup for this host.

You can only use MX lookup if you specified a hostname.

9 If you have more than one host listed, enter a number in the **Preference** field. Connections for lower numbered servers are attempted first.

10 Click **Save** to save the delivery settings or check **Apply above settings to all hosts** and then click **Save**.

To configure Scanner outbound mail delivery settings for an outbound-only Scanner

1 Perform the steps in the previous procedure.

See "To configure Scanner outbound mail delivery settings for an inbound and outbound Scanner" on page 94.

2 Locate the **Outbound Local Mail Delivery** section.

Note: The Outbound Local Mail Delivery section is available only on Scanners used for outbound mail filtering only. Ensure that **Outbound mail filtering only** in the **Mail filtering** section is selected to see this option.

---

3 Check a local host under **Relay local mail to** and click **Edit**

Alternatively, click **Add** to add delivery information about a local email host or **Delete** to delete a host.

4 In the host table, edit the **Host** (IP address) or **Port** (number) of the mail server to which you want the filtered outbound mail delivered.

5 Check **MX Lookup** if you want to enable MX lookup for this host.

You can only use MX lookup if you specified a hostname.

6 If you have more than one host listed, enter a number in the **Preference** field. Connections for lower numbered servers are attempted first.

7 Click **Save** to save the delivery settings or check **Apply above settings to all hosts** and then click **Save**.

Configuring SMTP advanced settings

Additional SMTP settings are available from the SMTP Advanced Settings page.

To configure SMTP advanced settings

1 In the Control Center, click **Administration > Hosts > Configuration**.

2 Click the underlined name of the host you want to configure.

3 Click the **SMTP** tab.

4 Scroll to the bottom of the page and click **Advanced Settings**.

5 Make the desired configurations changes.

6 Click **Continue**.

7 On the **Host Configuration** page, click **Save**.

You can configure advanced SMTP settings for:

- Inbound messages
- Outbound messages
- Delivering messages
- Delivery bindings
- SMTP authentication

See "SMTP advanced inbound settings" on page 98.

See "SMTP advanced outbound settings" on page 99.

See "SMTP advanced delivery settings" on page 101.

See "SMTP advanced settings for delivery bindings" on page 102.

See "SMTP advanced authentication settings" on page 96.

Specifying the **MTA host name** lets you define the HELO banner during the initial portion of the SMTP conversation. The **MTA host name** also appears in Received headers. This host name is not connected to the OS-level host name of the Scanner. If you change the host name on the Edit Host page, the **MTA host name** is not changed.

SMTP advanced authentication settings

Table 3-3 describes the advanced authentication SMTP settings that you can use to further define your SMTP configuration.

Table 3-3 SMTP Advanced Settings—SMTP Authentication Configuration

Item	Description
Maximum number of connections	Sets the maximum number of simultaneous SMTP authentication connections. The default is 2,000 connections.

Table 3-3 SMTP Advanced Settings—SMTP Authentication Configuration *(continued)*

Item	Description
Maximum number of connections from a single IP address	Sets the maximum number of simultaneous SMTP authentication connections that can be made from a single IP address. The default is 20 connections.
Maximum message size in bytes	Sets the maximum size allowable for a message before it is rejected. The default is 10,485,760 bytes.
Maximum number of recipients per message	Indicates the maximum number of recipients permitted to receive a message. The default is 1,024 recipients.
Insert RECEIVED header	Places a RECEIVED header in the message during outbound processing of messages sent using SMTP authentication when checked. When unchecked, no RECEIVED header is inserted during outbound SMTP processing. If this option and **Strip pre-existing RECEIVED headers** are both checked, the outbound SMTP RECEIVED header remains when the message goes to the delivery queue.
Strip RECEIVED headers	Removes all RECEIVED headers for outbound messages sent using SMTP authentication when checked. When headers are stripped, message looping can occur depending on the settings of other MTAs. When unchecked, RECEIVED headers remain in the message during outbound processing. The RECEIVED header for outbound SMTP processing remains in the message when both **Insert RECEIVED header** and **Strip pre-existing RECEIVED headers** are checked. **Warning:** Enabling this setting can reduce you ability to diagnose message flow issues.

Table 3-3 SMTP Advanced Settings—SMTP Authentication Configuration *(continued)*

Item	Description
Enable reverse DNS lookup	Causes the system to perform reverse DNS lookup on the SMTP client IP addresses to resolve the IP address to a name when checked. This is the default condition. When unchecked, reverse DNS lookup is not performed for outbound messages sent using SMTP authentication.
Enable optional SASL plain support	By default, Symantec Brightmail Gateway uses SASL login authentication. You can choose to also enable SASL plain authentication. For more information, see the following RFCs: 4954 4616
Session timeout	This setting controls how long the MTA waits for a request or response from the connecting MTA. If this limit is exceeded, the appliance will drop the connection.

Note: The SMTP authentication listener uses the outbound message queue, and inherits some behavior from the outbound listener.

SMTP advanced inbound settings

Table 3-4 describes inbound SMTP settings that you can configure to further define the Scanner SMTP connections.

Table 3-4 SMTP Advanced Settings—Inbound Configuration

Item	Description
Maximum number of connections	Sets the maximum number of simultaneous inbound connections. The default is 2,000 connections.

Table 3-4 SMTP Advanced Settings—Inbound Configuration *(continued)*

Item	Description
Maximum number of connections from a single IP address	Sets the maximum number of simultaneous inbound connections that can be made from a single IP address. The default is 20. If Connection Classification is enabled, the settings for each Connection Class override this setting.
Maximum number of recipients per message	Sets the maximum number of recipients for a message. The default is 1,024 recipients.
Maximum message size in bytes	Sets the maximum size of a message before it is rejected. The default is 10,485,760 bytes.
Maximum number of messages in inbound queue	Sets the maximum threshold. When this threshold is met or exceeded, alerts are sent (when enabled) and connections are deferred. The default is 5,000 messages.
Defer new connections when inbound queue is full	When the number of messages in the inbound queue exceeds the maximum, defers messages, by issuing an SMTP 4xx error. If unchecked, messages are accepted for as long as resources allow. See "Troubleshooting the message queue" on page 662.
Insert a RECEIVED header to inbound messages	Places a RECEIVED header in the message during inbound SMTP processing.
Enable reverse DNS lookup	Causes the system to perform reverse DNS lookup on the SMTP client IP addresses to resolve the IP address to a name when checked. This is the default condition. When unchecked, reverse DNS lookup is not performed for inbound messages.
Session Timeout	This setting controls how long the MTA waits for a request or response from the connecting MTA. If this limit is exceeded, the appliance will drop the connection.

SMTP advanced outbound settings

Table 3-5 describes the advanced outbound SMTP settings that you can use to further define your SMTP configuration.

Table 3-5 SMTP Advanced Settings—Outbound Configuration

Item	Description
Maximum number of connections	Sets the maximum number of simultaneous outbound connections. The default is 2,000 connections.
Maximum number of connections from a single IP address	Sets the maximum number of simultaneous outbound connections that can be made from a single IP address. The default is 20 connections.
Maximum message size in bytes	Sets the maximum size allowable for a message before it is rejected. The default is 10,485,760 bytes.
Maximum number of recipients per message	Indicates the maximum number of recipients permitted to receive a message. The default is 1,024 recipients.
Default domain for sender addresses with no domain	Sets a default domain when none can be found in the message.
Maximum number of messages in outbound queue	Sets the maximum threshold. When this threshold is met or exceeded, alerts are sent (when enabled) and connections are deferred. The default is 5,000 messages.
Defer new connections when outbound queue is full	When the number of messages in the outbound queue reaches or exceeds the maximum, defers messages, by issuing an SMTP 4xx error. If unchecked, messages are accepted for as long as resources allow.
Insert RECEIVED header	Places a RECEIVED header in the message during outbound SMTP processing when checked. When unchecked, no RECEIVED header is inserted during outbound SMTP processing. If this option and **Strip pre-existing RECEIVED headers** are both checked, the outbound SMTP RECEIVED header remains when the message goes to the delivery queue.
Strip pre-existing RECEIVED headers	Removes all RECEIVED headers for outbound messages when checked. When headers are stripped, message looping can occur depending on the settings of other MTAs. When unchecked, RECEIVED headers remain in the message during outbound processing. The RECEIVED header for outbound SMTP processing remains in the message when both **Insert RECEIVED header** and **Strip pre-existing RECEIVED headers** are checked. **Warning:** Enabling this setting can reduce your ability to diagnose message flow issues.

Table 3-5 SMTP Advanced Settings—Outbound Configuration *(continued)*

Item	Description
Enable reverse DNS lookup	Causes the system to perform reverse DNS lookup on the SMTP client IP addresses to resolve the IP address to a name when checked. This is the default condition. When unchecked, reverse DNS lookup is not performed for outbound messages.
Session Timeout	This setting controls how long the MTA waits for a request or response from the connecting MTA. If this limit is exceeded, the appliance will drop the connection.

SMTP advanced delivery settings

Table 3-6 describes SMTP delivery configuration message settings for your site.

Table 3-6 SMTP Advanced Settings—Delivery Configuration

Item	Description
Maximum number of external connections	Sets the maximum number of simultaneous external connections. The default is 100 connections.
Maximum number of external connections to a single IP address	Sets the maximum number of simultaneous connections made to a single IP address. The default is 50 connections.
Maximum number of connections to all internal mail servers	Sets the maximum number of connections that can be made to all defined internal mail servers. The default is 100 internal mail server connections.
Maximum number of connections per single internal mail server	Sets the maximum number of connections to one internal mail server. The default is 50 connections.
Maximum number of messages in delivery queue	Sets the maximum threshold. When this threshold is met or exceeded, alerts are sent (when enabled) or connections are deferred. The default is 150,000 messages.
Defer new connections when delivery queue is full	When the number of messages in the delivery queue reaches or exceeds the maximum, defers messages, by issuing an SMTP 4xx error. If unchecked, messages are accepted for as long as resources allow.
Minimum retry interval	Sets the smallest interval the SMTP server waits before trying to deliver a message again. The default is 15 minutes.

Table 3-6 SMTP Advanced Settings—Delivery Configuration *(continued)*

Item	Description
Sent message time-out	Sets the time after which an undelivered message times out and is rejected from the queue. The default is 5 days.
Bounce message time-out	Sets a time-out period for deletion of messages in your bounce queue. This can be particularly useful in environments where you cannot configure LDAP settings. The default is 1 day.
Message delay time in queue before notification	Sets the time a message waits in the mail queue before notification of nondelivery is sent. The default is 4 hours.
Attempt TLS encryption for all messages	Instructs the MTA to attempt TLS encryption for all messages delivered.
Offer TLS client certificate	Instructs the MTA to offer a client certificate with every TLS connection.

SMTP advanced settings for delivery bindings

Table 3-7 describes the settings available for delivery bindings. Delivery bindings allow you to specify the IP addresses from which messages are sent.

You can also set domain-specific delivery bindings for non-local messages. For each of your local domains, this feature enables you to define one or more IP addresses that messages from that domain are sent from.

Table 3-8 describes the settings available for per-domain delivery bindings for non-local messages.

Table 3-7 SMTP Advanced Settings—Delivery Bindings

Item	Description
Local messages	Sets the IP address that delivers messages locally. The drop-down menu provides a list of IP addresses for this Scanner from which you can choose: the inbound listener IP address, the outbound listener IP address, all virtual addresses, or Auto. If you choose Auto, Symantec Brightmail Gateway automatically chooses the best route based on current traffic flow.

Table 3-7 SMTP Advanced Settings—Delivery Bindings *(continued)*

Item	Description
Non-local messages	Sets the IP address that delivers non-local messages. The drop-down menu provides a list of IP addresses for this Scanner from which you can choose: the inbound listener IP address, the outbound listener IP address, all virtual addresses, or Auto. If you choose Auto, Symantec Brightmail Gateway automatically chooses the best route based on current traffic flow.
Dynamically routed messages	Sets the IP address that delivers messages to non-static routes. The drop-down menu provides a list of IP addresses for this Scanner from which you can choose: the inbound listener IP address, the outbound listener IP address, all virtual addresses, or Auto. If you choose Auto, Symantec Brightmail Gateway automatically chooses the best route based on current traffic flow. **Note:** If you are using multiple IP addresses and your system is provisioned for content encryption, **Dynamically routed messages** must be set to "Auto." See “Managing the host and port information for content encryption” on page 460.
Messages destined for the Control Center	Sets the IP address that delivers mail to the Control Center for storage in Spam Quarantine, Suspect Virus Quarantine, or content incident folders. The drop-down menu provides a list of IP addresses for this Scanner from which you can choose: the inbound listener IP address, the outbound listener IP address, all virtual addresses, or Auto. If you choose Auto, Symantec Brightmail Gateway automatically chooses the best route based on current traffic flow.

Table 3-8 SMTP Advanced Settings—Non-Local SMTP Delivery Bindings Per Domain

Item	Description
Domain	This list includes all domains defined for this Scanner, as local or non-local domains. To select multiple domains, hold down Ctrl.

Table 3-8 SMTP Advanced Settings—Non-Local SMTP Delivery Bindings Per Domain *(continued)*

Item	Description
IP addresses	This list includes all IP addresses available on this Scanner. To select multiple IP addresses, hold down Ctrl.
Add	Click to add the combination(s) of domains and IP addresses selected to the Domain/IP Adress list. By selecting more than one domain or IP address first, you can add multiple rows with one click.
Delete	Check the box next to a row in the Domain/IP Address list and click **Delete** to delete that row.
Delete All	Click to delete all rows in the Domain/IP Address list, including any rows not currently visible.
Entries per page	Set the number of entries to display per page.
Display	Select a range of entries to display.
⏮	Go to beginning of entries.
◀	Go to previous page of entries.
▶	Go to next page of entries.
⏭	Navigate to last page of members or 50 pages ahead if there are more than 50 pages.
Domain/IP Address list	Each row in this list shows only one domain and one IP address. You can use the Add button to add multiple rows simultaneously. Each row defines the delivery binding for non-local messages from one domain.

Specifying internal mail hosts for non-gateway deployments

Internal mail hosts are mail transfer agents (MTAs) that pass email from the Internet to a Scanner. If your Scanners are at the Internet gateway, you do not need to specify internal mail hosts. However, if your network is configured with one or more MTAs that are, with respect to inbound mail flow, upstream from your Scanners, you must specify the IP addresses of these MTAs as internal mail hosts.

If your network has MTAs that are upstream from Symantec Brightmail Gateway, it is important to specify these MTAs as internal mail hosts for the following reasons:

- Email from upstream MTAs to Scanners will likely contain some spam messages. Scanners will see all external email as coming from the IP addresses of the gateway MTAs. If you have enabled Connection Classification, this may result in all email arriving from the Internet being deferred.
- Scanners will not be able to determine the IP address of a sender. Sender groups that match IP addresses such as Local Bad Sender IPs will not function properly.

In addition to internal mail hosts you can add, Symantec Brightmail Gateway includes a series of IP address ranges in the internal hosts list.

See "Internal mail servers: non-gateway deployments" on page 106.

Follow these procedures to add or delete internal mail hosts from which the Scanner is always allowed to receive mail.

To add an internal mail host to the list of allowed hosts

1. From the Control Center, click **Administration > Hosts> Configuration**.
2. Check the Scanner that you want to configure.
3. Click **Edit**.
4. Click the **Internal Mail Hosts** tab.
5. Specify the IP address for an internal mail host.
6. Click **Add**.
7. Click **Save** to store the information.

To delete an internal mail host

1. From the Control Center, click **Administration > Hosts > Configuration**.
2. Check the Scanner you want to configure.
3. Click **Edit**.
4. Click the **Internal Mail Hosts** tab.
5. Select an internal mail host.
6. Click **Delete**.
7. Click **Save** to store the information.

Internal mail servers: non-gateway deployments

When deployed at the gateway, Symantec Brightmail Gateway obtains the physical or peer IP connection for an incoming message and compares it to entries in the good sender and bad sender groups. If a Scanner is deployed elsewhere in your network, for example, downstream from a gateway MTA that is not identified as an internal mail host, Symantec Brightmail Gateway may identify the IP address of your gateway server as a source of spam. You should accurately identify all internal mail hosts that are, with respect to inbound mail flow, upstream from your Symantec Brightmail Gateway appliance.

See "Specifying internal mail hosts for non-gateway deployments" on page 104.

In addition to internal mail hosts you can add, Symantec Brightmail Gateway includes a series of IP address ranges in the internal hosts list as follows:

- 0.0.0.0/255.0.0.0
- 10.0.0.0/255.0.0.0
- 127.0.0.0/255.0.0.0
- 169.254.0.0/255.255.0.0
- 172.16.0.0/255.240.0.0
- 192.168.0.0/255.255.0.0

Symantec Brightmail Gateway will exclude the IP addresses of internal mail hosts from the following verdicts:

- Local Good Sender IPs
- Third Party Good Senders
- Local Bad Sender IPs
- Third Party Bad Senders
- Directory Harvest Attacks
- Symantec Global Bad Senders
- Symantec Global Good Senders
- Connection Classification
- Email Virus Attacks
- Fastpass

Modifying Scanner configurations

You can modify a Scanner's configuration at any time. For example, you can suspend the flow of mail or enable different components and services.

Note: If you have provisioned content encryption for your messaging system and then add or change the IP address of a Scanner, you must inform your Symantec provisioning representative. For more information, see the Symantec Content Encryption Provisioning page by clicking **Content > Settings > Content Encryption** and then clicking the **Click here** link.

See "About encrypting messages with Symantec Content Encryption" on page 458.

You must have Full Administration rights or Manage Settings modify rights to modify Scanner settings.

To modify Scanner configurations

1. In the Control Center, click **Administration > Hosts > Configuration.**
2. Click the linked name of the Scanner that you want to edit.
3. Make any changes to the host or its included components and services.

   See "Adding Scanners" on page 79.

   See "About Scanner email settings " on page 86.

Enabling and disabling Scanners

You can also disable or enable a Scanner, or delete a Scanner. When you disable a Scanner, you stop the flow of statistics, logs, and configuration information between that Scanner and your Control Center. The Scanner can still process messages. The Control Center can still route Spam Quarantine mail to the Scanner. Message Audit Log queries omit the Scanner.

To disable or enable a Scanner

1. In the Control Center, click **Administration > Hosts > Configuration.**

   This page lists your Scanners. A black dash in the Enabled column indicates that the Scanner is disabled. A green check in the Enabled column indicates that the Scanner is enabled.
2. Check the Scanner that you want to change.
3. Click **Enable** to enable the Scanner or click **Disable** to disable the Scanner.

Deleting Scanners

When you delete a Scanner, you permanently remove that Scanner's services from the Control Center. Symantec recommends that you stop a Scanner before you delete it. Otherwise, you can loose the email messages that are in the Scanner email queues. You cannot delete the host on which the Control Center is running. You must have Full Administration rights or Manage Setting modify rights to delete Scanners.

Once you delete a Scanner, you cannot retrieve or access its configuration settings. If you are uncertain as to whether you want to delete a Scanner, you can stop the Scanner. When you stop a Scanner, it still exists but no longer scans messages.

See “Stopping and starting Scanners” on page 108.

See “Enabling and disabling Scanners” on page 107.

To delete Scanners

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Check the box next to the Scanner that you want to delete.
3. Click **Delete**.

Stopping and starting Scanners

You may have an occasion when you want to stop a Scanner. For example, you may want to temporarily stop the mail flow so that you can troubleshoot an issue. After you resolve the issue, you can restart the Scanner. Or you may want to stop a Scanner so that you can delete it. Otherwise, you can lose the email messages that are in the Scanner email queues. You must have Full Administration rights or Manage Setting modify rights to stop and start Scanners.

See “Adding administrators” on page 682.

Symantec recommends that you stop a Scanner before you delete it. A Scanner does not process mail when it is stopped.

Note: You cannot stop the host on which the Control Center is running.

If you have a Scanner that you want to stop permanently or remove, you can delete it.

See “Deleting Scanners” on page 108.

Note: The best procedure to stop a Scanner may vary based on your system parameters and message flow characteristics. You can design your own procedure for stopping a Scanner based on the impact of each of the settings.

See "MTA and message queue behavior" on page 613.

See "Managing services and MTA operations" on page 110.

To stop a Scanner

1 In the Control Center, click **Administration > Hosts > Configuration.**

2 Click the Scanner that you want to stop.

3 Click **Do not accept incoming messages**.

4 Click **Save**.

5 Click **Status > SMTP > Message Queues**.

6 In the **Host** drop-down list, choose a Scanner.

7 In the **Queue** drop-down list, choose a queue.

8 In the **List** drop-down list, click **All** .

Or, to proceed more quickly on a high-volume Scanner, click **10 in queue longest** instead.

9 Click **Display Filtered**.

10 Click **Flush All**.

11 Repeat steps 7 - 10 for the other queues.

12 Let the messages drain from the queue.

To check the message queue status, repeat steps 7 - 9 for each queue.

13 Click **Administration > Hosts > Configuration**.

14 Click the Scanner that you want to stop.

15 Check **MTA** and click **Stop**.

16 Click **Save** to save your changes and return to the **Host Configuration** page.

17 Check the box next to the Scanner that you want to stop and click **Disable**.

The Scanner list updates to reflect your change.

To start a Scanner

1. In the Control Center, click **Administration > Hosts > Configuration.**
2. To enable a Scanner that is currently disabled, check the box next to the Scanner and click **Enable**.

   You can check multiple boxes.

   The Scanner list updates to reflect your change.

Managing services and MTA operations

Table 3-9 lists the various settings that you use to start or stop services, view the status of services, and configure MTA operations.

Table 3-9 Edit Host Configuration page—Services tab

Item	Description
Start/Stop	Allows you to start or stop one or more services by checking the box next to the service name under Scanner Services and clicking Start or Stop. Stopping the MTA service will stop the inbound, outbound, and delivery listeners. To stop one or more listeners separately or in sequence, stop the corresponding message queues on the **Status > SMTP > Message Queues** page. **Note:** If you use SMTP authentication, Symantec Brightmail Gateway employs an additional listener for authentication. This listener is controlled by the outbound listener controls. If you stop the outbound listener, the authentication listener also stops. See "MTA and message queue behavior" on page 613. **Note:** If you stop the Brightmail Engine or the MTA on a host configured to receive alerts, and wish to continue receiving alerts, specify an operating MTA IP address under **SMTP Host** on the **Administration > Settings > Control Center** page.

Table 3-9 Edit Host Configuration page—Services tab *(continued)*

Item	Description
Scanner Services	Lists the Scanner services available to the selected host. These are: ■ **Conduit** ■ **LiveUpdate** ■ **Brightmail Engine** ■ **MTA** ■ **IM Relay** ■ **Directory Data Service** Check the name of the service you want to start or stop for this host. Check **Scanner Services** to select all the services.
Status	Indicates whether a particular service is either **Running** (in black) or **Stopped** (in red). If the service crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours.
MTA Operation	Use the radio buttons to determine how the host handles messages. Choices are: ■ **Accept and deliver messages normally** – Processes messages in accordance with defined policies ■ **Pause message scanning and delivery** – Accepts inbound and outbound messages; holds messages in queues for future scanning and delivery. This option can be useful if you want to pause incoming messages while waiting for new virus definitions. ■ **Do not accept incoming messages** – Rejects incoming messages; scanning and delivery of messages in message queues continues. This option is useful when you need to drain queues in order to remove a host from use. When a message is rejected, the SMTP server is sent a `service not available` (450) error message. Once this option is selected, all previously received messages are processed, but no new messages are accepted. See "MTA and message queue behavior" on page 613.

Use the following procedures from the Services tab to manage individual Scanner services and MTA operations.

To start and stop services

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Click the name of the Scanner on which you want to stop or start a service.
3. Check the services to be started or stopped.
4. Click **Stop** to stop a running service or **Start** to start a stopped service.

To manage a Scanner's MTA operations

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Click the name of the Scanner that you want to change.
3. On the MTA Operation portion of the page, perform one of the following actions:

   To pause message delivery, click **Pause message scanning and delivery**. Inbound and outbound messages are placed in a queue for future scanning and delivery.

   To reject incoming messages, check **Do not accept incoming messages**. All messages currently in message queues are scanned and delivered, but all new messages are rejected.

   To restore normal operation, click **Accept and deliver messages normally**.

   See "MTA and message queue behavior" on page 613.
4. Click **Save** to store your changes.

Working with Services

Use the Services tab to configure a Scanner to perform any of the following tasks:

- Enable or disable the following services on a Scanner using the Services tab on the Edit Host Configuration page. Each host runs several services that it uses to communicate with the Internet and other servers on your network.

Service	Description
Conduit	Retrieves new and updated email filters from Symantec Security Response through secure HTTPS file transfer. Conduit authenticates filters and alerts the Brightmail Engine to their presence. Conduit manages statistics for use by Symantec Security Response and by Symantec Brightmail Gateway to compile reports.

Service	Description
LiveUpdate	Automatically downloads virus definitions from Symantec Security Response to the Scanner. This information is used by the Scanner's Brightmail Engine to identify known security threats.
Brightmail Engine	Scans email and attachments, instant messages, and file transfers for viruses, spam, and content filtering according to filter polices that you have configured.
MTA	The mail transfer agent routes inbound and outbound messages to the Brightmail Engine for processing and delivers filtered messages to their internal destinations or to the Internet.
IM Relay	Retrieves new and updated virus definitions and Spim filters from Symantec Security Response. Uploads suspected Spim to Symantec Security Response and subsequently downloads filters to other Symantec Brightmail Gateway appliances that are configured to detect heuristic-based Spim.
Directory Data Service	The directory data service lets you use the information that is stored in your Lightweight Directory Access Protocol (LDAP) directories for features in the Symantec Brightmail Gateway.

Note: If you stop the Brightmail Engine or the MTA service on a host configured to receive alerts, you must specify another host to continue receiving alerts. To avoid an interruption in alerting, modify the SMTP Host and Port fields on the Control Center Settings page (Administration > Settings > Control Center) before stopping either of these services.

- Enable, disable, or pause incoming message scanning
  Enabling a Scanner to accept and deliver messages normally is the default behavior. However, if you have to take a Scanner offline, you can limit MTA operations in stages while you assign them to other Scanners.
  See “Managing services and MTA operations” on page 110.

Configuring email settings

This chapter includes the following topics:

About email domains

When inbound mail arrives at a Scanner, Symantec Brightmail Gateway verifies that the message is addressed to a valid local domain before accepting it to the appropriate message queue.

You must define any domain for which you want Symantec Brightmail Gateway to accept inbound email as a local domain. Symantec Brightmail Gateway only accepts inbound email that is addressed to local domains. However, a domain from which you send email can be a local domain or a non-local domain.

Defining more than one domain lets you assign different routing and delivery options to each domain, including TLS encryption for secure delivery, for local or non-local domains. By configuring TLS encryption options for non-local domains, you can secure connections for delivery by external servers.

You can also enable DKIM signing for a local or non-local domain, so that each message sent from that domain will include a DKIM signature that enables DKIM authentication by the receiving MTA.

You can also add non-local domains. You specify non-local domains primarily to route outbound email over established connections to external servers for non-local delivery. You can also define delivery options for non-local domains.

Several features of Symantec Brightmail Gateway make use of domain designations.

You can add or edit domains to:

- Create different email acceptance settings for each domain.
  See "About email domain acceptance settings" on page 122.
- Define different delivery options for each domain.
  You can route inbound email addressed to different local domains using an unlimited number of default relays.
  You can designate an directory data routing source for delivery of email that is addressed to specified non-local domains.
  MX Lookup and TLS encryptions options are available for delivery of email that is addressed to non-local domains.
  See "Adding or editing domains" on page 117.
- Configure Recipient Validation to validate recipients of incoming email against your LDAP directory.
  See "About invalid recipients" on page 129.

  Note: If you intend to set up probe accounts from invalid recipient email addresses you must check Enable Recipient Validation for the domains you plan to use for probe accounts.

  See "Creating probe accounts from invalid recipient email addresses" on page 253.

- Configure static routing, to route inbound or outbound email to specified mail servers for internal delivery.
  See "Adding or editing domains" on page 117.
- Configure domain-specific delivery bindings for non-local messages. This lets you define one or more IP addresses that messages from a domain are sent from.
  See "SMTP advanced settings for delivery bindings" on page 102.

You can import lists of domains that you then edit individually.

See "Importing a domains list" on page 120.

Typically, a domain is the part of the recipient's email address that follows the @ sign. For example, anywhere.com is the domain for someone@anywhere.com. Domains can include subdomains. For example, somewhere.anywhere.com is a subdomain of anywhere.com. Alternatively, you can specify a single email address as a domain.

If you want to include all subdomains with a domain, enter a period before the domain. For example, if you want to include all subdomains in example.com, enter **.example.com**. However, entering a period before the domain omits the domain itself. For example, to accept email that is addressed to example.com and all subdomains of example.com, you must specify both **example.com** and **.example.com**.

If you want to include only certain subdomains, you must specify each subdomain separately. For example, you must specify both **elsewhere.anywhere.com** and **somewhere.anywhere.com** as separate domains to accept email that is addressed to either subdomain but not **overthere.anywhere.com**.

Note: A domain can be a fully qualified domain name (FQDN), subdomain, or RFC5321-compliant email address. These levels of granularity allow you maximum control over what addresses are acceptable and how email that is addressed to them are routed.

Adding or editing domains

Local domains are domains and email addresses for which Symantec Brightmail Gateway accepts inbound messages. When adding or editing a local domain, you can assign routing behaviors based on MX Lookup, designated hosts, or directory data sources, and enable or disable recipient validation for messages accepted from the domain. You can also import lists of local domains.

Note: If you have provisioned content encryption for your messaging system and add a domain or change the address of a domain, you must inform your Symantec provisioning representative. For more information, see the Symantec Content Encryption Provisioning page by clicking **Content > Settings > Content Encryption** and then clicking the **Click here**. You might also need to replace your certificate or certificate bundle to include the domain.

See "About encrypting messages with Symantec Content Encryption" on page 458.

Configuring separate local and non-local domains provides you with different routing and delivery options. For example:

- You can route inbound email addressed to different local domains using an unlimited number of default relays.
- You can route outbound mail sent to non-local domains using a non-local default relay for delivery by an external mail host at a subsidiary or business partner.

- You can route outbound email that is addressed to a specified non-local domain to an internal mail server that you reserve for confidential communications.

Non-local domains are typically used to route outbound mail to external mail hosts with which you regularly exchange email. You can also route outbound email that is addressed to a non-local domain to an internal mail host. Because email addressed to non-local domains is first scanned for policy violations, adding non-local domains provides you with options to securely exchange mail with business partners or parts of your own organization whose networks reside behind separate firewalls.

Configure any non-local domains that you add to:

- Statically route outbound email addressed to non-local domains to any one of up to three mail hosts for non-local delivery

  You can optionally enable MX Lookup on any non-local domains for which you define a static route. Symantec recommends that you enable MX Lookup for delivery to multiple mail hosts.

- Enable TLS encryption options for specified domains

Note: You can enable routing for a domain once it has been imported by editing the domain name in the Domains list.

You can also query a routing data source to route outbound email addressed to recipients at remote domains.

During site setup, you specify one or more local domains for which the appliance accepts messages. Any domain that you designate during site setup is by default a local domain. Symantec Brightmail Gateway only accepts connections from a sender IP when an email is addressed to a local domain. After processing, the Scanner relays any email that does not violate policy conditions or setting limitations to the mail server that hosts the local domain for delivery to recipients.

A mail host that serves as a default relay can host more than one local domain. If only one mail host is specified for local domains, it acts as the default relay for delivery of all inbound mail. You can designate an unlimited number of additional destination servers as default relays for inbound email addressed to local domains.

On the **Protocols > SMTP > Domains** page you can:

- Add more local domains for which you want Symantec Brightmail Gateway to accept inbound email.
- Add probe account domains.

  If you intend to create probe accounts from unused or invalid recipient email addresses, the probe account domain must be added to your list of local domains.

If you intend to set up probe accounts from invalid recipient email addresses, you must also enable Recipient Validation for that domain.

See "Setting up probe accounts" on page 251.

- Limit inbound mail that is addressed to local domains to valid recipients

  Symantec Brightmail Gateway checks recipients of email addressed to specified local domains against LDAP directory data before allowing the sender IP to connect to the scanner to inbound email interface. You must configure an appropriate data source, enable the recipient validation function, and enable Invalid Recipient Handling to reject or drop inbound email that is addressed to invalid recipients.

- Define different delivery options for separate domains.

To add or edit a domain

1. In the Control Center, click **Protocols > SMTP > Domains**.
2. On the Domains page, click **Add** or click the name of a domain whose settings you want to edit.
3. In **Domain or email address for which to accept inbound mail**, enter a local domain, subdomain, or email address or edit the domain name.

   Placing a period in front of the domain enables Symantec Brightmail Gateway to accept all subdomains.
4. If you are adding or editing a non-local domain, deselect the **Local domain** checkbox.
5. If you wish to reject or drop invalid recipients, or if you are a probe account participant, check **Enable Recipient Validation for this domain**.

   You cannot enable recipient validation on an email address or a non-local domain.

   Further configurations are required if you want to accept or reject invalid recipient email.

   See "Setting up invalid recipient handling" on page 130.
6. To define delivery options for this domain, click the **Delivery** tab.
7. If you wish to have messages addressed to this domain routed via MX lookup, to specific destination servers via host-based routing, or based on directory data, check **Optionally route this domain or email address**.

   If you do not check this option, messages are routed to the default relay that you configure under Inbound Local Mail Delivery on the SMTP tab of the Edit Host Configuration page.
8. Choose one of the following options:

- Click **MX Lookup** to route messages to the host indicated by the recipient address using MX lookup.

  If you specified a fully qualified domain name or an email address, MX lookup is performed on the specified domain for messages to recipients in that domain. If you specified a subdomain parent, MX lookup is performed on the fully qualified domain name of the recipient address, for messages to recipients in any of the subdomains of the specified parent domain.

- Click **Destination hosts** to specify one or several specific hosts to route messages to, then click **Add** to add each host. Optionally, type a port to which messages addressed to this domain are routed. You can check **MX Lookup** to enable MX lookup for the destination host. You cannot use MX lookup if you specified the host using an IP address. Type a number between 1 and 100, inclusive, under Preference, to indicate the order in which hosts should be used. Lower-numbered hosts are used first. To perform load balancing, type the same number for each host.

  See "About Scanner email settings " on page 86.

- Click **Directory Data Source** to route messages using directory-based routing for delivery resolution. Choose the LDAP routing source from the accompanying drop-down list.

  You must have an LDAP directory data source defined for routing to use this option.

  See "Creating a data source" on page 491.

  See "Creating a routing data source" on page 516.

9 To deliver email to a non-local host using TLS encryption, click **Optional delivery encryption**, then click the radio button that best describes the basis for TLS encryption for this domain.

10 To enable DKIM signing for this domain, use the Domain Key Identified Mail section.

See "Enabling DKIM signing for a domain" on page 143.

11 Click **Save** to add the domain, subdomain, or email address to the list or to confirm your edits.

Importing a domains list

Lists of domain definitions and email addresses can be imported from an ASCII file, similar to the Sendmail `mailertable`. You can include optional routing information to default local destination hosts can be included as part of the definition.

After successfully importing a list of domains, each domain appears in the list on the **Domains** page. You can then edit any domain and, if desired, uncheck **Local domain** to use that domain as a non-local domain instead

Note: If you provision content encryption for your messaging system and later add a domain, you must inform your Symantec provisioning representative. For more information, see the Symantec Content Encryption Provisioning page by clicking **Content > Settings > Content Encryption** and then clicking the **Click here**. You may also need to update your certificate to cover the new domains.

See "About encrypting messages with Symantec Content Encryption" on page 458.

Note: You can enable directory-based routing for a domain once it has been imported by editing the domain name in the Domains list.

See "Adding or editing domains" on page 117.

You can also enable recipient validation for imported domains. To do this you must also enable Recipient Validation Handling and configure a data source for recipient validation.

See "Setting up invalid recipient handling" on page 130.

See "Creating a recipient validation data source" on page 512.

In the import file, place each domain definition on a line by itself. The domain definition can consist of the following tab-delimited attributes:

Domain name	Can be either a complete domain name, a subdomain name, or an email address. To include all subdomains within a domain, add a period at the beginning of the domain name.
Destination	Consists of destination type and destination host name. Only definitions with a destination type (Mailer) of SMTP or ESMTP are supported, and %backreferences are not supported. After import, ESMTP destination types convert to SMTP. When the host name is enclosed in brackets—smtp:[destination.domain.com]—MX lookup is not performed for the destination host.
Validation	Indicates whether or not recipient validation is enabled for the domain. The value VALIDATE_RCPTS in this column enables recipient validation. Any other value or no value will not enable recipient validation. Recipient validation can only be enabled for domains. It cannot be enabled for email addresses.

Here is a sample import file:

```
local1@domain.com  smtp:local1.com
local2@domain.com  smtp:local2.com:20
local3@domain.com  smtp:[local3.com]:30
local4@domain.com  smtp:[local4.com]
.local5.com        smtp:[192.168.248.105]
local6.com         smtp:[192.168.248.106]:60 VALIDATE_RCPTS
```

To import a list of local domains

1. In the Control Center, click **Protocols > SMTP > Domains**.
2. Click **Import**.
3. In the **Specify the import file** text box, enter the filename path or browse to the file containing the list of domain definitions.
4. Click **Import**.

   If entries in the import file do not match the required file format, an error message with a link appears. Click on the link to download a file containing the unprocessed entries.

Deleting domains

You can delete domains to which you no longer want email sent. Deleting a local domain means that the MTA no longer accepts inbound email messages that are addressed to that domain. Deleting a non-local domain means that any routing or TLS options that are configured for that domain no longer operate.

Note: If you delete a local domain, any probe accounts associated with that domain will no longer be valid. Check the active probe accounts before deleting the domain.

See "Adding or editing domains" on page 117.

See "Disabling a probe account" on page 256.

To delete one or more domains

1. In the Control Center, click **Protocols > SMTP > Domains**.
2. Select one or more domains in the Domains list or click **Delete All** to delete all listed domains.
3. Click **Delete** to delete only the domains selected.

About email domain acceptance settings

During site setup, you designate at least one local domain for which Symantec Brightmail Gateway accepts inbound email. You also define a single mail host to which the inbound scanner routes email addressed to local domains. Any domain that is added during site setup is by default a local domain and uses this static route as its default relay to deliver inbound email.

After your site setup is complete, you can further configure inbound email domain acceptance to:

- Reject or drop email addressed invalid recipients at specified local domains.
  See "About invalid recipients" on page 129.
- Detect directory harvest attacks directed against a domain.
  See "Configuring directory harvest attack recognition" on page 172.

You must enable recipient validation for any domain for which you want Symantec Brightmail Gateway to validate recipients or detect directory harvest attacks.

You configure a domain to validate recipient addresses within a local domain by configuring an LDAP recipient validation profile and enabling recipient validation on a per-domain basis. If recipient validation is enabled for a local domain, Symantec Brightmail Gateway checks with the data source to determine that the email address exists in the LDAP directory. If there is no match, any attempted connection is rejected or dropped according to your site's settings for handling invalid recipients.

See "Setting up invalid recipient handling" on page 130.

Symantec Brightmail Gateway accepts all email from internal mail hosts.

See "Changing Scanner outbound mail acceptance settings" on page 93.

About aliases and address masquerades

An alias translates an email address into one or more destination addresses. Windows users may understand this concept as a "distribution list." You can add an alias as a convenient shortcut for typing a long list of recipients. An alias can also translate addresses from one top-level domain to another. For example, you can create an alias source domain to translate example.com to a target domain, example-internetsecurity.com. Symantec Brightmail Gateway translates email addressed to someone@example.com delivered to someone@example-internetsecurity.com.

Address masquerading is a method of concealing email addresses or domain names behind the mail gateway by assigning replacement values to them. Symantec Brightmail Gateway lets you implement address masquerading on inbound mail, outbound mail, or both. A typical use of address masquerading is to hide the names of internal mail hosts so that outgoing mail appears to be coming from a different domain than that of the actual host.

Outbound address masquerades change the apparent sender of a message. Inbound address masquerades change the apparent recipient of a message.

Alias translation only applies to inbound or internal messages processed by Symantec Brightmail Gateway. Once the gateway allows email to connect to the inbound scanner, it determines whether the address in the SMTP envelope `To:` field is an alias and translates it into any destination addresses during the connection session. Transformed addresses are written back to the SMTP envelope `To:` before the scanner filters the message. The contents of the message `To:` and `Cc:` headers are ignored and not changed.

Note: If you employ probe accounts along with aliases and address masquerades you may lose one or the other account. Follow the guidelines for creating probe accounts.

See "About creating probe accounts" on page 250.

Inbound address masquerading takes precedence over the alias translation. If the same original email address or domain exists in both the address masquerading list and the aliases list, messages to the source domain are routed to the masqueraded address or domain. Symantec Brightmail Gateway does not route the message to the alias address or domain.

Except where address masquerading applies to a recipient address, you must add the source domain of an alias to the list of local domains for which Symantec Brightmail Gateway accepts inbound email.

See "Adding or editing domains" on page 117.

Alias translation does not apply to outbound messages routed to the Internet.

See "Alias addresses" on page 126.

See "Adding or editing aliases" on page 124.

See "Importing aliases" on page 125.

See "Adding or editing address masquerades " on page 127.

See "Importing an address masquerade list" on page 128.

Adding or editing aliases

An alias translates an email address into one or more destination addresses. Specify one and only one source email address or domain for each alias that you enter. For each destination address, you can enter a single email address or multiple email addresses separated by commas, semicolons, or spaces.

See "About aliases and address masquerades" on page 123.

See “Alias addresses” on page 126.

To add an alias

1. In the Control Center, click **Protocols > SMTP > Aliases**.
2. Click **Add**.
3. On the Add Aliases page, type the alias in the **Alias domain or email address** box:
4. Type a domain or one or more destination email addresses in the **Domain or email addresses for this alias** box:

   You can specify multiple email addresses, or a single domain.
5. Click **Save**.

To edit an alias

1. In the Control Center, click **Protocols > SMTP > Aliases**.
2. Click an alias.
3. In the Edit Aliases page, modify the text in the **Alias domain or email address** box as desired.
4. Modify the text in the **Domain or email addresses for this alias** box as desired.
5. Click **Save**.

Importing aliases

Aliases can be imported from a text file. Each address in the text file must be separated with one or more spaces or tabs, or a combination of spaces and tabs. Commas or semicolons are not valid delimiters. In the import file, each line must contain an alias address followed by one or more destination addresses.

Following is a sample import file:

```
oak@example.com quercus@symantec-internetsecurity.com
ops@example.com tla@example.com bmi@example.com
blockads.com noadsorspam.com
```

To import aliases

1. In the Control Center, click **Protocols > SMTP > Aliases**.
2. Click **Import**.

3 In the **Specify the import file** text box, enter or browse to the filename containing the list of aliases.

4 Click **Import**.

If entries in the import file are not specified correctly, do not match the required file format or are duplicates, an error message is displayed. You can click a link to download a file containing the unprocessed entries.

See “About aliases and address masquerades” on page 123.

See “Alias addresses” on page 126.

Click **Cancel** to ignore the unprocessed entries and return to the main Aliases page to review the valid imported entries.

Alias addresses

You can enter multiple destination email addresses for each alias that you enter as a source email address. You can only enter a single destination domain, however, for each alias that you enter as a source domain. Alias source addresses must have local domains. Destination domains can be local or non-local domains.

Alias source example	Destination examples
anyone@example.com	somebody@example.com
anybody@example.com	someone@elsewhere.com
help@example.com	anybody@example.com, anyone@example.com, someone@example.com, somebody@example.com
example.com	elsewhere.com

Note: Except where address masquerading applies to a recipient address, you must add the source domain of an alias to the list of local domains for which Symantec Brightmail Gateway accepts inbound email.

See “Adding or editing domains” on page 117.

Aliases are recursive. This means that an alias specified in the destination email address list is expanded as defined in the list of aliases up to 1000 addresses.

In the example shown below, a message addressed to it@example.com would be delivered to the destination addresses for both it@example.com and ops@example.com, because it@example.com includes ops@example.com.

Alias	Destination addresses
it@example.com	alro@example.com, oak@example.com, ops@example.com
ops@example.com	tla@example.com, bmi@example.com, map@example.com

You cannot add aliases that duplicate aliases already stored. You cannot create an email address alias for a domain. You cannot create a domain alias for an email address. You can only create an email address alias for one or more other email addresses. You can only create a domain alias for one or more other domains.

Both of the aliases shown below are invalid.

Alias	Destination addresses
it@example.com	alro@example.com, symantec.com
example.com	sample@symantec.com

Adding or editing address masquerades

Address masquerading is a method of concealing email addresses or domain names behind the mail gateway by assigning replacement values to them.

See "About aliases and address masquerades" on page 123.

Warning: If you masquerade a probe address or domain, the probe account will become invalid. Make sure there are no probe accounts associated with an account you masquerade. See "About creating probe accounts" on page 250.

Follow these steps to add or edit masqueraded entries.

To add a masqueraded entry

1. In the Control Center, click **Protocols > SMTP > Address Masquerading**.
2. Click **Add**.
3. In the **Original address or domain to masquerade** box, specify an address or domain to masquerade.
4. In the **New address or domain** box, specify a new name for the address or domain name.
5. From the **Apply to** drop-down list, specify the messages to which this masqueraded name applies: **Outbound messages**, **Inbound messages**, or **Inbound and outbound messages**.
6. Click **Save**.

To edit a masqueraded entry

1. In the Control Center, click **Protocols > SMTP > Address Masquerading**.
2. Click the masqueraded address or domain that you want to modify.
3. On the Edit Masqueraded Entry page, modify the masqueraded entry as desired.
4. Click **Save**.

Importing an address masquerade list

In addition to creating new masqueraded entries, you can import them from a text file similar to the Sendmail `virtusertable`. In the import file, place each masqueraded address definition on a line by itself. Each address in the file must be separated with one or more spaces or tabs, or a combination of spaces and tabs. Commas or semicolons are not valid delimiters.

Note: You cannot import a file with extended ASCII or non-ASCII characters; you can only import files encoded in US-ASCII format.

The masquerade address definition consists of the following elements:

Original entry
Specifies the original email address or domain name to be masqueraded.

Replacement entry
Specifies the replacement email address or domain name.

Apply to
Indicates the direction to which masquerading is applied. Available choices are:

- Inbound messages
- Outbound messages
- Inbound and outbound messages

Following is a sample import file:

```
orig1@domain.com  new1@domain.com  inbound
orig2@domain.com  new2@domain.com  outbound
orig3@domain.com  new3@domain.com  inbound/outbound
orig4@domain.com  new4.com         inbound
orig5@domain.com  new5.com         outbound
orig6@domain.com  new6.com         inbound/outbound
orig7.com         new7@domain.com  inbound
```

```
orig8.com           new8@domain.com   outbound
orig9.com           new9@domain.com   inbound/outbound
```

To import a list of masqueraded entries

1. In the Control Center, click **Protocols > SMTP > Address Masquerading**.
2. Click **Import**.
3. In the **Specify the import file** text box, enter or browse to the filename containing the list of masqueraded entries.
4. Click **Import**.

   If entries in the import file are not specified correctly, do not match the required file format, or are duplicates, an error message is displayed. You can click a link to download a file containing the unprocessed entries. Click **Cancel** to return to the Address Masquerading page to review the valid imported entries.

About invalid recipients

By default, when an email message arrives addressed to your domain that is not addressed to a valid user, Symantec Brightmail Gateway passes the message to the internal mail server. The internal mail server may either accept the message and generate a bounce message, or the internal mail server may reject the message. Upon receiving the bounce message, a legitimate sender can resend the original message with the correct address. However, messages with invalid recipients can also result from a spammer's directory harvest attack.

Note: The **Remove unresolved recipients** action on the Directory Harvest Attacks page only removes unresolved recipients when a directory harvest attack occurs. You can combine this action with your invalid recipient handling setting or enable the two settings individually.

See “Configuring directory harvest attack recognition” on page 172.

You can configure Symantec Brightmail Gateway to accept, reject, or drop any messages that are sent to invalid recipients, as follows:

- If you choose to accept all recipients, Symantec Brightmail Gateway accepts all messages, whether or not the recipients are valid. However, if the internal mail server rejects a recipient, Symantec Brightmail Gateway sends a bounce message. The internal mail server may also send bounce messages if it is configured to send them.

- If you choose to reject invalid recipients, Symantec Brightmail Gateway rejects any messages that are addressed to email addresses that do not exist in your LDAP directory. The sending MTA may generate a bounce message to the sender. You must have a data source configured for recipient validation. Recipients are rejected at the initial SMTP conversation with a 5xx SMTP error.

  See "About using the recipient validation function with your data source" on page 488.
- If you choose to drop invalid recipients, Symantec Brightmail Gateway drops from the mail stream any messages that are addressed to email addresses that do not exist in your LDAP directory. No bounce messages are returned to the sender.You must have a data source configured for recipient validation.
- If you choose to reject or drop invalid recipients, Symantec Brightmail Gateway applies your choice to each local domain that you configure to enable recipient validation. If you do not enable recipient validation for any local domains, no messages are dropped or rejected.
  See "Adding or editing domains" on page 117.

Warning: Dropping messages for invalid recipients is an extreme measure. Enabling this feature may prevent diagnosis of serious problems with your email configuration. Only enable this feature after you are sure that your email system is stable. Also, if enabled, accidentally mis-addressed messages are dropped, and no bounce messages are sent. You can instead reject invalid recipients, which allows the sending MTA to generate a bounce message if so configured.

See "To set up invalid recipient handling" on page 131.

Setting up invalid recipient handling

You can set up invalid recipient handling to accept all recipients, or to drop or reject invalid recipients. If you choose to drop or reject invalid recipients, you then must enable invalid recipient validation on a per-domain basis.

You must configure a directory data source for recipient validation before choosing to reject or drop invalid recipients.

See "Adding a data source " on page 492.

The **Remove unresolved recipients** action on the Directory Harvest Attacks page only removes unresolved recipients when a directory harvest attack occurs. You can combine this action with your invalid recipient handling setting or enable the two settings individually.

Warning: Dropping messages for invalid recipients is an extreme measure. Enabling this feature may prevent diagnosis of serious problems with your email configuration. Only enable it after you are sure that your email system is stable. Also, if enabled, accidentally mis-addressed messages are dropped, and no bounce message are sent. You can instead reject invalid recipients, which allows the sending MTA to generate a bounce message if so configured.

To set up invalid recipient handling

1. In the Control Center, click **Protocols > SMTP > Invalid Recipients**.
2. Do one of the following:
   - Click **Accept all recipients** to accept all messages, whether or not the recipients are valid. Bounce messages will be sent if your internal mail server is configured to send bounce messages or to reject invalid recipients.
   - Click **Reject invalid recipients** to reject any messages that are addressed to user names that do not exist in your LDAP directory. The sending MTA may generate a bounce message to the sender if configured to do so. You must configure a directory data source for recipient validation to reject invalid recipients. Recipients are rejected at the initial SMTP conversation with a 5xx SMTP error.
   - Click **Drop invalid recipients** to drop any messages that are addressed to user names that do not exist in your LDAP directory from the mail stream. No bounce messages are returned to the sender. You must have an LDAP source configured for recipient validation to use this setting.

     This setting is independent of action you specify using the Directory Harvest Attacks page, but it can be used with that action.
3. Click **Save**.

Note: If you chose to **Reject** or **Drop** invalid recipient mail, you must enable recipient validation for that domain.

See “Adding or editing domains” on page 117.

About email message flow

Understanding exactly what happens to an email message during processing can help you to configure your system optimally and troubleshoot any problems that arise.

Symantec Brightmail Gateway lets you manage two messaging protocols: SMTP and Instant Messaging. This section provides an overview of the message flow for SMTP (Simple Mail Transfer Protocol). Although you can filter instant messages on the same host as email, the two protocols operate through separate IP addresses and ports.

See “About IM” on page 289.

Email policies and SMTP settings can apply to both the inbound and outbound message flow. Some policies and settings, such as those covered by Brightmail Adaptive Reputation Management, address issues unique to inbound message flow. Content filtering policies, on the other hand, address the data loss prevention and regulatory compliance issues that most often affect outbound message flow.

Warning: Symantec Brightmail Gateway allows you to configure various MTA configuration parameters to manage your email message flow. If you relay messages to other MTAs, some settings for these MTAs may conflict with Symantec Brightmail Gateway settings. For example, you configure maximum message size for 10 MB, and your local relay MTA has a maximum of 1 MB. Such conflicts can result in errors that are difficult to diagnose.

See “About blocking and allowing messages at connection time” on page 163.

See “About content filtering” on page 331.

See “Email message flow phases” on page 132.

Note: You can vary the number of times Symantec Brightmail Gateway scans a potentially malformed message.

See “Configuring bad message handling” on page 701.

Email message flow phases

Symantec Brightmail Gateway processes an email message in the following phases:

- Phase 1, SMTP connection – During the connection phase Symantec Brightmail Gateway accepts connections from legitimate senders, unless a specific policy or setting requires rejection or deferral. All of the actions that are taken during this phase are based on the IP address alone. Symantec Brightmail Gateway rejects connection attempts from any IP addresses that are not permitted by a Scanner's SMTP settings. For inbound connections from IP addresses that appear in IP-based bad sender groups, Symantec Brightmail Gateway rejects the SMTP connection, unless a different action is configured for the group. Based on Connection Classification limits, some connections are deferred.

See “Changing Scanner inbound mail settings” on page 89.

See “About managing connection load at the gateway” on page 166.

- Phase 2, SMTP session – During the SMTP session, Symantec Brightmail Gateway accepts, rejects, or defers messages on the basis of the message envelope. It also checks Connection Classification settings and SMTP settings to determine whether accepting the message exceeds the configured limits.

  See “Changing Scanner inbound mail acceptance settings” on page 90.
- Phase 3, Message filtering – After the Scanner accepts a message, Symantec Brightmail Gateway evaluates the message content and renders verdicts. It evaluates the message content based on applicable policies and settings, including, if required, content filtering policies. Based on the verdicts, it applies the configured actions.

  See “About filtering” on page 717.
- Phase 4, Message routing – Symantec Brightmail Gateway routes messages that have not been quarantined or held for review to a mail host. The route is determined by Domains Delivery settings and the Scanner's Local Mail Delivery and Non-local Mail Delivery settings.

  See “About email domains” on page 115.
- Phase 5, Message delivery – Symantec Brightmail Gateway enforces limits on the number of connections to internal mail servers. If the address binding settings are configured to specify the IP addresses that deliver email, Symantec Brightmail Gateway uses the specified IP addresses.

  See “Configuring Scanner inbound email delivery settings” on page 91.

Figure 4-1 illustrates the inbound message flow.

Figure 4-1 Inbound Message Flow

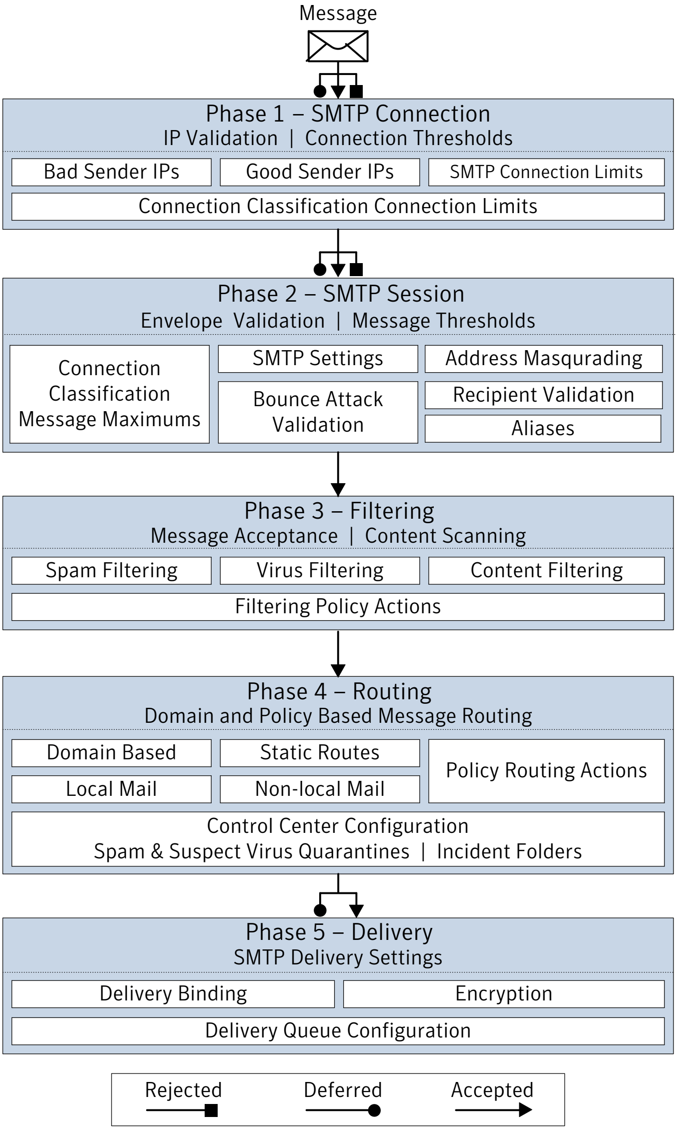

Chapter 5

Setting up email authentication

This chapter includes the following topics:

Enabling SPF and Sender ID authentication

Symantec Brightmail Gateway can authenticate a sender's IP address by checking it against the published DNS record for the named mail server. If the DNS record includes a hard outbound email policy (one that requires content filtering), and it does not include the sending IP address, Symantec Brightmail Gateway processes the inbound message according to the action that you specify on the Sender Authentication page. If the sender's IP address matches the IP address that is published in DNS record, or if the domain publishes only an informational policy or does not publish a policy at all, no action is taken.

Authenticating the IP addresses of senders can reduce spam because spammers often attempt to forge the mail server name to evade detection. Symantec Brightmail Gateway uses the Sender Policy Framework (SPF) or the Sender ID standard to authenticate sender IP addresses. If you specify domains whose IP addresses you want Symantec Brightmail Gateway to authenticate, the best practice is to specify the highest-level domain possible, such as example.com, because tests for compliance include all subdomains of the specified domain—for example, my.example.com and your.example.com.

Warning: Authenticating all domains can significantly increase processing load. Many domains do not publish an outbound email policy, or they publish only an informational policy. Attempting to authenticate the IP addresses belonging to such domains will not produce any action on mail sent from them and can unnecessarily expend processing resources, at times excessively. Authentication is most effective for domains that publish hard policies that are frequently spoofed in phishing attacks.

To enable SPF and Sender ID authentication

1. In the Control Center, click **Spam > Settings > Sender Authentication**.
2. Check **Enable Sender Authentication**.
3. Under **Authentication Types**, check **Sender Policy Framework (SPF)** or **Sender ID**.

   Choosing Sender ID also enables SPF because when you authenticate Sender ID with DNS, it also provides SPF authentication.
4. Under **Domain Authentication**, choose a domain authentication method.

   To initiate sender authentication on incoming messages from all domains, click **Authenticate all domains** and click **Save.**

   To select specific domains to authenticate, click **Authenticate only the following domains** and check the domains to authenticate.
5. Perform additional actions as needed.
   - To add a new domain to the list click **Add**. Type a domain name in the text field and click **Save**.
   - To edit the spelling of a domain click the domain name and click **Edit**. Make changes and click **Save**.
   - To delete a domain from the list, check the domain name and click **Delete**.
   - To change the default action, or to add additional actions, choose from the drop-down menu. Some action choices display additional fields where you can provide specifics for the action. By default, each failed message has the phrase [**sender auth failure**] prepended to its subject line.
     See " Verdicts and actions for email messages" on page 718.
6. Click **Save** to commit your changes.

Configuring DKIM authentication

Domain Key Identified Mail (DKIM) is a protocol that uses public-key cryptography to allow the sending MTA to electronically sign legitimate email messages in a way that can be verified by recipient MTAs. Symantec Brightmail Gateway can perform DKIM signing on outbound messages. This enables your recipients to identify messages as validly originating from you, and also to detect whether messages were modified after leaving your MTA. Symantec Brightmail Gateway can also perform DKIM validation on inbound messages, to verify the authenticity of a DKIM signature and detect whether a message has been modified.

You implement DKIM signing on a per-domain basis. Symantec Brightmail Gateway can add only one DKIM signature to an outbound message.

After enabling DKIM validation for all inbound messages, you can create a content filtering policy to choose the action that Symantec Brightmail Gateway takes when an inbound message from a specific domain or group of domains fails DKIM validation. Symantec Brightmail Gateway does not grant any type of enhanced processing to messages that pass DKIM validation.

Note: Enabling DKIM validation may have an adverse affect on mail processing performance.

To configure DKIM signing for outbound mail that is sent from a specific domain, follow the steps in Table 5-1.

To configure DKIM validation for inbound mail, follow the steps in Table 5-2.

Table 5-1 Configuring DKIM signing

Step	Action	Description
Step 1	Create a domain key	When you add a domain key, Symantec Brightmail Gateway generates an RSA key pair. If you prefer, you can generate your own RSA key and import that key. See "Adding a domain key" on page 138. See "Importing a domain key" on page 140.
Step 2	Enable DKIM signing for a specific domain	You can choose the domains for which you want to use DKIM signing. See "Enabling DKIM signing for a domain" on page 143.

Table 5-1 Configuring DKIM signing *(continued)*

Step	Action	Description
Step 3	Generate a DKIM text record from the domain key	When you enable DKIM signing for a domain, you can also generate a DKIM text record. See "Enabling DKIM signing for a domain" on page 143.
Step 4	Store the domain key in DNS	You must update your DNS record so that it can make use of the domain key. See "Enabling DKIM signing for a domain" on page 143.

Table 5-2 Configuring DKIM validation

Step	Action	Description
Step 1	Enable sender authentication and DKIM validation	You enable DKIM validation for incoming mail on a system-wide basis. See "Enabling DKIM validation" on page 140.
Step 2	Create a content filtering policy	Apply the content filtering policy to all inbound messages that contain a DKIM=fail message header. Optionally, you can limit the policy to specific recipient or sender domains, or in other ways. Choose an action to take on these messages. See "Creating a content filtering policy for DKIM validation" on page 141.

You may also want to perform the following additional DKIM-related tasks:

See "Deleting a domain key" on page 139.

See "Changing a self-signed certificate or domain key name" on page 195.

See "Viewing a domain key" on page 139.

Adding a domain key

You can add a domain key to enable Scanners to perform DKIM signing on outbound messages. When you add a domain key, Symantec Brightmail Gateway generates an RSA key pair that includes a public key and a private key. The private key is used for creating a signature that accompanies a message. The public key is used for reading the signature to validate the message.

If you want to generate your own RSA key using OpenSSL or another program, you can then import the key you generated.

See "Importing a domain key" on page 140.

To add a domain key

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **Domain Keys** tab.
3. Click **Add**.
4. In the **Domain key name** field, type a unique name for this domain key.
5. In the **Key length** drop-down list, choose a length, in bits, for the RSA key.

   The default key length is 1024 bits.
6. Click **Create**.

Deleting a domain key

You can view or delete a domain key. You cannot delete a domain key that is in use for DKIM signing on outbound messages.

See "Viewing a domain key" on page 139.

To delete a domain key

1. Click **Administration > Settings > Certificates**.
2. Click the **Domain Keys** tab.
3. Check the box next to the domain key that you want to delete.
4. Click **Delete**.

Viewing a domain key

You can view any of your domain keys. Each domain key consists of an RSA key pair that includes a public key and a private key. The private key is used for creating a signature that accompanies a message. The public key is used for reading the signature to validate the message. You can only view the public key.

See "Viewing existing CA certificates" on page 198.

To view a domain key

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **Domain Keys** tab.

3 Check the box next to the domain key you want to view.

4 Click **View**.

Importing a domain key

You can import an existing domain key. If you do not have a domain key, you must first generate your own private RSA key, using OpenSSL or another program. An RSA key pair includes a public key and a private key. The private key is used for creating a signature that accompanies a message. The public key is used for reading the signature to validate the message. When you import a domain key, Symantec Brightmail Gateway generates a public key based on the private key you imported.

To import a domain key, you need to have the domain key stored in a text file that you can access from your Control Center.

Instead of importing a domain key, you can add a domain key. When you add a domain key, Symantec Brightmail Gateway generates both the private and public RDA keys for you.

See "Adding a domain key" on page 138.

Note: The domain key you import must be in PEM format. If the domain key is not in PEM format, or is not acceptable to OpenSSL, Symantec Brightmail Gateway will attempt to convert the domain key to correct the issue.

See "PEM format requirements for certificates and domain keys" on page 196.

To import a domain key

1 In the Control Center, click **Administration > Settings > Certificates**.

2 Click the **Domain Keys** tab.

3 Click **Import**.

4 Next to the **File name** field, click **Browse** and select a text file containing the domain key you want to add.

5 In the **Domain key name** field, type a unique name for this domain key.

6 Click **Import**.

Enabling DKIM validation

You can enable DKIM validation on all inbound messages.

Note: Enabling DKIM validation may have an adverse impact on mail processing performance.

After you enable DKIM validation, you can create a content filtering policy for inbound messages that specifies an action to take on messages that fail DKIM validation.

See "Creating a content filtering policy for DKIM validation" on page 141.

See "Configuring DKIM authentication" on page 137.

To enable DKIM validation

1. In the Control Center, click **Spam > Settings > Sender Authentication**.
2. Ensure that **Enable Sender Authentication** is checked.
3. Click the **DKIM** tab.
4. Check **Enable DomainKeys Identified Mail (DKIM) validation for all incoming mail**.
5. Type your authentication identifier in the **Authentication Identifier** box.

   This is a site-specific string that Symantec Brightmail Gateway inserts into each message's DKIM authentication results header. The authentication identifier uses a syntax similar to the syntax of a domain name, as in the following examples of valid authentication identifiers:

   - example.com
   - mail.example.org
   - ms1.newyork.example.com
   - example-auth
6. Type an integer in the **Maximum number of DKIM signature validations** box.

   Type an integer between 1 and 20, inclusive. When the maximum number of signatures is exceeded for a single message, DKIM validation is aborted for that message. The default value is 10 signatures.
7. Click **Save**.

Creating a content filtering policy for DKIM validation

You can create a content filtering policy that applies a specific action to messages that fail DKIM validation. To use such a policy, you should also enable DKIM validation.

See "Enabling DKIM validation" on page 140.

See "Configuring DKIM authentication" on page 137.

For an explanation of the meaning of the results that can be returned in the DKIM Authentication-Results header, see RFC 5451, Section 2.41:

tools.ietf.org/html/rfc5451#section-2.4.1

To create a content filtering policy for DKIM validation

1 In the Control Center, click **Content > Policies > Email**.

2 Click **Add**.

3 Under **Policy Template**, ensure that the **Blank** radio button is selected.

4 Click **Select**.

5 Under **Policy Name**, type a name for your policy.

6 If desired, check **Track violations of this policy in the dashboard and reports**.

7 Under **Conditions**, click on the **Apply to** drop-down menu and choose **Inbound messages**.

8 Click **Add** to add a condition.

9 Click **Text in this specific part of the message**.

10 Click the drop-down menu next to **Text in this specific part of the message** and select **Message header**.

11 In the **Header name** field, type: **Authentication-Results**

12 Click **contains** and type **1** next to **or more occurences of**.

13 Type **fail** in the adjacent text entry field.

14 Click **Add condition** at the bottom of the page.

15 Under **Actions**, click the **Perform the following action** drop-down list and select an action to perform on messages that fail DKIM validation.

16 For some actions, additional fields appear. Fill in the information as needed.

17 Click **Add Action**.

18 Under **Apply to the following groups**, check one or more groups to which you want to apply this policy.

19 Click **Save**.

Enabling DKIM signing for a domain

You can enable DKIM signing for all outbound messages from a specific domain, using an existing domain key. Symantec Brightmail Gateway can add only one DKIM signature to an outbound message.

Note: If you have no domain keys, you must first add a domain key.

See "Configuring DKIM authentication" on page 137.

See "Adding a domain key" on page 138.

To enable DKIM signing for a domain

1. In the Control Center, click **Protocols > SMTP > Domains**.
2. Click on the underlined name of the domain you want to edit.
3. On the **Edit Domain** page, click on the **Delivery** tab.
4. Under **DomainKeys Identified Mail**, click **Enable DKIM signing for messages from this domain**.
5. In the **Base domain** box, type the name of the domain in the form: **example.com**
6. In the **Selector** box, type the selector string that receiving MTAs will use to perform DNS lookup to retrieve your public key.

   The selector string is composed of any number of abritrary strings of no more than 63 lower case alphanumeric characters (a-z or 0-9) each, with a period in between the strings. If only one string is included, no period is needed.

   You can use selectors to support multiple concurrent public keys per signing domain. For more information on the use of selectors, see RFC 4871, Section 3.1.

   http://tools.ietf.org/html/rfc4871#section-3.1
7. Click on the **Signing key** drop-down list and choose the domain key you want to use.

   You must have an existing domain key to use.

   See "Adding a domain key" on page 138.
8. In the **Signature expiration** box, type an integer between 1 and 9999, inclusive, then click either **Hours** or **Days**.

   The default value is 30 days.
9. If desired, click **Show Advanced** and complete the following optional fields:

Identity

An email address, with or without the portion before the @, that includes either the base domain or a subdomain of the base domain. For example, if your base domain is example.com, acceptable identity strings include:

- @example.com
- user@example.com
- @new.example.com
- user@old.example.com

Override default signed headers

Check this box to replace the default signed headers with headers of your own design, then type one or more headers separated by colons.

You can append any header with one of the following characters:

- ? - Sign a single copy of the header. Do not assert a non-existent header if the header does not exist.
- * - Sign all existing copies of the header. Assert a non-existent header if the header does not exist.
- + - Sign all existing copies of the header. Do not assert a non-existent header if the header does not exist.
- [No character] - Sign a single copy of the header. Assert a non-existent header if the header does not exist.

Example:

```
Received+:X-Example*:From:Subject?:Received
```

Whether or not you override the default signed headers, Symantec Brightmail Gateway includes the From: header.

Headers

You can choose the method used to prepare the signature for the message headers. Simple preparation bases the signature on the exact content of the headers, including such details as spacing. Relaxed preparation creates a signature based on a representation of the headers that includes minor changes, such as changes to white spaces. If minor alterations of the headers occur during transit, relaxed canonicalization in many cases still results in a matching signature.

The default for message headers is relaxed canonicalization.

Body

You can choose the method used to prepare the signature for the message body. Simple preparation bases the signature on the exact content of the message body, including such details as spacing. Relaxed preparation creates a signature based on a representation of the message body that includes minor changes, such as changes to white spaces. If minor alterations of the message body occur during transit, relaxed canonicalization in many cases still results in a matching signature.

The default for the message body is simple canonicalization.

For more information on canonicalization, see RFC 4871, Section 3.4.

RFC 4871, Section 3.4

10. Click on **Generate** to create a DKIM DNS text record using the base domain, selector, and signing key details you specified in the above steps.

    Symantec Brightmail Gateway does not publish DKIM DNS records.

11. Click **Save**.

    Store the public key in DNS.

    Receiving MTAs will access your DNS entry to retrieve your public key.

    You can use the Linux facility `dig` to confirm that you have configured your DNS correctly.

Using SMTP authentication

SMTP authentication allows an MTA to authenticate an email client before permitting it to send messages. You can use SMTP authentication to allow remote users to send email via Symantec Brightmail Gateway. A typical use of SMTP authentication is to allow authorized users to relay mail.

SMTP authentication is a service extension to the ESMTP protocol. For more information on SMTP authentication, see RFC 4954:

http://www.ietf.org/rfc/rfc4954.txt

Many email clients, also known as Mail User Agents (MUAs), support SMTP authentication. Supported clients allow users to provide appropriate credentials to enable SMTP authentication.

Symantec Brightmail Gateway has been tested against versions of the following MUAs for SMTP authentication:

- Outlook Express
- Outlook 2003

- Outlook 2007
- Foxmail (Chinese)
- Thunderbird 2
- Mail.app (for MacOS)

Symantec Brightmail Gateway provides two methods for authenticating user credentials supplied for SMTP authentication. You can use an LDAP authentication source, or you can forward the credentials supplied by the MUA to another SMTP server for authentication.

Symantec Brightmail Gateway supports SMTP authentication via LDAP using simple bind for all supported LDAP directory types. For SMTP authentication via LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported.

Note: Symantec Brightmail Gateway searches all of your authentication directory data sources for a user attempting to authenticate. If the user exists in more than one authentication directory data source, SMTP authentication fails.

See "About data sources and functions" on page 482.

For SMTP authentication via SMTP forwarding to an SMTP server, Symantec Brightmail Gateway has been tested against servers hosting versions of the following MTAs:

- Exchange
- Domino
- Sendmail

Warning: If not configured correctly, with appropriate security safeguards, use of SMTP authentication can expose your system to significant security threats. Be sure to take appropriate steps to protect your users, systems, and data when you configure SMTP authentication.

See "Best practices for using SMTP authentication" on page 152.

To use SMTP authentication, perform the steps listed inTable 5-3

Table 5-3 Using SMTP Authentication

Step	Action	Description
Step 1	Choose your authentication source.	Required You can use either an LDAP server or an SMTP server as your authentication source. If you choose an LDAP server, you must have or create a directory data source for authentication. See the links in step 5 for more information. If you choose SMTP forwarding, you must provide details for an SMTP server that supports SMTP authentication. This server cannot be another Symantec Brightmail Gateway appliance. Skip step 2 and see the link in step 3 for more information. **Note:** Using SMTP forwarding may have an adverse impact on mail processing performance.
Step 2	Choose your authentication method.	Required if you choose LDAP, skip if you choose SMTP forwarding You can authenticate user passwords using either simple bind or password fetching. If you do not create a custom LDAP query, Symantec Brightmail Gateway defaults to using simple bind. If you want to use password fetching you must create a custom LDAP query. See "Creating and testing a custom authentication and quarantine address resolution query" on page 505. Symantec Brightmail Gateway supports SMTP authentication via LDAP using simple bind for all supported LDAP directory types. For SMTP authentication via LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported

Table 5-3 Using SMTP Authentication *(continued)*

Step	Action	Description
Step 3	Configure SMTP authentication mail settings	Required Enable authentication and provide key authentication details, including whether you will authenticate client credentials via LDAP or via SMTP forwarding. You must use an IP address/port combination for SMTP authentication that is different from both your inbound and outbound IP address/port combinations. See "Configuring SMTP authentication mail settings" on page 149.
Step 4	Configure advanced SMTP authentication settings	Optional Set maximums and other advanced configuration parameters for SMTP authentication. See "Configuring SMTP advanced settings" on page 95. See "SMTP advanced authentication settings" on page 96.
Step 5	Configure LDAP authentication	Optional Required for SMTP authentication via LDAP. If you have not yet created a data source, create a data source and enable SMTP authentication. See "Adding a data source " on page 492. See "Enabling functions on a new data source" on page 498. See "Creating an authentication data source" on page 500. If you have already created a data source, enable SMTP authentication. See "Enabling or editing the authentication function" on page 540.
Step 6	Provide instructions to end users or configure MUAs	Required Configure MUAs to connect to the authentication listener port. The default port for the authentication listener is port 587. If you changed this in Step 3 use the changed value.

Table 5-3 Using SMTP Authentication *(continued)*

Step	Action	Description
Step 7	Configure SMTP authentication alerts	Optional Configure alerts to notify administrators of SMTP authentication login failures. See "Types of alerts" on page 615. See "Configuring alerts" on page 614.

Configuring SMTP authentication mail settings

You can set up SMTP authentication to enable users who connect remotely to send messages from Symantec Brightmail Gateway.

See "Using SMTP authentication" on page 145.

See "Best practices for using SMTP authentication" on page 152.

You can also configure alerts for SMTP authentication login failures.

See "Types of alerts" on page 615.

To configure SMTP authentication mail settings

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Check the Scanner whose settings you want to modify and click **Edit**.
3. On the **Edit Host Configuration** page, click the **SMTP** tab.
4. Under **Authentication Mail Settings**, click **Enable Authentication**.

   In order to view the Authentication Mail Settings section, you must enable either **Outbound mail filtering only** or **Inbound and Outbound mail filtering**.
5. In the **Authentication mail IP address** drop-down menu, select the IP address for which you want to authenticate users.

   The available choices are the Ethernet and virtual interfaces available on the selected Scanner.

6 Either leave the default port assignment of 587, or enter a new port in the **Port** field.

The port you assign here is the port that you either configure mail clients to access or instruct users to configure in their mail clients.

You must use an IP address/port combination for SMTP authentication that is different from both your inbound and outbound IP address/port combinations.

For more information, see the following RFC:

4954

7 Check **Accept TLS encryption** if you want the host to accept connections using TLS encryption.

If you leave this option unchecked, Symantec Brightmail Gateway will not advertise support for TLS encryption during the SMTP session.

Note: You must configure an MTA TLS certificate and assign it to this Scanner before it can accept TLS encrypted email from a connection.

8 Select the name of a certificate from the drop-down menu to authenticate the Scanner as a trusted source to clients sending over TLS-encrypted connections.

See "About certificates" on page 191.

9 Check **Request client certificate** if you want the scanner to request a TLS encryption certificate from a sender before accepting a TLS-encrypted connection.

10 Check **Require TLS encryption** to allow only TLS-encrypted connections.

Warning: Symantec strongly recommends that you require TLS encryption when enabling SMTP authentication.

See "Best practices for using SMTP authentication" on page 152.

11 Under **Authentication Source**, click either **LDAP server** or **SMTP forwarding**.

If you choose **LDAP server** you must have a directory data source defined for authentication.

See "Using SMTP authentication" on page 145.

If you chose **LDAP server**, skip the next step.

If you chose **SMTP forwarding**, you must have a host running an MTA that supports SMTP authentication.

12 If you chose **SMTP forwarding**, specify the following: SMTP server, host, port, and TLS services you want to use.

- Specify the IP address or hostname in the **Host** field, and specify a **Port**.
- Optionally, for **Binding**, you can either choose **Auto** or choose a specific IP address. Symantec Brightmail Gateway sends the message to the SMTP server from the IP address you choose.
- Under TLS, choose one of the following:

Do not attempt TLS encryption	Choose this option for unencrypted communication with the SMTP server.
Attempt TLS encryption	Choose this option to attempt TLS-encrypted communication. If the SMTP server does not support TLS, communication will be unencrypted.
Require TLS encryption and don't verify certificate	Choose this option to abort communication if the SMTP server does not support TLS, without verifying the SMTP server certificate.
Require TLS encryption and verify certificate	Choose this option to abort communication if the SMTP server does not support TLS, or if the SMTP server certificate cannot be successfully verified.

- If you are using TLS, optionally check **Offer TLS client certificate** and choose the certificate you want to use.

Warning: Symantec strongly recommends that you require TLS encryption when using SMTP forwarding.

Warning: Do not specify another Symantec Brightmail Gateway appliance as the SMTP server for SMTP forwarding.

See "Best practices for using SMTP authentication" on page 152.

13 Under Authentication Mail Connections, click either:

Accept authenticated mail connections from all IP addresses	This option may be best if the users who will be connecting via SMTP authentication frequently travel.
Accept authenticated mail connections from only the following IP addresses	If users consistently connect from the same IP addresses, this option provides better security. Click **Add** to add an IP address. You can also edit or delete an IP address.

14 Click **Save**.

Best practices for using SMTP authentication

You can use SMTP authentication to allow remote users to send email using your MTA. Using SMTP authentication can introduce increased security risks if appropriate steps are not taken to ensure secure communications.

Symantec strongly recommends that you implement the following best practices if you are using SMTP authentication:

- Require TLS encryption for SMTP authentication.
  See "Configuring SMTP authentication mail settings" on page 149.
- If you use SMTP forwarding, require TLS encryption for SMTP forwarding.
- Do not attempt to use another Symantec Brightmail Gateway appliance as your SMTP server for SMTP forwarding.
- When configuring SMTP authentication mail settings, do not choose to accept authenticated mail connections from all IP addresses. Instead, specify a list of IP addresses. This best practice may not be feasible for all organizations.
- Configure and enable outbound filtering policies that protect against spam.
  See "Creating email spam policies" on page 242.
- Configure alerts for SMTP authentication login failures.
  See "Types of alerts" on page 615.

Chapter 6

Managing sender groups

This chapter includes the following topics:

- Enabling or disabling good and bad sender groups
- Choosing actions for good and bad sender groups
- Adding senders to administrator and third party sender groups

Enabling or disabling good and bad sender groups

Follow the steps below to enable or disable sender groups.

To enable or disable good or bad sender groups

1 In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.

A black hyphen in the Enabled column indicates that the entry is currently disabled. A green check in the Enabled column indicates that the entry is currently enabled.

2 Check or uncheck the boxes next to the groups that you want to enable or disable.

3 Click **Enable** or **Disable**.

Choosing actions for good and bad sender groups

All sender groups have default actions. You can choose other actions for any sender group. The following procedure does not apply to Fastpass.

To choose actions for a good or bad senders group

1. In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.
2. Click one of the bad or good sender groups.
3. Under **Actions**, click on the **Perform the following action** drop-down list and choose the action you want to add.
4. Click **Add Action**.
5. To add more actions, repeat steps 3 and 4.

   Symantec Brightmail Gateway prevents you from combining contradictory actions. For example, if the action **Defer SMTP Connection** appears, you cannot add another action because no other action can be taken on a deferred connection. If you want to add a different action, first check the box next to **Defer SMTP Connection** and click **Delete**.

   See " Verdicts and actions for email messages" on page 718.
6. To delete an action, check an action in the actions list and click **Delete**.
7. Click **Save**.

Adding senders to administrator and third party sender groups

To prevent undesired messages from being delivered to inboxes, you can add specific email addresses, domains, and connections to your bad sender groups. To ensure that messages from specific email addresses, domains, and connections are not treated as spam, you can add them to your good sender groups. You cannot add senders to the Symantec Global Good Senders or Symantec Global Bad Senders groups. This procedure does not apply to directory harvest attacks, email virus attacks or Fastpass.

To add senders to administrator and third party sender groups

1. In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.
2. Check a good or bad sender group and click **Edit**.
3. On the **Local Good or Bad Sender Domains** page, click **Add**.
4. On the **Add Sender Group Members** page, enter the information appropriate to the sender list to add it to the current sender group.

   See "Supported methods for identifying senders" on page 160.

5 Click **Save**.

6 On the **Local Good or Bad Sender Domains** page, modify the default action for messages originating from senders in this sender group, if desired.

7 Click **Save**.

Editing good and bad sender group members

Follow these steps to change sender information. This procedure does not apply to Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

To edit good and bad sender group members

1 In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.

2 Check a sender group from the sender groups list and click **Edit**.

3 To modify the information on a sender, check the sender whose information you want to modify, and click **Edit**.

You can also click an underlined sender name to automatically jump to the corresponding edit page.

4 On the **Edit Sender Group Member** page, make any changes, and click **Save**.

5 Click **Save** to commit your changes.

Deleting good and bad sender group members

Follow the steps below to delete senders. This procedure does not apply to Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks or Fastpass.

To delete senders from a good or bad senders group

1 In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.

2 Click one of the bad or good sender groups.

3 In the Members list, check the box next to the sender that you want to remove from your list, and then click **Delete**.

4 Click **Save**.

Enabling or disabling good and bad sender group members

When you add a new sender to a Sender Group, Symantec Brightmail Gateway automatically enables the filter and puts it to use when evaluating incoming messages. You may need to periodically disable and then re-enable sender group members for troubleshooting or testing purposes. Symantec Brightmail Gateway treats mail from a sender that you have disabled as it would any other message. This procedure does not apply to Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

To enable or disable good and bad sender group members

1. In the Control Center, click **Reputation > Policies > Good Senders** or **Reputation > Policies > Bad Senders**.
2. Click one of the good or bad sender groups.

   A black dash in the **Enabled** column indicates that the entry is currently disabled. A green check in the **Enabled** column indicates that the entry is currently enabled.
3. In the **Members** list on the page for that sender group, perform one of the following tasks:
   - To reinstate a member that is currently disabled in a sender group, check the box adjacent to the sender information and click **Enable**.
   - To disable a member that is currently enabled in a sender group, check the box adjacent to the sender information and click **Disable**.
4. Click **Save**.

Importing good and bad sender group entries

Use the following procedure to import LDIF-formatted text files into good and bad sender groups. This procedure does not apply to Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

Be aware of the following limitations on the number of entries that can be imported into sender groups:

- The maximum number of total senders that can be stored, including good and bad senders, is 650,000.
- The maximum number of lines per file when importing senders is 500,000. To add more, divide senders into multiple files and import each file separately.
- No warning is displayed if you exceed these limits. Sender data is silently dropped.

To import good and bad sender group entries

1. In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.

2. Check one of the bad or good sender groups and click **Edit**.

   You can import entries for all good sender groups, or all bad sender groups in one import action, no matter which group you open. However, you cannot import entries for Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

3. Click **Import**.

4. In the **Import** dialog box, specify or browse to the location of the text file with the sender information that you want to import.

   The sender information must be formatted correctly.

   See "Sender group import file formats" on page 157.

5. Click **Import**.

   Symantec Brightmail Gateway merges data from the imported list with the existing sender information.

6. Click **Save**.

Sender group import file formats

Use the specifications in this section when importing sender information for your sender groups. You cannot import sender entries for Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

See "Importing good and bad sender group entries" on page 156.

The file that you import should be line-oriented and use a format similar to the Lightweight Directory Interchange Format (LDIF), which includes the following restrictions and characteristics:

- The file must have the required LDIF header. Do not change the first three uncommented lines from the following:

  ```
  dn: cn=mailwall, ou=bmi
  objectclass: top
  objectclass: bmiBlackWhiteList
  ```

- After the header, each line must contain exactly one attribute, along with a corresponding pattern.

- Empty lines or white spaces are not allowed.
- Lines beginning with # are ignored.
- Entries terminating with the colon-dash pattern (:-) are disabled; entries terminating with the colon-plus pattern (:+) are enabled; entries with neither set of terminating symbols are enabled.

To populate the list, specify an attribute, which is followed by a pattern. In the following example, a list of attributes and patterns follows the LDIF header. See below for an explanation of the attribute codes.

```
## Permit List
#
dn: cn=mailwall, ou=bmi
objectclass: top
objectclass: bmiBlackWhiteList
AC: 65.86.37.45/255.255.255.0
AS: grandma@example.com
RC: 20.45.32.78/255.255.255.255
RS: spammer@example.com
BL: sbl.spamhaus.org
# Example notations for disabled and enabled entries follow
RS: rejectedspammer@example.com:-
RS: rejectedspammer2@example.com:+
```

The following table lists the attributes and the syntax for the values.

Attribute	Description	Examples

AC:	Allowed (good) connection or network. Specify a numerical IP address, numerical IP address and network mask, or Classless Inter-Domain Routing (CIDR) IP address.	`AC: 76.86.37.45`
		`AC: 76.86.37.45/255.255.255.0`
		`AC: 76.86.37.00/18`

RC:	Rejected or blocked (bad) connection or network. Specify a numerical IP address, numerical IP address and network mask, or CIDR IP address.	`RC: 76.86.37.45`
		`RC: 76.86.37.45/255.255.255.0`
		`RC: 76.86.37.00/18`

Attribute	Description	Examples

AS:	Allowed (good) sender. Specify an email address or domain using alphanumeric and special characters, except the plus sign (+).	`AS: example.com`
		`AS: spammer@example.org`
		`AS: john?????@example.com`

RS:	Rejected or blocked (bad) sender. Specify an email address or domain using alphanumeric and special characters, except the plus sign (+).	`RS: example.com`
		`RS: spammer@example.org`
		`RS: john?????@example.com`

BL:	Third party blocked (bad) sender list. Use the zone name specified by the list provider.	`BL: sbl.spamhaus.org`
WL:	Third party allowed (good) sender list. Use the zone name specified by the list provider.	`WL: allowed.example.com`

Exporting sender group information

Occasions can arise when you want to export the data stored in your sender group for use in another application. Use the following procedure to export sender group entries to a text file. This procedure does not apply to Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.

To export sender group information to a text file

1. In the Control Center, click **Reputation > Policies > Bad Senders** or **Reputation > Policies > Good Senders**.
2. Click one of the bad or good sender groups.

   The entries for all good sender groups, or all bad sender groups are exported no matter which list you open. However, the you cannot export entries for Symantec Global Good Senders, Symantec Global Bad Senders, directory harvest attacks, email virus attacks, or Fastpass.
3. Click **Export**.

   Your browser will prompt you to open the file from its current location or save it to disk.

Supported methods for identifying senders

You can use the following methods to identify senders for your good sender groups and bad sender groups:

Method	**Notes**
IP-based	Specify IP connections. Symantec Brightmail Gateway checks the IP address of the mail server initiating the connection to verify if it is in your good sender groups or bad sender groups. Wildcards are not supported. Although you can use network masks to indicate a range of addresses, you cannot use subnet masks that define non-contiguous sets of IP addresses (for example, 69.84.35.0/255.0.255.0). The following notations are supported: - Single host: 128.113.213.4 - IP address with subnet mask: 128.113.1.0/255.255.255.0 - Classless Inter-Domain Routing (CIDR) IP address: 192.30.250.00/18
Third-party services	Supply the lookup domain of a third-party sender service. Symantec Brightmail Gateway can check the message source against third party DNS-based lists to which you subscribe—for example, list.example.org. **Note:** Be sure to confirm the quality of a third-party list before using it. Symantec is unable to resolve false positives that result from use of third-party lists.
Domain-based	Specify sender addresses or domain names. Symantec Brightmail Gateway checks the following characteristics of incoming mail against those in your lists: - MAIL FROM: address in the SMTP envelope. Specify a pattern that matches the value for localpart@domain in the address. You can use the * or ? wildcards in the pattern to match any portion of the address. - `From`: address in the message headers. Specify a pattern that matches the value for localpart@domain in the FROM: header. You can use wildcards in the pattern to match any portion of this value.

If you choose to identify senders by address or domain name, use the following examples to model entries when you add members to a sender group:

Example	Sample matches
example.com	chang@example.com, marta@example.com, john@bank.example.com
malcolm@example.net	malcolm@example.net
sara*@example.org	sara@example.org, sarahjane@example.org
jo??@example.corg	john@example.org, josh@example.org

Chapter 7

Enabling reputation filtering

This chapter includes the following topics:

About blocking and allowing messages at connection time

Scanning email for spam, viruses, and content filtering is a resource-intensive task. Any email that must be processed past the gateway taxes your mail infrastructure, resource capacity, and system performance. Symantec Brightmail Gateway features Brightmail Adaptive Reputation Management (Brightmail ARM). Brightmail ARM includes features designed to reduce unnecessary incoming email traffic, protect your network from attacks, and optimize the use of your processing resources.

Brightmail ARM includes technologies that can reject or defer incoming connection attempts based solely on the incoming IP address. To accomplish this, Brightmail ARM uses dynamic, self-learning local reputation data, global reputation data, and administrator-defined Bad Sender Policies and Good Sender Policies.

Brightmail ARM generates local reputation data based on good and bad verdicts rendered on messages in your mail stream. Brightmail ARM builds global reputation data by leveraging the extensive world-wide data collection capabilities of Brightmail IQ Services. Brightmail IQ Services includes the Probe Network, Symantec's collection of millions of honeypot emails that collect spam throughout the Internet, as well as the Global Intelligence Network. The Global Intelligence Network includes threat detection and response centers around the world, managed by Symantec Security Response.

Brightmail ARM uses these diverse technologies to achieve five goals:

- Reduce the volume of incoming email traffic by eliminating most spam messages at the gateway.
- Stop virus, malware, and directory harvest attacks at the gateway.
- Allow messages from senders with the best local reputation to bypass spam scanning.
- Provide uninterrupted connection abilities to your best senders, regardless of the volume of spam or attacks at any moment.
- Protect you from denial-of-service attacks by limiting the connection abilities of illegitimate senders.

Symantec Brightmail Gateway conserves, protects, and optimizes your physical assets, your message flow, and your vital data. Brightmail ARM is the first stage in the inbound protection process. By examining the incoming IP addresses, and in some cases also the message envelope, Brightmail ARM can take preventive action. By rejecting or deferring undesirable connections, Brightmail ARM restores valuable filtering cycles to the Scanner.

Brightmail ARM employs the following features and technologies to achieve these aims.

Table 7-1

Feature	Description
Connection Classification	Connection Classification provides the best connection abilities to your best senders, and progressively worse connection abilities to all other senders. Connection Classification ensures that your worst senders cannot degrade the connection experience of your best senders. Connection Classification automatically places every incoming sender IP into one of 10 classes based on local reputation. Class membership is determined based on how many legitimate and spam messages each IP has sent to the Scanner, and is constantly updated. New IPs are assigned to the Default class. Senders in Good Sender groups always use the best class (Class 1). Senders in Bad Sender groups always use the worst class (Class 9). See "About managing connection load at the gateway" on page 166. See "Configuring Connection Classification" on page 168.
Email virus attack prevention	If Symantec Brightmail Gateway detects a specified number of infected messages from an IP address, email virus attack prevention can then defer further connections. Or, you can choose other actions. See "Configuring email virus attack recognition" on page 170.
Directory harvest attack prevention	If Symantec Brightmail Gateway detects a specified number and percentage of invalid recipient from an IP address, directory harvest attack prevention can then defer further connections. Or, you can choose other actions. See "Configuring directory harvest attack recognition" on page 172.
Bad Sender Policies	You can add senders to administrator-defined groups and use Symantec Global Bad Senders to block email from bad senders, or choose other actions. See "About blocking and allowing messages using sender groups" on page 174.

Table 7-1 *(continued)*

Feature	Description
Good Sender Policies	You can add senders to administrator-defined groups and use Symantec Global Good Senders to deliver messages from good senders normally, or choose other actions. See "About blocking and allowing messages using sender groups" on page 174.
Fastpass	The Fastpass feature conserves resources by exempting senders with the best local reputation from spam scanning. Symantec Brightmail Gateway automatically collects local sender reputation data to support Fastpass determinations and regularly re-evaluates senders granted a pass. Symantec Brightmail Gateway grants and revokes passes based solely on how many messages from each sender it determines to be spam. You can exclude specific senders from ever receiving a pass. See "About conserving resources using Fastpass" on page 178. See "Configuring Fastpass" on page 179.

About managing connection load at the gateway

In most networks the great majority of email traffic today is spam. By intelligently managing connection load and distinguishing between connections from senders known to send spam and legitimate senders, you can significantly reduce processing costs. The Connection Classification feature dynamically manages connection load based on automatically collected local reputation data. Connection Classification is a self-learning feature. In response to the latest changes in local reputation, Connection Classification updates its management of connection load on a just-in-time basis.

Spammers routinely leverage vast networks of compromised client machines, known as botnets, to disseminate their attacks. This enables them to generate huge volumes of messages without sending enough messages from any single IP address to merit entry on a global blacklist. Connection Classification supplements global lists from Symantec and third parties and your own administrator-defined lists with an approach that is effective against botnet-driven spam and the huge overall volume of spam.

By reducing the system resources used by senders with poor local reputation, Connection Classification protects your legitimate mail flow from denial-of-service attacks. With Connection Classification enabled, spammers get fewer connections. As a result, more resources are available to your legitimate senders.

Note: To take advantage of Connection Classification, your Symantec Brightmail Gateway appliance must be deployed at the gateway.

Connection Classification works by assigning each connecting IP address to one of 10 classes, based on the amount of spam sent by that IP address. Connection Classification assigns new IP addresses to the default class. Connection Classification regularly changes the classifications of senders, as it continues to learn more about sender reputation in real time.

Connection Classification allows most connections for the best senders (class 1). As one moves from the best class to the worst class (Class 9), the network resources allowed a sender decrease. For Class 9, Connection Classification defers most connections.

Senders in the Symantec Global Good Senders, Local Good Sender IPs, and Third Party Good Senders groups are always assigned to the best class (Class 1). Senders in the Symantec Global Bad Senders, Local Bad Sender IPs, and Third Party Bad Senders groups are always assigned to the worst class (Class 9).

Symantec Brightmail Gateway determines class membership separately for each Scanner in your system. The same sending IP can be in Class 3 on one Scanner and Class 4 on another Scanner. Based on the amount of spam sent from each IP address, the classifications can change constantly, to dynamically reflect the latest local, per-Scanner reputation.

The restrictions placed on a sender's ability to consume system and network resources correlate directly with the sender's reputation for spamming your organization. Senders with a clean history are placed in the best class and allowed more frequent connections than those with poor records. Conversely, an IP address with a poor reputation can improve its class over time by sending less spam and more legitimate email.

Connection Classification uses the data collected in the reputation database to determine the probability that a message sent from a given IP is spam. As the appliance collects more data over time, the probabilistic determination yields more accurate results.

The only action Symantec Brightmail Gateway takes based on Connection Classification is to defer some SMTP connections. Connection deferral is also known as soft rejection or a 450 SMTP error. Connection Classification defers connections during the connection phase of the inbound message flow and also during the SMTP session phase.

See “About email message flow” on page 131.

Symantec Brightmail Gateway does not take any action based on Connection Classification until the appliance has recorded enough data to make accurate predictions. Immediately after the initial installation of a Scanner, Connection Classification is in learning mode. Learning mode ends when 50,000 messages have been received and the statistics gathered from them have been added to the database. At that point, if Connection Classification is enabled, connection management begins. If you have multiple Scanners, a newly installed Scanner is initially be in learning mode, while your other Scanners may already be managing connection load.

Note: If you disable Connection Classification, the Scanner continues to record sender reputation information. This means that you can disable this feature temporarily and not miss any sender data during that time.

You can edit the connection parameters for each class.

See “Configuring Connection Classification” on page 168.

You can query the status of an IP's reputation.

See “Researching IP address reputation” on page 186.

Configuring Connection Classification

Using Connection Classification ensures that the most abusive senders cannot degrade the connection ability of your best senders. Connection Classification automatically classifies every incoming IP address into one of 10 classes. Symantec Brightmail Gateway automatically gathers local reputation data to inform the classification. Senders in the best class, because they rarely if ever send spam, benefit from the best connection parameters. Senders in the worst class are subject to the worst connection parameters. New IP addresses are initially placed into the default class.

Upon initial installation, Connection Classification is in learning mode for the first 50,000 messages. During learning mode no messages are deferred based on their connection class.

Connection Classification is designed to work without any configuration. However, you can configure Connection Classification to customize the parameters for your message flow. Use the procedures in this section to enable, disable, or configure Connection Classification.

To configure Connection Classification

1 In the Control Center, click **Reputation** > **Policies** > **Connection Classification**.

2 On the **Connection Classification** page, check **Enable Connection Classification**.

To disable **Connection Classification**, uncheck the box.

3 To configure Connection Classification parameters, click **Edit**.

The fields in the table become editable, and the **Edit** button changes to a **Load Defaults** button.

See "Connection class default settings" on page 169.

4 To change the maximum connections for each class, type new values in the 10 fields on the **Maximum connections** row.

Each value is the percent of total available connections that are allocated to that class. The total of all 10 values must equal 100%. In each field, you can enter a value between 0.1 and 99.1, inclusive. As you edit the fields, the current total of the amounts you entered appears in red below the **Maximum connections** label on the left.

5 To change the connections that are allowed on a per-IP address basis, type new values in the **Maximum Connections per IP** fields.

6 To vary the number of messages that are allowed on a per-connection basis, type new values in the **Messages per Connection** fields.

7 To vary the time, in seconds, before a sender IP is allowed to reconnect, type new values in the **Reconnect Timeout** fields.

8 To vary the proportion of connections that are deferred for each class, type new values in the **Deferred Messages** fields.

Each value on this row represents the percentage of the total messages for sender IP addresses in that class that must be deferred. You can type any integer between 0 and 100, inclusive. The values do not need to add up to 100%.

9 To abandon your changes and return to the default values, click **Load Defaults**.

10 To commit your changes, click **Save**.

Connection class default settings

Table 7-2 shows the default values for each connection class. A value of zero (0) indicates that there is no limit.

Table 7-2 Default values for each connection class

Field	Default Class	Class 9 (worst)	Class 8	Class 7	Class 6	Class 5	Class 4	Class 3	Class 2	Class 1 (best)
Maximum connections (must total 100%)	10.0%	0.2%	0.4%	1.0%	5.0%	10.0%	10.0%	10.0%	19.0%	34.4%
Maximum connections per IP	1	1	1	1	1	1	25	50	100	200
Messages per connection	1	1	1	1	1	5	10	20	40	0
Reconnect timeout	10 sec	60 sec	30 sec	30 sec	15 sec	5 sec	2 sec	1 sec	1 sec	0 sec
Deferred messages	10%	95%	80%	60%	30%	10%	5%	0%	0%	0%

See “About managing connection load at the gateway” on page 166.

Configuring email virus attack recognition

In an email virus attack, a specified quantity of infected email messages has been received from a particular IP address. By default, any connections that are received from violating senders are deferred.

Set up email virus attack recognition as described in the following procedure. Email virus attack recognition is disabled by default, and must be enabled to be activated.

To enable or disable email virus attack recognition

1. In the Control Center, click **Reputation > Policies > Bad Senders**.
2. To enable or disable email virus attack recognition on this page, click **Email Virus Attacks**, then click **Enable** or **Disable**.

   Or, continue with the next step.
3. Click **Email Virus Attacks**.
4. Check **Enable Email Virus Attack detection** to enable email virus attack recognition, or uncheck **Enable Email Virus Attack detection** to disable email virus attack recognition.

To configure email virus attack recognition

1. In the Control Center, click **Reputation > Policies > Bad Senders**.
2. Click **Email Virus Attacks**.

3 Accept the defaults or modify the values under Email Virus Attack Configuration:

Minimum percentage of virus messages	Percentage of virus messages from a single server that must be exceeded to trigger the specified action. The minimum number must also be exceeded.
Minimum number of virus messages	Number of virus messages from a single server that must be exceeded to trigger the specified action. The minimum percentage must also be exceeded.
Qualification time window	Time period in which the specified percentage and number of virus messages must be exceeded to trigger the specified action.
Penalty box time	Period of time during which Symantec Brightmail Gateway performs the specified action against all messages from the sending SMTP connection.

4 Under **Actions**, you can:

- Accept the default, recommended action of Defer SMTP Connection with a message of "try again later."
- Edit the action to enter a new message and click **Update Action**.
- Or, select another action from the drop-down list under **If an email virus attack occurs**

Other actions may provide additional options for you to configure. For instance, if you choose the **Archive the message** action, the **Email Virus Attacks** page displays an **Optional archive tag** text box and an **Encoding** drop-down list.

5 Click **Add Action** to add the action to the list of actions Symantec Brightmail Gateway takes upon recognizing a virus attack.

Symantec Brightmail Gateway prevents you from combining contradictory actions. For example, you cannot add another action to the default action because no other action can be taken on a deferred connection. If you want to add a different action, first check the box next to **Defer SMTP Connection** and click **Delete**.

See “ Verdicts and actions for email messages” on page 718.

6 To change the settings for an existing action, check the box next to the action and click **Edit**.

Any available options for that action appear. Click **Update Action** after configuring the options

7 Click **Save**.

Configuring directory harvest attack recognition

Spammers employ directory harvest attacks to find valid email addresses at the target site. A directory harvest attack works by sending a large quantity of possible email addresses to a site. An unprotected mail server rejects any messages that are sent to invalid addresses. This behavior allows spammers to tell which email addresses are valid by checking the rejected messages against the original list.

When directory harvest attack recognition is enabled, any connections that are received from violating senders are deferred by default. Deferring a connection slows down the progress of a possible attack and discourages spammers from maintaining the connection.

Set up directory harvest attack recognition as described in the following procedures. Directory harvest attack recognition is disabled by default. You must enable directory harvest attack recognition to activate it.

Note: Before enabling directory harvest attack recognition, you must perform the following actions:

- Create and enable a data source with recipient validation enabled.
  See "Creating a recipient validation data source" on page 512.
- Set up your local domains. Symantec Brightmail Gateway accepts inbound messages only for the domains you specify.
  See "Adding or editing domains" on page 117.
- Enable invalid recipient handling, configured to reject invalid recipients.
  See "Setting up invalid recipient handling" on page 130.

To enable or disable directory harvest attack recognition

1 In the Control Center, click **Reputation > Policies > Bad Senders**.

2 To enable or disable directory harvest attack recognition on this page, check **Directory Harvest Attack** and click **Enable** or **Disable**.

Or, continue with the next step.

3 Click **Directory Harvest Attack**.

4 Check **Enable DHA detection** to enable directory harvest attack recognition, or uncheck **Enable DHA detection** to disable directory harvest attack recognition.

5 Click **Save**.

To configure directory harvest attack recognition

1 In the Control Center, click **Reputation > Policies > Bad Senders**.

2 Click **Directory Harvest Attack**.

3 Accept the defaults or modify the values under Directory Harvest Attack Configuration:

Minimum percentage of bad recipients	Percentage of bad recipient messages from a single server that must be exceeded to trigger the specified action. The minimum number must also be exceeded. Bad recipient messages are messages sent to addresses in your local domains that do not exist.
Minimum number of bad recipients	Number of bad recipient messages from a single server that must be exceeded to trigger the specified action. The minimum percentage must also be exceeded.
Qualification time window	Time period in which the specified percentage and number of bad recipient messages must be exceeded to trigger the specified action.
Penalty box time	Period of time during which Symantec Brightmail Gateway performs the specified action against all messages from the sending SMTP connection.

4 Under **Actions**, you can:

- Accept the default, recommended action of Defer SMTP Connection with a message of "try again later."
- Edit the action to enter a new message and click **Update Action**.
- Or, select another action from the drop-down list under **If a directory harvest attack occurs**.

Other actions may provide additional options for you to configure. For instance, if you choose the **Archive the message** action, the **Directory Harvest Attack** page displays an **Optional archive tag** text field and an **Encoding** drop-down list.

5. Click **Add Action** to add the action to the list of actions Symantec Brightmail Gateway takes upon recognizing a directory harvest attack.

   Symantec Brightmail Gateway prevents you from combining contradictory actions. For example, you cannot add another action to the default action because no other action can be taken on a deferred connection. If you want to add a different action, first check the box next to **Defer SMTP Connection** and click **Delete**.

   See “ Verdicts and actions for email messages” on page 718.

6. To change the settings for an existing action, check the box next to the action name and click **Edit**.

   Any available options for that action appear. Click **Update Action** after configuring the options.

7. Click **Save**.

About blocking and allowing messages using sender groups

Filtering email based on the sender's domain, IP address, or email address provides administrators and end users a powerful way to reduce spam and malware.

Note: This section describes administrator-defined and global sender groups, which are applied at the server level for your organization. To allow end users to maintain individual sender lists, enable personal good and bad sender lists by going to Administration > Users > Groups.

See “Enabling and disabling end user settings for policy groups” on page 327.

Symantec Brightmail Gateway lets you customize spam detection in the following ways:

Define good senders	Symantec Brightmail Gateway treats mail coming from an address or connection in the Local Good Sender Domains and Local Good Sender IPs groups as legitimate mail. The good sender groups reduce the small risk that messages sent from trusted senders will be treated as spam or filtered in any way. By default messages from these senders are delivered normally.

Define bad senders	Symantec Brightmail Gateway supports a number of actions for mail from a sender or connection in the Local Bad Sender Domains and Local Bad Sender IPs groups. By default, messages from senders in the Local Bad Sender Domains group are deleted. By default, SMTP connections from senders in the Local Bad Sender IPs and Third Party Bad Senders groups are rejected. However, you can instead choose other actions.
Use global sender groups	By default, Symantec Brightmail Gateway is configured to use Symantec Global Good Senders and Symantec Global Bad Senders. Symantec monitors hundreds of thousands of email sources to determine how much email sent from these IP addresses is legitimate and how much is spam. Symantec Global Good Senders consists of IP addresses known as legitimate senders based on reputation data collected by Symantec. Symantec Global Bad Senders consists of IP addresses that have sent large amounts of spam to mail servers protected by Symantec. Both groups are continuously compiled, updated, and incorporated into Symantec Brightmail Gateway filtering processes at your site. No configuration is required for these lists. You can choose to disable either of these lists. By default, messages from senders in the Symantec Global Good Senders group are delivered normally. By default, SMTP connections from senders in the Symantec Global Bad Senders group are rejected. However, you can instead choose other actions.
Incorporate lists managed by other parties	Third parties compile and manage lists of desirable or undesirable IP addresses. These lists are queried using DNS lookups. You can add third-party sender lists to your Third Party Bad Senders or Third Party Good Senders groups. By default, SMTP connections from bad senders in these groups are rejected, and message from good senders in these groups are delivered normally. However, you can instead choose other actions. **Note:** Be sure to confirm the quality of a third party list before using it. Symantec is unable to resolve false positives that result from third-party lists.

Table 7-3 describes why you might want to maintain lists of good or bad senders for your organization and gives examples of patterns that you might use to match the sender.

Table 7-3 Use cases for good and bad sender groups

Problem	Solution	Pattern example
Mail from an end-user's colleague is occasionally flagged as spam.	If personal good and bad sender lists are enabled for end users, the user can add the colleague's email address to their Good Senders list. To enable this capability for an end user, go to **Administration > Users > Policy Groups**, edit the policy group containing the end user, and click on the **End User** tab. The user can then add colleague@trustedco.com to their Good Senders list. See "Enabling and disabling end user settings for policy groups" on page 327.	colleague@trustedco.com
Desired newsletter from a mailing list is occasionally flagged as spam.	Add newsletter.com to the Local Good Sender Domains group. See "Adding senders to administrator and third party sender groups" on page 154.	latest@newsletter.com
An individual is sending unwanted mail to people in your organization.	Add Joe.unwanted@getmail.com to the Local Bad Sender Domains group. See "Adding senders to administrator and third party sender groups" on page 154.	Joe.unwanted@getmail.com
Numerous people from a specific range of IP addresses are sending unsolicited mail to people in your organization.	After analyzing the received headers to determine the sender's network and IP address, add 218.187.0.0/255.255.0.0 to the Local Bad Sender IPs group. See "Adding senders to administrator and third party sender groups" on page 154. See "Supported methods for identifying senders" on page 160.	218.187.0.0/255.255.0.0

When evaluating domain name matches, Symantec Brightmail Gateway automatically expands the specified domain to include subdomains. For example, Symantec Brightmail Gateway expands example.com to include biz.example.com and jenny@foo.example.com, to ensure that any possible subdomains are allowed or blocked as appropriate.

See "Supported methods for identifying senders" on page 160.

You cannot have the exact same entry in both a good sender group and a bad sender group. If an entry already exists in one group, you see an error message when you try to add the same entry to the other group. If you prefer that an entry in one group appear as an entry on the other, first delete the entry from the group where it currently resides, then add it to the other group.

Incorporating third-party lists adds additional steps to the filter process. For example, similar to a typical DNS query, the IP address of the sending mail server for each incoming message is checked against a DNS list maintained in the third-party database. If the sending mail server is on the list, the mail is flagged as spam. If your mail volume is sufficiently high, running incoming mail through a third-party database could hamper performance because of the requisite DNS lookups. Symantec recommends that you use the Symantec Global Good Senders and Symantec Global Bad Senders groups instead of enabling third-party lists.

When deployed at the gateway, Symantec Brightmail Gateway obtains the physical or peer IP connection for an incoming message and compares it to entries in the good sender and bad sender groups. If a Scanner is deployed elsewhere in your network, for example, downstream from a gateway MTA that is not identified as an internal mail host, Symantec Brightmail Gateway may identify the IP address of your gateway server as a source of spam. You should accurately identify all internal mail hosts that are upstream relative to inbound mail flow from your Symantec Brightmail Gateway appliance.

See "Specifying internal mail hosts for non-gateway deployments" on page 104.

In addition to internal mail hosts you can add, Symantec Brightmail Gateway includes a series of IP address ranges in the internal hosts list as follows:

- 0.0.0.0/255.0.0.0
- 10.0.0.0/255.0.0.0
- 127.0.0.0/255.0.0.0
- 169.254.0.0/255.255.0.0
- 172.16.0.0/255.240.0.0
- 192.168.0.0/255.255.0.0

Symantec Brightmail Gateway will exclude the IP addresses of internal mail hosts from the following verdicts:

- Local Good Sender IPs
- Local Good Third Party Senders
- Local Bad Sender IPs
- Local Bad Third Party Senders
- Directory Harvest Attacks
- Symantec Global Bad Senders
- Symantec Global Good Senders
- Connection Classification
- Email Virus Attacks

- Fastpass

About conserving resources using Fastpass

Scanning email messages for spam is a resource-intensive process. Fastpass conserves resources by providing a temporary exemption from spam scanning for senders with a demonstrated history of sending no spam messages. A "pass" is granted to such a message source if that source has sent a specified number of consecutive legitimate messages, 15 by default.

Once a source has received a pass, the amount of antispam processing applied to messages from that source decreases over time. The number of messages allowed to bypass antispam filtering increases as more and more legitimate email comes from the source. This reduces the processing time required for messages from legitimate sources. This may also decrease effectiveness in detecting spam. The Fastpass feature is designed to deliver a significant increase in performance, at the cost of a minimal decrease in effectiveness.

If a message source holding a pass subsequently sends a spam message that is sampled, the pass is immediately revoked from all IP addresses in the /24 range of the offending IP address. Full antispam analysis is performed on all messages from sources in that range. The source remains eligible to receive another pass however, by once again meeting all the configured criteria.

Fastpass uses a database to store and categorize message source IP addresses. The database consists of two tables:

Fastpass trial table	Contains entries being evaluated for possible inclusion in the Fastpass table. A determination is made to move an IP address from the trial table to the pass table based on successful testing for legitimate messages for the IP address. All messages from IP addresses in the Fastpass trial table are scanned for spam.
Fastpass table	Contains entries that have been granted a pass by Fastpass based on no spam coming from the IP address for a specified number of messages.

You can configure:

- The size of the Fastpass database. The default size is 250,000 IP addresses; you can change this to any integer between 1,000 and 1,000,000, inclusive. 25% of the database is reserved for the Fastpass table. The remaining 75% reserved for the Fastpass trial table.
- The rate of growth of the Fastpass trial table, by specifying the probability that an unknown IP address is added to the Fastpass trial table when it sends a legitimate message. You can specify any (integer) probability from a 1/1

probability to a 1/100 probability, inclusive. You only type the denominator of the fraction to indicate the probability. The default value is 3, which indicates a 1/3 or 1 in 3 probability. Note that the number is a probability, not a certainty. If set to 1/3, for example, there may be occasions when 5 sequential different unknown IPs are not added, or when 2 IPs in a row are added.

- The rate of growth of the Fastpass table, by specifying the number of sequential legitimate messages required before an IP address in the Fastpass trial table is added to the Fastpass table. You can specify any integer between 1 and 1,000, inclusive, without commas. The default setting is 15.
- The initial probability that a message is scanned for spam after a pass is issued to an IP address. As the IP address continues to send legitimate messages, this sampling rate decreases from this rate. You can specify and integer between 2 and 50, inclusive. You specify only the denominator of the fraction that indicates the probability. The default value is 5, which indicates a 1/5 or 1 in 5 probability.
- The list of IP addresses that cannot be granted a pass. You can create this list by adding individual IP addresses or importing a list of IP addresses.

When either table reaches the configured limit, Fastpass discards the least recently used entry to make room for the next entry.

Note that Fastpass only exempts senders from spam scanning. Messages from senders with passes are scanned for viruses and content filtering. However, because these messages are not scanned for spam, they cannot receive Suspected Spam or Spam verdicts.

Senders who are members of the Local Good Sender IPs or Symantec Global Good Senders sender groups cannot receive a pass. However, those senders are already exempted from spam scanning.

The Fastpass database only takes spam and suspected spam verdicts into account. Virus verdicts, content filtering verdicts, spim verdicts, sender authentication verdicts, and sender group verdicts do not affect the granting or revoking of passes.

Configuring Fastpass

Fastpass conserves processing resources by exempting sending IP addresses with the best local reputation from spam scanning. Symantec Brightmail Gateway automatically collects data on the level of spam sent by each IP address, and uses this data to grant or revoke passes. After Symantec Brightmail Gateway grants a pass to a sender, it still scans a small sample of the email from that sender for spam. If Symantec Brightmail Gateway identifies a spam message, it immediatly revokes the pass for that sender.

By default, the Fastpass feature is not enabled. Enabling Fastpass can yield significant savings in processing resources. Use the procedure below to enable or disable Fastpass. Fastpass is designed to work without any custom configuration. Advanced users can customize Fastpass using the procedures in this section.

To enable or disable Fastpass

1. In the Control Center, click **Reputation** > **Policies** > **Good Senders**.
2. Click **Fastpass**.
3. Check **Enable Fastpass** to enable Fastpass, or uncheck **Enable Fastpass** to disable Fastpass.

To configure Fastpass

1. In the Control Center, click **Reputation** > **Policies** > **Good Senders**.
2. Click **Fastpass**.
3. Click **Show Advanced**.
4. To change the database size, type an integer between 1,000 and 1,000,000, inclusive in the **Maximum number of sending IPs tracked in the database** field.

   Do not type commas. The default value is 250,000. The database includes two tables. The Fastpass trial table includes sending IPs being evaluated for a pass, and is limited to 75% of the maximum database size. The Fastpass table includes sending IPs that currently have a pass, and is limited to 25% of the maximum database size.
5. To vary the rate of growth of the Fastpass trial table, type an integer between 1 and 100, inclusive, in the **Chance that a new IP will be added to the Fastpass trial table** field.

   A smaller value will result in more frequent sampling. For example, a value of 1 means that every new IP that sends a legitimate message will be added to the table. A value of 5 means that a new sending IP not in the Fastpass trial table has a 20% (1/5 or 1 in 5) chance of being added to the table. The default value is 3, which indicates a 1 in 3 probability. All messages from IP addresses in the Fastpass trial table are scanned for spam. If a spam message is received, the entire /24 range of sending IP addresses is removed from both the Fastpass trial table and the Fastpass table.

6 To vary the rate of growth of the Fastpass table, type in integer between 1 and 1,000, inclusive, in the **Minimum required legitimate messages before granting fastpass** field.

Do not type commas. The default value is 15. After an IP enters the Fastpass trial table, this value is the number of sequential legitimate messages that must be received from the IP before the IP moves to the Fastpass table. Any spam message received will cause the entire /24 range of IP addresses to be dropped from both tables.

7 To vary the rate at which IPs that have passes are checked for current behavior, type an integer between 2 and 50, inclusive, in the **Initial message sample rate after pass is issued** field.

A smaller value will result in more frequent sampling. For example, a value of 2 means that 1 of every 2 messages for an IP just granted a pass is scanned for spam. A value of 25 means that 1 of every 25 messages is scanned for spam. The default value is 5, which indicates a 1 in 5 probability. As additional legitimate messages are received from an IP address, the initial sampling rate is adjusted so that fewer messages are sampled. The sampling rate cannot fall to less than 5 times the initial sampling rate. For example, an initial sampling rate of 8 would gradually decrease as additional legitimate messages are processed, until the sampling rate is 1 message out of 40.

To exclude IPs from receiving Fastpasses

1 In the Control Center, click **Reputation** > **Policies** > **Good Senders**.

2 Click **Fastpass**.

3 Click **Show Advanced**.

4 Under **Fastpass Exclusions**, type the IP addresses you want to exclude from receiving passes in the **IP addresses** field.

You can type multiple IP addresses, separated by commas. You can type fully qualified IP addresses or multiple IP addresses using CIDR notation. Wildcards are not supported. If you specify hostnames, some of the perfomance benefit of Fastpass is lost, as Symantec Brightmail Gateway then needs to look up the hostname of the IP for every sampled message to ensure that it does not match a hostname you have specified to exclude.

5 To import a list of IP addresses, click **Import**.

Imported files must be plain text files containing a single entry per line.

6 To delete and entry from the list, check the box next to the IP address and click **Delete**.

7 To delete all entries on the current page of the list, click **Delete All**.

8 To export the currently saved list of excluded IPs in a text file, click **Export**.

About defending against bounce attacks

A bounce attack occurs when a spammer obscures message origins by using one email server to bounce spam to an address on another server. The spammer does this by inserting a target address into the “Mail From” value in the envelope of their messages then sending those messages to another address.

If the initial recipient finds the message undeliverable, that mail server recognizes the forged "Mail From" value as the original sender, and returns or "bounces" the message to that target. When the targeted system recognizes the server from which the message was bounced as a legitimate sender, it accepts the message as a legitimate non-deliverable receipt (NDR) message.

Bounce attacks can be used to leverage the initial recipient's "good" reputation when sending spam, pollute the initial recipient's IP reputation, or create denial of service attacks at the target's server.

Symantec Brightmail Gateway uses bounce attack prevention to eliminate NDRs that are a result of such redirection while still delivering legitimate NDRs.

To set up bounce attack prevention for your mail system, you must:

- Provide a Bounce attack prevention seed value in your Control Center.
  See “Configuring the Control Center for bounce attack prevention” on page 183.
- Determine and configure the policy groups to which you want the system to apply bounce attack prevention.
  See “Configuring policy groups for bounce attack prevention” on page 184.
- Assign a policy (a default policy is provided) to the policy group that determines the actions to be taken for NDRs that do not pass bounce attack prevention validation.
  See “Creating an email spam policy for bounce attack prevention” on page 185.

Note: For successfull processing you must also ensure that all of your applicable outbound mail is routed through the appliance.

Once your system is configured for bounce attack prevention, Symantec Brightmail Gateway calculates a unique tag that uses the provided seed value as well as the current date. Your Scanner attaches this tag to outbound messages sent by users in your defined policy groups.

If the message is then returned as undeliverable, the envelope's return address will contain this tag.

When the system receives a message that appears to be a message returned as undeliverable, the system will compare the inbound message's recipient with the policy group configuration to see if the user's policy group is configured for bounce attack prevention. If the policy group is configured, the system calculates a new tag that includes the seed value and current date, then uses that new tag to validate the tag in the email.

A valid tag on an inbound NDR will include the following:

- The correct tag format
- A seed value that matches the seed value in the new calculated tag
- A date that falls within a week of the new calculated tag

Based on this evaluation, Symantec Brightmail Gateway will do the following:

- If the system determines that the tag is valid, the system strips the tag from the envelope and sends the message forward for filtering and delivery per your mail filtering configuration.
- If there is no tag, or the tag content is found to not match the tag that is calculated for validation, the address will be rewritten without tag information then managed per your bounce attack prevention policy configuration. An error will be logged and this message will be accounted for in your message statistics as a message with a "single threat." The message is also included in your system spam statistics as a "bounce threat."
  If your policy defines an action other than "reject" when the message fails validation, the message can acquire more threats and could then be counted in your statistics as a "multiple threat."
- If, due to an unrecognizable format, validation cannot be performed by the system, the system will not strip the tag and will keep the tag as part of the address. The system will then act upon this message based on the actions you define in your spam policy configuration.

Note: Bounced messages over 50k are truncated. Attachments in truncated messages may be unreadable.

Configuring the Control Center for bounce attack prevention

You must configure bounce attack prevention in the Control Center by providing a seed value that will be used to calculate a tag that will be appended to outgoing messages and later used to validate that message if it is returned.

See “About defending against bounce attacks” on page 182.

To create a seed value to be used when creating validation tags for outgoing messages

1. In the Control Center, select **Administration > Settings > Control Center**.
2. Click the **Certificates** tab.
3. Under **Control Center Certificate**, enter a **Bounce attack prevention seed**.

   This seed value should consist of eight alphanumeric characters.
4. Click **Save**.

Warning: If you are running your inbound and outbound messages on different Scanners with different Control Centers, repeat steps 1 through 3 for each Control Center, using the same seed value. This ensures that all inbound and outbound servers are calculating the same tags for validation.

Note: For successfull processing you must ensure that all of your applicable outbound mail is routed through the appliance.

You must now enable bounce attack prevention for your policy groups and assign a spam policy that describes the actions to be taken when a message does not pass bounce attack validation. If you do not enable at least one policy group for bounce attack prevention, bounce attack prevention will be disabled and your system will not be protected from bounce attacks.

See “Configuring policy groups for bounce attack prevention” on page 184.

A default spam policy is provided, called "Failed Bounce Attack Validation: Reject message". You can use this policy as is, edit it, or create your own policy.

See “Creating an email spam policy for bounce attack prevention” on page 185.

Configuring policy groups for bounce attack prevention

Once you configure bounce attack prevention in the Control Center Settings page, you must enable the policy groups to which the system should apply validation and assign a bounce attack prevention policy.

See “About defending against bounce attacks” on page 182.

To configure policy groups for bounce attack prevention

1. In the Control Center, select **Administration > Users > Policy Groups**.
2. Select the policy group you want to edit, or create a new one, then select the **Spam** tab for that policy group.
3. Under **Email**, check **Enable bounce attack prevention for this policy group**.
4. For the **Bounce attack prevention policy**, select the policy you want to apply to bounced messages.

   This policy must contain the condition, "If a message fails bounce attack validation" and actions to be taken should bounce address tag validation fail. Symantec Brightmail Gateway provides a default policy: "Failed Bounce Attack Validation: Reject message." This default policy is configured to reject messages that fail tag validation.

   You can also edit this policy or create a new one.

   See "Creating an email spam policy for bounce attack prevention" on page 185.
5. Click **Save**.

Note: For successfull processing you must ensure that all of your applicable outbound mail is routed through the appliance.

Creating an email spam policy for bounce attack prevention

In order to enable bounce attack prevention, you must enable your policy groups for bounce attack prevention and assign a spam policy that describes the actions to be taken when a message does not pass bounce attack validation.

See "About defending against bounce attacks" on page 182.

Symantec Brightmail Gateway provides you with a default bounce policy called "Failed Bounce Attack Validation: Reject message". This default policy provides one action, which is to reject all messages that fail tag validation. You can edit this policy to change the actions taken, or you can create a new policy to suit your specific needs.

Create an email spam policy for bounce attack prevention conditions

1. In the Control Center, click **Spam > Policies > Email**.
2. Click **Add** to create a new policy.

3. Enter a name for the new policy, and for **If the following condition is met:** select "If a message fails bounce attack validation".

   The **apply to** field will automatically be set to "inbound messages" and disabled. You can only configure an inbound policy for this condition. The outbound policy is static and cannot be modified.

4. Select the actions that should be applied if a bounce message fails validation. An action "Reject messages failing bounce attack validation" is provided, or you can select any other action as desired.

   Be sure to consider your existing spam policies and how they might affect your overall email configuration.

5. Under **Apply to the following policy groups**, select the policy groups to which you want to apply this policy.

6. Click **Save**.

Researching IP address reputation

Use the IP Reputation Lookup page to research historical and current statistical information about a particular IP address. You can view the sender groups (if any) that currently include the IP address, add the IP address to your Local Good Sender IPs or Local Bad Sender IPs sender groups, or clear the current sender policy of the IP address. Clearing the sender policy removes the IP address from the Local Good Sender IPs or Local Bad Sender IPs group.

You can also reset the local reputation of this IP address (thereby clearing the data used to manage the connections given to that IP address). The page displays data collected since the last time the spam reputation was reset for the specified IP address.

The IP reputation functionality is designed to render verdicts when traffic crosses your organization's gateway, before it enters your network. The full benefit of this feature comes from the ability to reject or defer bad connections before the traffic enters your network and consumes resources. This works best when your Symantec Brightmail Gateway is deployed at the network edge. However, if you have deployed Symantec Brightmail Gateway behind relays or elsewhere within your network, verdicts can still be rendered based on the contents of the received headers.

Note: Historical data for this page is not available unless you have configured Symantec Brightmail Gateway to store this data on an ongoing basis. You can do this on the Administration > Settings > Reports page. Under **Reports Data Storage**, check **Sender IP connections**, and click **Save**.

To check current and historical information for an IP address

1. In the Control Center, click **Reputation > Reputation Tools > IP Reputation Lookup**.
2. Specify the IP address that you want to query.
3. Click **Find**.

The following information is displayed for an IP address that you specify:

Table 7-4 Reputation Status

Item	Description
Add to Local Good Sender IPs	Click to add this IP address to your Local Good Sender IPs group. Connections from IP addresses in Local Good Sender IPs are allowed by default, and the messages bypass spam filtering.
Add to Local Bad Sender IPs	Click to add this IP address to your Local Bad Sender IPs group. Connections from IP addresses in Local Bad Sender IPs are rejected by default.
Clear Sender Policy	Click to remove this IP address from either the Local Bad Sender IPs or Local Good Sender IPs group.
Bad Reputation - Global Bad Senders	If this IP address is in the Symantec Global Bad Senders group, a green checkmark appears. If not, a black dash appears. If the Symantec Global Bad Senders group is disabled on the Bad Senders page, this column is grayed out.
Request Removal	This link appears only if the IP address is in the Symantec Global Bad Senders group. Click the link to request that Symantec remove this IP address from Symantec Global Bad Senders.
Bad Reputation - Local Bad Sender IPs	If this IP address is in the Local Bad Sender IPs group, a green checkmark appears. If not, a black dash appears. If the Local Bad Sender IPs group is disabled on the Bad Senders page, this column is grayed out.

Table 7-4 Reputation Status *(continued)*

Item	Description
Good Reputation - Global Good Senders	If this IP address is in the Symantec Global Good Senders group, a green checkmark appears. If not, a black dash appears. If the Symantec Global Good Senders group is disabled on the Good Senders page, this column is grayed out.
Good Reputation - Local Good Sender IPs	If this IP address is in the Local Good Sender IPs group, a green checkmark appears. If not, a black dash appears. If the Local Good Sender IPs group is disabled on the Good Senders page, this column is grayed out.

Note: An IP address can be in the Local Good Sender IPs group or in the Local Bad Sender IPs group, but not in both. Therefore, if the IP address is currently in one of the two groups, that add button appears grayed out. You can click the button for the other group to switch the IP address from one group to another.

Table 7-5 Local status

Item	Description
Reset Status	After you click and confirm this action, Symantec Brightmail Gateway "forgets" all previous reputation data on this sender. This action clears the **Current Action**, **Message Volume**, and **Total % Spam** columns. However, if the sender is in any of the Local Good Sender IPs, Local Bad Sender IPs, or Fastpass sender groups, the sender remains in those groups. To remove the sender from any of these sender groups, click **Clear Sender Policy**. Clicking **Reset Status** also moves the sender to the default Connection Classification class. See "About managing connection load at the gateway" on page 166.
Scanners	The Scanners that have seen traffic from this IP address. Each Scanner in your installation has its own line.

Table 7-5 Local status *(continued)*

Item	Description
Current Action	This line displays the action that will be taken on connections from this IP address based on the information in the IP Reputation database. Connections from a given IP address can be assigned one of the following actions: ■ Reject: This IP address is in a Bad Sender Group and the action for that group is to reject the SMTP connection ■ Defer: This IP address is in a Bad Sender Group and the action for that group is to defer the SMTP connection. ■ Skip AS Filtering: This IP address is in a Good Sender Group and the action for that group is to deliver the message normally. Although Symantec Brightmail Gateway will not filter messages from this IP address for spam, the messages will undergo all other filtering, including antivirus filtering. ■ Filter Partially: This IP address was granted a pass by the Fastpass feature. In most cases, Symantec Brightmail Gateway does not filter messages from this IP address for spam. However, Symantec Brightmail Gateway does filter a sample of the messages for spam. All the messages undergo all other filtering, including antivirus filtering. ■ Filter Normally: This IP address has no negative local or global reputation, and has not been granted a pass by the Fastpass feature.
Message Volume	The quantity of messages from this IP address. One of the following four values appears: ■ None (this is the default value) ■ Low ■ Medium ■ High
Total % Spam	The percentage of messages from this IP address identified as spam
Connection Class	The connection class (1-9 or Default) assigned to this IP address on this Scanner. Connection Classification defers a higher percentage of connections from IP addresses in higher (worse) classes.
Fastpass	If this IP address currently has a pass granted by the Fastpass feature on this Scanner, a green check appears in this column. If not, a black dash appears.
DHA	If this IP address is under DHA restriction based on your settings in **Reputation > Bad Senders > Directory Harvest Attack**, a green check appears. If not, a black dash appears.

Table 7-5 Local status *(continued)*

Item	Description
Virus Attack	If this IP address is under Virus Attack restriction based on your settings in **Reputation > Bad Senders > Email Virus Attacks**, a green check appears. If not, a black dash appears.
Last Message	The last time traffic was seen from this IP address.

Table 7-6 Local connection history

Item	Description
Time range	Choose the time range for the data displayed in the Local Connection History table. You can choose the past hour, the past day, or the past week.
Attempted	The number of connections from this IP address attempted within the specified time range.
Accepted	The number of connections from this IP address accepted within the specified time range.
Rejected	The number of connections from this IP address rejected within the specified time range.
Deferred	The number of connections from this IP address deferred within the specified time range.

To check sender group membership for a domain, email address, or IP address

1. In the Control Center, click **Reputation > Reputation Tools > Find Sender**.
2. Type a domain, email address, or IP address.
3. Click **Find Sender**.

   If the sender is any of the following groups, the name of the group appears:

   - Local Bad Sender Domains
   - Local Bad Sender IPs
   - Third Party Bad Senders
   - Local Good Sender Domains
   - Local Good Sender IPs
   - Third Party Good Senders

Chapter 8

Setting up certificates

This chapter includes the following topics:

About certificates

Certificates secure and authenticate communications between client and server IP addresses or domains.You can generate a self-signed certificate or import a signed certificate that a Certificate Authority (CA) issues.

For successful HTTPS authentication using a CA certificate, there must be a complete "path" or "chain" from the client certificate to a CA certificate. Additionally, both participants in the negotiation must recognize the signing authority. Symantec Brightmail Gateway includes pre-installed certificates for the most common Certificate Authority vendors. The Certificate Authority tab on the Certificate Settings page lists the pre-installed CA certificates. You can add additional root or intermediate CA certificates. Some certificate issuers require and provide an intermediate CA certificate for the certificates that they issue for additional security.

For SMTP/TLS authentication using a CA certificate, Symantec Brightmail Gateway allows you to use a certificate even if there is not a complete path or chain from the client certificate to a CA certificate.

Symantec Brightmail Gateway supports the following types of certificates:

MTA TLS certificate

The inbound and the outbound mail flows in each Scanner use the TLS certificate that is assigned to them to accept TLS-encrypted messages for scanning.

See "Assigning an MTA TLS certificate to a Scanner" on page 200.

User interface HTTPS certificate

The Control Center uses the HTTPS certificate to secure communications for its Web-based management tools.

See "Assigning a user interface HTTPS certificate to the Control Center" on page 201.

Note: When you purchase or generate a certificate, you must specify whether you intend to use it for SMTP/TLS or HTTPS. A Certificate Authority may require you to import an intermediate CA certificate in addition to the certificate itself. Make sure that you install both the certificate and any intermediate certificate that you receive from the Certificate Authority.

You can add certificates to the list of available certificates in one of the following ways:

- Generate a self-signed certificate by completing the Add Certificate page. The self-signed certificate is immediately available as an HTTPS certificate for the Control Center and for Scanner MTAs for accepting TLS encryption.
- Add a Certificate Authority signed certificate by submitting a certificate request to a Certificate Authority. When you receive the certificate back from the Certificate Authority, save it locally and import it to the Control Center to add it to the list of available certificates.

After you add a certificate, assign it to the Control Center to secure Web-based communications or to a Scanner MTA to support TLS encryption.

Adding a self-signed certificate

A self-signed certificate for HTTPS communication does not offer the same level of security as a CA-signed certificate. A self-signed certificate is not recommended for SMTP/TLS.

See “Adding a CA certificate” on page 193.

See “About certificates” on page 191.

To add a self-signed certificate

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **TLS & HTTPS Certificates** tab.
3. Click **Add**.
4. In the **Certificate name** box, type a name for the certificate.
5. In the **Certificate type** drop-down list, click **Self Signed**.
6. Complete the remainder of information on the page.
7. Click **Create**.

Adding a CA certificate

Symantec Brightmail Gateway includes pre-installed certificates for the most common Certificate Authority vendors. Add a CA certificate if the CA issues you an SMTP/TLS or HTTPS certificate that is not already in the Control Center. Another reason to add a CA certificate is if your certificate requires an intermediate CA certificate. When you add a CA certificate, you complete the certificate chain to permit authentication of the new certificate. All of your configured Scanners can access the CA certificates in the Control Center for SMTP/TLS and HTTPS authentication.

Ensure that you have the CA certificate before you proceed. The CA certificate may have been included when you received the certificate from the CA. Alternatively, you may be able to download the CA certificate from the Certificate Authority's Web site. The file that contains the CA certificate must be in PEM format.

When you receive a certificate from a Certificate Authority (CA), you must import it to make it available in the Control Center.

See “Importing a Certificate Authority-signed certificate” on page 197.

See “About certificates” on page 191.

See “Requesting a Certificate Authority-signed certificate” on page 194.

See “PEM format requirements for certificates and domain keys” on page 196.

To add a CA certificate bundle

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **Certificate Authority** tab.
3. Click **Update**.
4. On the Update CA Certificates page, click **Browse**.
5. Locate and select the file that contains the CA certificate.
6. Click **Update**.

   A status message appears at the top of the page to indicate success or failure.

Requesting a Certificate Authority-signed certificate

A Certificate Authority (CA)-signed certificate provides more security than a self-signed certificate and is appropriate for HTTPS and SMTP/TLS communication. Before you proceed, determine the CA from which you want to purchase your certificate. Some possible CAs to use are listed on the Certificate Authority tab in the Control Center. However, other CAs are also supported.

For the common name, use the domain name or the fully qualified domain name of the computer where the certificate will be installed. Some CAs may not support certificates that are created using an IP address instead of a domain name for the common name. Check with your CAs.

See “Viewing existing CA certificates” on page 198.

See “About certificates” on page 191.

Each CA has its own set of procedures to request certificates and issue certificates. Consult your CA for details and follow the instructions that are appropriate for your installation.

See “Adding a CA certificate” on page 193.

To request a Certificate Authority-signed HTTPS or TLS certificate

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **TLS & HTTPS Certificates** tab.
3. Click **Add**.

4 In the **Certificate name** box, type a name for this certificate.

5 In the **Certificate type** drop-down list, click **Certification Authority Signed**.

6 Fill in the information on the remainder of the page as appropriate.

7 Click **Create**.

8 Copy the block of text that appears, paste it into a text file, and save it.

Save the generated text as a text file. You can copy and paste the information from the text file into a CA request form at a later time.

9 Submit the Certificate Authority-Signed Request (CSR) to a CA, using the method that the CA requires.

Changing a self-signed certificate or domain key name

You can change a self-signed certificate or domain key name, but you cannot modify any other part of a certificate. To change another part of a certificate, you must create a new certificate.

See “Adding a self-signed certificate” on page 193.

See “Adding a domain key” on page 138.

You cannot change the name of a Certificate Authority-signed certificate.

To change a self-signed certificate name

1 Click **Administration > Settings > Certificates**.

2 Click the **TLS & HTTPS Certificates** tab.

3 Check the box beside the certificate that you want to modify.

4 Click **Edit**.

5 On the **Edit Certificate** page, type the new name of the certificate in the **Certificate name** field.

6 Click **Save**.

To change a domain key name

1 Click **Administration > Settings > Certificates**.

2 Click the **Domain Keys** tab.

3 Check the box beside the domain key that you want to modify.

4 Click **Edit**.

5 On the **Edit Domain Key** page, type the new name of the certificate in the **Domain Key name** field.

6 Click **Save**.

PEM format requirements for certificates and domain keys

When you add a certificate, whether self-generated or Certificate Authority-signed, and when you import a domain key, ensure that the certificate or domain key meets the following requirements:

- The certificate or domain key must be stored in a file in PEM format with the certificate or domain key included as Base64-encoded text between the following markers:
  For a certificate, `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----`.
  For a domain key, `-----BEGIN PUBLIC KEY-----` and `-----END PUBLIC KEY-----`.
  Any text outside of the begin and end markers is ignored.
  The formats for certifcates and domain keys are identical, except for the beginning and ending markers.
  Base64 text consists of only uppercase and lowercase Roman alphabet characters (A–Z, a–z), the numerals (0–9), and the "+" and "/" symbols.
- The file must be encoded as US-ASCII or UTF. The file cannot contain extended ASCII or non-ASCII characters.
- When you add or replace CA certificates (Update or Restore), a file can contain multiple certificates.
- The extension of the file that contains the certificate or domain key does not matter. The .txt or .crt extension are typically used for certificates, and the .key extension is typically used for domain keys.
- The file that contains the certificate or domain key must be accessible from the browser that you use to access the Control Center.

The following is a sample PEM format CA certificate:

```
Text before Begin Certificate is ignored.
-----BEGIN CERTIFICATE-----
MIICPTCCAaYCEQDNun9W8N/kvFT+IqyzcqpVMA0GCSqGSIb3DQEBAgUAMF8xCzAJ
BgNVBAYTAlVTMRcwFQYDVQQKEw5WZXJpU2lnbiwgSW5jLjE3MDUGA1UECxMuQ2xh
```

```
c3MgMSBQdWJsaWMgUHJpbWFyeSBDZXJ0aWZpY2F0aW9uIEF1dGhvcml0eTAeFw05
NjAxMjkwMDAwMDBaFw0yODA4MDEyMzU5NTlaMF8xCzAJBgNVBAYTAlVTMRcwFQYD
VQQKEw5WZXJpU2lnbiwgSW5jLjE3MDUGA1UECxMuQ2xhc3MgMSBQdWJsaWMgUHJp
bWFyeSBDZXJ0aWZpY2F0aW9uIEF1dGhvcml0eTCBnzANBgkqhkiG9w0BAQEFAAOB
jQAwgYkCgYEA5Rm/baNWYS2ZSHH2Z965jeu3noaACpEO+jglr0aIguVzqKCbJF0N
H8xlbgyw0FaEGIeaBpsQoXPftFg5a27B9hXVqKg/qhIGjTGsf7A01480Z4gJzRQR
4k5FVmkfeAKA2txHkSm7NsljXMXg1y2He6G3MrB7MLoqLzGq7qNn2tsCAwEAATAN
BgkqhkiG9w0BAQIFAAOBgQBMP7iLxmjf7kMzDl3ppssHhE16M/+SG/Q2rdiVIjZo
EWx8QszznC7EBz8UsA9P/5CSdvnivErpj82ggAr3xSnxgiJduLHdgSOjeyUVRjB5
FvjqBUuUfx3CHMjjt/QQQDwTw18fU+hI5Ia0e6E1sHslurjTjqs/OJ0ANACY89Fx
lA==
-----END CERTIFICATE-----
Text after End Certificate is ignored.
```

When you add a domain key, Symantec Brightmail Gateway generates the domain key in a way that meets PEM format requirements.

See “Adding a CA certificate” on page 193.

See “Adding a self-signed certificate” on page 193.

See “Importing a domain key” on page 140.

See “Adding a domain key” on page 138.

Importing a Certificate Authority-signed certificate

When you receive a certificate from a Certificate Authority (CA), you must import it to make it available in the Control Center. In addition to the certificate, the CA might have sent you an intermediate certificate that you also need to install in the Control Center.

See “Adding a CA certificate” on page 193.

To import a Certificate Authority signed certificate

1. When you receive the certificate file from the CA, save the file to a location that you can access from the Control Center.

   In some cases you may need to store more than one file, depending on your Certificate Authority's requirements.

2. In the Control Center, click **Administration > Settings > Certificates**.

3. Click the **TLS & HTTPS Certificates** tab.

4. Click **Import**.

5. On the **Import Certificate** page, type the full path and the file name of the certificate or click **Browse** and choose the file.
6. Click **Import**.

Replacing existing CA certificates

You can replace existing CA certificates in the Control Center with another set of CA certificates. All existing CA certificates are removed and replaced with the CA certificates in the file that you specify.

Ensure that you have the CA certificates before you proceed. The file that contains the CA certificates must be in PEM format.

See "PEM format requirements for certificates and domain keys" on page 196.

To replace existing CA certificates

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **Certificate Authority** tab.
3. Click **Restore**.
4. On the **Restore CA Certificates** page, click **Browse** and locate the file that contains the CA certificates.
5. Click **Restore**.

   A status message appears at the top of the page to indicate success or failure.

Viewing existing CA certificates

You can view a list of the currently installed CA certificates.

To view existing CA certificates

1. In the Control Center, click **Administration > Settings > Certificates**.
2. Click the **Certificate Authority** tab.

   The currently installed CA certificates appear.

Backing up CA certificates

You can back up the CA certificates that are installed in the Control Center.

See "Viewing existing CA certificates" on page 198.

To back up CA certificates

1 In the Control Center, click **Administration > Settings > Certificates**.

2 Click the **Certificate Authority** tab.

3 Click **Backup**.

4 In the browser **File Download** dialog box, click **Save** the file and then specify the file location.

The file may be saved to your default browser download directory or in a location that you specify.

Deleting an SMTP/TLS or HTTPS certificate

You can view or delete a certificate. You cannot delete a certificate that is in use for SMTP/TLS or HTTPS authentication.

See “Viewing an SMTP/TLS or HTTPS certificate” on page 199.

To delete an SMTP/TLS or HTTPS certificate

1 Click **Administration > Settings > Certificates**.

2 Click the **TLS & HTTPS Certificates** tab.

3 Check the box next to the certificate that you want to delete.

4 Click **Delete**.

Viewing an SMTP/TLS or HTTPS certificate

You can view any of your SMTP/TLS or HTTPS certificates.

See “Viewing existing CA certificates” on page 198.

To view an SMTP/TLS or HTTPS certificate

1 In the Control Center, click **Administration > Settings > Certificates**.

2 Click the **TLS & HTTPS Certificates** tab.

3 Check the box beside the certificate that you want to view.

4 Click **View**.

Assigning an MTA TLS certificate to a Scanner

You can assign an MTA TLS certificate to a Scanner. You need to do this if you want the Scanner to accept TLS-encrypted messages for scanning of inbound or outbound mail.

See "About certificates" on page 191.

You can also assign a certificate for SMTP authentication.

See "Configuring SMTP authentication mail settings" on page 149.

You can use a self-signed or CA-signed certificate. You may also need to install an intermediate CA certificate. For SMTP/TLS authentication, Symantec Brightmail Gateway allows you to use a CA certificate even if there is not a complete path or chain from the client certificate to a CA certificate.

See "Requesting a Certificate Authority-signed certificate" on page 194.

See "Adding a CA certificate" on page 193.

To assign an MTA TLS certificate to a Scanner

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Check the box beside the host that you want, and click **Edit**.
3. Click the **SMTP** tab.
4. Under either**Inbound Mail Settings** or **Outbound Mail Settings**, check **Accept TLS encryption** if you want this Scanner to scan for inbound or outbound TLS-encrypted email, respectively.
5. In the adjacent drop-down list, choose the MTA TLS certificate that is appropriate to the inbound mail flow or outbound mail flow.

   You can assign the same certificate to both inbound and outbound TLS-encrypted email filtering.
6. Check **Request client certificate** if you want the connecting client to present a TLS certificate for authentication.

   This step is required for the inbound mail setting only.

   You may need to install an intermediate CA certificate to authenticate the connecting client's TLS certificate.
7. Click **Save**.

Assigning a user interface HTTPS certificate to the Control Center

You can assign a user interface HTTPS certificate to the Control Center. You can use either a self-signed certificate or a CA-signed certificate. If you use a CA-signed certificate, you may need to install an intermediate CA certificate.

See "Adding a CA certificate" on page 193.

To assign a user interface HTTPS certificate to the Control Center

1. In the Control Center, click **Administration > Settings > Control Center**.
2. Under **Control Center Validation**, click the **User interface HTTPS certificate** drop-down list and select a certificate.
3. Click **Save**.

Chapter 9

Detecting viruses and malicious attacks

This chapter includes the following topics:

About detecting viruses and malicious attacks

Viruses and other types of malicious attacks can wreak havoc in an organization. The damage can range from email server crashes to network downtime and the compromise and destruction of company data. Given the damage that can result from viruses and other types of malicious attacks, it is essential to employ virus protection as early in the mail flow process as possible. Virus and malicious threat detection is optional.

See "How to detect virus and malicious threat detection" on page 208.

Create virus policies to protect your server from the following types of attacks:

Viruses

Symantec Brightmail Gateway detects viruses, worms, and Trojan horses in all major file types (for example, Microsoft Word files), including compressed file formats.

See "Product technologies that detect viruses and malicious attacks" on page 205.

Mass-mailer worms

Symantec Brightmail Gateway detects that an email message is a mass-mailer worm or virus. It can automatically delete the infected email message and any attachments.

Suspicious attachments

Symantec Brightmail Gateway detects the email messages that exhibit virus-like signs. It also detects the messages that contain a suspicious new pattern of message flow that involves email message attachments.

Encrypted attachments

Infected files can be intentionally encrypted. Encrypted files cannot be decrypted and scanned without the appropriate decryption tool. You can configure how you want Symantec Brightmail Gateway to process encrypted container files.

Adware and spyware

Symantec Brightmail Gateway detects the security risks that do any of the following:

- Provide unauthorized access to computer systems
- Identity theft or fraud by logging keystrokes
- Capture email and instant messaging traffic
- Harvest personal information, such as passwords and logon identifications
- Present some type of disruption or nuisance

See "Spyware or adware verdict details" on page 207.

See "Creating email virus policies" on page 210.

Symantec Brightmail Gateway must be able to decompose and scan a container file to detect viruses. You can specify the maximum size and scanning depth levels of container files to reduce your exposure to zip bombs or denial-of-service attacks.

See "Setting limits on nested files" on page 220.

When Symantec Brightmail Gateway scans a message and detects a virus policy violation, it takes the verdict that you specify in that policy.

See "Selecting virus policies for a policy group" on page 322.

You must have a valid antivirus license to perform antivirus scanning functions and to obtain updated virus definitions.

See "Licensing your product" on page 678.

Product technologies that detect viruses and malicious attacks

Table 9-1 describes the technologies that Symantec Brightmail Gateway uses to detect viruses and malicious attacks.

Table 9-1 Technologies that detect viruses and malicious attacks

Technology	Description
Antivirus engine	The antivirus engine provides rapid and reliable virus protection through a multi-threaded scanning system. It scans incoming and outgoing email traffic. It identifies and cleans the messages that contain viruses and related malicious executables. It also attempts to repair viruses within email attachments. The antivirus engine itself cannot be modified.
Heuristics technology	The product uses Symantec Bloodhound heuristic technology to detect virus-like behavior to identify and repair unknown viruses. You can adjust heuristic settings for more or less aggressive identification of viruses. The technology detects up to 90 percent of new macro viruses and up to 80 percent of new and unknown executable file viruses. You can modify the heuristics detection level. See "Modifying the heuristic level" on page 219.

Table 9-1 Technologies that detect viruses and malicious attacks *(continued)*

Technology	Description
Virus definitions	Virus definitions are available every hour to protect against the latest, fast-spreading threats. Symantec LiveUpdate is the process by which the appliance receives current virus definitions from Symantec Security Response. By default, the appliance downloads certified virus definitions. However, you can obtain more frequent, less tested Rapid Response definitions. You can also obtain certified daily Platinum definitions for faster response to emerging threats. You can configure how and when you want to obtain updated definitions. See "About updating virus definitions" on page 221.
Antivirus policies	You can create policies to detect viruses or malicious attacks. When you create a policy, you specify the action that you want Symantec Brightmail Gateway to take if the policy is violated. For example, you can clean infected attachments, but delete spyware attachments entirely. You can create as many antivirus policies as needed. See "Creating email virus policies" on page 210.
Day-zero detection	This feature leverages the Symantec view of email threats as well as heuristic analysis to identify a suspicious attachment before antivirus definitions are available. Messages that contain suspicious attachments can be moved to the Suspect Virus Quarantine. Symantec Brightmail Gateway holds the message in the quarantine for the period of time that you specify (up to 24 hours). It then releases the message to be scanned again with updated virus definitions. You can create the virus policies that contain verdicts to quarantine suspect message attachments. You can also configure how long an attachment remains in the Suspect Virus Quarantine. See "About quarantining suspected viruses" on page 229.

See "About detecting viruses and malicious attacks" on page 204.

See "Product technologies that detect viruses and malicious attacks" on page 205.

What you can do with suspicious attachments

When you create a policy and choose the condition, "If a message contains a suspicious attachment," additional options become available as follows:

Hold message in Suspect Virus Quarantine	Select this option to quarantine the message and all attachments.
Strip and Delay in Suspect Virus Quarantine	Select this option to delete the suspicious attachment and quarantine the message. When you select this option, the suspicious attachment cannot be retrieved.

Both of these actions move the message to quarantine. After the amount of time that you specify, the messages are rescanned. This time, however, the messages are scanned with the newest definitions that are available.

See "About quarantining suspected viruses" on page 229.

See "Creating email virus policies" on page 210.

Spyware or adware verdict details

Symantec Brightmail Gateway can detect security risks.

Security risks are the programs that do any of the following:

- Provide unauthorized access to computer systems
- Compromise data integrity, privacy, confidentiality, or security
- Present some type of disruption or nuisance

Symantec Brightmail Gateway applies the spyware or adware verdict to all of the security risks that it detects.

See "About detecting viruses and malicious attacks" on page 204.

Table 9-2 lists the categories of security risks that Symantec Brightmail Gateway detects.

Table 9-2 Security risk categories included in spyware or adware verdict

Category	Description
Adware	Standalone or the appended programs that gather personal information through the Internet and relay it back to a remote computer without the user's knowledge. Adware might monitor browsing habits for advertising purposes. It can also deliver advertising content.

Table 9-2 Security risk categories included in spyware or adware verdict *(continued)*

Category	Description
Hacking tools	Programs that are used to gain unauthorized access to a user's computer. For example, a keystroke logger tracks and records individual keystrokes and sends this information to a remote computer. The remote user can perform port scans or vulnerability scans. Hacking tools can also be used to create viruses.
Dialers	Programs that use a computer, without the user's permission or knowledge, to dial out through the Internet to a 900 number or FTP site. These programs typically to accrue charges.
Joke programs	Programs that alter or interrupt the operation of a computer in a way that is intended to be humorous or bothersome. For example, a joke program might move the Recycling Bin away from the mouse when the user tries to click on it.
Remote access programs	Programs that let a remote user gain access to a computer over the Internet to gain information, attack, or alter the host computer.
Spyware	The Standalone programs that can secretly monitor system activity. These programs can and detect passwords and other confidential information and then relay the information back to a remote computer.

How to detect virus and malicious threat detection

Table 9-3 describes the tasks that you can perform to detect viruses and malicious threats; you can perform any or all of the tasks in any order.

Table 9-3 Detecting virus and malicious threat detection

Task	Description
Email virus attack recognition.	In an email virus attack, a specified quantity of infected email messages has been received from a particular IP address. By default, any connections that are received from violating senders are deferred. Email virus attack recognition is disabled by default and must be enabled to be activated. See "Configuring email virus attack recognition" on page 170.

Table 9-3 Detecting virus and malicious threat detection *(continued)*

Task	Description
Create and enable email virus policies.	Symantec Brightmail Gateway comes with the pre-configured virus policies that are automatically enabled. You can modify these polices and create your own custom policies. See "Default email virus policies" on page 52. See "Creating email virus policies" on page 210.
Set the heuristic detection level.	Symantec Brightmail Gateway contains Symantec Bloodhound heuristics technology. This technology scans for unusual behaviors (such as self-replication) to target potentially infected message bodies and attachments. The default setting is Medium. However, you can modify this setting or turn Bloodhound off. Bloodhound heuristics involve a trade-off between higher virus detection rates and the speed with which Symantec Brightmail Gateway processes mail. Lower heuristic levels may miss more viruses but require less processing power. Higher heuristic levels may catch more viruses but consume more processing power. See "Modifying the heuristic level" on page 219.
Specify the file types that can bypass antivirus scanning.	You can specify the file types that can bypass antivirus scanning. For example, certain file types typically do not contain viruses, such as .mpg files. File types that you feel confident do not contain viruses can bypass virus scanning, which saves system resources. Symantec Brightmail Gateway provides a default list of file type categories. But you must create Exclude Scanning Lists, select the categories that you want to include, and enable the list. You can also add and remove file types from Exclude Scanning Lists. See "Excluding file types from virus scanning" on page 214.

Table 9-3 Detecting virus and malicious threat detection *(continued)*

Task	Description
Configure the Suspect Virus Quarantine.	You can create virus policies to quarantine suspicious message attachments in the Suspect Virus Quarantine. Symantec provides default values for the following Suspect Virus Quarantine settings; however, you can change these settings as needed: ■ Maximum amount of the time that messages are held in the quarantine The default setting is 6 hours. ■ Disk space available for the quarantine The default setting is 10 GB. See "About quarantining suspected viruses" on page 229.
Enable definition updates.	By default, LiveUpdate is enabled. Platinum definition updates are scheduled to occur every 10 minutes from Monday through Friday. However, you modify when and how you want to obtain updates. See "About updating virus definitions" on page 221.
Configure outbreak notification alerts.	Set up alert notifications to let you know any of the following virus-related events occur: ■ An outbreak is detected ■ Virus filters are older than the time period that you specify ■ New virus filters are available ■ The antivirus license has expired See "Types of alerts" on page 615.
Monitor reports.	Monitor reports to determine how effective virus detection and policies are. Reports also indicate the volume of threats that your organization receives. This information can help you fine-tune your antivirus detection and threat detection settings. See "About working with reports" on page 568.

Creating email virus policies

Symantec Brightmail Gateway installs with several preconfigured virus policies that are enabled by default. In addition to these policies, you can create your own custom policies. Content filtering, spam, and virus policy names must be unique.

For example, if you have a content filtering policy called XYZ, you cannot have a spam policy or virus policy called XYZ. Email virus policies are enabled by default when you create them.

See "Default email virus policies" on page 52.

See "Modifying email virus policies" on page 212.

See "Enabling or disabling email virus policies" on page 213.

See "Deleting email virus policies" on page 213.

See "Copying email virus policies" on page 214.

To create email virus policies

1. In the Control Center, click **Virus > Policies > Email**.
2. Click **Add**.
3. In the **Policy name** box, type a name for the virus policy.
4. Under **Conditions**, click the **Apply to** drop-down list and choose to which type of message the virus policy should apply:
   - Inbound messages
   - Outbound messages
   - Inbound and Outbound messages
     This option specifies where this virus policy is available on the **Virus** tab when you configure a policy group.
     For example, assume that you choose **Inbound messages** and the **If the message contains a mass-mailing worm** condition. This virus policy is only available in the **Inbound mass-mailing worm policy** list when you configure a policy group.
5. Click the **If the following condition is met** drop-down list to select a condition.

   See "What you can do with suspicious attachments" on page 206.
6. Under **Actions**, click the **Perform the following action** drop-down list and select an action.

   See "Selecting virus policies for a policy group" on page 322.

   For some actions you need to specify additional information in the fields beneath the action.

   For example, assume that you select the action to **Forward a copy of the message**. A box appears in which you can type of the email address of the person to which you want to forward the message.

7 Click **Add Action** to add more actions, if needed.

See "User interface action combinations" on page 728.

8 Under **Policy Groups**, check one or more groups to which this policy should apply.

9 Click **Save**.

Modifying email virus policies

You can modify default and custom email virus policies to fine-tune them or to expand or reduce their scope.

Table 9-4 describes the setting that you can modify for default email virus policies and custom email virus policies.

Table 9-4 Modifiable email virus policy settings

Setting	Modifiable in default policies	Modifiable in custom policies
Policy name	No	Yes
Apply to	No	Yes
If the following condition is met	Yes	Yes
Perform the following action	Yes	Yes
Apply to the following groups	Yes	Yes

See "Default email virus policies" on page 52.

See "Enabling or disabling email virus policies" on page 213.

See "Deleting email virus policies" on page 213.

See "Copying email virus policies" on page 214.

To modify email virus policies

1 In the Control Center, click **Virus > Policies > Email**.

2 Check the box beside the policy that you want to modify, and then click **Edit**.

3 Make the modifications that you want.

4 Click **Save**.

Enabling or disabling email virus policies

Default email virus policies are enabled by default. When you create a new email virus policy, it is enabled by default. You can disable any policy that you do not want Symantec Brightmail Gateway to use when it scans email messages.

You can disable a virus policy to troubleshoot email virus scanning issues. You can also create custom email virus policies when an outbreak occurs and then disable the policies when the outbreak ends. You can turn on the policy in the event of another outbreak.

You can also disable the policies that you no longer want to use but do not want to delete yet.

See "Creating email virus policies" on page 210.

See "Default email virus policies" on page 52.

See "Modifying email virus policies" on page 212.

See "Deleting email virus policies" on page 213.

See "Copying email virus policies" on page 214.

To enable or disable email virus policies

1. In the Control Center, click **Virus > Policies > Email**.
2. Check the box beside the policy that you want to enable or disable.
3. Click one of the following options:

Enable	When you enable a policy, a green check mark appears in the Enabled column.
Disable	When you disable a policy, a horizontal line appears in the Enabled column.

Deleting email virus policies

You can delete the email virus policies that you no longer need. However, when you delete a policy, the policy cannot be retrieved. If you are unsure if you want to permanently delete a policy, you can disable it instead.

See "Enabling or disabling email virus policies" on page 213.

See "Modifying email virus policies" on page 212.

See "Copying email virus policies" on page 214.

To delete email virus policies

1. In the Control Center, click **Virus > Policies > Email**.
2. Check the box beside the policy that you want to delete.
3. Click **Delete**.
4. Click **OK** in the confirmation dialog box.

Copying email virus policies

You may have instances in which you create a complicated email antivirus policy and want to create a similar policy with only a few variances. Symantec Brightmail Gateway lets you copy email virus policies.

When you copy an email virus policy, the new policy must have a unique name. For example, if you have a content filtering policy called XYZ, you cannot have a spam policy or virus policy called XYZ. Email virus policies are enabled by default when you create them.

See "Creating email virus policies" on page 210.

See "Default email virus policies" on page 52.

See "Modifying email virus policies" on page 212.

See "Enabling or disabling email virus policies" on page 213.

See "Deleting email virus policies" on page 213.

To copy email virus policies

1. In the Control Center, click **Virus > Policies > Email**.
2. Check the box beside the policy that you want to copy.
3. Click **Copy**.
4. On the **Add Email Virus Policies** page, type a new name for the policy.
5. Make any other changes you want.
6. Click **Save**.

Excluding file types from virus scanning

Viruses and malicious threats are typically contained in executable file types. You can improve scanning performance by letting the file types that do not contain executables bypass scanning.

Symantec Brightmail Gateway lets you create the lists that contain the file types that can bypass antivirus scanning. Multiple lists can help you categorize the file types.

When you create a list, select from the following preconfigured lists:

File classes

Symantec Brightmail Gateway provides preconfigured file classes from which to choose.

An example of some file classes are as follows:

- Movie file
- Library format
- Desktop publishing

The All File Classes file class contains all of the file types in all of the file classes.

File types

When you select a file class, the preconfigured file types for that class appear.

You can select all of the file types for a file class to include in the **Exclude scanning** list. Or you can select only the file types that you want to include in the list.

The list must be enabled if you want it to be used during scans. When you create a new list, it is enabled by default.

See “Modifying the file types to exclude from scanning” on page 216.

See “Enabling or disabling the lists of file types to exclude from scanning” on page 217.

See “Deleting lists of file types to exclude from scanning” on page 218.

See “Exporting lists of the file types to exclude from scanning” on page 218.

To exclude file types from virus scanning

1. In the Control Center, click **Virus > Settings > Scan Settings**.
2. Click the **Exclude Scanning** tab.
3. Click **Add** to create a definition of files for exclusion from virus scanning.
4. In the **Exclude scanning list name** box, type a name for the list.
5. In the **File classes** list, select the file class that you want to exclude from scanning.

   To select multiple classes, hold down the **Ctrl** key while you click the names of file classes.

6 Perform one of the following tasks:

To exclude all file types in a class or group of classes	Click **Add File Classes**.
To exclude only certain file types	Click the file classes from which you want to exclude specific file types. Then select the file types from the **File Types** list. Hold down **Ctrl** to select multiple file types. Click **Add File Types**.

7 Click **Save**.

The names of the file types appear in the **Description** list.

Modifying the file types to exclude from scanning

You can modify the file types that you want to exclude from scanning as you fine-tune your virus scanning policies. For example, you may find that some malicious executables can be modified to look like harmless file types. So you want to remove the file types that can be spoofed from the list of files to be excluded from scanning.

See "Excluding file types from virus scanning" on page 214.

See "Enabling or disabling the lists of file types to exclude from scanning" on page 217.

See "Deleting lists of file types to exclude from scanning" on page 218.

See "Exporting lists of the file types to exclude from scanning" on page 218.

To modify the file types to exclude from scanning

1 In the Control Center, click **Virus > Settings > Scan Settings**.

2 Click the **Exclude Scanning** tab.

3 Check the box beside the Exclude Scanning List that you want to modify.

4 Click **Edit**.

5 Do any of the following tasks:

To add a file type to the list	In the **File classes** list, select the file class that you want to exclude from scanning. To select multiple classes, hold down the **Ctrl** key while you click the names of file classes. Click the file classes from which you want to exclude specific file types. Then select the file types from the **File Types** list. Hold down **Ctrl** to select multiple file types. Click **Add File Types**.
To remove a file type from the list	Under **Description**, check the box beside the file type that you want to delete, and then click **Delete**.

6 Click **Save**.

Enabling or disabling the lists of file types to exclude from scanning

You can have multiple lists of file types to exclude from scanning. However, the file types that are included in the list do not bypass antivirus scanning unless the list is enabled. A list is enabled by default when you create it. Disable any list that you do not want Symantec Brightmail Gateway to use when it scans email messages.

You can disable a list to troubleshoot email virus scanning issues. You can also create custom a list to use for an outbreak and disable the list when the outbreak ends. You can enable the list again in the event of another outbreak.

You can also disable the lists that you no longer want to use but do not want to delete yet.

See “Excluding file types from virus scanning” on page 214.

See “Modifying the file types to exclude from scanning” on page 216.

See “Deleting lists of file types to exclude from scanning” on page 218.

See “Exporting lists of the file types to exclude from scanning” on page 218.

To enable or disable the lists of file types to exclude from scanning

1 In the Control Center, click **Virus > Settings > Scan Settings**.

2 Click the **Exclude Scanning** tab.

3 Check the box beside the policy that you want to enable to disable.

4 Select one of the following options:

Enable	When you enable a policy, a green check mark appears in the **Enabled** column.
Disable	When you disable a policy, a horizontal line appears in the **Enabled** column.

Deleting lists of file types to exclude from scanning

You can delete a list of file types to exclude from scanning that you no longer need. However, when you delete a list, the contents of the list cannot be retrieved. If you are unsure if you want to permanently delete a list, you can disable it instead. Before you delete the list, you may want to export it and maintain the list for your records.

See "Excluding file types from virus scanning" on page 214.

See "Modifying the file types to exclude from scanning" on page 216.

See "Enabling or disabling the lists of file types to exclude from scanning" on page 217.

See "Exporting lists of the file types to exclude from scanning" on page 218.

To delete lists of file types to exclude from scanning

1 In the Control Center, click **Virus > Settings > Scan Settings**.

2 Click the **Exclude Scanning** tab.

3 Check the box beside the list that you want to delete.

4 Click **Delete**.

5 Click **Delete** in the confirmation dialog box.

Exporting lists of the file types to exclude from scanning

You can export lists of the file types to exclude from scanning to a text file.

You may want to export file types for any of the following reasons:

- To have a list that you can distribute throughout your organization to those who need to know the file types that bypass scanning.

- To have an archive copy of a file type list before you copy the list, modify the list, or delete the list.

See “Excluding file types from virus scanning” on page 214.

See “Modifying the file types to exclude from scanning” on page 216.

See “Enabling or disabling the lists of file types to exclude from scanning” on page 217.

See “Deleting lists of file types to exclude from scanning” on page 218.

To export lists of file types to exclude from scanning

1. In the Control Center, click **Virus > Settings > Scan Settings**.
2. Click the **Exclude Scanning** tab.
3. Put a check beside the list that contains the file types that you want to export.
4. Click **Export**.
5. In the confirmation dialog box, specify whether you want to open the file or save it.

Modifying the heuristic level

The heuristic level determines the way in which the system uses heuristics to detect viruses. Symantec Brightmail Gateway uses Symantec Bloodhound heuristics technology to scan for threats for which no known definitions exist. Bloodhound heuristics technology scans for unusual behaviors (such as self-replication) to target potentially infected message bodies and attachments. Bloodhound technology can detect up to 80 percent of new and unknown executable file threats. Bloodhound-Macro technology detects and repairs over 90 percent of new and unknown macro viruses.

Bloodhound requires minimal overhead because it examines only message bodies and the attachments that meet stringent prerequisites. In most cases, Bloodhound determines in microseconds whether a message or attachment is likely to be infected. If it determines that a file is not likely to be infected, it moves to the next file.

Bloodhound heuristics involve a trade-off between higher virus-detection rates and the speed with which Symantec Brightmail Gateway processes mail. Lower heuristic levels may miss more viruses but require less processing power. Higher heuristic levels may catch more viruses but consume more processing power.

See “About detecting viruses and malicious attacks” on page 204.

See “How to detect virus and malicious threat detection” on page 208.

See “Product technologies that detect viruses and malicious attacks” on page 205.

To modify the heuristic level

1. Click **Virus > Settings > Scan Settings**.
2. Click the **General** tab.
3. Under **Bloodhound Level**, click the level that you want.

   The default setting is Medium.
4. Click **Save**.

Setting limits on nested files

When Symantec Brightmail Gateway processes certain compressed files, these files can expand to the point where they deplete system memory. Such container files are often referred to as “zip bombs.” Symantec Brightmail Gateway can handle such situations by automatically sidelining large attachments and stripping the attachments. It assumes that such a file can be a zip bomb and should not be allowed to deplete system resources. Action is taken on the file only because of its size, not because of any indication that it contains a virus or other violation.

You can specify this size threshold and the maximum extraction level that Symantec Brightmail Gateway processes in memory. You can also specify a time limit for scanning containers. If a configured limit is reached, Symantec Brightmail Gateway performs the action that you specify for the **Unscannable for viruses** category.

The following table describes at what threshold a container is unscannable for each option that you can configure:

Maximum container scan depth	The nested depth in a container file (such as a .zip file or email message) exceeds the number specified. Do not set this value too high. You can be vulnerable to denial-of-service attacks or zip bombs, which contain many levels of nested files.
Maximum time to open container	The specified time elapses during a scan of container attachments (such as .zip files). Use this setting to detect the containers that do not exceed the other container settings, but include container nesting, many files, large files, or a combination of these.
Maximum individual file size when opened	Any individual component of the container exceeds the size that is specified when unpacked.

Maximum accumulated file size when opened	The total size of all the files in a container exceeds the size that is specified when unpacked.

See "About detecting viruses and malicious attacks" on page 204.

See "How to detect virus and malicious threat detection" on page 208.

See "Product technologies that detect viruses and malicious attacks" on page 205.

To set limits on nested files

1. In the Control Center, click **Protocols > SMTP > Settings**.
2. Under **Container Limits**, in the **Maximum container scan depth** box, type the maximum number of container depths.
3. In the **Maximum time to open container** box, type a value, and then click the drop-down menu to specify the **Seconds**, **Minutes**, or **Hours**.
4. In the **Maximum individual file size when opened** box, type the maximum file size, and then click the drop-down menu to select **KB**, **MB**, or **GB**.
5. In the **Maximum accumulated file size when opened** box, type the maximum accumulated file size, and then click the drop-down menu to select **KB**, **MB**, or **GB**.
6. Click **Save**.

About updating virus definitions

Symantec Brightmail Gateway relies on up-to-date information to detect viruses and threats. One of the most common reasons that problems occur is that virus definition files are not up-to-date. Symantec regularly supplies the updated virus definition files that contain the necessary information about all newly discovered viruses and threats. Regular updates of that information maximize security and guard your organization against infections and the downtime that is associated with an outbreak.

Every 10 minutes, Symantec Brightmail Gateway polls Symantec servers to see if updates are available. If they are, that information appears on the **Status** page. However, updates are only downloaded based on the schedule that you specify. Or you can download definition updates manually.

Table 9-5 lists the methods that you can use to obtain updated virus definitions from Symantec.

Table 9-5 Methods to obtain updated virus definitions from Symantec

Method	Description
LiveUpdate	You can use LiveUpdate to automatically update your protection. When LiveUpdate runs, it downloads and installs available definitions. You can configure LiveUpdate to run on a scheduled basis, or you can run it on demand. LiveUpdate is enabled by default to update definitions Monday through Friday every 10 minutes. See "Scheduling automatic virus definition updates" on page 223. See "Initiating virus definition updates on demand" on page 225.
Rapid Response	You can use Rapid Response when you need quick responses to emerging threats. Rapid Response definitions are most useful for a perimeter defense to mitigate quickly spreading threats. Rapid Response is an alternative to LiveUpdate. Rapid Response definitions are created when a new threat is discovered. Rapid Response definitions undergo the basic quality assurance testing by Symantec Security Response. They do not undergo the intense testing that is required for a LiveUpdate release. Symantec updates Rapid Response definitions as needed. They respond to high-level outbreaks and might be made available before the LiveUpdate definitions quality assurance process is complete. They can be augmented later on by more robust detection capabilities in certified definitions. Rapid Response definitions are automatically updated every 10 minutes. You cannot schedule Rapid Response updates. **Warning:** Rapid Response definitions do not undergo the same rigorous quality assurance testing as LiveUpdate definitions. Symantec encourages users to rely on the full quality-assurance-tested definitions whenever possible. Ensure that you deploy Rapid Response definitions to a test environment before you install them on your network. See "Obtaining definitions when a new, emerging threat is discovered" on page 226.

You can select the source from where you want to obtain virus definitions. If your organization has several appliances, you can obtain definitions on an internal server. Then you can disseminate the definitions to all of your Symantec Brightmail Gateway appliances. This configuration lets you limit the amount of Internet traffic that accesses Symantec LiveUpdate. In this scenario, you must specify the information for the LAN host and proxy, if required.

See "Specifying from where to obtain virus definitions" on page 225.

You must have a valid content license to install definition files. A content license is a grant by Symantec Corporation for you to update Symantec corporate software with the latest associated content, such as new definitions. When you do not have a content license or your license expires, your product does not receive the most current definitions. Your environment is vulnerable to attacks.

See "Licensing your product" on page 678.

Viewing the status of your virus definitions

Symantec Brightmail Gateway provides details about the status of your virus definitions.

The LiveUpdate status provides the following details:

Last LiveUpdate attempt	The day, date, and time that Symantec Brightmail Gateway last attempted a virus definition update.
Last virus definitions LiveUpdate status	Whether the last attempted virus definition update was successful.
Virus definitions version (revision)	The version of the current virus definitions.
Current virus definition manifest	A list of virus definitions that are contained in this update.

See "About updating virus definitions" on page 221.

See "Scheduling automatic virus definition updates" on page 223.

See "Disabling automatic virus definition updates" on page 224.

See "Initiating virus definition updates on demand" on page 225.

See "Obtaining definitions when a new, emerging threat is discovered" on page 226.

See "Specifying from where to obtain virus definitions" on page 225.

To view the status of your virus definitions

◆ In the Control Center, click **Virus > Settings > LiveUpdate**.

The LiveUpdate status appears under the **LiveUpdate Settings** label.

Scheduling automatic virus definition updates

When you schedule automatic virus definition updates, you ensure that you obtain the most current definitions available from Symantec. You can specify how often you want Symantec Brightmail Gateway to attempt to obtain virus definitions.

You can only schedule automatic updates through LiveUpdate. Rapid Response does not support scheduled updates.

See "Specifying from where to obtain virus definitions" on page 225.

See "Disabling automatic virus definition updates" on page 224.

To schedule automatic virus definition updates

1. In the Control Center, click **Virus > Settings > LiveUpdate**.
2. Under **LiveUpdate Schedule**, click **Enable automatic updates on the following schedule**.
3. Specify a day or days of the week and time at which to begin LiveUpdate.
4. Specify the time that you want the first LiveUpdate attempt to begin.
5. Specify the frequency (in minutes) with which LiveUpdate runs after the first time.

   Enter a value between 3 and 60. The default value is 10.
6. Specify how long Symantec Brightmail Gateway should attempt the LiveUpdate before it times out.
7. Click **Save**.

Disabling automatic virus definition updates

You may want to disable automatic updates. When you disable automatic updates, you must perform on-demand updates to obtain new virus definitions.

Warning: When you disable automatic updates, you run the risk of allowing viruses and malicious attacks into your environment. Automatic virus definition updates ensures that your network always has the latest protection available to defend against threats.

See "Initiating virus definition updates on demand" on page 225.

See "Scheduling automatic virus definition updates" on page 223.

See "Initiating virus definition updates on demand" on page 225.

See "Obtaining definitions when a new, emerging threat is discovered" on page 226.

See "Specifying from where to obtain virus definitions" on page 225.

To disable automatic virus definition updates

1. In the Control Center, click **Virus > Settings > LiveUpdate**.
2. Click **Disable automatic updates**.
3. Click **Save**.

Initiating virus definition updates on demand

You can initiate a LiveUpdate at any time, even if you schedule automatic updates.

See "Viewing the status of your virus definitions" on page 223.

See "Obtaining definitions when a new, emerging threat is discovered" on page 226.

See "Specifying from where to obtain virus definitions" on page 225.

To initiate virus definition updates on demand

1. Click **Virus > Settings > LiveUpdate**.
2. Click **LiveUpdate Now**.

Specifying from where to obtain virus definitions

You can specify the source from where you want to obtain virus definitions as follows:

Symantec Web site	Downloads the virus definitions directly from the Symantec LiveUpdate server. This option is the default setting.
LAN host	If your organization has several appliances, you can obtain definitions on an internal server. Then you can disseminate the definitions to all of your Symantec Brightmail Gateway appliances. This configuration lets you limit the amount of Internet traffic that accesses Symantec LiveUpdate. In this scenario, you must specify the information for the LAN host and proxy, if required.

See "Scheduling automatic virus definition updates" on page 223.

See "Initiating virus definition updates on demand" on page 225.

See "Obtaining definitions when a new, emerging threat is discovered" on page 226.

To obtain virus definition updates from Symantec LiveUpdate server

1 In the Control Center, click **Virus > Settings > LiveUpdate**.

2 Under **Source**, click **Download Platinum virus definitions from the Symantec Web site**.

LiveUpdate uses the proxy that is defined on the **Proxy** tab of the **Administration > Hosts > Configuration > Edit** page.

3 In the **Rapid response and automatic update timeout** box, specify the number of minutes that Symantec Brightmail Gateway waits to connect to the Symantec server before timing out.

The default setting is 20 minutes.

4 Click **Save**.

To obtain virus definition updates from a LAN host

1 In the Control Center, click **Virus > Settings > LiveUpdate**.

2 Under **Source**, click **Download virus definitions from a LAN host**.

If you download virus definitions from a LAN host, LiveUpdate uses a proxy only if you specify one under **Use a proxy**.

Refer to the *LiveUpdate Administrator's Guide* for more information about how to set up a LAN host.

3 In the **Address** field, type the address of the LAN host.

Use a URL, not a host name.

4 In the **Username** field and **Password** field, type the user name and password, if required to access the LAN host.

5 If you use a proxy server, check **Use a proxy**.

6 In the **Proxy** host field, type a valid host name.

7 In the **Proxy port** field, type a valid port number.

8 In the **Username** field and **Password** field, type the user name and password if they are required to access the proxy host.

9 Click **Save**.

Obtaining definitions when a new, emerging threat is discovered

You can use Rapid Response when you need quick responses to emerging threats. Rapid Response definitions are most useful for a perimeter defense to mitigate quickly spreading threats.

Warning: Rapid Response definitions do not undergo the same rigorous quality assurance testing as LiveUpdate definitions. Symantec encourages users to rely on the full quality-assurance-tested definitions whenever possible. Ensure that you deploy Rapid Response definitions to a test environment before you install them on your network.

When you enable Rapid Response virus definition updates, Symantec Brightmail Gateway uses the Symantec Web site as the source for definition updates by default. But you can modify the source to a LAN host.

See “Specifying from where to obtain virus definitions” on page 225.

You cannot schedule definition updates through Rapid Response. If you want to schedule automatic virus definition updates, use LiveUpdate.

See “Scheduling automatic virus definition updates” on page 223.

To obtaining definitions when a new, emerging threat is discovered

1 Click **Virus > Settings > LiveUpdate**.

2 Click **Enable Rapid Response updates**.

   Symantec Brightmail Gateway checks for updated definitions every 10 minutes after this setting is saved.

3 Click **Save**.

Chapter 10

Quarantining suspected viruses

This chapter includes the following topics:

About quarantining suspected viruses

Symantec Brightmail Gateway can quarantine the suspicious messages that might contain viruses. Messages are held in Suspect Virus Quarantine for the period of time that you specify (6 hours by default). Then the messages are released and rescanned. This delay provides time for Symantec to release updated virus definitions with which to scan the suspicious messages. If Symantec Brightmail

Gateway rescans the message, is unable to detect a virus, but still deems the attachment suspicious, the message is returned to Suspect Virus Quarantine.

See "Specifying how long suspect virus messages are retained in quarantine" on page 237.

To use the quarantine, your virus policy must use one of the following actions for suspicious attachments:

- Hold message in Suspect Virus Quarantine
- Strip and Delay in Suspect Virus Quarantine

See "Creating email virus policies" on page 210.

You can do the following with the messages that are in Suspect Virus Quarantine:

- View messages
  See "Viewing suspect virus messages in quarantine" on page 230.
- Sort messages
  See "Sorting suspect virus messages in quarantine" on page 233.
- Search messages
  See "Searching quarantined virus messages" on page 233.
- Delete messages
  See "Deleting suspect virus messages in quarantine" on page 236.
- Release messages
  See "Releasing suspect virus messages from quarantine" on page 236.

Quarantined messages are stored on the Control Center. You can specify the amount of disk space that you want to allocate to Suspect Virus Quarantine.

See "Modifying the disk space allotted for Suspect Virus Quarantine" on page 238.

Viewing suspect virus messages in quarantine

You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Suspect Virus Quarantine.

See "About navigating Suspect Virus Quarantine" on page 232.

See "Specifying the number of suspect virus message entries to view per page" on page 231.

See "Sorting suspect virus messages in quarantine" on page 233.

See "Searching quarantined virus messages" on page 233.

To view suspect virus messages in quarantine

- Do one of the following:

If you are not on the Suspect Virus Message Quarantine page	In the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
If you are on the Suspect Virus Message Quarantine page and want to see the newly arrived messages	On the **Virus > Quarantine > Email Suspect Virus** page, if the **Display All** icon is not visible, click **Show Filters** , and then click **Display All**.

Choosing the language encoding for suspect virus messages

In most cases, the Auto-detect setting properly determines the language encoding for a message in Suspect Virus Quarantine. However, the Control Center may not be able to determine the proper language encoding for some messages. If the message is garbled, select the language encoding most likely to match the encoding that is used in the message.

See "Viewing suspect virus messages in quarantine" on page 230.

See "About navigating Suspect Virus Quarantine" on page 232.

Only the administrators that have Full Administration rights or Manage Quarantine modify rights can choose language encoding for messages in quarantine.

To choose the language encoding for suspect virus messages

1. In the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click on the subject line of the message that you want to view.
3. On the message details page, select the language encoding in the drop-down list.

Specifying the number of suspect virus message entries to view per page

You can specify how many quarantined message entries appear per page.

You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Suspect Virus Quarantine.

See “Viewing suspect virus messages in quarantine” on page 230.

See “About navigating Suspect Virus Quarantine” on page 232.

See “Sorting suspect virus messages in quarantine” on page 233.

See “Searching quarantined virus messages” on page 233.

To specify the number of suspect virus message entries to view per page

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click the **Entries per page** drop-down list, and select a number.

About navigating Suspect Virus Quarantine

The following navigation icons help you navigate through messages on the Suspect Virus Message Quarantine page as follows:

⏮	Go to beginning of messages
⏭	Navigate to last page of messages or 50 pages ahead if there are more than 50 pages.
◀	Go to previous page of messages
▶	Go to next page of messages
1-10 ▾	Choose up to 500 pages before or after the current page of messages

When you navigate to a different page of messages, the status of the check boxes in the original page is not preserved. For example, assume that you select three messages on the first page of messages and then move to the next page. When you return to the first page, all of the message check boxes are cleared.

See “Viewing suspect virus messages in quarantine” on page 230.

See “Specifying the number of suspect virus message entries to view per page” on page 231.

See “Sorting suspect virus messages in quarantine” on page 233.

See “Searching quarantined virus messages” on page 233.

Sorting suspect virus messages in quarantine

You can sort messages in Suspect Virus Quarantine by date to make it easier to categorize the messages or locate a specific message. By default, messages appear in date descending order. The newest messages are listed at the top of the page.

You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Suspect Virus Quarantine.

See "Viewing suspect virus messages in quarantine" on page 230.

See "Specifying the number of suspect virus message entries to view per page" on page 231.

See "About navigating Suspect Virus Quarantine" on page 232.

See "Searching quarantined virus messages" on page 233.

To sort suspect virus messages in quarantine

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click on the **Date** column heading to sort by ascending order or descending order.

Searching quarantined virus messages

You can search for messages in Suspect Virus Quarantine. The ability to search messages lets you more easily find a specific message that you want to view, delete, or release. You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Suspect Virus Quarantine.

See "Viewing suspect virus messages in quarantine" on page 230.

See "Specifying the number of suspect virus message entries to view per page" on page 231.

See "About navigating Suspect Virus Quarantine" on page 232.

See "Sorting suspect virus messages in quarantine" on page 233.

To search quarantined virus messages

1. In the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. On the message list page, click **Show Filters**.

3 Do any of the following to perform a search:

To search message envelope "To" recipient	Type a name or address in the **To** box to search the message envelope `RCPT TO:` header. You can search for a display name, the user name portion of an email address, or any part of a display name or email user name. If you type a full email address in the To box, Symantec Brightmail Gateway searches only for the user name portion of `user_name@example.com`. The search is limited to the envelope `To:`, which may contain different information than the header `To:` that appears on the message details page.
To search "From" headers	Type a name or address in the **From** box to search the `From:` header in all messages for a particular sender. You can search for a display name, email address, or any part of a display name or email address. The search is limited to the visible message `From:` header, which is usually forged in spam messages. The visible message `From:` header may contain different information than the message envelope.
To search subject headers	Type in the **Subject** box to search the `Subject:` header for all messages about a specific topic.
To search a time range	Click the **Time range** drop-down list to display all of the messages that were received during the time range that you specify.

4 Click **Display Filtered**.

Suspect virus message search criteria and tips

The search function is optimized for searching a large number of messages. However, this can lead to unexpected search results.

Consider the following tips and information to help you conduct searches in Suspect Virus Quarantine:

Tokens

Tokens are matched with substring semantics. Searching for a subject with the search target <in> will match "Lowest rate in 45 years," "RE: re: Sublime Bulletin (verification)," "Up to 85% off Ink Cartridges + no shipping!," and "Re-finance at todays super low rate."

Multiple word searches

If any word in a multiple word search is found in a message, that message is considered a match. For example, searching for red carpet match "red carpet," "red wine," and "flying carpet."

Case sensitivity

All text searches are case-insensitive. For example, assume you type emerson in the From box. Messages with a From header that contains emerson, Emerson, and eMERSOn all appear in the search results.

Exact phrases

To search for an exact phrase, enclose the phrase in " " (double quotes).

Wildcards

You can use * (asterisk) to perform wildcard searches. It also functions as a logical AND character. In addition, you can search on special characters such as & (ampersand), ! (exclamation point), $ (dollar sign), and # (pound sign).

Single characters

Even a single character is treated as a substring target.

Special characters

You can search on special characters such as & (ampersand), ! (exclamation point), $ (dollar sign), and # (pound sign).

Multiple characteristics

If you search for multiple characteristics, only the messages that match the combination of characteristics are listed in the search results. For example, assume you type LPQTech in the From box and Inkjet in the Subject box. Only the messages that contain LPQTech in the `From:` header and Inkjet in the `Subject:` header appear in the search results.

Forged header information

Spammers usually "spoof" or forge some of the visible messages headers such as From and To and the invisible envelope information. Sometimes they forge header information using the actual email addresses or domains of innocent people or companies.

Time to perform a search

The amount of time it takes to perform the search depends on how many search boxes you use and the number of messages in the mailbox. Searching in the administrator mailbox takes longer than searching in a user's mailbox.

See “Searching quarantined virus messages” on page 233.

Deleting suspect virus messages in quarantine

You can delete messages in Suspect Virus Quarantine one at a time, or you can delete several messages at once. When you delete a message, it is no longer accessible. Only the administrators that have Full Administration rights or Manage Quarantine Modify rights can delete messages in quarantine.

See "Viewing suspect virus messages in quarantine" on page 230.

You can also use a purge utility to automatically delete messages from Suspect Virus Quarantine. This utility frees you from having to manually delete messages from the quarantine to free up space. The utility purges messages based on the schedule that you specify.

See "Specifying how long suspect virus messages are retained in quarantine" on page 237.

To delete individual messages in quarantine

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click on the check box beside each message that you want to delete.
3. Click **Delete**.

To delete all messages in quarantine

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click **Delete All** to delete all of the messages in Suspect Virus Quarantine, including those on other pages.

Releasing suspect virus messages from quarantine

You can release messages from the Suspect Virus Quarantine to be rescanned with the latest virus definitions. If Symantec Brightmail Gateway is unable to repair the virus and your policy is to quarantine suspect viruses, the message is returned to the Suspect Virus Quarantine.

You can release one message at a time, or you can release all messages at once. You can also set up automatic releases and rescanning of suspect virus messages.

See "Specifying how long suspect virus messages are retained in quarantine" on page 237.

Releasing messages requires access to the IP address of the Control Center. If you limit inbound or outbound SMTP access, check the Inbound Mail Settings and Outbound Mail Settings.

Only administrators with Full Administration rights or Manage Quarantine modify rights can release messages from the quarantine.

See “Deleting suspect virus messages in quarantine” on page 236.

See “About Scanner email settings ” on page 86.

To release individual messages from quarantine

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Check the box beside each message that you want to release.
3. Click **Release**.

To release all messages from quarantine

1. From the Control Center, click **Virus > Quarantine > Email Suspect Virus**.
2. Click **Release All** to release all the messages in Suspect Virus Quarantine, including those on other pages.

Specifying how long suspect virus messages are retained in quarantine

You can choose the maximum amount of time a message can be held in Suspect Virus Quarantine, up to 24 hours. After the period of time that you specify, messages are automatically released from Suspect Virus Quarantine and rescanned with updated virus definitions.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See “About scheduled tasks” on page 621.

See “Modifying the disk space allotted for Suspect Virus Quarantine” on page 238.

You can also delete messages manually from Suspect Virus Quarantine.

See “Deleting suspect virus messages in quarantine” on page 236.

You must have Full Administration or Modify rights to change Suspect Virus Quarantine settings.

To specify how long suspect virus messages are held in quarantine

1. On the Control Center, click **Virus > Settings > Suspect Virus Settings**.
2. Under **Message Release**, click the **Automatically release messages older than** drop-down list and select the number of hours to hold messages in quarantine.

   The default is six hours. The maximum setting is 24 hours.
3. Click **Save**.

Modifying the disk space allotted for Suspect Virus Quarantine

You can specify the amount of disk space that Suspect Virus Quarantine uses. The default disk space is 10 GB.

Only administrators with Full Administration rights or Manage Settings modify rights can modify Suspect Virus Quarantine settings.

See “Specifying how long suspect virus messages are retained in quarantine” on page 237.

See “Deleting suspect virus messages in quarantine” on page 236.

See “Releasing suspect virus messages from quarantine” on page 236.

To modify the disk space allotted for Suspect Virus Quarantine

1. Click **Virus > Settings > Suspect Virus Settings**.
2. Check **Maximum size of the Suspect Virus Quarantine** to enable the threshold.
3. Specify the amount of disk space you want to allot for Suspect Virus Quarantine.

   The default is 10 GB.
4. Click **Save**.

Chapter 11

Filtering spam

This chapter includes the following topics:

About filtering spam

Symantec Brightmail Gateway comes with default spam policies that are enabled by default. However, you can modify these policies or create your own.

See "Default email spam policies" on page 50.

See "Creating email spam policies" on page 242.

You can also modify the settings that further detect spam.

See "Configuring the threshold for suspected spam identification" on page 246.

You can specify how you want spam or suspected spam handled. For example, you can send suspected spam messages to the Spam Quarantine. You can also prepend the subject line to let users know that the message is suspected spam.

See “ Verdicts and actions for email messages” on page 718.

See “About quarantining spam” on page 258.

You must have a valid antispam license to perform antispam scanning functions.

See “Licensing your product” on page 678.

Configuring spam detection

Table 11-1 describes the tasks that you can perform to detect spam and enhance spam detection capabilities.

See “About filtering spam” on page 239.

Table 11-1 Configure spam detection

Task	Description
Create and enable email spam policies.	Symantec Brightmail Gateway comes with preconfigured spam policies that are automatically enabled. You can modify these polices or create your own custom policies. See “Default email spam policies” on page 50. See “Creating email spam policies” on page 242. See “Modifying spam policies” on page 244.
Set the suspected spam level.	When Symantec Brightmail Gateway evaluates whether messages are spam, it calculates a spam score from 1 to 100 for each message. This score is based on the results of a wide variety of antispam technologies. You can define a range of scores from 25 to 89 that constitutes “suspected spam.” Through policies, you can specify different actions for the messages that are identified as suspected spam and messages that are identified as spam. Symantec recommends that you not adjust the spam threshold until you have some exposure into the filtering patterns at your site. Gradually move the threshold setting down 1 point to 5 points per week until the number of false positives is at an acceptable level. See “Configuring the threshold for suspected spam identification” on page 246.

Table 11-1 Configure spam detection *(continued)*

Task	Description
Enable the features that provide Symantec with information that helps us create better spam filters.	You can help Symantec create better spam filters when you enable the following features: ■ Uniform Resource Identifiers (URI) reporting When you enable URI reporting, Symantec Brightmail Gateway sends a report to Symantec Security Response. The report contains URIs that appear in the messages that Symantec Brightmail Gateway scans for spam. Symantec uses this information to develop new URI-based filters. See "Enabling or disabling URI reporting to Symantec" on page 248. ■ Probe accounts You can forward unused email addresses or invalid email addresses to the Symantec Probe Network. Symantec uses these email addresses to attract spammers. Then Symantec uses the spam messages it receives at these addresses to create better spam filters. See "Enabling probe participation" on page 253. See "Creating probe accounts from invalid recipient email addresses" on page 253. See "Enabling probe accounts" on page 255.

Table 11-1 Configure spam detection *(continued)*

Task	Description
Configure the Spam Quarantine.	You can route spam, suspected spam, or both to Spam Quarantine. Users access Spam Quarantine to determine if messages in the quarantine are false positives. If a message is marked as spam or suspected spam, but is legitimate, users can release the messages to their inboxes. Users can notify you of false positives so that you can continue to adjust your spam settings and spam policies accordingly. You can also set up summary notifications to be delivered to users inboxes. See "Before you use Spam Quarantine" on page 258. See "Configuring Spam Quarantine for administrator-only access" on page 260. See "Specifying who to notify of false positive messages" on page 274. See "About configuring the user and distribution list notification digests" on page 275. You can configure thresholds to control the space that is allocated for Spam Quarantine. See "Modifying Spam Quarantine thresholds" on page 272. See "Specifying when and how often Spam Quarantine is expunged" on page 284.
Monitor reports.	Monitor reports to determine how effective spam policies are. Reports also indicate the volume of spam that your organization receives. This information can help you fine-tune your spam detection settings. See "About working with reports" on page 568.

Creating email spam policies

You can create the email spam policies that determine the actions to be taken for specific spam conditions.

Symantec Brightmail Gateway installs with several preconfigured policies. You can use these policies as is, edit them to suit your needs, or create your own custom policies.

See "Default email spam policies" on page 50.

When you create a new spam policy, it is enabled by default.

See “Enabling and disabling spam policies” on page 244.

See “Modifying spam policies” on page 244.

See “Copying spam policies” on page 245.

See “Deleting spam policies” on page 246.

To create the policies that detect spam

1. In the Control Center, click **Spam > Policies > Email**.
2. Click **Add**.
3. On the **Email Spam Policy** page, in the **Policy name** box type a name for the spam policy.

   Content filtering, spam, and virus policy names must be unique. For example, if you have a content filtering policy called "XYZ," you cannot have a spam policy or virus policy called "XYZ."
4. Under **Conditions** click the **Apply to** drop-down list and choose whether the policy is applied to inbound messages only, outbound messages only, or both.
5. Click the **If the following condition is met** drop-down list and select one of the following options:

If a message fails bounce attack validation	Performs the specified action if a message is a Non-Delivery Receipt message that does not pass validation. This condition is only available for Inbound messages.
If a message is spam	Performs the specified action if a message is spam.
If a message is spam or suspected spam	Perform the specified action if a message is either spam or suspected spam.
If a message is suspected spam	Perform the specified action if a message might be spam. The suspected spam level is adjustable on the Spam > Settings > Scan Settings page.

6. Under **Actions**, click the **Perform the following action** drop-down list and select the action that you want to take on a message that meets the specified spam condition.

   For some actions, you need to specify additional information.

   See “ Verdicts and actions for email messages” on page 718.

   See “About defending against bounce attacks” on page 182.
7. Click **Add Action**.

8 If you want, add more actions.

See Table A-5 on page 728.

9 Under **Apply to the following groups**, check one or more policy groups to which this policy should apply.

See "Creating a policy group" on page 316.

10 Click **Save**.

Enabling and disabling spam policies

Default and newly created spam policies are enabled by default. When you create a new spam policy, it is enabled by default. You can disable any policy that you do not want Symantec Brightmail Gateway to use when it scans email messages.

You can disable a virus policy to troubleshoot antispam scanning issues. For example, you can disable the policies that you no longer want to use but do not want to delete yet.

See "About filtering spam" on page 239.

See "Creating email spam policies" on page 242.

See "Modifying spam policies" on page 244.

See "Copying spam policies" on page 245.

See "Deleting spam policies" on page 246.

To enable or disable spam policies

1 In the Control Center, click **Spam > Policies > Email**.

2 Check the box beside the policy that you want to enable to disable.

3 Select one of the following options:

Enable	When you enable a policy, a green check mark appears in the **Enabled** column.
Disable	When you disable a policy, a horizontal line appears in the **Enabled** column.

Modifying spam policies

You can modify spam policies to fine-tune them or to expand or reduce their scope.

Table 11-2 describes the setting that you can modify for default spam policies and custom spam policies.

Table 11-2 Modifiable spam policy settings

Setting	Modifiable in default policies	Modifiable in custom policies
Policy name	No	Yes
Apply to	No	Yes
If the following condition is met	Yes	Yes
Perform the following action	Yes	Yes
Apply to the following policy groups	Yes	Yes

See "About filtering spam" on page 239.

See "Creating email spam policies" on page 242.

See "Enabling and disabling spam policies" on page 244.

See "Copying spam policies" on page 245.

See "Deleting spam policies" on page 246.

To modify spam policies

1. In the Control Center, click **Spam > Policies > Email**.
2. Check the box beside the policy that you want to modify, and then click **Edit**.
3. Make the modifications that you want.
4. Click **Save**.

Copying spam policies

You may have instances in which you create a complicated spam policy and want to create a similar policy with only a few variances. Symantec Brightmail Gateway lets you copy spam policies.

When you copy a spam policy, the new policy must have a unique name. For example, if you have a content filtering policy called XYZ, you cannot have a spam policy or virus policy called XYZ. Spam policies are enabled by default when you create them.

See "About filtering spam" on page 239.

See “Creating email spam policies” on page 242.

See “Enabling and disabling spam policies” on page 244.

See “Modifying spam policies” on page 244.

See “Deleting spam policies” on page 246.

To copy spam policies

1. In the Control Center, click **Spam > Policies > Email**.
2. Check the box beside the policy that you want to copy.
3. Click **Copy**.
4. On the **Email Spam Policies** page, type a new name for the policy.
5. Make any other changes you want.
6. Click **Save**.

Deleting spam policies

You can delete the spam policies that you no longer need. However, when you delete a policy, the policy cannot be retrieved. If you are unsure if you want to permanently delete a policy, you can disable it instead.

See “Enabling and disabling spam policies” on page 244.

See “Copying spam policies” on page 245.

See “Modifying spam policies” on page 244.

To delete spam policies

1. In the Control Center, click **Spam > Policies > Email**.
2. Check the box beside the policy that you want to delete.
3. Click **Delete**.
4. Click **Delete** in the confirmation dialog box.

Configuring the threshold for suspected spam identification

When Symantec Brightmail Gateway evaluates whether messages are spam, it calculates a spam score from 1 to 100 for each message. This score is based on techniques such as pattern matching and heuristic analysis.

Symantec Brightmail Gateway categorizes the spam scores as follows:

90 - 100	If an email receives a score in the range of 90 to 100, it is defined as spam. Symantec determines which messages are spam. This determination cannot be adjusted.
89 - 25 suspected spam threshold 75 by default	You can define a range of scores from 25 to 89. The messages that score within this range are considered "suspected spam." Through policies, you can specify different actions for the messages that are identified as suspected spam and messages that are identified as spam. For example, assume that you configure your suspected spam scoring range to encompass scores from 80 through 89. If an incoming message receives a spam score of 83, Symantec Brightmail Gateway considers this message to be suspected spam. It applies the action you specify for suspected spam messages, such as Modify the Message (tagging the subject line).
Less than the suspected spam threshold	If a message receives a score that is less than the suspected spam threshold, Symantec Brightmail Gateway considers it to be non-spam email.

Symantec recommends that you not adjust the spam threshold until you have some exposure into the filtering patterns at your site. Gradually move the threshold setting down 1 point to 5 points per week until the number of false positives is at an acceptable level. One way to test the effects of spam scoring is to configure your spam policy to hold suspected spam in Spam Quarantine. Then set up a designated mailbox to receive false positives complaints and monitor their numbers as you change the spam-score threshold.

Note: Symantec does not consider the legitimate messages that receive a suspected spam verdict to be false positives. Messages that are submitted to Symantec Security Response that receive suspected spam verdicts are not reviewed. The reason is that Symantec cannot control how organizations configure the Suspect Spam threshold value. So Symantec does not create filters or modify filters based on suspected spam verdicts. Filters that are created based on the suspected spam threshold values that are set too low can impact spam effectiveness for all Symantec customers.

See "About filtering spam" on page 239.

To configure the threshold for suspected spam identification

1. In the Control Center, click **Spam > Settings > Scan Settings**.
2. Click the **Email** tab.
3. Under **Do you want messages to be flagged as suspected spam**, click **Yes** or **No**.

4. Under **Select a Suspected Spam Threshold between 25 and 89**, click and drag the slider to increase or decrease the lower limit of the range for suspected spam. You can also type a value in the box.
5. Click **Save**.

Enabling or disabling URI reporting to Symantec

You can help Symantec create better spam filters that block messages based on Uniform Resource Identifiers (URI). When you enable URI reporting, Symantec Brightmail Gateway sends a report to Symantec Security Response. The report contains URIs that appear in the messages that Symantec Brightmail Gateway scans for spam. Symantec uses this information to develop new URI-based filters. You receive these updated filters through the Conduit.

See "Working with Services" on page 112.

The URI reporting feature has no measurable impact on product performance.

To enable or disable URI reporting to Symantec

1. In the Control Center, click **Spam > Settings > Scan Settings**.
2. On the **Email** tab, under **Uniform Resource Identifier Reporting**, do either of the following tasks:

To enable URI reporting	Check **Report URIs to Symantec Security Response**. This feature is enabled by default for new installations.
To disable URI reporting	Uncheck **Report URIs to Symantec Security Response**. This feature is disabled by default for migrated installations.

3. Click **Save**.

Participating in the Symantec Probe Network

About the Symantec Probe Network

The Probe Network is crucial to Symantec's ongoing effort to fight spam. As spammers find new ways to bypass filters, the Symantec Probe Network helps Symantec stay one step ahead by monitoring spamming methods through the use of probe accounts.

The Probe Network is effective for the following reasons:

Drives early detection of spam attacks	Probe accounts are the first step in the real-time detection and analysis of spam. The structure of the Probe Network essentially provides Symantec Security Response with a stream of real-time spam being disseminated over the Internet. This virtual net of numerous accounts spread all over the Internet makes it easy for Symantec to verify that a given message was sent using bulk methods.
Speeds the development of accurate filters	A key marketplace differentiator for Symantec Brightmail Gateway is the near perfect accuracy rate of its spam filtering technology. The antispam capability is largely due to core filters that are based on actual spam. The probe network also provides key data that is used to develop Symantec's more predictive filters, such as heuristics. What makes all this possible is the volume, quality, and timeliness of data that flows in real time from the probe network to Symantec Security Response.
Aids ongoing trend research	Spammers are constantly changing their tactics and dissemination methods to evade filtering software. Symantec's Customer Response and AntiSpam Systems teams mine the data from the probe network to advance Symantec's AntiSpam technology. Examples include staying abreast of the latest spam trends, evaluating the spam-catching differences between product versions and monitoring detection rates in different languages.

See “About probe accounts” on page 249.

See “Enabling probe participation” on page 253.

About probe accounts

Symantec Brightmail Gateway provides options to convert your invalid recipient email addresses into probe accounts which can be used in the Symantec Probe Network. Probe accounts help Symantec track spam and learn from it. The intelligence that Symantec gains from probe accounts enables continuous improvement of the rules that govern spam filters. Better filters mean fewer spam intrusions on your network.

See “About the Symantec Probe Network” on page 248.

The tools for creating probe accounts are available from the **Spam > Settings > Probe Accounts** page.

See “Enabling probe participation” on page 253.

See “Setting up probe accounts” on page 251.

See “Creating probe accounts from invalid recipient email addresses” on page 253.

See “Creating probe accounts manually” on page 254.

See “Enabling probe accounts” on page 255.

See “Disabling a probe account” on page 256.

See “Deleting a probe account” on page 256.

You can track the effectiveness of your probe accounts by viewing the reports that track the top probe accounts.

See “Report types” on page 571.

About creating probe accounts

When you create probe accounts, you may wonder which email addresses make the best probe accounts. You have the option to select as many invalid recipient email addresses as you want. You can also create any number of invalid or unused email addresses for use in the probe network.

However, when you upload email addresses for use as probe accounts, it is important to consider the following guidelines:

Invalid Recipients

You should select the invalid recipient addresses that receive the most amount of email or those that you believe receive high percentages of spam email. You should not select the invalid recipients addresses that are former employees' addresses or common misspellings of public addresses (for example, support@symantecexample.com). These addresses likely receive mostly valid emails.

To see your top invalid recipients, use the **Reporting** feature to view the Invalid Recipient report.

See “Selecting the data to track for reports” on page 569.

See “Report types” on page 571.

See “Creating probe accounts from invalid recipient email addresses” on page 253.

Manual entering of addresses

Use this method to add other addresses that you believe receive only spam content. These may include the addresses that you have seeded on the Internet that you expect to receive only spam messages. Seeding is the deliberate publishing of email addresses on the Internet in order for spammers to harvest and target these addresses.

See “Creating probe accounts manually” on page 254.

In some cases you may employ **Alias** addresses. Probe accounts always override aliases.

Remember the following when you create probe accounts and alias addresses:

- If you create a probe account for an address that already has an alias, then the probe feature overwrites that alias.
- You cannot create an alias for an address that is already a probe account.

In some cases you may employ **Masqueraded** addresses. Masqueraded accounts override probe accounts.

For example, you use the reports to find that jay@symantecexample.com is a top invalid recipient. You make jay@symantecexample.com a probe address.

Sometime later, the company changes its name from symantecexample to symantecdomain, and all of the email addresses change to @symantecdomain.com.

You use address masquerading to masquerade domain ‘symantecexample’ to ‘symantecdomain’ so the mails that are sent to someone@symantecexample.com gets sent to the new someone@symantecdomain.com address. In doing so, you invalidate the jay@symantecexample.com probe account. All mail to that address is seen as going to jay@symantecdomain.com, which is not a probe address.

See “Adding or editing address masquerades ” on page 127.

Setting up probe accounts

The following table describes the steps to create probe accounts and add them to the Symantec Probe Network.

Table 11-3 Creating probe accounts

Step	Action	Description
Step 1	Enable probe participation.	Probe participation must be enabled to activate probe accounts. See “Enabling probe participation” on page 253.
Step 2	Verify domain and enable recipient validation.	The probe account domain must match one of your local domains. Verify that there is a matching domain and that it is enabled for recipient validation. If there is no matching domain you can add one using the add domain task. See “Adding or editing domains” on page 117.

Table 11-3 Creating probe accounts *(continued)*

Step	Action	Description
Step 3	Specify invalid recipient handling.	You must configure invalid recipient handling on the domain that you want to use for a probe account. You can choose any of the three handling options: accept, reject, or drop; but you must choose one. See "Setting up invalid recipient handling" on page 130.
Step 4	Add a data source (optional) and enable recipient validation. **Note:** Skip this step if you plan to create probe accounts using the manual method only.	A directory data source that is enabled for recipient validation is necessary only if you want to upload invalid recipient email addresses captured in the directory data service filters. See "Enabling or editing the recipient validation function" on page 556. If you do not already have a directory data source, you need to create one and enable recipient validation. See "Adding a data source " on page 492.
Step 5	Enable reporting.	To track top invalid recipients and top probe accounts, enable the Invalid Recipients report. See "Selecting the data to track for reports" on page 569. See "Creating and configuring reports" on page 569.
Step 6	Create probe accounts.	The probe feature provides two methods for creating probe accounts. You can manually create probe accounts by entering email addresses or uploading them from a file. Or, you can create probe accounts using the invalid recipient method. See "Creating probe accounts manually" on page 254. See "Creating probe accounts from invalid recipient email addresses" on page 253.

Table 11-3 Creating probe accounts *(continued)*

Step	Action	Description
Step 7	Enable or disable one or more probe accounts.	Although probe accounts are enabled by default, you may want to temporarily disable a probe account and enable it again later. See "Enabling probe accounts" on page 255. See "Disabling a probe account" on page 256.

Enabling probe participation

To begin adding email addresses for use in the Symantec Probe Network, you must first enable the feature.

To enable probe participation

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. Under **Probe Email Address**, check **Enable probe participation**.

   Once you have enabled probe participation, you can select individual probe accounts to enable or disable.

   See "Enabling probe accounts" on page 255.

Creating probe accounts from invalid recipient email addresses

You must first make a few minor system configurations before completing this task.

See "Setting up probe accounts" on page 251.

For assistance on determining which accounts make good probe accounts, refer to the probe account create guidelines.

See "About creating probe accounts" on page 250.

Once you create a probe account, you can track its effectiveness on the **Reports** tab by running a **Top Probe Accounts** report.

See "Creating and configuring reports" on page 569.

Note: To view **Top Probe Account** reports, you must first enable the **Invalid Recipient** setting on the **Administration > Settings > Reports** page.

See "Selecting the data to track for reports" on page 569.

To create probe accounts from invalid recipient email addresses

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. Check **Enable probe participation** if it is not already checked.
3. Under **Add Probe Addresses**, click **Invalid Recipients**.

   Ensure that you have configured the system to track invalid recipient data, otherwise an error message appears.

   See “Selecting the data to track for reports” on page 569.
4. Click the **Time range** drop-down list, select a time range, and click **View**.

   In the **Invalid Recipient** address table, use the drop-down menus to specify the number of **Entries** (email addresses) you want to see per page and which ones to **Display**. You can select 1-10 to see the top 10 invalid recipient accounts, or you can select 10-20 to see the next 10 accounts. Use the arrows to quickly navigate through your results.
5. Check the box beside each address that you want to make into a probe account and click **Add Selected to Probes.** Click **Add All to Probes** to add all the addresses to the Symantec Probe Network .

   This returns you to the initial probe accounts page where you see the selected addresses in the **Email Address** table with a status of **Enabled**.

Creating probe accounts manually

You must make a few minor system configurations before completing this task.

See “Setting up probe accounts” on page 251.

For assistance on determining which accounts make good probe accounts, refer to the probe account create guidelines.

See “About creating probe accounts” on page 250.

Once you create a probe account you can track its effectiveness on the **Reports** tab by running a **Top Probe Accounts** report.

See “Creating and configuring reports” on page 569.

Note: To view the Top Probe Account reports, you must first enable the **Invalid Recipient** setting on the **Administration > Settings > Reports** page.

See “Selecting the data to track for reports” on page 569.

To manually create a probe account

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. Check the **Enable probe participation** box if it is not already checked.
3. In the **Manually add probe email addresses** field, type one or more email addresses, separated by commas, to add to the probe network.

   You can alternatively add email addresses from an external text file by clicking **Browse**. When the **Browse** window opens, select the file that contains the email addresses that you want to add as probe accounts and click **Upload**.

   Email addresses uploaded from a text file must be formatted correctly. Apply one email address per line; no commas.
4. Click **Add**.

   When the screen refreshes, the probe accounts appear in the **Email Address** table with an Enabled status.

Enabling probe accounts

By default a probe account is enabled when you create it. In some cases you may want to temporarily disable a probe account. When you are ready to reactivate the account, you use the **Enable** option.

Before enabling individual probe accounts, make sure the **Enable Probe Participation** check box is checked at the top of the **Spam > Settings > Probe Accounts** page. If it is unchecked, it overrides individually enabled probe accounts, making them inactive.

See "Enabling probe participation" on page 253.

To enable probe accounts

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. From the probe accounts list, check each probe account that you want to enable.

   You can also check the topmost box to select all the accounts that are listed on the page.

   You can sort the list of probe accounts by clicking the column title **Email Address** , which sorts the accounts alphabetically. To sort the probe addresses by status, click the table title **Enabled**.
3. Click **Enable** or **Enable All**.

   The selected probe accounts are now active in the Symantec Probe Network.

Disabling a probe account

When you want to temporarily disable a probe account in the Symantec Probe Network, you use the **Disable** option.

Note: To permanently remove a probe account, use the **Delete** option.

See "Deleting a probe account" on page 256.

To disable a probe account

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. In the **Email Addresses** table, check each probe account that you want to disable, or click the topmost box to select all accounts listed on the page.
3. Click **Disable** or **Disable All**.

   The disabled account is no longer active, but it remains in the Symantec Probe Network where it can be reactivated.

Deleting a probe account

When you no longer want a probe account in the Symantec Probe Network, you can permanently remove it using the **Delete** option.

Note: If you want to save the probe account but temporarily remove it from the Symantec Probe Network, use the **Disable** option.

See "Disabling a probe account" on page 256.

To delete a probe account

1. In the Control Center, click **Spam > Settings > Probe Accounts**.
2. In the **Email Addresses** table, check the check box beside each probe account that you want to delete. You can also check the topmost box to select all the accounts that are listed on the page.
3. Click **Delete** or **Delete All**.

   The deleted accounts are removed from the Symantec Probe Network and no longer appear in the list of probe accounts.

Chapter 12

Quarantining spam

This chapter includes the following topics:

- Specifying when and how often Spam Quarantine is expunged
- Troubleshooting Spam Quarantine

About quarantining spam

You can route spam, suspected spam, or both to Spam Quarantine so that administrators and users can access the messages to check for false positives, if necessary. Spam Quarantine can help you fine-tune your spam settings and spam policies. Too many false positives can indicate that spam settings are too stringent and should be modified. Use of Spam Quarantine is optional.

See "Creating email spam policies" on page 242.

See "Before you use Spam Quarantine" on page 258.

See "Viewing spam and suspected messages in quarantine" on page 261.

See "About navigating Spam Quarantine" on page 264.

Spam Quarantine stores spam messages, and it provides Web-based user and administrator access to those messages. Users access Spam Quarantine with their LDAP user names and authentication. If a message is marked as spam or suspected spam, but is legitimate, users can release the messages to their inboxes. Users can notify you of false positives so that you can continue to adjust your spam settings and spam policies accordingly. You can also set up summary notifications to be delivered to users inboxes.

See "Configuring Spam Quarantine for administrator-only access" on page 260.

See "Specifying who to notify of false positive messages" on page 274.

See "About configuring the user and distribution list notification digests" on page 275.

You can configure thresholds to control the space that is allocated for Spam Quarantine.

See "Modifying Spam Quarantine thresholds" on page 272.

See "Specifying when and how often Spam Quarantine is expunged" on page 284.

Before you use Spam Quarantine

If you intend to permit users to access Spam Quarantine, before you use Spam Quarantine, ensure that you have done all of the following:

Create and enable the spam policies that quarantine spam and suspected spam.	One or more policy groups must have an associated filter policy that quarantines messages. For example, you can create a spam policy that quarantines inbound suspected spam messages for the Default group. See "Creating email spam policies" on page 242.
Configure your LDAP server and ensure it works properly.	Control Center access to your LDAP server using authentication must work properly for users to logon to Spam Quarantine to check their quarantined messages. You also need LDAP authentication to expand LDAP email aliases and for the **Delete Unresolved Email** setting. See "About using the authentication function with your data source" on page 487.
Ensure that you have an SMTP mail server available.	Spam Quarantine does not require a separate SMTP mail server to send notifications and resend misidentified messages. However, an SMTP mail server must be available to receive notifications and the misidentified messages that Spam Quarantine sends. The SMTP server that you choose should be downstream from the Scanner, as notifications and misidentified messages do not require filtering. See "Configuring SMTP advanced settings" on page 95.

See "About quarantining spam" on page 258.

See "Enabling users to bypass Control Center login credentials" on page 674.

Forwarding spam messages for non-existent addresses to quarantine

If LDAP is configured, the messages that are sent to non-existent email addresses (based on LDAP lookup) are deleted by default. If you uncheck the check box for **Delete messages sent to unresolved email addresses**, these messages are stored in the Spam Quarantine `postmaster` mailbox. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify this setting.

See "Viewing spam and suspected spam messages sent to the postmaster mailbox" on page 263.

The deletion of unresolved email is disabled by default if LDAP is not configured. In that case, the option **Delete messages sent to unresolved email addresses** is unchecked and is grayed out.

Note: If an LDAP server connection fails or LDAP settings have not been configured correctly, then quarantined messages addressed to non-existent users are consigned to the Scanner's deferred queue and the SMTP connection between Control Center and Scanner is closed, whether the Delete unresolved email check box is checked or unchecked. Once the Scanner's deferred retry or timeout limit is reached, the message bounces back to the sender.

See “About quarantining spam” on page 258.

To forward spam messages for non-existent addresses to quarantine

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. Uncheck **Delete messages sent to unresolved email addresses**.
3. Click **Save**.

Configuring Spam Quarantine for administrator-only access

If you do not have an LDAP directory server configured or do not want users in your LDAP directory to access Quarantine, you can configure Quarantine so that only administrators can access the messages in Quarantine.

When administrator-only access is enabled, you can still perform all the administrator tasks available for sites with LDAP integration enabled. These tasks include redelivering misidentified messages to local users, whether or not you use an LDAP directory at your organization. However, notification of new spam messages is disabled when administrator-only access is enabled.

The **Administrator-only Quarantine** option is enabled by default if LDAP is not configured. In that case, this option is checked and grayed out.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “How Spam Quarantine differs for administrators and users” on page 262.

To configure Spam Quarantine for administrator-only access

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. If it is unchecked, check **Administrator-only Quarantine**.
3. Click **Save**.

Viewing spam and suspected messages in quarantine

View messages that are in Spam Quarantine to determine if they are spam or false positives.

Administrators with Full Administration rights or Manage Quarantine view rights can view spam messages in quarantine. However, these administrators cannot release messages or delete messages in Spam Quarantine. Administrators with Full Administration rights or Manage Quarantine modify rights can view, delete, and release spam messages from Spam Quarantine.

When you click on the subject line of a message on the Message List page, the contents appear on the Message Details page. When you finish viewing the details of that spam message, you can return to the Message List page.

Note the following Message Details page behavior:

Graphics appear as gray rectangles	The original graphics in messages are replaced with graphics of gray rectangles. The purpose is to suppresses offensive images and prevents spammers from verifying your email address. If you release the message by clicking Release, the original graphics are viewable by the intended recipient. Users cannot view the original graphics within Spam Quarantine.
Attachments cannot be viewed	The names of attachments are listed at the bottom of the message, but the actual attachments cannot be viewed from within Spam Quarantine. However, if you redeliver a message by clicking Release, the message and attachments are accessible from the inbox of the intended recipient.

Note: The "To” column in the Message List page indicates the intended recipient of each message as listed in the message envelope. Use caution when considering this information, since spammers oftentimes forge this header.

See “About navigating Spam Quarantine” on page 264.

See “Configuring Spam Quarantine for administrator-only access” on page 260.

See “How Spam Quarantine differs for administrators and users” on page 262.

To view spam and suspected messages in quarantine

- Do one of the following:

If you are not on the Spam Message Quarantine page	In the Control Center, click **Spam > Quarantine > Email Spam**.
If you are on the Spam Message Quarantine page and want to see newly arrived messages	Click **Show Filters** if the Display All option is not visible, and click **Display All**.

To view the contents of a spam and suspected message

- Click on the subject line of the message.

  The Message Details page appears.

To return to the Message List page from the Message Details page

- To return to the message list, click **Back to Messages**.

How Spam Quarantine differs for administrators and users

Table 12-1 describes the differences between the Spam Quarantine for administrators and users.

Table 12-1 Spam Quarantine differences between administrators and users

Page	Differences
Message List page	The Message List page has the following differences: ■ Users can only view and delete their own quarantined messages. Quarantine administrators can view and delete all users' quarantined messages, either one by one, deleting all messages, or deleting the results of a search. ■ When users click Release, the message is delivered to their own inbox. When a Quarantine administrator clicks Release, the message is delivered to the inbox of each of the intended recipients. ■ The administrator Message List page includes a "To" column that contains the intended recipient of each message. Users can only see their own messages, so the "To" column is unnecessary. ■ Users only have access to Spam Quarantine, not the rest of the Control Center. **Note:** Users access Spam Quarantine by logging into the Control Center. They use the user name and password that your LDAP server requires.
Message Details page	Users can only view and delete their own quarantined messages. Quarantine administrators can view and delete messages for all users.
Search filters	Quarantine administrators can search for recipients. In the search results, users can only delete their own quarantined messages. Quarantine administrators can delete all users' quarantined messages.

See “Configuring Spam Quarantine for administrator-only access” on page 260.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

Viewing spam and suspected spam messages sent to the postmaster mailbox

If Spam Quarantine cannot determine the proper recipient for a message that it receives and it is configured not to delete such messages, it delivers the message to a postmaster mailbox accessible from Spam Quarantine. Your network may also have a postmaster mailbox that you access with a mail client that is separate from Spam Quarantine postmaster mailbox. Spam messages may also be delivered to the Spam Quarantine postmaster mailbox if there is a problem with the LDAP configuration.

No notification messages are sent to the postmaster mailbox.

You must have Full Administration Rights or Manage Quarantine view or modify rights to view the messages in the postmaster mailbox.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

See “Configuring Spam Quarantine for administrator-only access” on page 260.

See “How Spam Quarantine differs for administrators and users” on page 262.

To view spam and suspected spam messages sent to the postmaster mailbox

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click **Show Filters**.
3. In the To box, type **postmaster**.
4. Specify additional filters as needed.
5. Click **Display Filtered**.

About navigating Spam Quarantine

The following icons show how to navigate through the Spam Quarantine message list page:

⏮	Go to beginning of messages
⏭	Navigate to last page of messages or 50 pages ahead if there are more than 50 pages.
◀	Go to previous page of messages
▶	Go to next page of messages
1-10 ▾	Choose up to 500 entries per page before or after the current page of messages

The following icons show how to navigate through the Spam Quarantine message details page:

◀	Go to the previous message
▶	Go to the next messages

When you navigate to a different page of messages, the status of the check boxes in the original page is not preserved. For example, assume that you select three messages on the first page of messages and then move to the next page. When you return to the first page, all of the message check boxes are cleared.

See “About quarantining spam” on page 258.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “How Spam Quarantine differs for administrators and users” on page 262.

Specifying the number of entries to appear on the Spam Quarantine Message list page

You can specify the number of entries that appear at a time on the message list page. You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Spam Quarantine.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

To specify the number of entries to appear on the Spam Quarantine message list page

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. On the **Entries per page** drop-down list, click a number.

Sorting spam and suspected messages in quarantine by date

You can sort messages in Spam Quarantine to make it easier to categorize the messages or locate a specific message. By default, messages appear in date descending order. The newest messages are listed at the top of the page.

A triangle appears in the date column that indicates ascending or descending sort order. Click on the column heading to toggle between ascending and descending sort order. By default, messages are listed in date descending order, meaning that the newest messages are listed at the top of the page.

You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Spam Quarantine.

See "Viewing spam and suspected messages in quarantine" on page 261.

See "About navigating Spam Quarantine" on page 264.

See "Searching quarantined spam messages" on page 268.

See "Viewing spam message headers" on page 271.

To sort spam and suspected messages in quarantine by date

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click the **Date** column heading to sort messages by date.

Releasing false-positive messages from quarantine

Occasionally you may see messages in Spam Quarantine that are not spam. You can redeliver these messages to the intended recipient. When you redeliver a message, it also removes the message from Spam Quarantine. Depending on how you configure Spam Quarantine, a copy of the message can also be sent to an administrator, Symantec, or both. This configuration lets the email administrator or Symantec monitor the effectiveness of the spam settings and filters.

See "Specifying who to notify of false positive messages" on page 274.

Only the administrators that have Full Administration rights or Manage Quarantine modify rights can release messages from quarantine.

See "Viewing spam and suspected messages in quarantine" on page 261.

See "About navigating Spam Quarantine" on page 264.

See "Modifying Spam Quarantine thresholds" on page 272.

See "Specifying when and how often Spam Quarantine is expunged" on page 284.

See "Deleting spam messages in quarantine" on page 267.

To release false-positive messages from the quarantine message list page

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the check box to the left of a misidentified message and then click **Release** to redeliver the message to the intended recipient.

To release false-positive messages from the quarantine message details page

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the subject line of the spam message that you want to review and possibly redeliver.
3. On the message details page, click **Release**.

Deleting spam messages in quarantine

Delete spam messages from Spam Quarantine to free up disk space. When you delete a message in the administrator's Spam Quarantine, you also delete it from the user's Spam Quarantine. For example, assume that you delete spam messages in the administrator's Spam Quarantine. The users to whom those messages are addressed cannot view the messages in their Spam Quarantine.

Users remove messages from the Quarantine when they release them or delete them. When you or a user deletes a message, it is no longer accessible.

See “Releasing false-positive messages from quarantine” on page 266.

You can delete messages from the message list page or from the message details page. Only the administrators that have Full Administration rights or Manage Quarantine modify rights can delete messages in quarantine. Users do not need special permissions to delete messages from their own quarantine.

You can also use an Expunger to automatically delete messages from Spam Quarantine. The Expunger frees you from having to manually delete messages from Spam Quarantine to free up space. The Expunger purges messages based on the schedule that you specify.

See “Specifying when and how often Spam Quarantine is expunged” on page 284.

To delete individual messages from the message list page

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the check box to the left of each message that you want to delete.
3. Click **Delete**.

To delete all messages from the message list page

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click **Delete All** to delete all the messages in Spam Quarantine, including those on other pages.

   This task deletes all of the spam messages in the users' Spam quarantine. Users see no mail in their quarantine.

To delete spam messages from the message details page

1 In the Control Center, click **Spam > Quarantine > Email Spam**.

2 Click on the subject line of the message that you want to view.

3 To delete the message that you are currently viewing, click **Delete**.

Searching quarantined spam messages

You can search for messages in Suspect Virus Quarantine. The ability to search messages lets you more easily find a specific message that you want to view or delete. You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Spam Quarantine.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

See “Sorting spam and suspected messages in quarantine by date” on page 265.

To search quarantined spam messages

1 In the Control Center, click **Virus > Quarantine > Email Spam**.

2 On the message list page, click **Show Filters**.

3 Do any of the following to perform a search:

To search message envelope "To" recipient	Type a name or address in the To box to search the message envelope `RCPT TO:` header. You can search for a display name, the user name portion of an email address, or any part of a display name or email user name. If you type a full email address in the To box, Symantec Brightmail Gateway searches only for the user name portion of `user_name@example.com`. The search is limited to the envelope `To:`, which may contain different information than the header `To:` that appears on the message details page. You can search for the domain portion of an email address by typing the domain.
To search "From" headers	Type a name or address in the From box to search the `From:` header in all messages for a particular sender. You can search for a display name, email address, or any part of a display name or email address. The search is limited to the visible message `From:` header, which is usually forged in spam messages. The visible message `From:` header may contain different information than the message envelope.
To search the Message ID header	Type in the Message ID box to search the message ID in all messages. You can view the message ID on the message details page in Spam Quarantine by clicking Display Full Headers. In addition, most email clients can display the full message header, which includes the message ID. For example, in Outlook 2000, double click on a message to show it in a window by itself, click **View** and then click **Options**. See "Viewing spam message headers" on page 271.
To search subject headers	Type in the Subject box to search the `Subject:` header for all messages about a specific topic.
To search a time range	Select a time range from the Time Range drop-down list to display all of the messages that were received during that time range.

4 Click **Display Filtered**.

Spam message search criteria and tips

The search function is optimized for searching a large number of messages but can lead to unexpected search results.

Consider the following tips and information to help you conduct searches in Spam Quarantine:

Tokens

Tokens are matched with substring semantics. Searching for a subject with the search target <in> will match "Lowest rate in 45 years," "RE: re: Sublime Bulletin (verification)," "Up to 85% off Ink Cartridges + no shipping!," and "Re-finance at todays super low rate."

Multiple word searches

If any word in a multiple word search is found in a message, that message is considered a match. For example, searching for red carpet match "red carpet," "red wine," and "flying carpet."

Case sensitivity

All text searches are case-insensitive. For example, assume you type emerson in the From box. Messages with a From header that contains emerson, Emerson, and eMERSOn all appear in the search results.

Exact phrases

To search for an exact phrase, enclose the phrase in " " (double quotes).

Wildcards

You can use * (asterisk) to perform wildcard searches. It also functions as a logical AND character.

Single characters

Even a single character is treated as a substring target.

Special characters

You can search on special characters such as & (ampersand), ! (exclamation point), $ (dollar sign), and # (pound sign).

Multiple characteristics

If you search for multiple characteristics, only the messages that match the combination of characteristics are listed in the search results. For example, assume you type LPQTech in the From box and Inkjet in the Subject box. Only the messages that contain LPQTech in the `From:` header and Inkjet in the `Subject:` header appear in the search results.

Forged header information

Spammers usually "spoof" or forge some of the visible messages headers such as From and To and the invisible envelope information. Sometimes they forge header information using the actual email addresses or domains of innocent people or companies.

Time to perform a search

The amount of time it takes to perform the search depends on how many search boxes you use and the number of messages in the mailbox. Searching in the administrator mailbox takes longer than searching in a user's mailbox.

See “Searching quarantined spam messages” on page 268.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

Viewing spam message headers

Viewing headers of spam messages may provide clues about the origin of a message. Keep in mind, however, that spammers usually forge some of the message headers. You must have Full Administration rights or Manage Quarantine view or modify rights to view messages in Spam Quarantine.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

See “Releasing false-positive messages from quarantine” on page 266.

See “Sorting spam and suspected messages in quarantine by date” on page 265.

To view full spam messages headers

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the subject line of the message that you want to view.
3. To display all headers available to Spam Quarantine, click **Display Full Headers**.

To view brief spam messages headers

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the subject line of the message that you want to view.
3. To display only the `From:`, `To:`, `Subject:`, and `Date:` headers, click **Display Brief Headers**.

Choosing the language encoding for spam messages

In most cases, the Auto-detect setting properly determines the language encoding for a message in Spam Quarantine. However, the Control Center may not be able to determine the proper language encoding for some messages. If the message is garbled, select the language encoding most likely to match the encoding that is used in the message.

Only the administrators that have Full Administration rights or Manage Quarantine modify rights can choose language encoding for messages in quarantine.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “About navigating Spam Quarantine” on page 264.

See “Sorting spam and suspected messages in quarantine by date” on page 265.

To choose the language encoding for spam messages

1. In the Control Center, click **Spam > Quarantine > Email Spam**.
2. Click on the subject line of the message that you want to view.
3. On the message details page, select the language encoding in the drop-down list.

Modifying Spam Quarantine thresholds

Spam Quarantine thresholds let you control the maximum size for Spam Quarantine. You can use the Expunger to enforce Spam Quarantine threshold settings. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

Note: Since the Expunger maintains the maximum size of Spam Quarantine, the quarantine can exceed maximum thresholds until the Expunger runs at the next scheduled interval.

See “Specifying how long spam messages are retained in quarantine” on page 283.

See “Specifying when and how often Spam Quarantine is expunged” on page 284.

Before you modify the Spam Quarantine thresholds, ensure that you understand the implications and considerations.

See “Spam Quarantine threshold considerations” on page 273.

Table 12-2 describes the Spam Quarantine thresholds that you can configure.

Table 12-2 Spam Quarantine Thresholds

Threshold	Description
Maximum size of quarantine	Maximum amount of the disk space that is used for quarantined messages for all users. The maximum Quarantine size reflects the actual size on disk of each message file in the message store. The actual disk usage may be slightly higher due to other unaccounted disk usage, such as database tables and indexes.

Table 12-2 Spam Quarantine Thresholds *(continued)*

Threshold	Description
Maximum size per user	Maximum amount of the disk space that is used for quarantine messages per user.
Maximum number of messages	Maximum number of messages for all users (a single message sent to multiple recipients counts as one message).
Maximum number of messages per user	Maximum number of quarantine messages per user.

To modify Spam Quarantine thresholds

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. Under **Thresholds**, for each type of threshold that you want to configure, check the box and enter the size threshold or message threshold.

   You can configure multiple thresholds.
3. Click **Save**.

Spam Quarantine threshold considerations

Table 12-3 describes the issues that you should consider before you modify Spam Quarantine thresholds.

Table 12-3 Issues to consider before you modify Spam Quarantine thresholds

Issues to consider	Details
Thresholds are enforced when the Expunger runs. The Expunger deletes older messages to enforce thresholds.	Thresholds may be exceeded temporarily until the next Expunger run.
Per-user thresholds are the most processing intensive to enforce.	Per-user thresholds are not recommended for larger deployments, such as those with over 5000 users. When the Expunger runs, per-user thresholds are checked and enforced before the other thresholds.

Table 12-3 Issues to consider before you modify Spam Quarantine thresholds *(continued)*

Issues to consider	Details
The "Maximum size of quarantine database threshold" and "Maximum number of messages" threshold provide the most precise control over disk usage and message count.	Spam Quarantine searches run faster with fewer messages.
Shortening the Spam Quarantine message retention period can also limit the size of Spam Quarantine.	A more efficient method to manage Spam Quarantine size is to conserve disk space rather than using Spam Quarantine thresholds.
No alert or notification occurs if the specific Spam Quarantine thresholds are exceeded.	You can configure an alert for **Usage of the maximum configured disk space for Spam Quarantine exceeds**. See "Types of alerts" on page 615.

See "Modifying Spam Quarantine thresholds" on page 272.

Specifying who to notify of false positive messages

If users or administrators find false positive messages in Spam Quarantine, they can click Release. Clicking Release redelivers the selected messages to the user's normal inbox. You can also automatically send a copy to a local administrator, Symantec, or both. These messages should be sent to an administrator who monitors misidentified messages at your organization to determine the effectiveness of Symantec Brightmail Gateway.

Symantec Security Response analyzes message submissions to determine if filters need to be changed. However, Symantec Security Response does not send confirmation of the misidentified message submission to the administrator or the user submitting the message. Nor is there any guarantee that filters are altered based on those submissions.

Note: Symantec does not consider the legitimate messages that receive a suspected spam verdict to be false positives. Messages that are submitted to Symantec Security Response that receive suspected spam verdicts are not reviewed. The reason is that Symantec cannot control how organizations configure the Suspect Spam threshold value. So Symantec does not create filters or modify filters based on suspected spam verdicts. Filters that are created based on suspected spam threshold values that are set too low can impact spam effectiveness for all Symantec customers.

Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

See "Releasing false-positive messages from quarantine" on page 266.

To specify who to notify of false positive messages

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. To report misidentified messages to Symantec, under **Misidentified Messages**, click **Symantec Security Response**.

   This option is selected by default.
3. To send copies of misidentified messages to a local administrator, under **Misidentified Messages**, click **Administrator** and type the appropriate email address.

   Type the full email address including the domain name, such as `admin@symantecexample.com`. The administrator email address must not be an alias, or a copy of the misidentified message is not delivered to the administrator email address.
4. Click **Save**.

About configuring the user and distribution list notification digests

By default, a notification process runs at 4 A.M. every day. The process determines if users have new spam messages in Spam Quarantine since the last time the notification process ran. If so, it sends a message to users who have new spam to remind them to check their spam messages in Spam Quarantine. The process can also send notification digests to users on distribution lists.

Note: Notification messages and notification settings are disabled if LDAP is not configured or if administrator-only access is enabled.

By default, the notification templates for standard quarantined messages and quarantined distribution list messages are different. Separate templates let you customize the notification templates for each type of quarantined message.

See "About how spam is handled when addressed to distribution lists" on page 276.

See "Specifying when to notify users of spam messages in their quarantine" on page 277.

See "Modifying the spam notification message digest templates" on page 278.

See "Enabling notification digests for distribution lists" on page 281.

See "Selecting the notification digest format" on page 282.

About how spam is handled when addressed to distribution lists

If Spam Quarantine is enabled, a spam message that is sent to an alias with a one-to-one correspondence to a user's email address is delivered to the user's normal quarantine mailbox. For example, if "tom" is an alias for "`tomevans`," the quarantined messages that are sent to "tom" or to "`tomevans`" all arrive in the Spam Quarantine account for "`tomevans`."

Note: An "alias" on UNIX or "distribution list" on Windows is an email address that translates to one or more other email addresses. In this text, distribution list is used to mean an email address that translates to two or more email addresses.

Symantec Brightmail Gateway does not deliver a spam message that is sent to a distribution list in the intended recipients' Spam Quarantine mailboxes. Instead, the message is delivered in a special Spam Quarantine mailbox for that distribution list. However, you can configure Spam Quarantine to send notification digests about the messages in a distribution list mailbox to the recipients of that distribution list. You configure this option by selecting the Notify distribution lists check box on the Quarantine Settings page.

If the Include View link box is selected, a list of the quarantined distribution list messages is included in the notification digest. Each message has a View link that users can click to view that message in Spam Quarantine. If the Include Release link box is selected, each message that is listed in the digest has a Release link. Users can click this link to release that distribution list message without accessing Spam Quarantine. If any one recipient clicks the **Release link** for a message in the quarantined distribution list mailbox, the message is delivered to the normal inboxes of all distribution list recipients. The View link and Release link do not appear if the notification format is text only.

Table 12-4 provides an example of how messages are routed to members of distribution lists.

Table 12-4 Distribution list notification and delivery examples

Scenario	Result
A distribution list that is called mktng contains Ruth, Fareed, and Darren	Spam sent to mktng and configured to be quarantined is not delivered to the Spam Quarantine inboxes for Ruth, Fareed, and Darren.
The Notify distribution lists check box on the Quarantine Settings page is selected	Ruth, Fareed, and Darren receive email notifications about the quarantined mktng messages.
The Include View link box is selected on the Quarantine Settings page	Ruth, Fareed, and Darren can view the quarantined mktng messages by clicking on the View link in the notification digests.
The Include Release link box is also selected	Ruth, Fareed, and Darren can redeliver any quarantined mktng message by clicking on the **Release** option in the notification digest.
Ruth clicks the Release option for a quarantined mktng message	The message is delivered to the normal inboxes of Ruth, Fareed, and Darren.

See “About configuring the user and distribution list notification digests” on page 275.

See “Specifying when to notify users of spam messages in their quarantine” on page 277.

See “Modifying the spam notification message digest templates” on page 278.

See “Enabling notification digests for distribution lists” on page 281.

See “Selecting the notification digest format” on page 282.

Specifying when to notify users of spam messages in their quarantine

You can change the frequency at which notifications are automatically sent to users. The default frequency is every day. To not send notification messages, change the notification frequency to `NEVER`.

If you modify the notification frequency, keep in mind the potential impact of frequent notifications. If you have a large number of users, notifications that occur more than once daily could become overwhelming for your users. And frequent notifications can impact network performance. Symantec recommends that for larger deployments, notifications should not occur more frequently than daily.

See "About configuring the user and distribution list notification digests" on page 275.

See "About how spam is handled when addressed to distribution lists" on page 276.

See "Modifying the spam notification message digest templates" on page 278.

See "Enabling notification digests for distribution lists" on page 281.

See "Selecting the notification digest format" on page 282.

Note: You must have Full Administration or Modify rights to change notification settings.

See "Administrator rights" on page 684.

To specify when to notify users of spam messages in their quarantine

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. Under **Notification Settings**, in the **Maximum summary entries per notification** box, specify how many items to include in the summary notification message.

   The default setting is 100.
3. Click the **Notification frequency** drop-down list and select how often you want notifications sent.
4. Click the **Notification start time** drop-down lists and select hour and minute that you want notifications sent.
5. Click **Save**.

Modifying the spam notification message digest templates

The notification digest templates determine the appearance of notification messages that are sent to users as well as the message subjects and send from addresses. The default notification templates are similar to the text that is listed as follows. The distribution list notification template lacks the information about how to logon. In your browser, the text does not wrap, so you need to scroll horizontally to view some of the lines. This layout prevents unusual line breaks or extra lines if you choose to send notifications in HTML format.

```
Spam Quarantine Summary for %USER_NAME%


There are %NEW_MESSAGE_COUNT% new messages in your Spam Quarantine
```

```
since you received your last Spam Quarantine Summary. These messages
will automatically be deleted after %QUARANTINE_DAYS% days.

To review the complete text of these messages, go to
%QUARANTINE_URL%
and log in.
===================== NEW QUARANTINE MESSAGES =====================
%NEW_QUARANTINE_MESSAGES%
===================================================================
```

Table 12-5 describes the variables that you can use in spam notification messages.

You can reposition each variable in the template or remove it. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

See "About configuring the user and distribution list notification digests" on page 275.

See "About how spam is handled when addressed to distribution lists" on page 276.

See "Specifying when to notify users of spam messages in their quarantine" on page 277.

See "Enabling notification digests for distribution lists" on page 281.

See "Selecting the notification digest format" on page 282.

To modify the spam notification message digest templates

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. Under **Notification Settings**, click **Edit** next to Notification template.
3. In the **Encoding** drop-down list, select the character encoding for the notification message.

   ISO-8859-1 and UTF-8 are appropriate for European languages. Windows-31j, EUC-JP, and ISO-2022-JP are appropriate for Japanese.
4. In the **Send from** box, type the email address from which the notification digests appear to be sent.

   Since users can reply to the email address that you provide, type an address where you can monitor users' questions about the notification digests. Specify the full email address, which includes the domain name, such as `admin@symantecexample.com`.

5 In the **Subject** box, type the text that should appear in the `Subject:` header of notification digests, such as "Your Suspected Spam Summary."

Use of message variables in the subject box is unsupported.

The Send from settings and Subject settings are the same for both the user notification template and distribution list notification template.

6 Edit the user notification and distribution list notification as necessary.

See Table 12-5 on page 280.

Refrain from using manually insert line breaks if you plan to send notifications in HTML.

7 Click one of the following icons:

Save	Saves and applies your changes.
Default	Erase the current information and replace it with default settings.
Cancel	Discard your changes to the notification template and close the template editing window.

8 Click **Save** on the Quarantine Settings page.

Spam notification message variables

Table 12-5 lists the spam notification message variables.

Table 12-5 Notification Message Variables

Variable	Description
%USER_NAME%	User name of user receiving the notification message.
%NEW_MESSAGE_COUNT%	Number of new messages in the user's Spam Quarantine since the last notification message was sent.
%QUARANTINE_DAYS%	Number of days messages in Spam Quarantine are kept. After that period, messages are purged.
%QUARANTINE_URL%	URL that the user clicks on to display the Spam Quarantine logon page.

Table 12-5 Notification Message Variables *(continued)*

Variable	Description
%NEW_QUARANTINE_MESSAGES%	List of messages in the user's Spam Quarantine since the last notification was sent. For each message, the contents of the `From:`, `Subject:`, and `Date:` headers are printed. View and Release links are displayed for each message if they are enabled and you've chosen a Multipart or HTML notification format.

See "About configuring the user and distribution list notification digests" on page 275.

See "About how spam is handled when addressed to distribution lists" on page 276.

See "Specifying when to notify users of spam messages in their quarantine" on page 277.

See "Modifying the spam notification message digest templates" on page 278.

See "Enabling notification digests for distribution lists" on page 281.

See "Selecting the notification digest format" on page 282.

Enabling notification digests for distribution lists

You can configure Spam Quarantine to send notification digests about the messages in a distribution list mailbox to the recipients in a distribution list. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

See "About configuring the user and distribution list notification digests" on page 275.

See "About how spam is handled when addressed to distribution lists" on page 276.

See "Specifying when to notify users of spam messages in their quarantine" on page 277.

See "Modifying the spam notification message digest templates" on page 278.

See "Selecting the notification digest format" on page 282.

To enable notification digests for distribution lists

1 In the Control Center, click **Spam > Settings > Quarantine Settings**.

2 Under **Notification Settings**, check **Notify distribution lists**.

3 Click **Save**.

Selecting the notification digest format

The notification digest template determines the MIME encoding of the notification message that is sent to users. It also determines as whether **View** and **Release** links appear in the message.

Details about the View and Release links are as follows:

View	When a user clicks on the View link in a notification digest message, the selected message appears in Spam Quarantine in the default browser. This check box is only available if you choose Multipart (HTML and text) or HTML only notification format. If you remove the `%NEW_QUARANTINE_MESSAGES%` variable from the notification digest template, the new message summary (including the View links) are not available.
Release	The **Release** link is for misidentified messages. When a user clicks on the **Release** link in a notification digest message, the adjacent message is released from Spam Quarantine and sent to the user's normal inbox. This check box is only available if you choose Multipart (HTML and text) or HTML only notification format. If you remove the `%NEW_QUARANTINE_MESSAGES%` variable from the notification digest template, the new message summary (which includes the **Release** links) are not be available.

See “About configuring the user and distribution list notification digests” on page 275.

See “About how spam is handled when addressed to distribution lists” on page 276.

See “Specifying when to notify users of spam messages in their quarantine” on page 277.

See “Modifying the spam notification message digest templates” on page 278.

See “Enabling notification digests for distribution lists” on page 281.

Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

To select the notification digest format

1. In the Control Center, click **Spam > Settings > Quarantine Settings**.
2. Under **Notification Settings**, click one of the following items in the **Notification format** drop-down list:

Multipart (HTML and Text)	Send notification messages in MIME multipart format. Users see either the HTML version or the text version depending on the type of email client they use and the email client settings. The View and Release links do not appear next to each message in the text version of the summary message.
HTML only	Send notification messages in MIME type text/html only.
Text only	Send notification messages in MIME type text/plain only. If you choose Text only, the View and Release links do not appear next to each message in the summary message.

3. Check the **Include View link** box to include a **View** link next to each message in the notification digest message summary.
4. Check the **Include Release link** box to include a **Release** link next to each message in the notification digest message summary.
5. Click **Save**.

Specifying how long spam messages are retained in quarantine

You can change the amount of time spam messages are kept before being deleted. You may want to shorten the retention period if quarantined messages use too much disk space. However, a shorter retention period increases the chance that users may have messages deleted before they have a chance to check them. The default retention period is 7 days. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

By default, the Expunger runs at 1 A.M. every day to delete messages older than the retention period. For example, if you have a retention period of 7 days, when the Expunger runs it deletes all messages older than 7 days. The Expunger also deletes messages as necessary to enforce the Spam Quarantine message and size thresholds.

See “Specifying when and how often Spam Quarantine is expunged” on page 284.

See “Modifying Spam Quarantine thresholds” on page 272.

You can also delete messages manually from Spam Quarantine.

See "Deleting spam messages in quarantine" on page 267.

To specify how long spam messages are retained in quarantine

1 In the Control Center, click **Spam > Settings > Quarantine Settings**.

2 Under **Spam Quarantine Expunger**, in the **Days to store in Quarantine before deleting** field type the number of days.

3 Click **Save** on the Quarantine Settings page.

Specifying when and how often Spam Quarantine is expunged

You can specify the time that the Quarantine Expunger begins the purge process and how frequently the purge process occurs. The Expunger lets you keep Spam Quarantine at a manageable size. Messages that are purged cannot be retrieved. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify quarantine settings.

See "Specifying how long spam messages are retained in quarantine" on page 283.

See "Modifying Spam Quarantine thresholds" on page 272.

See "Deleting spam messages in quarantine" on page 267.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See "About scheduled tasks" on page 621.

To specify when and how often Spam Quarantine is purged

1 In the Control Center, click **Spam > Settings > Quarantine Settings**.

2 Click the **Quarantine Expunger frequency** drop-down list to specify how often the Expunger runs.

3 In the **Quarantine Expunger start time** drop-down lists, specify the time that you want the Expunger to start.

4 Click **Save**.

Troubleshooting Spam Quarantine

Table 12-6 lists some problems that may occur with Spam Quarantine.

Table 12-6 Spam Quarantine issues

Issue	Description/solution
Error in log file "error.mail.transport. connect=Cannot release mail, cannot connect to any available MTA service"	This error can occur if the IP address of the Control Center is not specified for inbound and outbound mail settings on the **Administration > Hosts > Configuration Add or Edit page, SMTP** tab. See "About Scanner email settings " on page 86.
Users do not see distribution list messages in their Spam Quarantine	A Scanner does not deliver a spam message that is sent to a distribution list in the intended recipients' Quarantine mailboxes. Instead, the message is delivered to a special Spam Quarantine mailbox for that distribution list. See "About how spam is handled when addressed to distribution lists" on page 276.
Undeliverable quarantined messages go to Spam Quarantine postmaster	If Spam Quarantine cannot determine the proper recipient for a message that is received by Symantec Brightmail Gateway, it delivers the message to a postmaster mailbox accessible from Spam Quarantine. Alternatively you can specify Delete message sent to unresolved email addresses in the Quarantine Settings page. Your network may also have a postmaster mailbox you access using a mail client that is separate from the Spam Quarantine postmaster mailbox. If the LDAP server fails or has been improperly configured, however, spam messages to non-LDAP-recognized addresses are held in the Scanner's deferred queue (up to the delivery retry and timeout limits) and not in the Spam Quarantine postmaster mailbox. **Note:** No notification messages are sent to the postmaster mailbox. See "Viewing spam and suspected spam messages sent to the postmaster mailbox" on page 263.
Users receive notification messages, but cannot access messages	If users who cannot access their messages are in a different Active Directory domain from users who can access their messages, configure Directory Integration in the Control Center to use a Global Catalog. Alternatively, ensure that you have a directory data source for each active directory domain. To configure access to an Active Directory Global Catalog, specify the port for the Global Catalog, usually 3268, on the **Administration > Settings > LDAP/Edit** page. See "Creating a data source" on page 491.

Table 12-6 Spam Quarantine issues *(continued)*

Issue	Description/solution
Duplicate messages appear in Spam Quarantine	You may notice multiple copies of the same message when logged into Spam Quarantine as an administrator. When you read one of the messages, all of them are marked as read. This behavior is intentional. If a message is addressed to multiple users at your company, Spam Quarantine stores one copy of the message in its database, although the status (read, deleted) of each user's message is stored per-user. Because the administrator views all users' messages, the administrator sees every user's copy of the message. If the administrator clicks on Release, a copy of the message is redelivered to each affected user mailbox.
Maximum number of messages in Spam Quarantine	The total Quarantine size is calculated by summing the block size (size on disk) of each message file in the message store. Be aware that the actual disk usage will be higher due to other unaccounted disk usage such as database tables and indexes. See "Modifying Spam Quarantine thresholds" on page 272.
Message "Cannot release the message" appears	This message may occur if there is a problem with message traffic on your inbound or outbound mail flow. It occurs when the message cannot be delivered to any of the configured non-local relays (default) or cannot be delivered to the SMTP host configured on the **Adminstration > Control Center > SMTP** page. It may also occur if a Scanner is not installed on the same appliance as the Control Center and the SMTP host setting has not been set for a host that has an MTA. This causes delivery problems for releasing messages from Spam Quarantine. See "Configuring Control Center SMTP settings for alerts and reports" on page 673.
Quarantining spam messages and suspected spam messages takes longer than expected	This issue might be the result of slow access to the LDAP server. Try the following tasks: ■ If the LDAP server is configured for Active Directory as a global catalog server, ensure that you use port 3268 instead of 389. See "Creating a data source" on page 491. ■ Check the base DN query. A more specific DN query returns faster than a more general one. For example, "ou=quarantine-test,dc=brightmail,dc=com" is a more specific query than "dc=brightmail,dc=com."

Table 12-6 Spam Quarantine issues *(continued)*

Issue	Description/solution
When an administrator clicks on one spam message, multiple messages are marked as 'read'	This situation occurs when the same message is sent to different recipients. When an administrator accesses one instance of the message, then all instances of that message (regardless of the recipient) are marked as read. The reason this situation occurs is because the administrator read flag is associated to the message itself. This situation does not occur for end-user Spam Quarantine. That read flag is associated to each message/recipient combination.

Chapter 13

Filtering IM

This chapter includes the following topics:

About IM

Symantec Brightmail Gateway offers enterprises a gateway-based instant messaging (IM) solution. Along with its email security solutions, Symantec Brightmail Gateway provides threat protection solutions to your enterprise for IM through the features that are described in Table 13-1.

You can install Instant Messaging filtering on a Scanner by itself or on a Scanner with email filtering. The IM settings that you configure in the Control Center apply to all the Scanners on which IM is filtered.

After you deploy a Scanner with Instant Messaging filtering, configure your enterprise DNS to route IM messages from your IM users to that Scanner. You must then configure the DNS for the Scanner to route IM messages to their public IM networks over the Internet.

The IM messages that are routed through your Instant Message–filtering Scanner are subject to the settings that you define in the Control Center.

Table 13-1 describes the features that Symantec Brightmail Gateway uses to detect, block, and monitor certain IM-related activity.

Table 13-1 IM Features

Feature	Description
Network Access Control	You can create policies to block access to the IM networks that you do not use. You can also specify whether file transfers and extended features are enabled for the networks that you do use. When you block access to an IM network, each IM-filtering Scanner is prevented from connecting to that network's public IM network servers. The IM client notifies IM users who attempt to sign on to a blocked IM network that the connection attempt failed. See "Blocking access to an IM network" on page 307.
Network Status	You can view the connection status of each IM network that you support from each Scanner that is in your corporate network. Viewing the network connection status may help you to troubleshoot IM connectivity problems after you configure a Scanner and your DNS. See "Viewing the connection status of your IM networks" on page 665.
Registered Users	You can register your IM users with Symantec Brightmail Gateway to ensure that only qualified users within your organization can sign on. Unregistered IM users can sign on to their IM networks but cannot send IM messages or transfer files to other IM users. Unregistered IM users who attempt to send IM messages are given the option to self-register at that time if they qualify. See "About registering IM users" on page 297.

Table 13-1 IM Features *(continued)*

Feature	Description
Active Users	You can view a list of both the registered and unregistered IM users that are currently signed on. IM users that are currently signed on are known as active IM users. You can view all active IM users, or you can create a filter to display only the active IM users that you want to view. See "Viewing IM users that are signed on" on page 662.
IM Virus Policies	You can create policies to allow or block your IM users from performing file transfers. Transferred files can contain viruses that may pose a security threat to your corporate network when opened. If you allow file transfers, Symantec Brightmail Gateway scans each file for known viruses, worms, and other threats and, optionally, blocks infected files from delivery. In addition to known threats, Symantec Brightmail Gateway also detects unscannable and encrypted files. Because these files cannot be scanned, Symantec Brightmail Gateway considers them a security threat and treats them as it would a virus. See "Creating IM virus policies" on page 309.

Table 13-1 IM Features *(continued)*

Feature	Description
IM Spim Policies	Symantec Brightmail Gateway detects spim through the following features: ■ Known spim detection Symantec Brightmail Gateway periodically downloads the latest virus definitions, worm signatures, and spim signatures from the Symantec Security Response. You can configure the appliance to scan and, optionally, block the IM messages that contain these known threats. ■ Heuristic-based spim detection Symantec Brightmail Gateway can detect heuristic-based IM activity, such as a URL that is sent in multiple IM messages in rapid succession. You can configure the appliance to detect these threats and share them with Symantec Security Response for download to other organizations that use Symantec Brightmail Gateway. See "Creating IM spim policies" on page 311.
Reports	Based on your IM settings, you can generate reports that detail IM-related activity, such as spim and virus detection and file transfer blocking. See "About working with reports" on page 568.

You must have a valid license to perform antispim scanning functions.

See "Licensing your product" on page 678.

Enabling IM filtering

You can enable instant messaging for a host when you install Symantec Brightmail Gateway or configure a Scanner to route instant messages and file transfers. You must configure an external DNS server to enable IM scanning.

See "Specifying DNS server addresses" on page 80.

For Yahoo IM, you can specify whether you are using the standard Yahoo IM network or the Japanese IM network. However, you cannot specify both.

Use the following procedure to enable instant messaging for a host.

To configure instant messaging for a host

1. Click **Administration > Hosts > Configuration**.
2. Check the Scanner for which you want to configure instant messaging.
3. Click **Edit**.
4. Click the **IM** tab.
5. Click **Enable IM on this Scanner**.
6. Under Outbound IM Interface, from the **Ethernet** drop-down list, select the Ethernet network interface that IM clients use to communicate internally.

   Note: You must have an Ethernet interface already configured for this Scanner. Use the Ethernet tab to configure an outbound Ethernet network if necessary.

   See "Configuring Ethernet settings and routes" on page 83.

7. Under Outbound IM Interface, select outbound IP addresses from the **Outbound IP address** and **Secondary IM IP address** drop-down lists.

   The outbound IP address listens for incoming instant messages. You can use the same outbound IP address for outbound IM as you use for incoming email.

   The secondary IP address routes file transfers through the Scanner. The secondary IP address must be a different address from the outbound IP address. You can, however, use the same inbound IP address for inbound IM as you use for outgoing mail.

8. Under Inbound IM Interface, from the **Ethernet** drop-down list, select the Ethernet network that IM clients use to communicate with public servers.

   The same Ethernet network card can be used for both the internal IM interface and the inbound IM interface.

9 Under Inbound IM Interface, select the inbound IP address from the **Inbound IP address** drop-down list.

The inbound IP address can be the same as the outbound IP address. You can, however, assign an inbound IP address for your Inbound IM Interface to a different card from the one that you use for the Outbound IP address of your Outbound Interface. If you are using different Ethernet networks for incoming and outgoing email traffic, it is best to assign one Ethernet network to the Outbound IP address and the other Ethernet network card to the Inbound IP address.

10 Under Network Options, click the version of Yahoo IM in use on your network. The Standard version is selected by default.

11 Check or uncheck **Limit AOL file transfers to known ports** as desired.

12 Click **Save**.

About spim

Symantec Brightmail Gateway is integrated with Symantec Security Response to protect your network against the threats posed by spim.

Spim is unsolicited commercial IM messages that typically contain links to the Web sites that a spimmer wants to market. Bots that simulate IM users are used to generate spim. These bots send spim to a pre-determined set of IM screen names that are generated randomly, or are harvested off of the Internet. Ill-intentioned programmers also generate spim to infect other computer systems with malware (malicious software), such as viruses, worms, and Trojan Horses. This malware is designed to infiltrate or damage computer systems without the owners' informed consent.

About how spim works

Spim is often generated from a computer system that is infected with malware, such as viruses, worms, and Trojan horses.

A virus is a computer program that attaches itself to another computer program. When the infected program is run, the virus is activated and attaches itself to another program. For example, an activated virus can attach itself to a messenger service on your computer. Once the virus detects your IM screen name, it can then send IM messages to users on your contact list using your screen name. These messages typically contain URLs that, when clicked, direct the user to a Web site that the spimmer wants to market. These URLs can also contain links that, when clicked, automatically download viruses through which the computer can be accessed.

Spim takes advantage of the social aspect of instant messaging. Most IM users regard spim as safe content because they are under the impression that it was sent to them from a trusted source. Most spim typically contains friendly, inviting text—such as "Hey, check this site out!"—along with the targeted URL. Such messages can lead spim recipients to believe that the URLs point to interesting Web sites that contain inoffensive online pictures or games.

The malware that generates spim can spread from computer to computer in a number of different ways. Viruses and Trojan horses typically spread to other computers by users who unknowingly share infected files, most often through email. Worms typically spread to other computers by using the originating computer's transport features.

About detecting spim

Symantec Brightmail Gateway is integrated with Symantec Security Response to protect your network against the threats that spim poses through Instant Messaging providers. Spim detection does not detect the malicious software that generates spim. Instead spim detection blocks known spim that an external IM user or bot sends to your IM users. It also blocks spim that an internal IM user can unknowingly spread from a computer that is infected with a virus. Finally, it blocks the IM messages that are suspected of containing spim based on a configured set of heuristic-based rules.

Symantec Brightmail Gateway detects spim through the following features:

- Known spim detection
  Symantec Brightmail Gateway periodically downloads the latest virus definitions, worm signatures, and spim signatures from the Symantec Security Response. By default, known spim detection is enabled.
- Heuristic-based spim detection
  Symantec Brightmail Gateway can detect heuristic-based IM activity, such as a URL that is sent in multiple IM messages in rapid succession. These threats are then shared with Symantec Security Response for download to other organizations that use Symantec Brightmail Gateway. By default, heuristic-based spim detection is enabled.

After you configure your IM spim settings, you can create spim policies to scan and optionally block the IM messages that contain spim. If you allow a suspected IM message to be delivered, you can append it with an annotation to alert the recipient. You can also send a notification to the sender, recipient, or both, to indicate that a suspected IM message has been deleted or delivered.

See “Creating IM spim policies” on page 311.

You can generate reports that provide spim-related statistics.

See “About working with reports” on page 568.

About detecting known spim

When you create a spim policy, each Scanner that is configured to filter IM downloads a group of rule sets from Symantec Security Response. These rule sets include the following:

Settings	A list of predefined settings that determine the blocking and reporting behavior of spim. Symantec Security Response defines these settings.
Known Spim	A list of words, phrases, and URLs that are known to constitute spim.
Known Spimmers	A list of IM users who are known to send spim.
White Lists	A list of URLs that are not detected as spim, even if they appear in the Known Spim list or violate heuristic-based settings. IM messages that contain URLs that appear in white lists are not blocked.

Symantec Security Response updates its list of known spim and spimmers as it identifies new instances of these threats. Symantec Brightmail Gateway updates your network with these lists by downloading them several times an hour and sending them to each IM-enabled scanner. In addition to new spim definitions, these lists also identify spim that is no longer considered a threat by Symantec Security Response.

About detecting suspected spim

Symantec Brightmail Gateway uses heuristic-based technology to detect and block the IM messages that contain content that is suspected of being spim.

Symantec Brightmail Gateway uses predefined settings from Symantec Security Response to scan each IM message for the content that is characteristic of spim. These settings define the number of times that particular content (such as a URL) appears in multiple IM messages during a specified number of seconds. For example, the URL `www.geocities.com/some_recent.pictures` is suspected of being spim if it is detected 5 times within a 75-second interval.

Suspected spim is uploaded to Symantec Security Response and subsequently downloaded by other Symantec Brightmail Gateway systems that are configured to detect heuristic-based spim. However, suspected spim is blocked only for a pre-configured length of time. (The default is 4 hours.) If the suspected spim does not continue to violate the heuristic-based spim settings during this time, it is no longer suspected of being spim. Its new status is then uploaded to Symantec Security Response and subsequently downloaded to the other Symantec Brightmail Gateway systems.

Additionally, you also receive suspected spim from other Symantec Brightmail Gateway systems that are configured to detect heuristic-based spim.

Enabling spim detection

By default, Symantec Brightmail Gateway is configured to detect both known and heuristic-based spim. You can configure Symantec Brightmail Gateway to enable or disable spim detection.

To enable or disable spim detection

1. In the Control Center, click **Spam > Settings > Scan Settings**.
2. Click the **Instant Messaging** tab.

   By default, both known spim and heuristic-based spim detection are enabled.
3. Do one or both of the following:
   - To disable known spim detection, uncheck **Enable detection of known spim attacks**.
   - To disable heuristic-based spim detection, uncheck **Enable heuristic based spim detection**.
4. Click **Save**.

About registering IM users

User registration is the process of associating an IM user's corporate email address with the user's IM screen names in Symantec Brightmail Gateway. You can register an IM user if the user's email address is included in your LDAP directory or contains one of your local domains (for example, symantec.com). You can then configure Symantec Brightmail Gateway to allow only the internal IM users that are registered to use IM.

If registration is required, unregistered IM users can sign on to their IM networks. However, unregistered users cannot send IM messages or transfer files to other IM users. Unregistered IM users who attempt to do so are given the option to self-register at that time by using IM. Symantec Brightmail Gateway notifies the unregistered IM users that they cannot use IM.

User registration is also required if you want to create IM-related policies. If user registration is required, the IM users that belong to a group must adhere to that group's policies. Default policies apply to a user that does not belong to a group, or belongs to a group that has no IM-related policies. If user registration is not required, all IM users must adhere to the default policies.

By default, user registration is not required.

About the IM user registration process

As an administrator, you can register IM users in the Control Center.

See "Registering an IM user as an administrator" on page 303.

As an IM user, you can register yourself when you attempt to send an IM message for the first time as an unregistered user. This process is known as self-registration.

See "Self-registering an IM user" on page 304.

Using a single email address, you can register an IM user in the following ways:

- Multiple IM screen names for the same IM network
  For example, you can register both john_smith and j_smith for the same AOL IM user.
- Same IM screen for multiple IM networks
  For example, you can register john_smith for both AOL and Yahoo IM for the same IM user.

Note: You cannot register the same IM screen name for multiple email addresses. For example, you cannot register john_smith for both john_smith@symantec.com and j_smith@symantec.com.

To ensure that only a valid IM user can register, Symantec Brightmail Gateway validates the user's email address with your Control Center's LDAP database. If the user's email address is not found in your LDAP database, Symantec Brightmail Gateway tries to register the user by validating the domain of the user's email address with your local domains. If the domain cannot be validated, the user cannot be registered.

You can view a list of both the registered and unregistered IM users that are currently signed on.

See "Viewing IM users that are signed on" on page 662.

Enabling IM user registration

If you want to allow only registered IM users to send IM messages and transfer files, or if you want to create IM-related policies, you must enable IM user registration. By default, IM user registration is not required.

To enable IM user registration, you must have at least one local domain configured. If you do not have any local domains configured when you enable user registration, you are directed automatically to the Domains page where you are required to add one.

See "Adding or editing domains" on page 117.

To enable user registration

1 In the Control Center, click **Protocols** > **Instant Messaging** > **Registered Users**.

2 Under Registration Settings, click **Require users to register**.

If you do not have a local domain configured for Symantec Brightmail Gateway, the following message appears: User registration requires at least one local email domain to be configured.

You are directed automatically to the Domains page.

3 (Optional) To edit the text of the registration notifications, click **Notification Text**.

See "To edit registration notification text" on page 299.

4 Click **Save**.

Editing IM user registration text

Unregistered IM users receive a series of default notifications that instruct them through the process of self-registration.

See "Self-registering an IM user" on page 304.

You can edit these notifications to be more suitable to your organization or rewrite them in another language. These notifications can be 256 characters in length and can contain a hyperlink. However, they cannot contain HTML characters.

To edit registration notification text

1 In the Control Center, click **Protocols** > **Instant Messaging** > **Registered Users**.

2 Click **Notification Text**.

3. Under Registration IM Text, edit one or more of the following notifications:

Start registration text

The IM notification that appears when an unregistered IM user tries to send an IM message.

This notification should inform the user to send the user's email address in an IM message by responding to this notification.

Default text: You must enter your screen name with your admin before you can use IM. Please specify your email address in this window.

Valid email entered

The IM notification that appears when an unregistered IM user sends a valid email address.

This notification should inform the user that a registration key was emailed to the user's email address. It should also inform the user to send the registration key in an IM message by responding to this notification.

Default text: A registration key will be sent to your email address. When you receive it, enter the registration key into this window to complete registration.

Successful registration

The IM notification that appears when an IM user successfully registers.

Default text: Registration complete. You may now use IM.

4 Under Registration Error Messages, edit one or more of the following notifications:

Invalid email address

The IM notification that appears when an IM user sends an invalid email address.

An invalid email address is an address that does not contain the @ character, contains spaces, or exceeds its maximum length. It also includes email addresses that are created for distribution lists (for example, everyone@symantec.com).

This notification should instruct the user to resend a valid email address.

Default text: The email address you entered is invalid.

Invalid Domain

The IM notification that appears when an IM user sends an email address during self-registration that does not exist in your LDAP database and does not contain a valid local domain (for example, symantec.com).

This notification should instruct the user to resend an existing email address (i.e., the user's corporate email address), or one that contains a valid local domain.

Default text: You must use your internal email address.

No local domain

The IM notification that appears when an IM user sends an email address that does not exist in your LDAP database and when you do not have any local domains configured.

This notification should instruct the user to resend an existing email address, or contact a system administrator to configure a local domain.

Default text: Unable to verify email address. No internal emails domain has been configured.

Invalid registration key

The IM notification that appears when an IM user sends an invalid registration key.

This notification should instruct the user to resend the registration key, or to type the restart command that you specify in the Restart command field.

Default text: The registration key you entered is invalid. Try again. If you want to start over, type restart.

Restart Command

The command that an IM user types to restart the user registration process.

Default restart command: Restart.

5 Under External Notification Text, edit the following notifications:

Unregistered user

The IM notification that an IM user receives when that user tries to send an IM message to an unregistered user.

Default text: This user is not permitted to send or receive messages via IM.

6 Under Registration Key Email Template, edit one or more of the following notifications:

File Encoding

From the Encoding drop-down list, select the character encoding that is used to send the email message that contains the registration key to the IM user who is self-registering.

- Unicode (UTF-8) – default
- Western European (ISO-8859-1)
- Japanese (Shift_JIS, ISO-2022-JP, or EUC-JP)
- Simplified Chinese (B2312 or GB18030)
- Traditional Chinese (Big5)
- Korean (KS_C_5601-1987)

Sent from

The email address that appears in the From field of the email message that contains the registration key.

Subject	The text that appears in the Subject field of the email message that contains the registration key.
Content	The text that appears in the body of the email message that contains the registration key. This text should instruct the user to send the enclosed registration key in an IM message. Default text: You are receiving this email because you have specified this email ID for Symantec IM Security registration. To confirm registration, please copy and paste the registration key specified below in your IM window.

7. Do one of the following:
   - Click **Default** to restore the default text for each notification.
   - Click **Save** to save your customized notifications.

Registering an IM user as an administrator

As an administrator, you can register IM users in the Control Center. For a single corporate email address, you can register as many IM screen names as you want for each IM network that you support. When you register an IM user in the Control Center, that user can sign on and use IM immediately.

To register an IM user as an administrator in the Control Center

1. In the Control Center, click **Protocols** > **Instant Messaging** > **Registered Users**.
2. Under Users, click **Add**.
3. Under Add IM User, in the **Email address** field, type the IM user's corporate email address.
4. Under IM Network Accounts, type the IM user's IM screen name for each IM network that you support.

   Separate multiple IM screen names for the same IM network with a comma.
5. Click **Save**.

Self-registering an IM user

As an unregistered IM user, you can register yourself if your email address contains a local domain that is configured for Symantec Brightmail Gateway.

See “Adding or editing domains” on page 117.

The following sequence describes the self-registration process for an unregistered IM user:

An unregistered IM user sends an IM message to another IM user.	Symantec Brightmail Gateway sends an IM message to the unregistered user that requests the user's email address. The IM message sent to the user is blocked.
The unregistered IM user sends an IM message with the requested email address to Symantec Brightmail Gateway.	If the email address (or its domain) is valid, Symantec Brightmail Gateway sends an email to the user that contains a registration key. This email instructs the user to send the registration key to Symantec Brightmail Gateway in an IM message. If the email address (or domain) is invalid, Symantec Brightmail Gateway sends an IM message to the user that requests the email address again.
The unregistered IM user sends an IM message with the requested registration key to Symantec Brightmail Gateway.	If the IM message contains a valid registration key, Symantec Brightmail Gateway sends an IM message to the user that states that the registration process is complete. The user is able to use IM. If the IM message contains an invalid registration key, Symantec Brightmail Gateway sends an IM message to the user that requests the registration key again. If the user wants to resend the email address instead, the user can type **restart**.

When you self-register, you can register only one IM screen name at a time. If you want to register an additional screen name for the same email address, you must sign on using that screen name.

Note: If the self-registration process is disrupted (for example, the IM user logs off or is disconnected), the process starts over the next time the user attempts to use IM.

To self-register as an IM user

1. Sign on to your IM network using your IM client.
2. Send an IM message to another IM user.

   The following IM message appears: You must register your screen name before you can use IM. To register, please specify your email address.
3. Type your corporate email address in the IM message window, and then click **Send**.

   The following IM message appears: A registration key will be sent to your email address. When you receive it, enter the registration key into this window to complete your registration.

   If your email address is valid, Symantec Brightmail Gateway sends you an email that contains a registration key.

   See Step 4.

   If your email address is invalid (for example, it does not contain the @ character), you receive the following IM message: The email address you entered is invalid.

   If your email address contains an invalid domain, you receive the following IM message: You must use your internal email address. To register, please specify your email address.

   Retype your corporate email address in the IM message window.
4. Copy the registration key into the message window, and then click **Send**.

   If your registration key is valid, the following IM message appears: Registration complete. You may now use IM.

   If your registration key is invalid, the following IM message appears: The registration key you entered is invalid. Try again. If you want to start over, type "restart".

   Recopy the registration key into the message window.

Editing and deleting registered IM users

As an administrator, you can add, edit, or delete the IM screen names associated with a registered IM user. You cannot edit the email address for a registered IM user. If an IM user's email address changes, you must delete and re-register the user with the new email address.

You can also delete a registered IM user, which deletes the user's email address and the IM screen names associated with it. Based on the results of your search, you can delete one or more specific IM users, or all IM users.

When you edit or delete an IM user, the changes that you make are not applied for 30 minutes. Until that time, the IM user's previous settings are in effect. For example, an IM user that you delete can temporarily sign on and use IM.

You can search for one or more of the registered IM users that you want to edit or delete by specifying one of the following:

- Email address
  The email address of the IM user. You can specify the complete email address of a single IM user, or you can use a wildcard (*) to represent one or more characters within an email address to search for multiple IM users. For example, if you specify **john_smith@*.com**, IM users john_smith@hotmail.com and john_smith@gmail.com appear.
- IM account
  The screen name of the IM user. You can specify the complete screen name of a single IM user, or you can use a wildcard (*) to represent one or more characters within a screen name to search for multiple IM users. For example, if you specify **jsmith***, IM users jsmith1 and jsmith2 appear.
- IM network
  The IM network of the IM user. If you search by IM network, all of the registered IM users for the network that you specify appear.

You can specify the number of registered IM users that you want to appear on each page of your search results based on increments of 10, 25, 50 or 100. Based on the increment that you specify, you can navigate immediately to any page that contains additional search results. For example, if you specify that you want 10 users to appear on each page, and your search results yield 100 users, you can navigate to the page that contains users 21-30. Using the control buttons, you can also navigate to the first page, the previous page, the next page, or the last page at any time.

To edit a registered IM user

1. In the Control Center, click **Protocols** > **Instant Messaging** > **Registered Users**.
2. Under Users, under Search for user by, select one of the following from the drop-down list:
   - Email address
   - IM account
   - IM network
3. Do one of the following:
   - If you selected Email address or IM account, type the value.

- If you selected IM network, select an IM network from the drop-down list.

4 Click **Search**.

5 Check the IM user that you want to edit, and then click **Edit**.

6 From the Edit IM User page, under IM Network Accounts, add or edit the user's IM screen names.

7 Click **Save**.

To delete a registered IM user

1 In the Control Center, click **Protocols** > **Instant Messaging** > **Registered Users**.

2 Under Users, under Search for users by, select one of the following options from the drop-down list:

- Email address
- IM account
- IM network

3 Do one of the following:

- If you selected Email address or IM account, type the value.
- If you selected IM network, select an IM network from the drop-down list.

4 Click **Search**.

5 Do one of the following:

- Check the specific IM users that you want to delete.
- Check **Email** to delete all of the IM users that are displayed.

6 Check **Delete**, and then click **OK**.

Blocking access to an IM network

You can create a policy to determine which IM networks that members of a particular group can access. You can also enable or disable file transfers and extended features for each IM network to which you allow access. Most IM clients provide extended features that allow IM users to communicate with each other by a means other than IM, such as audio or video. Extended features also include such features as application sharing and games. If you do not want your IM users to use these features, you can disable them. You can also send a notification to an IM user who attempts to use an extended feature that is disabled.

You can view the connection status of each IM network that you support from each Scanner that is in your corporate network.

See "Viewing the connection status of your IM networks" on page 665.

Symantec Brightmail Gateway is installed with a default Network Access Control policy that allows access to the following public IM networks:

- AOL
- Yahoo IM
- MSN Messenger
- Google Talk

This policy also enables file transfers and extended features for each IM network.

You can configure Symantec Brightmail Gateway to block access to an IM network that you do not support. When you block access to an IM network, each IM filtering Scanner is prevented from connecting to that network's public IM network servers. IM users that attempt to sign on to a blocked IM network are notified by the IM client that the connection attempt failed.

If you block access to an IM network on which IM users are currently signed on, those users remain signed on until they purposely sign off.

Note: Before you create a Network Access Control policy, you must first create any notifications that you want to select for that policy.

See "Creating policy violation notifications" on page 401.

Note: When you block access to an IM network, you prevent your IM users from signing on to that network from their client workstations. However, some networks allow their users to sign on by using a Web-based IM client that is available on that network's public Web site. IM conversations that occur in this manner are not directed through Symantec Brightmail Gateway; instead, they are directed through your corporate network, and may therefore pose a security threat.

To prevent unauthorized IM conversations, you must block access to Web-based IM clients. See the *Symantec Brightmail Gateway Installation Guide*.

To create a Network Access Control policy

1. In the Control Center, click **Protocols** > **Instant Messaging** > **Network Access Control**.
2. Click **Add**.
3. In the **Policy name** box, type a name for this policy.

4 Under Enabled Networks and Features, check each IM network that you want to enable for this policy.

If you enable an IM network, file transfers and extended features are automatically enabled for that network.

5 (Optional) Under each IM network, uncheck **File Transfers** if you want to disable file transfers for that network.

6 (Optional) Under each IM network, uncheck **Extended Features** if you want to disable extended features for that network.

7 Under Actions, check **If a blocked network feature is detected send the following notification** if you want to send a notification to an IM user who is blocked from sending a file or using an extended feature.

This option is only available if you disabled extended features for one of the IM networks.

8 Select a notification from the **Notification** drop-down list.

9 Under Apply to the following groups, check each group to which this policy applies.

10 Click **Save**.

Creating IM virus policies

Most IM networks allow their IM users to send files to each other. However, these files, such as EXE and BAT files, can contain viruses that may pose a security threat to your corporate network when opened. Using the Network Access Control settings, you can allow or block file transfers for each IM network that you support.

See "Selecting Network Access Control policies for a policy group" on page 325.

If you allow file transfers, you can use the Instant Messaging Virus Policies page to add, edit, copy, delete, and enable or disable IM virus policies. Symantec Brightmail Gateway scans each file for known viruses, worms, and other threats, and optionally blocks infected files from being delivered. If an infected outbound file is blocked, you can notify the sender, recipient, or both that the file was blocked.

Note: Symantec Brightmail Gateway does not allow file transfers for Google Talk.

In addition to viruses, you can also create policies for encrypted and unscannable files. Encrypted files are files that require a password to open. Unscannable files are files that have a file type that cannot be identified by the scan engine. Because these files cannot be scanned, Symantec Brightmail Gateway considers them a potential security threat.

You can specify whether to scan files that are sent in the following ways:

- Inbound
  Files that are sent from external IM users
- Outbound
  Files that are sent by your internal IM users
- Both inbound and outbound

You can generate a report that details file transfer activity.

See "Report types" on page 571.

Note: If you allow file transfers, and use a load balancer to distribute traffic to multiple Scanners, you must configure your load balancer for single affinity mode. In this mode, all TCP/IP requests that originate from the same client (IP address) are routed to the same Scanner.

Note: Before you create an IM virus policy, you must first create any notifications that you want to select for that policy.

See "Creating policy violation notifications" on page 401.

To add an instant messaging virus policy

1. In the Control Center, click **Virus > Policies > Instant Messaging**.
2. Click **Add**.
3. In the **Policy name** box, type a name for this policy.

   This name appears on the **Instant Messaging Virus Policies** page, and on the **Virus** tab when you configure a group. Content filtering, spam, and virus policy names must be unique. For example, you cannot have both a spam policy and a virus policy called XYZ.
4. Under **Conditions**, click the **Apply to** drop-down list and select the file transfers for which this policy applies.
5. Click the **If the following condition is met** drop-down list and select the type of threat to which this policy applies.

6 Under **Actions**, click the **Perform the following action** drop-down list and select the action that this policy should take when a threat is detected.

You can select more than one action provided the actions are not contradictory. For example, you can block a file transfer as well as send a notification. However, you cannot both allow and block a file transfer within the same policy.

If you select **Send notification**, you must select a notification from the **Notification** drop-down list.

7 Click **Add Action**.

8 Optionally, repeat steps 6 and 7 to add more actions.

9 Under **Apply to the following policy groups**, do one of the following:

- Check each group to which this policy applies.
- Check **Policy Groups** to apply this policy to all of your groups.

10 Click **Save**.

Creating IM spim policies

You can add, edit, copy, delete, and enable or disable IM spim policies.

You can create IM spim policies to detect whether IM messages contain known or heuristic-based spim.

See "About how spim works" on page 294.

You can configure your IM spim policies to perform one or more of the following actions when spim is detected:

Table 13-2 SPIM policy actions

Action	Description
Add annotation	Prepends the IM message with a customized annotation so that the recipient is alerted.
Delete the message	Deletes the IM message so that the recipient does not receive it.
Deliver message normally	Delivers the IM message to the recipient even though the message was detected as spim.
Send notification	Sends a customized notification to the sender of the IM message, the recipient of the IM message, or both.

You can specify the IM messages to which you want your IM spim policy to apply:

- **Inbound instant messaging spim**

  Applies to IM messages that are received from an internal or external IM user.

  IM messages, including those that are sent by your internal IM users, are sent to their IM network servers over the Internet before they are delivered to their recipients. This means that all IM messages are subject to any inbound IM spim policies that you create.
- **Outbound instant messaging spim**

  Applies to IM messages that are sent from an internal IM user.
- **Inbound and outbound instant messaging spim**

  Applies to IM messages that are sent from an internal IM user, or received from another internal IM user or an external IM user.

To add an instant messaging policy

1. In the Control Center, click **Spam > Policies > Instant Messaging**.
2. Click **Add**.
3. In the **Policy name** box, type a name for this policy.
4. Under **Conditions**, click the **Apply to** drop-down list and select one of the IM message types.
5. Under **Actions**, click the **Perform the following action** drop-down list and select an action to apply to the policy.

   You can select more than one action; however, the combination should be logical. For example, you can't both delete the message and deliver the mail. Nor can you add an annotation and delete the IM message within the same policy.

   Depending on the actions that you select, you may have to provide additional information. For example, if you select the action **Forward a copy of this message**, you must type the email address of the person to whom you want to forward the message.
6. Click **Add Action**.
7. Optionally, repeat steps 5 and 6 to add more actions.

   The actions that you add appear in the **Actions** list in the order in which you add them.
8. Under **Apply to the following policy groups**, do one of the following:
   - Check each group to which this policy applies.

- Check **Policy Groups** to apply this policy to all of your groups.

9 Click **Save**.

Chapter 14

Creating policy groups

This chapter includes the following topics:

About policy groups

You can specify configurable message management options for an unlimited number of policy groups that you define. Policy groups collect the spam, virus, and content filtering policies for a set of users.

When Symantec Brightmail Gateway processes a message, the results for each sender or recipient can differ, based on policy group membership. If an inbound message goes to recipients in more than one policy group, the message is processed for the recipients in each policy group according to the filtering policies assigned to that policy group. For an outbound message, it is the sender's policy group that determines processing.

By default, policy groups that you create are assigned the default filter policies for spam, spim, and viruses (there is no default for content filtering policies).

After you create your own filter policies, you can assign different filter policies to different policy groups.

For a particular recipient or sender who is a member of more than one policy group, only the group with the highest group precedence applies. Policy group precedence is determined by the order of groups on the Policy Groups page.

The Default policy group is always the last group in the list. You cannot change the precedence of the Default policy group.

Note: If you have enabled probe participation, the Probe Account policy group is the first group in the list. You cannot change the precedence of the Probe Account policy group. When you choose to edit the Probe Account policy group, the **Probe Accounts** page appears.

See "Enabling probe participation" on page 253.

See "Creating a policy group" on page 316.

See "Importing and exporting policy group members" on page 319.

See "Selecting virus policies for a policy group" on page 322.

See "Selecting spam and spim policies for a policy group" on page 324.

See "Selecting Network Access Control policies for a policy group" on page 325.

See "Selecting content filtering policies for a policy group" on page 326.

The policy group management options let you do the following:

- Set policy group precedence, which determines the filter policies applied to each message.
- Edit policy group membership and actions.
- Enable and disable policy groups.
- Delete policy groups.
- View policy group information for particular users.

See "Editing, deleting, enabling, or disabling a policy group" on page 320.

See "Setting policy group precedence" on page 329.

See "Researching policy group membership for a user" on page 320.

Creating a policy group

The Policy Groups page lists each policy group. The default policy group, which contains all users and all domains, always appears last. Although you can add or modify actions for the default policy group, you cannot add members to the default policy group. You cannot delete, disable, or change the precedence of the default policy group.

Note: If you have enabled probe participation, the Probe Account policy group is the first group in the list. You cannot change the precedence of the Probe Account policy group. When you choose to edit the Probe Account policy group, the **Probe Accounts** page appears.

To create a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the **Policy Groups** page, click **Add**.
3. Enter a name in the **Policy group name** text box.
4. Click **Save**.

   Once you have created a policy group you can add members.

   See "Adding members to a policy group" on page 317.

Adding members to a policy group

You can assign members to a policy group based on email addresses, domain names, or LDAP groups for the purpose of applying policies. Once you have created your policy group and added members you can select the policies that you want your group to have.

See "Selecting virus policies for a policy group" on page 322.

Note: There is no edit button for editing policy group members. Use the add / remove method for making changes to member names. For example, if you want to correct a typo in a member's name, you must delete the member then add the member again.

If you use distribution lists or groups stored in an LDAP directory as part of your policy group membership, and you make changes to the structure of your directory that causes changes to the distinguished name (DN) of any of the LDAP groups, distribution lists, or the users that are members of these LDAP entities, do the following to ensure that the policy group is applied consistently.

- If a policy group member is an LDAP DN, update the DN in the policy group if it has changed.

- Clear the cache on the directory data source that is configured to perform address resolution on those LDAP groups and distribution lists and their user members. This step should be taken even if you reference LDAP distribution lists by email address in your policy groups.
  See "About the directory data cache" on page 526.

To add a member to a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. Click the Policy Group you want to edit.
3. Ensure that the **Members** tab is displayed, and click **Add**.
4. Specify members using one of the following methods:
   - Type email addresses, domain names, or both in the box. To specify multiple entries, separate each with a comma, semicolon, or space. Use * to match zero or more characters and ? to match a single character. To add all recipients of a particular domain as members, type any of the following:

     ```
     domain.com
     @domain.com
     *@domain.com
     ```

   - Type in the LDAP distinguished name of a group in the field provided, for example: cn=some ldap group,dc=domain,dc=com.
     Any modification that would change the distinguished name string of direct members added this way must be manually updated in the policy group. To protect mailflow behavior, the new distinguished name string membership should be added to the policy prior to modifying the group and the old distinguished name string should be removed after the modifications have been completed.
   - Check the box next to one or more LDAP groups.
     The LDAP groups listed on this page are loaded from your LDAP server. At least one address resolution function must be enabled and group listing query must be configured to use the LDAP groups list. See "About using the address resolution function with your data source" on page 490.
5. Click **Add members** to add the new member(s).
6. Click **Save** on the Edit Policy Group page.

Managing policy group members

Importing and exporting policy group members

You can import policy group members from a file, and you can export group members to a file.

Note: You cannot import or export LDAP group members described by distinguished names.

To import policy group members from a file

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. Click the underlined name of the Policy Group that you want to edit.
3. On the **Members** tab of the Edit Policy Group page, click **Import**.
4. Enter the appropriate path and filename (or click **Browse** to locate the file on your hard disk), and then click Import.

   Separate each domain or email address in the plain text file with a newline. Below is a sample file:

   ```
   ruth@example.com
   rosa@example.com
   ben*@example.com
   example.net
   *.org
   ```

   The email addresses in the samples behave as follows:

   - ruth@example.com and rosa@example.com match those exact email addresses.
   - ben*@example.com matches ben@example.com and benjamin@example.com, etc.
   - example.net matches all email addresses in example.net.
   - *.org matches all email addresses in any domain ending with .org.
5. Click **Save**.

To export policy group members to a file

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. Click the underlined name of the policy group you want to edit.

3 In the Members tab of the Edit Policy Group page, click **Export**.

4 Complete your operating system's save file dialog box as appropriate.

Researching policy group membership for a user

You can identify all of the policy groups to which a user is assigned.

To research policy group membership for a user

1 Do one of the following:

- In the Control Center, click **Administration > Users > Policy Groups**, click on the name of a group, and click **Find User**.
- In the Control Center, click **Administration > Users > Find User**.

2 In the **Email address** box, type the user's email address.

3 Click **Find User**.

The Control Center lists the enabled policy groups in which the specified user exists, in order by policy group precedence.

Editing, deleting, enabling, or disabling a policy group

The following sections describe common administrative tasks for policy groups.

To edit an existing policy group

1 In the Control Center, click **Administration > Users > Policy Groups**.

2 Click the policy name or check the box next to a policy group, and then click **Edit**.

Add or delete members or change filtering actions for this policy group as you did when you created it.

See "Creating a policy group" on page 316.

To enable a policy group

1 In the Control Center, click **Administration > Users > Policy Groups**.

2 Check the box next to a policy group, and then click **Enable**.

To disable a policy group

1 In the Control Center, click **Administration > Users > Policy Groups**.

2 Check the box next to a policy group, and then click **Disable**.

Note: You cannot disable or delete the Default policy group.

To delete a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the Policy Groups page, check the box next to a policy group, and then click **Delete**.

Selecting policies for policy groups

About assigning filter policies to policy groups

By default, groups that you create are assigned the default filter policies for spam, spim, and viruses (there is no default for content filtering policies). Follow the steps in the sections below to assign different filter policies to groups.

See "Selecting virus policies for a policy group" on page 322.

See "Selecting spam and spim policies for a policy group" on page 324.

See "Selecting content filtering policies for a policy group" on page 326.

See "Selecting Network Access Control policies for a policy group" on page 325.

Virus categories and default actions

Virus policies determine what to do with inbound and outbound email and IM messages that contain any of the threat categories listed in the table below.

See "Selecting virus policies for a policy group" on page 322.

See "Default email virus policies" on page 52.

Table 14-1 Virus categories and default actions

Category	Default email action	Default IM action
Viruses	Clean the message	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender
Mass-mailing worms	Delete the message	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender

Table 14-1 Virus categories and default actions *(continued)*

Category	Default email action	Default IM action
Unscannable messages	Delete the message	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender
Encrypted attachments	Prepend `[WARNING ENCRYPTED ATTACHMENT NOT VIRUS SCANNED]` to Subject: header.	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender
Spyware or adware	Prepend `[SPYWARE OR ADWARE INFECTED]` to Subject: header.	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender
Suspicious attachments	Inbound message: Strip and Delay in Suspect Virus Quarantine. Outbound message: Hold message in Suspect Virus Quarantine.	Inbound message: Block the file transfer Outbound message: Block the file transfer and send a notification to the sender

Selecting virus policies for a policy group

For each policy group, you can specify virus policies. Symantec Brightmail Gateway comes with preconfigured virus policies that are enabled by default. When you select virus policies for a policy group, you select from these preloaded policies.

See “Default email virus policies” on page 52.

See “Default IM virus policies” on page 56.

You can, however, also make your own custom virus policies and apply them to your policy groups. Or, you can modify the default virus policies to fine tune them or to expand or reduce their scope.

See “Creating email virus policies” on page 210.

See “Modifying email virus policies” on page 212.

Note: By default inbound and outbound email messages containing a mass-mailing worm, and unscannable messages, including malformed MIME messages, will be deleted. If you are concerned about losing important messages, you may want to create a different filter policy for unscannable messages and apply that new filter policy to some or all of your groups.

To select virus policies for a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the Policy Groups page, click the group for which you want to select virus policies.
3. Click the **Virus** tab.
4. Optionally, click **View** next to any policy to view the details of that policy.
5. If desired, under Email, check **Enable inbound email virus scanning for this policy group**, and then select the desired policy from each of the following drop-down lists:
   - **Inbound email antivirus policy**
   - **Inbound email mass-mailing worm policy**
   - **Unscannable inbound email message policy**
   - **Inbound encrypted email attachment policy**
   - **Inbound suspicious email attachment policy**
   - **Inbound email spyware/adware policy**
6. If desired, under Email, check **Enable outbound email virus scanning for this policy group**, and then select the desired policy from each of the following drop-down lists:
   - **Outbound email antivirus policy**
   - **Outbound email mass-mailing worm policy**
   - **Unscannable outbound email message policy**
   - **Outbound encrypted email attachment policy**
   - **Outbound suspicious email attachment policy**
   - **Outbound email spyware/adware policy**
7. If desired, under Instant Messaging, check **Enable inbound instant messaging virus scanning for this policy group**. You can view the following policy:
   - **Inbound instant messaging antivirus policy**

This policy includes the default worm, spyware/adware, and suspicious attachment policies.

- **Inbound encrypted instant messaging attachment policy**
- **Unscannable inbound instant messaging message policy**

8 If desired, under Instant Messaging, check **Enable outbound instant messaging virus scanning for this policy group**. You can view the following policy:

- **Outbound instant messaging antivirus policy**
  This policy includes the default worm, spyware/adware, and suspicious attachment policies.
- **Outbound encrypted instant messaging attachment policy**
- **Unscannable outbound instant messaging message policy**

9 Click **Save**.

Selecting spam and spim policies for a policy group

Spam and spim policies determine what to do with inbound and outbound email messages that contain spam or suspected spam and IM messages that contain spim or suspected spim.

Symantec Brightmail Gateway installs with preconfigured default email spam and IM spim policies which you apply to policy groups.

See "Default email spam policies" on page 50.

See "Default IM spim policies" on page 55.

Note: You can also apply custom email and IM spam policies. You must first create these policies to make them available from the selection menus.

See "Creating email spam policies" on page 242.

See "Creating IM spim policies" on page 311.

To select spam or spim policies for a policy group

1 In the Control Center, click **Administration > Users > Policy Groups**.

2 On the Policy Groups page, click the policy group for which you want to select spam or spim policies.

3 Click the **Spam** tab.

4 Optionally, click **View** next to any policy to view the details of that policy.

5. If desired, under Email, check **Enable inbound email spam scanning for this policy group**, and then select the desired policy from each of the following drop-down lists:
   - **Inbound email antispam policy**
   - **Inbound email suspected spam policy**
6. If desired, under Email, check **Enable outbound email spam scanning for this policy group**, and then select the desired policy from each of the following drop-down lists:
   - **Outbound email antispam policy**
   - **Outbound suspected email spam policy**
7. If desired, under Email, check **Enable bounce attack prevention for this policy group**, and then select the desired policy from **Bounce attack prevention policy** drop-down list.

   Symantec Brightmail Gateway provides a default policy: Failed Bounce Attack Validation: Reject message. You can also edit this policy or create a new policy, which must contain the condition, **If a message fails bounce attack validation** and actions to be taken should bounce address tag validation fail.

   See "About defending against bounce attacks" on page 182.
8. If desired, under Instant Messaging, check **Enable inbound instant messaging spim scanning for this policy group**. You can view the following policy:
   - **Inbound instant messaging spim policy**
9. If desired, under Instant Messaging, check **Enable outbound instant messaging spim scanning for this policy group**. You can view the following policy:
   - **Outbound instant messaging spim policy**
10. Click **Save**.

    You cannot change spam policy details from the Edit Policy Group page.

Selecting Network Access Control policies for a policy group

Network Access Control (or NAC) policies determine which public IM networks members of a particular group can access.

By default, all of the public IM networks that are supported by Symantec Brightmail Gateway can be accessed by all of the IM users that are in your network. If you want to block access to a particular IM network for a group, enable a Network Access Control policy for that group.

See “Blocking access to an IM network” on page 307.

To select a Network Access Control policy for a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the Policy Groups page, click the policy group for which you want to select a Network Access Control policy.
3. Click the **IM NAC** tab.
4. Check **Enable Network Access Control policies for this policy group**.
5. Select the desired policy from the **Network access control policy** drop-down list.
6. Optionally, click **View** next to any policy to view the details of that policy.
7. Click **Save**.

Selecting content filtering policies for a policy group

By associating an appropriate content filtering policy with a group, you can check messages for attachment types, keywords, or match regular expressions. Depending on the message content, you can add annotations, send notifications, or copy messages to an email address. You can also use your content filtering policies to check for content filtering with statutory regulations or organizational policies.

Note: Because there are no default content filtering policies, the drop-down list on the **Edit Policy Group** page is initially blank. Before you select content filtering policies for a policy group, you must first create at least one content filtering policy.

See “Creating content filtering policies” on page 334.

To select content filtering policies for a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the **Policy Groups** page, click the policy group for which you want to select content filtering policies.
3. Click the **Content Filtering** tab.
4. Check **Enable inbound Content Content Filtering for this policy group**.

5 Select the desired policy from the **Content Filtering Policies** drop-down list.

You must already have applied the policy to the group on the **Edit Email Content Filtering Policy** page for it to appear in the drop-down list.

If you want, click **View** to see a summary of the content filtering policy, and then click **OK** to return to the **Edit Policy Group** page. As you add content filtering policies from the drop-down list, they appear in the bottom list and become unavailable in the drop-down list.

6 Click **Add**.

7 If you want, add additional policies from the **Content Filtering Policies** drop-down list.

8 To configure outbound content filtering policies for the group, check **Enable outbound Content Filtering for this policy group** and follow 5 through 7 again.

9 Click **Save**.

You cannot change content filtering policy details (such as conditions and actions) from the **Edit Policy Group** page. Although you can add existing policies to the lists on this page, you cannot add new content filtering policies from this page.

Enabling and disabling end user settings for policy groups

End-user settings determine whether end users in a policy group can log into the Control Center to perform either of the following tasks:

- Configure personal Good and Bad Senders lists.
- Block or allow email in specified languages.

Note: You must have a data source configured for address resolution for this page to be enabled, and both authentication and address resolution data sources are required for the system to execute your settings.

See “About data sources and functions” on page 482.

See “Requirements for enabling end-user settings for policy groups” on page 328.

To log in, users access the same URL in their browser as Control Center administrators: https://<hostname>. The login and password for end users is the same as their LDAP login and password. For information about supported browsers, see the *Symantec Brightmail Gateway Installation Guide*.

Note: End users are limited to a total of 200 entries in their combined Good Senders and Bad Senders lists.

Note: Although the language identification technology employed by Symantec Brightmail Gateway to identify the language of a message has a high rate of accuracy, false language identifications can occur. Note that messages identified to be in a disallowed language are deleted.

To select end user settings for a policy group

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. On the Policy Groups page, click the group for which you want to select end user policies.
3. Click the **End Users** tab.
4. Check **Enable end user settings for this policy group**.
5. If desired, check **Create personal Good and Bad Senders Lists**.
6. If desired, check **Specify language settings**.
7. Click **Save**.

Requirements for enabling end-user settings for policy groups

The following requirements must be satisfied before end users can configure their own personal Good and Bad Senders Lists and block or allow email in specified languages:

- At least one data source is enabled for authentication.
- At least one data source is enabled for address resolution.
- End user preference replication frequency must be set.
- End-user preferences must be enabled for the given policy group on the End Users tab on the Edit Policy Group page.
- Members of the policy group can only be LDAP users, not a locally defined user (that is, an email address you typed manually).

Allowing or blocking email based on language

Using the language identification offered by Symantec Brightmail Gateway, you can block or allow messages written in specified languages for a group. For example, you can choose to only allow English and Spanish messages, or block messages in English and Spanish and allow messages in all other languages.

To allow or block email based on language for a group

1. In the Control Center, click **Administration > Users > Groups**.
2. On the Groups page, click the group for which you want to select language policies.
3. Click the **Language** tab.
4. Click the desired setting.
5. If you chose **Only receive mail in the following languages** or **Do not receive mail in the following languages**, check the box for each desired language.

   Available language settings are: Chinese, Dutch, English, French, German, Italian, Japanese, Korean, Portugese, Russian, and Spanish.
6. Click **Save**.

   Although the language identification technology employed by Symantec Brightmail Gateway to identify the language of a message has a high rate of accuracy, false language identifications can occur. Note that messages identified to be written in a disallowed language are deleted.

Setting policy group precedence

The Policy Groups page lists policy groups in a specific order. Policy groups higher in the list have a higher precedence. If a user is a member of multiple groups, the policy group with higher precedence applies in determining how messages are processed for that user.

Note: The Default policy group is always the last group in the list. You cannot change the precedence of the Default policy group.

Note: If you have enabled probe participation, the Probe Account policy group is the first group in the list. You cannot change the precedence of the Probe Account policy group. When you choose to edit the Probe Account policy group, the **Probe Accounts** page appears.

See “Enabling probe participation” on page 253.

See “About policy groups” on page 315.

To set policy group precedence

1. In the Control Center, click **Administration > Users > Policy Groups**.
2. Click on the group that you want to move, and drag it up or down to the location that you want.

Chapter 15

Filtering content for violations

This chapter includes the following topics:

- About content filtering
- Creating content filtering policies
- About preventing data loss with structured data
- About policy violation notifications
- About content filtering dictionaries
- About patterns
- Annotating messages that contain violations
- About attachment lists
- Specifying where to save archived messages
- About content incident folders
- About monitoring and acting on incidents
- About encrypting messages with Symantec Content Encryption
- About Symantec Network Prevent

About content filtering

Content filtering policies determine how Symantec Brightmail Gateway evaluates email message content, their attachments, and attributes. Symantec Brightmail

Gateway scans message content for conditions and applies the actions that you specify for the groups that you select.

Some reasons to use content filtering policies are as follows:

- Block email from the marketing lists that generate user complaints or use excessive bandwidth.
- Block or redirect messages or attachments with specific content or specific file attachment types or file names.
- Block oversized messages to control message volume and preserve disk space.
- Prevent confidential or sensitive information from leaving your organization.
- Protect sensitive customer data from being sent to unauthorized individuals and organizations.
- Limit the ability of email users to communicate or conduct the activities that are contrary to your organization's values and policies.
- Ensure that employees do not send or receive any messages that violate state and federal regulations.

A content filtering policy consists of the following components:

Template

Symantec provides the predefined templates that you can use to create content filtering policies for specific scenarios. For example, the Credit Card template provides predefined conditions to detect a credit card number in incoming or outgoing email messages. You can also use a blank template to create a policy that does not relate to any of the predefined templates.

See "About content filtering policy templates" on page 336.

Conditions

A condition is a statement that Symantec Brightmail Gateway uses to evaluate messages. When a particular condition is met, Symantec Brightmail Gateway takes the actions that you specify.

See "Considerations for content filtering policy conditions" on page 347.

Actions

When you define conditions for a policy, you also define the actions to take when a message meets those conditions.

See "Specifying content filtering policy actions" on page 365.

Groups

When you create a content filtering policy, you specify the groups for whom the policy applies.

See "Specifying the policy groups for which content filtering policies apply" on page 371.

See "Creating content filtering policies" on page 334.

Symantec Brightmail Gateway provides the following resources that you can use to create a content filtering policy:

Annotations

You can append text to a message that has violated a policy. The text that you choose depends on the policy. For example, it may advise the recipient that the accompanying email violates company norms and policies for corporate governance.

See "Annotating messages that contain violations" on page 421.

Archive

Some regulations require that you archive any messages that might violate corporate policies. You can send copies of the message that you want to archive to a designated email address on a regular mail server. You can also send copies of the message to an archive server.

See "Specifying where to save archived messages" on page 437.

Attachment lists

Attachment lists are predefined lists based on an attachment's true file type, MIME-type, or file name extension. An attachment list contains the file extensions and the file application types that you want Symantec Brightmail Gateway to detect. You use an attachment list as a condition of a content filtering policy. When Symantec Brightmail Gateway detects an attachment that has an extension that is on the file extension list, it applies the action that you specify. Symantec Brightmail Gateway can determine the email attachment type based on the application that created it, regardless of its file extension.

See "About attachment lists" on page 424.

Dictionaries

Dictionaries provide lists of the predefined words that you can use when you create policy conditions. For example, you can use dictionaries to detect vulgar language or the language that might suggest a job search.

See "About content filtering dictionaries" on page 404.

Notifications

Symantec Brightmail Gateway can send customized notifications to the sender or recipient of an email whenever a policy's conditions are met. Notifications can be sent with or without the accompanying message. You can also configure the policy to notify the content filtering officer, manager, or the administrator that is charged with enforcing the policy of the violation.

See "About policy violation notifications" on page 400.

Patterns	Patterns are predefined lists of the character patterns that are associated with an object type or data type that you may want to restrict. For example, you can use patterns to screen outgoing messages for credit card numbers by searching for the standard credit card patterns. See "About patterns" on page 417.
Records	A record consists of the structured data that your organization provides. Structured data contains the company-specific, delimited data that you want to protect. You can create views of records to use in content filtering policies that detect whether content in your data source file is in messages. Content can be detected in both incoming email messages and outgoing email messages. If it is, you can specify the action that you want Symantec Brightmail Gateway to take. See "About preventing data loss with structured data" on page 375.

You can specify actions in your content filtering policies to create incidents in content incident folders. Content incident folders help you organize, monitor, and manage the incidents of content filtering policy violations.

See "About content incident folders" on page 438.

See "About monitoring and acting on incidents" on page 444.

Symantec Brightmail Gateway integrates with Symantec Network Prevent to deliver, route, hold, or block email traffic. Symantec Network Prevent is a component of Symantec Data Loss Prevention, which discovers, monitors, and protects confidential data wherever it is stored or used. With Symantec Data Loss Prevention, you can create the policies that extend across endpoint, network, and storage systems.

See "About Symantec Network Prevent" on page 462.

Creating content filtering policies

Table 15-1 describes the process to create a content filtering policy.

Table 15-1 How to create content filtering policies

Step	Description
1	Configure the resources that you want to use in your policy so that they are available when you are ready to create the policy. See "About content filtering" on page 331.

Table 15-1 How to create content filtering policies *(continued)*

Step	Description
2	Select the template that you want to use for your policy. You can use a Structured Data policy template, a Described Content policy template, or a blank template. See "About content filtering policy templates" on page 336. See "Selecting the content filtering policy template" on page 338.
3	Define the policy. Specify a name for your policy and whether to track violations on the dashboard or in reports. See "Defining the content filtering policy" on page 346.
4	Create the content filtering policy conditions. Conditions are based on the content of messages. Conditions can specify whether they apply to inbound messages, outbound messages, or both. They can indicate that a violation occurs either when all conditions are met, any condition is met, or a combination of conditions are met. Conditions can apply to specific areas of the message. They can contain keywords and regular expressions. And you can use specific content filtering resources to create conditions (such as dictionaries, patterns, or records). See "Considerations for content filtering policy conditions" on page 347.
5	Configure the policy actions. Specify the actions that you want Symantec Brightmail Gateway to take if the policy is violated. See "Specifying content filtering policy actions" on page 365.
6	Select the groups for which the policy applies. See "Specifying the policy groups for which content filtering policies apply" on page 371.

Table 15-1 How to create content filtering policies *(continued)*

Step	Description
8	Specify the order that you want the policy evaluated. Policies are evaluated in the order that you specify. The same message can violate more than one policy. Depending on the higher-precedence policy action, the actions for the lower-precedence policy may not apply. For example, assume that you have a higher-precedence policy whose action is to delete a message. Also assume that you have a lower-precedence policy whose action is to quarantine the message. If a message triggers both policies, Symantec Brightmail Gateway deletes the message, which is the action for the higher-precedence policy. However, since the message is deleted, it cannot be quarantined as specified in the action for the lower-precedence policy. See "Specifying the order that content filtering policies are evaluated" on page 375.

To create a content filtering policy, you must have Full Administration privileges or privileges to modify policies.

See "About content filtering" on page 331.

See "Editing content filtering policies" on page 372.

See "Copying content filtering policies" on page 373.

See "Deleting content filtering policies" on page 373.

See "Creating a content filtering policy for DKIM validation" on page 141.

About content filtering policy templates

Table 15-2 lists the templates that you can use to create content filtering policies.

Table 15-2 Template types

Template name	Description
Predefined templates	Symantec provides the predefined templates that can help you create content filtering policies for specific scenarios. For example, the Credit Card template provides predefined conditions to detect a credit card number in incoming or outgoing email messages. You can modify conditions in predefined templates to meet your specific needs. And you must configure certain policy settings, such as the policy name and the action to take if the policy is violated. Predefined templates consist of the following types: ■ Described Content Described Content templates let you create the policy conditions that use dictionaries, attachment lists, and patterns. See "Described Content policy templates" on page 338. ■ Structured data Structured Data policy templates let you create the policy conditions that reference proprietary data sets that you upload to Symantec Brightmail Gateway as a record. Messages that are filtered for your proprietary data help you protect your organization from data loss or theft. In addition to records, Structured data templates also let you create the policy conditions that use dictionaries, attachment lists, and patterns. See "About preventing data loss with structured data" on page 375. See "Structured Data policy templates" on page 344. Predefined policy templates may not have both a Described Content and a Structured Data type. For example, the Canadian Social Insurance Number template only uses the Described Content type. However, the Caldicott Report uses both Described Content and Structured Data. If both template types are available, you can only select one.
Blank template	Use a blank template to create a policy that does not relate to any of the predefined templates.

Symantec Brightmail Gateway does not support the ability to create a custom template that appears in the list of templates. However, you can create a policy to use as a template. When you create and save the policy, it appears in your list of content filtering policies along with any other policies that you have created. You can save a policy and give it a name that is easily recognizable as a template. Ensure that you disable the policies that you want to use as templates.

See "Enabling and disabling content filtering policies" on page 374.

Selecting the content filtering policy template

Select the type of template that you want to use to create your content filtering policy.

See "About content filtering policy templates" on page 336.

See "Creating content filtering policies" on page 334.

To select the content filtering policy template

1. In the Control Center, click **Content > Policies > Email**.
2. Click **Add**.
3. Click the radio button beside the policy template that you want to use.
4. Click **Select**.

   After you specify the template that you want to use, define the content filtering policy.

   See "Defining the content filtering policy" on page 346.

Described Content policy templates

Table 15-3 lists the Described Content policy templates and the resources that they use.

See "About content filtering policy templates" on page 336.

Table 15-3 Described Content policy templates

Template name	Policy type	Associated resources
Export Administration Regulations (EAR)	U.S. regulatory policies See "U.S. regulatory policy templates" on page 819.	Dictionaries: EAR Country Codes; EAR CCL Keywords
Gramm- Leach-Bliley		Dictionaries: US SSN Keywords; Credit Card Number Keywords; ABA Routing Number Keywords Patterns: Valid Credit Card; Valid Social Security Number Regex rule: US Social Security Numbers
HIPAA (including PHI)		Dictionaries: US SSN Keywords; Prescription Drug Names; Medical Treatment Keywords; Disease Names; TPO Email Addresses Patterns: Valid Social Security Number Regex rule: Drug Codes
International Traffic in Arms Regulations (ITAR)		Dictionaries: ITAR Country Codes; ITAR Munition Names
NASD Rule 2711 and NYSE Rules 351 and 472		Dictionaries: Analysts' Email Addresses (user-defined); NASD 2711 Keywords
NASD Rule 3010 and NYSE Rule 342		Dictionaries: NASD 3010 Stock Keywords; NASD 3010 Buy/Sell Keywords; NASD 3010 General Keywords
NERC Security Guidelines for Electric Utilities		Dictionaries: Sensitive Keywords; Vulnerability Keywords

Table 15-3 Described Content policy templates *(continued)*

Template name	Policy type	Associated resources
Office of Foreign Assets Control (OFAC)	U.S. regulatory policies	Dictionaries: SDN List; OFAC SDN Country Codes; OFAC Country Codes
Payment Card Industry Data Security Standard		Dictionary: Credit Card Number Keywords Pattern: Valid Credit Card
Sarbanes-Oxley		Dictionaries: SEC Fair Disclosure Keywords; Company Name Keywords (user-defined); Financial Keywords; Confidential/Proprietary Words Attachment List: SEC Fair Disclosure Regulation
SEC Fair Disclosure Regulation		Dictionaries: SEC Fair Disclosure Keywords; Company Name Keywords (user-defined) Attachment List: SEC Fair Disclosure Regulation
State Data Privacy		Dictionaries: US SSN Keywords; ABA Routing Number Keywords; Credit Card Number Keywords; California Keywords; New York Keywords; Letter/12 Num. DLN State Words; Illinois Keywords; New Jersey Keywords; Affiliate Domains Patterns: Valid Social Security Number; Valid Credit Card Regex rules: ABA Routing Numbers; Drivers License Keywords; CA Drivers License Numbers; NY Drivers License Numbers; IL Drivers License Numbers; NJ Drivers License Numbers; Letter + 12 Digits Drivers License Numbers

Table 15-3 Described Content policy templates *(continued)*

Template name	Policy type	Associated resources
Confidential Documents	Confidential data protection See “Confidential data-protection policy templates” on page 831.	Dictionaries: Confidential Keywords; Proprietary Keywords; Internal Use Only Keywords; Not For Distribution Words Attachment Lists: Confidential Documents; Documents Nor for Distribution; Internal Use Only Documents; Proprietary Documents
Defense Message System (DMS) GENSER Classification		Dictionaries: Top Secret; Secret; Classified or Restricted; Other Sensitive Information
Design Documents		Dictionary: Design Document Extensions Attachment List: Design Documents
Encrypted Data		Dictionaries: GPG Encryption Keywords; PGP file extensions; PGP8 Keywords Attachment List: Password Protected Files; PGP Files Regex rule: S/MIME
Financial Information		Dictionary: Financial Keywords; Confidential/Proprietary Words Attachment Lists: Financial Information
Mergers and Acquisitions Data		Dictionary: M & A Project Code Names (user-defined)
Publishing Documents		Dictionary: Publishing Document Extensions Attachment List: Publishing Documents
Project Data		Dictionary: Sensitive Project Code Names (user-defined)
Resumes		Dictionaries: Job Search Keywords, Education; Job Search Keywords, Work; Job Search Keywords, General Attachment Lists: Resumes, All
Source Code		Dictionary: Source Code Extensions Regex rules: C Source Code; VB Source Code; Java Import Elements; Java Class Files; PERL indicator; PERL variable Regex rule: PERL Keywords
		Dictionaries: Top Secret; Secret; Classified or Restricted

Table 15-3 Described Content policy templates *(continued)*

Template name	Policy type	Associated resources
US Intelligence Control Markings (CAPCO) & DCID 1/7		
Competitor Communications	Acceptable use enforcement See "Acceptable use policy templates" on page 838.	Competitor Domains (user-defined)
Gambling		Dictionaries: Gambling Keywords, Confirmed; Gambling Keywords, Suspect
Illegal Drugs		Dictionaries: Street Drug Names; Manufd. Controlled Substances
Media Files		Dictionary: Media Files Extensions Attachment List: Media Files
Offensive Language		Dictionaries: Offensive Language, Explicit; Offensive Language, General
Sexually Explicit Language		Dictionaries: Sex. Explicit Words, Confirmed; Sex. Explicit Words, Suspect; Sex. Explicit Words, Possible
Racist Language		Dictionary: Racist Language
Restricted Files		Attachment List: MSAccess files and Executables
Restricted Recipients		Dictionary: Restricted Recipients (user-defined)
Violence & Weapons		Dictionaries: Violence Keywords; Weapons Keywords

Table 15-3 Described Content policy templates *(continued)*

Template name	Policy type	Associated resources
Canadian Social Insurance Number	Customer and employee data protection See "Customer and employee data-protection templates" on page 842.	Dictionary: Canadian Social Ins. No. Words Pattern: Canadian Social Insurance Numbers
Credit Card Numbers		Dictionary: Credit Card Number Keywords Pattern: Valid Credit Card
Customer Data Protection		Dictionaries: US SSN Keywords; Credit Card Number Keywords; ABA Routing Number Keywords Patterns: Valid Credit Card; Valid Social Security Number Regex rule: US Social Security Numbers
Employee Data Protection		Dictionaries: US SSN Keywords; Credit Card Number Keywords; ABA Routing Number Keywords Patterns: Valid Credit Card; Valid Social Security Number Regex rule: US Social Security Numbers
Individual Taxpayer Identification Numbers (ITIN)		Dictionary: US ITIN Keywords Regex rule: US ITIN
SWIFT Codes		Dictionary: SWIFT Code Keywords Regex rule: SWIFT Code Regex
UK Drivers License Numbers		Dictionary: UK Keywords Regex rule: UK Drivers License Numbers, Drivers License Keywords
UK Electoral Roll Numbers;		Dictionaries: UK Keywords; UK Electoral Roll Number Words Regex rule: UK Electoral Roll Numbers
UK National Insurance Number		Dictionary: UK NIN Keywords Regex rule: UK National Insurance Number
UK Passport Numbers		Dictionary: UK Passport Keywords Regex rules: UK Passport Numbers (Old Type); UK Passport Numbers (New Type)
UK Tax ID Numbers		Dictionary: UK Tax ID Number Keywords Regex rule: UK Tax ID Numbers

Table 15-3 Described Content policy templates *(continued)*

Template name	Policy type	Associated resources
US Social Security Numbers		Dictionary: US SSN Keywords Pattern: Valid Social Security Number
Network Security	Network security enforcement See "Network security policy templates" on page 848.	Dictionaries: Hacker Keywords; Keylogger Keywords Regex rule: GoToMyPC Activity
Network Diagrams		Attachment List: Network Diagrams with IP Address Keyword; Valid IP Address
Password Files		Dictionary: Password Filename Regex rules: /etc/passwd Format; /etc/shadow Format; SAM password
Caldicott Report	UK and international regulatory enforcement See "UK and international regulatory policy templates" on page 849.	Dictionaries: Prescription Drug Names; Disease Names; Medical Treatment Keywords; UK NIN Keywords Regex rule: UK National Insurance Number
Data Protection Act 1998		Dictionaries: UK NIN Keywords; UK Tax ID Number Keywords; UK Keywords; UK Passport Keywords; UK Electoral Roll Number Words Regex rules: UK Electoral Roll Numbers; UK National Insurance Number; UK Tax ID Numbers; UK Drivers License Numbers; Drivers License Keywords; UK Passport Numbers (Old Type); UK Passport Numbers (New Type)
Human Rights Act 1998		Dictionaries: UK Personal Data Keywords; UK Keywords; UK Electoral Roll Number Words Regex rules: UK Electoral Roll Numbers
PIPEDA		Dictionaries: Canadian Social Ins. No. Words; ABA Routing Number Keywords; Credit Card Number Keywords Patterns: Valid Credit Card Regex rule: ABA Routing Numbers; Canadian Social Insurance Numbers

Structured Data policy templates

Table 15-4 lists the Structured Data policy templates and the resources that they use.

See "About content filtering policy templates" on page 336.

Table 15-4 Structured Data policy templates

Template name	Policy type	Associated resource
Export Administration Regulations (EAR)	U.S. regulatory policies See "U.S. regulatory policy templates" on page 819.	Dictionary: EAR Country Codes Record resource with Stock Keeping Unit (SKU) data
Gramm-Leach-Bliley		Record resource with following fields: Account number; Bank card number; Email address; First name; Last name; PIN; Phone number; Social Security Number; ABA Routing Number; Canadian Social Insurance Number; UK National Insurance Number; Date of Birth
HIPAA (including PHI)		Dictionaries: Prescription Drug Names; Medical Treatment Keywords; Disease Names; TPO Email Addresses Regex rule: Drug Codes Record resource with following fields: Last name; Tax payer ID (SSN); Email address; Account number; ID card number; Phone number
International Traffic in Arms Regulations (ITAR)		Dictionary: ITAR Country Codes Record resource with SKU data
Payment Card Industry Data Security Standard		Record resource view
State Data Privacy		Dictionary: Affiliate Domains Record resource with following fields: First name; Last name; Tax payer ID (SSN); Bank card; Account; PIN; State ID; Driver's license; Password; ABA number; Date of birth; SSN
Price Information	Confidential data protection See "Confidential data-protection policy templates" on page 831.	Record resource with fields for SKU numbers and prices
Resumes		Dictionaries: Job Search Keywords, Education; Job Search Keywords, Work; and Job Search Keywords, General Attachment list: Resumes, Employee Record resource with fields for employee first and last names

Table 15-4 Structured Data policy templates *(continued)*

Template name	Policy type	Associated resource
Customer Data Protection	Customer and employee data protection See "Customer and employee data-protection templates" on page 842.	Record resource with the following fields: SSN; Phone; Email; First Name; Last Name; Bank Card number; Account Number; ABA Routing Number; Canadian Social Insurance Number; and UK National Insurance Number; Date of Birth
Employee Data Protection		Record resource with the following fields: SSN, Phone, Email, First Name, Last Name, Bank Card Number, Account Number, ABA Routing Number, Canadian Social Insurance Number, and UK National Insurance Number, employee number, medical insurance number, salary, direct deposit account, and Date of Birth
Caldicott Report	UK and international regulatory enforcement See "UK and international regulatory policy templates" on page 849.	Dictionaries: Prescription Drug Names; Disease Names; Medical Treatment Keywords Record resource with the following fields: NIN (National Insurance Number), account number, last name, ID card number, email, phone, and UK NHS (National Health Service) number
Data Protection Act 1998		Record resource with the following fields: NIN (National Insurance Number), account number, PIN, bank card number, first name, last name, drivers license, password, tax payer ID (SSN), UK NHS number, date of birth, mother's maiden name, email address, and phone number
EU Data Protection Directives		Dictionary: EU Country Codes Record resource with the following fields: last name, bank card number, driver's license, account number, pin, medical account number, and ID card number, username, password, ABA routing number, email, phone, and mother's maiden name
Human Rights Act 1998		Dictionary: UK Personal Data Keywords Record resource with last name and electoral roll number fields
PIPEDA		Record resource with following fields: last name; bank card; medical account number; medical record; agency number; account number; PIN, username; password; SIN; ABA routing number; email; phone; mother's maiden name

Defining the content filtering policy

After you select the content filtering policy template that you want to use, name your policy. By default, Symantec Brightmail Gateway tracks your policy's violations on the dashboard and in reports. But you can disable this feature.

See "Selecting the content filtering policy template" on page 338.

See “Creating content filtering policies” on page 334.

To define the content filtering policy

1 On the **Add Email Content Filtering Policy** page, in the **Policy name** box, type the name for your content filtering policy.

  Select a unique name for the policy that describes what the policy is configured to detect.

2 Uncheck **Track violations of this policy in the dashboard and reports** if you do not want to track this policy's violations in these locations.

  This option is checked by default.

  See “About the Dashboard” on page 604.

  See “About working with reports” on page 568.

  After you define the content filtering policy, create the policy conditions.

  See “Considerations for content filtering policy conditions” on page 347.

  See “Configuring content filtering policy conditions” on page 348.

Considerations for content filtering policy conditions

Keep in mind the following suggestions and requirements as you create content filtering policy conditions:

Test new content filtering policies	Initially, you may want to set your policies so that the messages that violate policy conditions are quarantined instead of deleted. When you are sure that the policies work as designed, you can adjust the actions accordingly.
Sieve scripts	Sieve scripts cannot be imported, including those that were created in previous versions of Symantec or Brightmail software.
Antispam content filtering policies	Spammers usually "spoof" or forge some of the visible headers and the invisible envelope information. Sometimes they forge header information with actual email addresses or domains of innocent people or companies. Use care when you create filters for spam.
File attachment size limits	Compressed files are evaluated based on their size after they are decompressed for the conditions that are based on attachment size limits. For example, assume that a policy is configured to strip message attachments over 10 MB. Symantec Brightmail Gateway strips a 5-MB compressed attachment file from a message because when the attachment is open, it expands to 35 MB. The 35-MB file size exceeds the 10-MB limit.

Compressed files — A compressed file attachment is not scanned if the uncompressed attachment exceeds the file size that is set for container settings. Ensure that any limit you set on attachment file size does not exceed the limit for the container files.

See "Setting limits on nested files" on page 220.

Case sensitivity — All tests for words and phrases are not case sensitive. Lowercase letters in your conditions match both lower- and uppercase letters in messages. Uppercase letters in your conditions match lower- and uppercase letters in messages.

Multiple white spaces — Multiple white spaces in an email header or body are treated as a single-space character.

String matches

If you tested that the subject contains this string	Then any message subject that contains these strings are matched

`inkjet`	`inkjet`
	`Inkjet`
	`INKJET`

`INKJET`	`inkjet`
	`Inkjet`
	`INKJET`

`inkjet cartridge`	`inkjet cartridge`
	`inkjet   cartridge`

`inkjet   cartridge`	`inkjet cartridge`
	`inkjet   cartridge`

`I n k j e t c a r t r I d g e`	`inkjet cartridge`
	`inkjet   cartridge`

Configuring content filtering policy conditions

After you define your content filtering policy, specify the conditions that create a violation.

See "Defining the content filtering policy" on page 346.

You can add, modify, and delete conditions as needed. If you chose a predefined template when you create your policy, the conditions that relate to that template already appear in the **Conditions** list.

You can have as many conditions per content filtering policy as needed, but you can only create one condition at a time. For example, assume that you want to create a policy that evaluates for message size and uses a record. Create a condition for the message size and add it to the **Conditions** list. Then create a separate condition that uses a record view and add it to the **Conditions** list.

After you create the policy conditions, you can create a set of compound conditions. You can also create compound conditions from individual policy conditions. Each compound condition contains the conditions that are linked by AND. You can then link the compound conditions using ALL or ANY. All of the combined conditions must be met to constitute a violation if you select ALL.

See "Considerations for content filtering policy conditions" on page 347.

See "Creating content filtering policies" on page 334.

To configure content filtering policy conditions

1. On the **Add Email Content Filtering Policy** page under **Conditions**, click the **Apply to** drop-down list to specify whether the policy applies to inbound messages, outbound messages, or both.
2. Click the **Which of the following conditions must be met** drop-down list. Specify whether all of the conditions in the **Conditions** list must be met or if any condition that is met creates a violation.
3. Do any of the following tasks:

To add a new condition	Click **Add**.
To modify an existing condition	Check the box beside the condition that you want to modify, and then click **Edit**.
To remove an existing condition	Check the box beside the condition that you want to remove, and then click **Delete**. No additional steps are required to remove an existing condition. Proceed to step 7.

4. On the **Content Filtering Policy Condition** page, do all of the following tasks:
   - Select the condition based on the message part.
     See "Content filtering policy conditions" on page 350.
     See "Message parts used in conditions" on page 352.
   - Select an appropriate resource from the available drop-down list or type a string or number for the condition to match message content against.
     See "Expressions for content filtering policy conditions" on page 362.
     See "Predefined regular expressions" on page 358.
     See "Putting predefined regular expressions into content filtering policy conditions" on page 357.
     See "Perl-compatible regular expressions" on page 435.

- Indicate the match criteria.
  See "Content filtering condition match criteria" on page 356.

5 Click **For all messages** to apply to action that you specify to all messages.

For example, assume that you want to archive all incoming messages. On the **Add Email Content Filtering Policy** page, in the **Apply to** drop-down list, you would select **Inbound messages**. For conditions, you would select **For all messages**. For actions, you would select **Archive the message**.

6 Click **Add Condition**.

7 To create a compound condition, check the conditions in the list that you want to combine, and click **(X & Y)**.

To disassociate compound conditions, click the conditions that you want to unbind and click **(X), (Y)**.

8 Repeat steps 3 through 7 to configure additional conditions for your policy.

9 After you configure all of the policy conditions, specify the actions that Symantec Brightmail Gateway takes if the policy is violated.

See "Specifying content filtering policy actions" on page 365.

Content filtering policy conditions

Table 15-5 describes the content filtering policy conditions that you can use to create your content filtering policy.

See "Considerations for content filtering policy conditions" on page 347.

Table 15-5 Content filtering policy conditions

Condition	Description
Text in the Subject, Body or Attachments	Applies any of the following rules to match words, expressions, patterns, or structured data in the message subject, body, or attachments: ■ Contains *n* or more words from the dictionary that you specify in the drop-down list. ■ matches/does not match the regular expression that you type in the text box. ■ matches/does not match the pattern that you select from the drop-down list. ■ Matches data in the record that you specify in the drop-down list according to the view and the **Minimum number of occurrences** required that you specify.

Table 15-5 Content filtering policy conditions *(continued)*

Condition	Description
Text in this specific part of the message	The policy condition applies one of the following rules to the specified message part and uses it to match the word, expression, or pattern that you specify: ■ Text in the part of the message that you select from the drop-down list. ■ contains/does not contain *n* or more occurrences of the text string that you specify in the text box. ■ starts with/does not start with/matches exactly/does not match exactly /ends with/does not end with the string that you enter in the text box. Regular expressions are designed to match on a line-by-line basis when configured to match against a message headers. But they only match against the entire body text when configured to match against the message body or attachments. ■ matches/does not match the regular expression that you type in the text box. ■ matches/does not match the pattern that select from the drop-down list. If you select **Message header** from the **Text in this specific part of the message** drop-down list, you have the option to type the header name and select exists/does not exist from the message header drop-down list.
Text in the specific part of the message header	Matches the messages that contain or do not contain an email address, domain, or country code from a specific dictionary to either the envelope recipient or the envelope sender. ■ Text in the part of the header that you select from the drop-down list ■ contains the email address that you select from drop-down list ■ from the dictionary that you select from drop-down list
Message size	Applies the policy according to whether the size of the message and its attachments are less than, greater than, or equal to the specified number of bytes, kilobytes, or megabytes. ■ Message size/Attachment size ■ is equal/is greater then/is less than ■ the number entered in the text box ■ bytes/KB/MB

Table 15-5 Content filtering policy conditions *(continued)*

Condition	Description
File metadata	This condition applies the policy to attachments whose file names match one of the specified criteria. ■ Is in the Attachment List that you select from the drop-down list ■ Has the file name that you type in the text box ■ Is MIME type that you type in the text box ■ contains/does not contain a file name from the dictionary that you select from the drop-down list ■ contains/does not contain a file extension from the dictionary that you select from the drop-down list
For all messages	For all of the messages that are not filtered by a higher-precedence policy.

Message parts used in conditions

Table 15-6 describes the parts of a message for which you can create conditions in content filtering policies. It also specifies the message part or resource that Symantec Brightmail Gateway uses to evaluate that condition.

See "Considerations for content filtering policy conditions" on page 347.

Table 15-6 Content filtering policy conditions message parts

Condition	What is evaluated	Example
Any part of the message	Any of the following resources: ■ Dictionary See "About content filtering dictionaries" on page 404. ■ Regular expression See "Predefined regular expressions" on page 358. See "Perl-compatible regular expressions" on page 435. See "Expressions for content filtering policy conditions" on page 362. ■ Pattern See "About patterns" on page 417. ■ Records See "About preventing data loss with structured data" on page 375.	Profanity ABA Routing Numbers Social Security number Customer phone numbers
Attachment content	Text within an attachment file	Finds all attachments that contain the word "discount" more than three times
Attachment type	An attachment list, file name, or MIME type	script.vbs application/octet-stream
Bcc: address	`Bcc:` (blind carbon copy) message header	jane symantecexample.com jane@symantecexample.com
Body	Contents of the message body This component test requires more resources to process. So you may want to add it as the last condition in a filter to optimize the filter.	You already may have won
Cc: address	`Cc:` (carbon copy) message header	jane symantecexample.com jane@symantecexample.com

Table 15-6 Content filtering policy conditions message parts *(continued)*

Condition	What is evaluated	Example
Envelope HELO	SMTP HELO domain in message envelope	symantecexample.com
Envelope recipient	Recipient in message envelope	jane symantecexample.com jane@symantecexample.com
Envelope sender	Sender in message envelope	jane symantecexample.com jane@symantecexample.com
From: address	`From:` message header	jane symantecexample.com jane@symantecexample.com
From/To/Cc/Bcc Address	`From:`, `To:`, `Cc:`, and`Bcc:` message headers	jane symantecexample.com jane@symantecexample.com
Message header	Message header that is specified in the accompanying text field A header is not case sensitive. Do not type the trailing colon in a header.	Reply-To reply-to Message-ID
Subject	`Subject:` message header	$100 FREE. Please Play Now!
To: address	`To:`message header	jane symantecexample.com jane@symantecexample.com
To/Cc/Bcc Address	`To:`, `Cc:`, and `Bcc:` message headers	jane symantecexample.com jane@symantecexample.com

Table 15-6 Content filtering policy conditions message parts *(continued)*

Condition	What is evaluated	Example
Recipient	Email addresses, domain names, and country codes in dictionaries	jane symantecexample.com jane@symantecexample.com
Sender	Email addresses, domain names, and country codes in dictionaries	jane symantecexample.com jane@symantecexample.com
Message size	Size of the message in bytes, kilobytes, or megabytes, including the header and body is less than or greater than the specified value	2000 bytes 200 KB 2 MB
Is in the attachment list	Archive files Document files Executable files Image files Multimedia files	fy1998.arc mystorysofar.msw startmeup.exe lookatme.jpg seenontv.mov
Has the filename	Use DOS wildcard characters if needed to allow for variations	Find all attachment file names that include attach*.exe (attached.exe, attachment.exe, attachments.exe)
Is MIME type	Filters the attachment according to Multipurpose Internet Mail Extension type for non-US ASCII characters sets, non-textual message bodies, and multipart message bodies	MIME-type begins with video/
contains/does not contain a filename from dictionary	Dictionaries See "About content filtering dictionaries" on page 404.	ABA Routing Number Keywords
contains/does not contain a file extension from	Dictionaries See "About content filtering dictionaries" on page 404.	California Keywords

Table 15-6 Content filtering policy conditions message parts *(continued)*

Condition	What is evaluated	Example
For all messages	All of the email that is not filtered by another content filtering policy	(Not applicable)

Content filtering condition match criteria

Table 15-7 describes how Symantec Brightmail Gateway evaluates the match criteria (for example, match/does not match) in content filtering policy conditions.

Table 15-7 Condition match criteria

Match criteria	Description
Contains	Tests the message content against the words that are found in selected dictionary.
Matches regular expression/Does not match regular expression	Tests the message content against user-specified regular expressions. Regular expressions combine alphanumeric characters with metacharacter variables to identify pattern variations.
Matches pattern/Does not match pattern	Tests the message content against known regular-expression patterns such as those found in Social Security numbers or credit card numbers patterns.
Matches data in the Record Resource	Compares the message content with the data that the selected record defines.
Contains/Does not contain	Tests for the supplied text within the component that is specified. Sometimes called a substring test. You can in some cases test for frequency - the number of instances of the supplied text that appear.
Starts with/Does not start with	The regular expression equivalent to ^text.*
Ends with/Does not end with	The regular expression equivalent to *text.$
Matches exactly/does not match exactly	The equivalent of match exactly or does not match exactly for the supplied test.
Is equal to/Is greater than/Is less than	Compares the message size with the specified number of bytes, kilobytes (KB), or megabytes (MB).

Note: All text tests are not case sensitive. Some tests are not available for some components.

Putting predefined regular expressions into content filtering policy conditions

You can copy an expression from a policy template and modify it as needed. You can put the modified expression directly into a content filtering policy condition. Or you can also put the expression in a new pattern or an existing pattern so that you can reuse it.

See "About content filtering policy templates" on page 336.

See "Configuring content filtering policy conditions" on page 348.

Table 15-8 describes the tasks to take to copy an existing regular expression from a template and use it in a content filtering policy condition.

Table 15-8 Process to access, copy, and paste expressions from templates

Step	Task
1	Identify the policy template that contains the regular expression that you want to use. **Note:** Policy template conditions do not identify regular expressions by name. See "Predefined regular expressions" on page 358.
2	Follow the procedure to set up a new content filtering policy from a template. This task lets you access the expressions in the policy template that you want to copy. See "Selecting the content filtering policy template" on page 338. In policy templates where more than one regular expression appears, identify the regular expression that is appropriate for the condition that you want to define. If you are unsure about which regular expression to use, check the conditions for the particular policy template in which the regular expression appears.
3	Copy the regular expression that you want to use from the template.

Table 15-8 Process to access, copy, and paste expressions from templates *(continued)*

Step	Task
4	Paste the regular expression to any of the following locations, and then make any desired modifications: ■ Pattern You can copy the expression to an existing pattern or to a new pattern. When you use an expression in a pattern, the regular expression appears when you select **Matches pattern** or **Does not match pattern** on the **Content Filtering Policy Condition** page. Create custom patterns for regular expressions you intend to use in more than one policy. See "Creating your own custom patterns" on page 419. See "Editing patterns" on page 420. ■ Content filtering policy condition You can copy the expression to a new content filtering policy condition or a policy condition that already exists. Copy the text to the text box beside the regular expression drop-down list on the **Content Filtering Policy Condition** page, and modify it as needed. Do this task if you only need to use the regular expression for a condition once. The regular expression is saved as a condition of the policy that you create. But it does not appear as a pattern resource for use in other policies. See "Configuring content filtering policy conditions" on page 348.

Predefined regular expressions

Table 15-9 describes the regular expressions that you can use in content filtering policy conditions.

See "Expressions for content filtering policy conditions" on page 362.

Table 15-9 Regular expressions and associated templates

Regular expression rule name	Description	Example	Associated template
ABA Routing Numbers	Any digit other than a 4, 5 or 9 followed by any 8 digits.	0123465789 222222222 841471235	Customer Data Protection; Employee Data Protection; PIPEDA; State Data Privacy

Table 15-9 Regular expressions and associated templates *(continued)*

Regular expression rule name	Description	Example	Associated template
CA Drivers License Numbers	Any sequence of a letter followed by 7 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	a1234567 A1234567	State Data Privacy
Canadian Social Insurance Numbers	Matches numbers in the form DDD-DDD-DDD or DDD DDD DDD	123-456-789 123 456 789	PIPEDA; Canadian Social Insurance Number

C Source Code	The sequence #include, followed by any number of spaces, followed by a '<' or ", followed by any number of letters, digits or underscores, followed by a < or ".	#include lt;stdio.h>
		#include "fdfdfd"
			Source Code

Drivers License Keywords	Looks for combinations of driv followed by characters followed by lic followed by characters OR dl # or dl# OR lic # or lic# all case insensitive.	Driver License Driver License dl # DL#	State Data Privacy
Drug Codes	Matches NDC Drug Code format http://www.fda.gov/cder/ndc/	12345-1234-12	HIPAA (including PHI)
/etc/passwd Format	Looks for sequences that match an /etc/passwd format.	bob:fdshfjhf:78978:45:fdsf:	Password Files
/etc/shadow Format	Looks for sequences that match an /etc/shadow format.	bob:fdsfd:343:45454: 4343: 122:343:545	Password Files
GoToMyPC Activity	Looks for a case insensitive match of the sequence 'jedi?request=' or '/erc/Poll?machinekey'	jedi?request= /erc/Poll?machinekey	Network Security
IL Drivers License Numbers	Any sequence of a letter followed by 11 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	a12346578901 A12345678901	State Data Privacy

Table 15-9 Regular expressions and associated templates *(continued)*

Regular expression rule name	Description	Example	Associated template
IP Address	Looks for a case insensitive match of the sequence 'IP Address.'	IP Address ip address iP aDdress	Network Diagrams
Java Class Files	Optional public, private or protected followed by class or interface followed a class name by an optional extends or implements followed by a class name.	public class Foo protected class Bar implements Foo	Source Code

Java Import Statements	Looks for import followed by a class name followed by a semicolon.	import
		java.util.Collection;
			Source Code

Java Source Code	Looks for the file name extensions "java"	stop_spam.java	Source Code
NY Drivers License Numbers	Any sequence of 9 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	123456789 987654321	State Data Privacy
Letter + 12 Digits Drivers License Numbers	Any sequence of a letter followed by 12 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	a123456789012 A123456789012	State Data Privacy
NJ Drivers License Numbers	Any sequence of a letter followed by 13 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	a1234567890123 A1234567890123	State Data Privacy
SAM Passwords	Looks for sequences that match a SAM passwords format.	bob:3434:12345678901234567890123456798012:12345678901234567890123456798012:	Password Files
PERL Keywords	A case sensitive match of the sequence 'perl'	perl	Source Code
PERL indicator	Looks for the sequence #!	#!	Source Code

Table 15-9 Regular expressions and associated templates *(continued)*

Regular expression rule name	Description	Example	Associated template
PERL variable	Looks for sequences that start with a $ followed by any letter or underscore then followed by any number of letters, digits or underscore.	$something; $another;	Source Code
S/MIME	Looks for MIME-type that begins with application/pkcs7-mime or application/x-pkcs7-mime or application/pkcs7-signature or application/x-pkcs7-signature	example.p7c example.p7m example.p7s example.p7a	Encrypted Data
SWIFT Code Regex	4 characters followed by a dash followed by 2 characters followed dash followed by 2 characters followed by an optional dash and 2 characters.	ABCD-EF-GH ABCD-EF-GH-IJK	Swift Codes
UK Drivers License Numbers	5 characters or digits followed by a digit followed by one of the inner groupings followed by a digit followed by 3 characters or digits followed by 2 characters.	A1234501023ABCDE	UK Drivers License Numbers
UK Electoral Roll Numbers	Two or three upper case letters, followed by one to four digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	AB1 AB12 AB123 AB1234 ABC1 ABC1234	UK Electoral Roll Numbers; Human Rights Act 1998
UK National Insurance Number	Two upper case letters from the group A-CEGHJ-NOPR-TW-Z followed by 6 digits, followed by an optional upper case letter (A through D), immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	AB123456 AB132456C	UK National Insurance Number; Caldicott Report

Table 15-9 Regular expressions and associated templates *(continued)*

Regular expression rule name	Description	Example	Associated template
UK Passport Numbers (Old Type)	An upper case letter or digit, followed by 4 digits, followed by an upper case letter or digit, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	A1234B 12345C 132456 A12345	UK Passport Numbers; Data Protection Act 1998
UK Passport Numbers (New Type)	Any sequence of 9 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	A1234B 12345C 132456 A12345	UK Passport Numbers; Data Protection Act 1998
UK Tax ID Numbers	Any sequence of 10 digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	AB123456 AB132456C	UK Tax ID Numbers
US ITIN	Looks for a 9 followed by 2 digits, followed by a dash, followed by a 7 or 8 followed by a digit, followed by a dash, followed by any four digits, immediately preceded and succeeded by a character that is not a letter, a digit, or underscore (_).	912-71-1234 902-81-0234	Individual Taxpayer Identification Numbers (ITIN)
VB Source Code	A case sensitive match of the sequence Attribute followed by a space character followed by VB_Name.	Attribute VB_Name	Source Code
Filename	Looks for a true file type that mataches the type identified in the associated template.	True file type is PGP Secret Keyring	Password Files; Encrypted Data

Expressions for content filtering policy conditions

Table 15-10 lists the types of expressions that you can use in content filtering policy conditions.

Table 15-10 Expression types

Expression type	Description	How messages are matched
Regular expressions	Symantec Brightmail Gateway provides the templates that you can use when you create a new content filtering policy. You can also use these templates to copy their predefined regular expressions into any of the following locations: ■ New pattern or existing pattern ■ New or existing content filtering policy condition These templates contain the regular expressions that are applicable to the type of rule that you can create with them. For example the Canadian Social Insurance Number template contains the regular expressions that can identify Canadian Social Insurance numbers. See "Predefined regular expressions" on page 358. See "Putting predefined regular expressions into content filtering policy conditions" on page 357.	When matched against a message subject line, matches on a line-by-line basis. When matched against the message body, matches against the entire body text, not on a line-by-line basis.

Table 15-10 Expression types *(continued)*

Expression type	Description	How messages are matched
Perl-compatible regular expressions	You can use the regular expressions that behave like Perl regular expressions when you create conditions for a content filtering rule. To use Perl expressions, you must use either **matches regular expression** or **does not match regular expression** as a policy condition. Symantec Brightmail Gateway supports PCRE version 6.6. See "Perl-compatible regular expressions" on page 435. See "Perl-compatible regular expression examples" on page 436. **Note:** To use a pattern to match certain special characters (including forward slashes), you must precede each character with backslash (\). For more information about Perl-compatible regular expressions, see: http://www.perl.com/doc/manual/html/pod/perlre.html.	When matched against a message header, matches on a line-by-line basis. When matched against the message body or attachments, matches against the entire body text, not on a line-by-line basis.

Symantec Brightmail Gateway uses different analyses when it scans messages for expression matches. If you specify a condition that uses a regular expression, Symantec Brightmail Gateway performs a regular expression analysis. If you specify a condition that uses a keyword or dictionary, Symantec Brightmail Gateway performs a text search.

See "Creating content filtering policies" on page 334.

See "Configuring content filtering policy conditions" on page 348.

See “Expressions for content filtering policy conditions” on page 362.

Specifying content filtering policy actions

After you define the conditions for a content filtering policy, define the actions that you want taken when a message meets those conditions.

See “Configuring content filtering policy conditions” on page 348.

You can select multiple actions to occur when the policy is violated. Symantec Brightmail Gateway applies the actions in the order that they are listed in the **Actions** list.

See “Creating content filtering policies” on page 334.

To specify content filtering policy actions

1. On the **Add Email Content Filtering Policy** page under **Actions**, click the drop-down list and select the action that you want taken if the conditions of the policy are violated.

   For some actions, you need to specify additional information.

   See “Content filtering policy actions” on page 365.

2. Click **Add Action**.
3. Repeat steps 1 and 2 to add as many actions as needed.
4. After you specify the actions that you want taken, specify the groups for which this policy applies.

   See “Specifying the policy groups for which content filtering policies apply” on page 371.

Content filtering policy actions

Table 15-11 describes the actions that you can take for content filtering policies.

See “Specifying content filtering policy actions” on page 365.

Table 15-11 Content filtering policy actions

Action	Description and additional options
Add a header	Adds a header to the message. You must also specify the following options: ■ **Header** Type the header that you want to add to the message. For example: X-Cfilter. ■ **Value** Specify the value of the header. For example: Data violation incident. ■ **Encoding** Click the drop-down list and select the encoding to use for the header.
Add annotation	Adds an annotation to the message body. You must also specify the following options: ■ **Annotation** Click the drop-down list and select the annotation that you want to use. ■ Specify whether you want to add the annotation before the message body (prepend) or after the message body (append). See "Annotating messages that contain violations" on page 421.
Add BCC recipients	Lets you send a blind copy of the message to the email recipients that you specify. You must also specify the following option: ■ **BCC recipients** Specify the email addresses of the recipients that you want to receive a blind copy of the message. Separate multiple entries with commas.

Table 15-11 Content filtering policy actions *(continued)*

Action	Description and additional options

Archive the message	Lets you specify where and how you want to archive the message.
	You must also specify the following options:
	■ **Archive email address**
	If you have already determined an archive email address on the **Archive** page, that address is the default. If you have not already configured an address, this field is blank and you must provide an address.
	■ **Optional archive tag**
	Specify an archive tag to add an `X-archive:` header to archived messages, followed by your text. The `X-archive:` header may be useful to sort archived messages when you view them with an email client. However, Symantec Brightmail Gateway itself does not use the `X-archive:` header.
	If multiple policies result in archiving the same message, each unique `X-archive:` header is added to the message. For example, the following archive tag:
	`Docket 53745`
	adds the following header to the message when it is archived:
	`X-archive: Docket 53745`
	Type any character except carriage return, line feed, or semicolon.
	■ **Encoding**
	Specify the encoding for the archive tag.
	■ **Archive server host**
	Type the host name or IP address for the archive email address.
	■ **Archive server port**
	If you specified an archive server host, also type the archive server port.
	■ **Enable MX Lookup**
	Check this box if you want to route archive messages with MX Lookup to locate the information that corresponds to the archive server host.
	See "Specifying where to save archived messages" on page 437.

Bypass spam scanning	The message is not scanned for spam. See "About filtering spam" on page 239.

Table 15-11 Content filtering policy actions *(continued)*

Action	Description and additional options
Create an informational incident	The violation incident is created in the Informational Incidents folder that you specify. You must also specify the following option: ■ **In content informational incident folder** Click the drop-down list and select the Informational Incidents folder in which you want the incident created. If you do not specify an Informational Incidents folder, the incident is created in the Informational Incidents default folder. See "About content incident folders" on page 438.
Create a quarantine incident	The violation incident is created in the Quarantine Incidents folder that you specify. You must also specify the following option: ■ **In content quarantine incident folder** Click the drop-down list and select the Quarantine Incidents folder in which you want the incident created. If you do not specify a Quarantine Incidents folder, the incident is created in the Quarantine Incidents default folder. After you add the action to create an incident in an incident folder, do the following tasks: ■ Click the **Actions** drop-down list again. ■ Select the action that you want taken if the content filtering officer approves the incident, and click **Add to Approved Actions**. ■ Select the action that you want taken if the content filtering officer rejects the incident, and click **Add to Rejected Actions**. These actions are initiated once a content filtering officer approves or rejects a message from the folder's **Incident Management** page. You can only apply this action once per policy. See "About content incident folders" on page 438.
Delete message	The message is permanently deleted and cannot be retrieved.
Deliver message normally	Delivers the message to the intended recipients without any modifications to the message.
Deliver message with content encryption	Encrypts all outbound messages. This action is only available for the content filtering policies that apply to outbound mail. See "About encrypting messages with Symantec Content Encryption" on page 458.

Table 15-11 Content filtering policy actions *(continued)*

Action	Description and additional options
Deliver message with TLS encryption	Sends the message over an encrypted channel. Specify one of the following options: ■ **Attempt TLS encryption** ■ **Require TLS encryption and don't verify certificate** ■ **Require TLS encryption and verify certificate**
Deliver the message to the recipient's spam folder	Delivers the message to the end user's spam folder. This action requires use of the Symantec Spam Folder Agent for Exchange or the Symantec Spam Folder Agent for Domino. **Note:** The Symantec Spam Folder Agent is no longer supported. See "About filtering spam" on page 239.
Forward a copy of the message	Forwards the message to the people that you specify. You must also specify the following option: ■ **Forward to** Type the email addresses of the people to whom you want to forward a copy of the message. Separate multiple entries with commas.
Hold message in Spam Quarantine	Places the message in the end user's Spam Quarantine. See "About quarantining spam" on page 258.
Modify the subject line	Modifies the message's `Subject:` line. You must also specify the following options: ■ **Modification** Type the modification to the subject line. ■ **Encoding** Click the drop-down list and select the encoding for message. ■ Specify whether to put the modified content before the subject line (prepend) or after the subject line (append).

Table 15-11 Content filtering policy actions *(continued)*

Action	Description and additional options
Route the message	Delivers the message to the location that you specify. You must also specify the following options: ■ **Host** Specify the host where you want to route the message. ■ **Port** Specify the host's port number. ■ **Use MX Lookup** Check this box if you want to route the message with MX Lookup to locate the information that corresponds to the SMTP host.
Send a bounce message	Returns the message to its `From:` address with a custom response, and delivers the message to the recipient. You must also specify the following options: ■ **Explanation** Optionally, type the explanation message that you want sent to the sender. ■ **Encoding** Specify the encoding that you want to use for the bounce message. ■ **Include original message** Check this box to include the original message with the bounce message.
Send notification	Delivers the original message to the intended recipient and sends a predefined notification. You must also specify the following option: ■ **Notification** Click the drop-down list and select the notification that you want to send. See "About policy violation notifications" on page 400.
Strip attachments	Lets you specify which types of attachments that you want to delete. Specify one of the following options: ■ **Strip all attachments** Deletes all attachments, regardless of whether the attachments violated the policy. ■ **Strip attachment lists** Click the drop-down list and select the type of attachments that you want to delete, regardless of whether any of the attachments violate the policy. ■ **Strip offending attachments** Only deletes the attachments that violated the policy.

Table 15-11 Content filtering policy actions *(continued)*

Action	Description and additional options
Treat as bad sender	Processes the message based on the action that you specify in the Local Bad Sender Domains group. This action applies even if the Local Bad Sender Domains group is disabled. See "About blocking and allowing messages using sender groups" on page 174.
Treat as good sender	Processes the message based on the action that you specify in the Local Good Sender Domains group. This action applies even if the Local Good Sender Domains group is disabled. These messages are not scanned for spam. See "About blocking and allowing messages using sender groups" on page 174.
Treat as a mass-mailing worm	Processes the message based on the action in the associated mass-mailing worm policy. See "About detecting viruses and malicious attacks" on page 204. See "Default email virus policies" on page 52.
Treat as a virus	Considers the message a virus and takes the actions that you specify in your virus policies. See "About detecting viruses and malicious attacks" on page 204. See "Default email virus policies" on page 52.
Treat as spam	Considers the message spam and takes the action that you specify in your spam policies. See "About filtering spam" on page 239. See "Default email spam policies" on page 50.
Treat as suspected spam	Considers the message suspected spam and takes the action that you specify in your suspected spam policies. See "About filtering spam" on page 239. See "Default email spam policies" on page 50.

Specifying the policy groups for which content filtering policies apply

After you define the actions for your content filtering policy, specify the policy groups for which the policy applies.

See "Specifying content filtering policy actions" on page 365.

You must create separate policies for the messages that meet similar conditions but require separate actions for different policy groups. Similarly, inbound messages may be treated differently from the outbound messages that otherwise meet the same conditions for a policy group.

For example, assume that separate actions are required to route the email that contains sensitive human resource data. Create one policy in which the messages that are sent to executives do not include attachments. Create a separate policy that contains the same conditions, but the messages that are sent to managers are held for review.

You can apply content filtering policies to policy groups in either of the following ways:

- Create the policy group first and then apply the policy to that group when you create your content filtering policy.
  The process to create a content filtering policy assumes that the policy group is created first, and then the policy is applied to the group.
  See "Creating a policy group" on page 316.
  See "To specify the policy groups for which content policies apply" on page 372.
- Create the content filtering policy and then create a policy group and apply the content filtering policy to the group.
  Use this method in situations in which new employees or policy groups are added after policies are created.
  See "Creating content filtering policies" on page 334.
  See "Selecting content filtering policies for a policy group" on page 326.

A content filtering policy is automatically enabled when you create it. If you do not want to use the content filtering policy yet, disable it.

See "Enabling and disabling content filtering policies" on page 374.

To specify the policy groups for which content policies apply

1. Under **Policy Groups**, check one or more groups to which this policy should apply.
2. Click **Save**.
3. After you create your content filtering policy, specify the order in which you want the policy evaluated.

   See "Specifying the order that content filtering policies are evaluated" on page 375.

Editing content filtering policies

After you create a content filtering policy, you may want to modify it to fine-tune it or update it for changing or new conditions.

To edit a content filtering policy, you must have Full Administration privileges or privileges to modify policies.

To edit content filtering policies

1. In the Control Center, click **Content** > **Policies** > **Email**.
2. Check the box beside the policy that you want to edit.
3. Click **Edit**.
4. Make the desired changes.

   See "Creating content filtering policies" on page 334.
5. Click **Save**.

Deleting content filtering policies

You can delete content filtering policies that you no longer need. When you delete unnecessary policies, the list of remaining policies is easier to manage. If you are uncertain as to whether you might want to use the policy again, disable it instead.

See "Enabling and disabling content filtering policies" on page 374.

To delete a content filtering policy, you must have Full Administration privileges or privileges to modify policies.

To delete content filtering policies

1. In the Control Center, click **Content > Policies > Email**.
2. Check the box beside the policy that you want to delete.

   Check **Email Content Filtering Policies** to check the boxes beside all of the policies.
3. Click **Delete**.

Copying content filtering policies

Instead of creating a new content filtering policy from scratch, you can copy an existing policy that is similar to the new one that you want to create. After you copy the policy, you can save it under a new name and modify it as needed.

You may also have a policy that you created as a template. Copy the policy template and create your new content filtering policy.

See "Creating content filtering policies" on page 334.

To copy a content filtering policy, you must have Full Administration rights or rights to modify policies.

To copy content filtering policies

1. In the Control Center, click **Content > Policies > Email**.
2. Check the box beside the policy that you want to copy.
3. Click **Copy**.
4. Modify the policy.

   See "Editing content filtering policies" on page 372.
5. Click **Save**.

Enabling and disabling content filtering policies

You must enable a content filtering policy for Symantec Brightmail Gateway to evaluate the policy when it scans messages. A content filtering policy is enabled by default when you create it.

See "Creating content filtering policies" on page 334.

You may want to disable a content filtering policy for any of the following reasons:

- You are uncertain as to whether you want to permanently delete the policy.
- You are testing and fine-tuning a new content filtering policy.
- You created the policy for a condition that has ended (such as an email virus attack). However, you want to retain the policy in case you need it again.
- The policy is a template to aid you when you create new policies of similar types.

The **Email Content Filtering Policies** page shows the status of each content filtering policy using the following symbols:

Green check mark: The policy is enabled.

Black dash: The policy is disabled.

To enable or disable a content filtering policy, you must have Full Administration rights or rights to modify policies.

To enable or disable content filtering policies

1. In the Control Center, click **Content > Policies > Email**.
2. Check the box beside the policy that you want to enable or disable.
3. Click **Enable** or **Disable**.

Specifying the order that content filtering policies are evaluated

You can change the order in which Symantec Brightmail Gateway applies content filtering policies to messages as they are scanned. Changing the order in which the content filtering policies are evaluated lets you prioritize the actions to be taken with messages.

For example, assume that a message triggers both of the following policies:

- Sexual Language
  This policy detects the messages that contain predefined sexual language. If this policy is violated, the action is to create an incident in an incident folder.
- Security Violations
  This policy detects the messages that contain security violations. If this policy is violated, the action is to forward a copy of the message to an administrator.

Assume that the policy for Sexual Language has a higher priority than the policy for Security Violations. When Symantec Brightmail Gateway scans the message, it detects the Sexual Language violation first. Symantec Brightmail Gateway applies the action for the Sexual Language policy (it creates an incident in an incident folder). In this example, the message can be approved and released for delivery, in which case, the administrator is not notified of the security violation.

Content filtering policies are applied in the order that they appear on the **Email Content Filtering Policies** page from top (the first policy to be applied) to bottom (the last policy to be applied).

To modify the order of content filtering policies, you must have Full Administration rights or rights to modify policies.

To specify the order that content filtering policies are evaluated

1. In the Control Center, click **Content > Policies > Email**.
2. Click on the policy that you want to move, and drag it up or down to the location that you want.

About preventing data loss with structured data

You can use structured data to protect customer, employee, patient, or other data that cannot be identified through regular expressions or keywords. For example, you can filter outbound email for sensitive, proprietary customer or employee data.

Table 15-12 describes the process to use structured data in a content filtering policy.

Table 15-12 How to use structured data in a policy

Step	Description
1	Decide which Structured Data policy template you intend to use to create your content filtering policy. The pattern fields in your data source file must correspond to the data types that your Structured Data policy template uses. For example, assume that you want to create a policy with the Customer Data Protection template. You should use the pattern fields that correspond to Social Security number, credit card number, phone, and email. See "About content filtering policy templates" on page 336. See "Structured Data policy templates" on page 344.
2	Obtain a data source file from your database administrator. Structured data is comprised of a data source file that your company provides. This file consists of columns of the company-specific, delimited data that you want to protect. See "About your data source files" on page 378.
3	Define a record. When you define a record, you specify a name for the record, a description (optional), data source attributes, and the processing error threshold. See "Defining records" on page 381.
4	Map the columns in your data source file to corresponding field names in Symantec Brightmail Gateway. See "About mapping your data source file columns" on page 383. See "Mapping data source file columns to fields in Symantec Brightmail Gateway" on page 388. See "System patterns" on page 384.
5	Upload your data source file. Symantec Brightmail Gateway validates the file when you upload it and lets you know if the upload is successful. If the upload is successful, Symantec Brightmail Gateway indexes the file. If the upload is unsuccessful, an error message appears that indicates the reason for failure. See "Uploading data source files" on page 389.

Table 15-12 How to use structured data in a policy *(continued)*

Step	Description
6	Replicate the record. You must replicate the record to all attached and enabled Scanners before you can use it in content filtering policies. Replicating data source files to your Scanners can be time consuming. You might want to wait to replicate the record during off-peak times. See "Replicating records" on page 390.
7	Define one or more views for a record. A view lets you specify which fields in a record you want Symantec Brightmail Gateway to search for when it evaluates a message. You also specify how many of those fields must be in a message to constitute an occurrence. You can have one or more views for a record. For example, one view might contain the Credit Card number pattern, the Email address pattern, and the IP address pattern. Another view for that same record might consist of the Credit Card number pattern, the US Phone pattern, and the US ZIP Code pattern. Multiple views let you create multiple content filtering policies from a single record or multiple conditions for the same policy. You must define at least one view before you can use the record in content filtering policies. See "Creating views" on page 397. See "Editing views" on page 398.
8	Create a content filtering policy with a Structured Data policy template. You must use a Structured Data policy template or a blank template to use a record view as a policy condition. See "About content filtering policy templates" on page 336. See "Structured Data policy templates" on page 344. See "Creating content filtering policies" on page 334.
9	Configure the policy to use a record view as a condition. In the content filtering policy, you specify how many occurrences must take place to violate the policy. See "Configuring content filtering policy conditions" on page 348.

About your data source files

To use structured data in content filtering policies, you must obtain a data source file from your database administrator.

Table 15-13 describes the requirements and considerations that you should know about your data source file.

Table 15-13 Data source file requirements and considerations

Requirement or consideration	Description
Only one word per delimited field consisting of at least two alphanumeric characters	Each entry must contain at least two alphanumeric characters. Single character entries in a field are unsupported. Also, a data source file can only have a single word per delimited field. For example, if you want to match the data "1st Street", place "1st" in one delimited field. Place "Street" in a separate (but following) field. If you place two or more words in a field, a match is less likely. For example, if you place the string "1st Street" in a single delimited field, you have placed multiple "words" in the same cell. A match is then unlikely since the only match that is expected would be when the data that is examined is in a tabular format. In a tabular format, the two strings (1st) and (Street) are evaluated as one string (1st Street). Similar behaviors exist when you try to match any languages that recognize white spaces differently, such as Korean or Chinese.

Table 15-13 Data source file requirements and considerations *(continued)*

Requirement or consideration	Description
Only specific separators are recognized for credit card and number patterns	Symantec Brightmail Gateway recognizes only certain separator characters when it attempts to match record entries in credit card and number pattern fields. The recognized separator characters (other than space) for credit card and number pattern fields are as follows: ■ Tab ■ Comma (,) ■ Pound sign (#) ■ Hyphen (-) ■ Plus sign (+) ■ Pipe (|) ■ Semicolon (;) ■ Colon (:) Symantec Brightmail Gateway interprets any numbers that contain an unrecognized separator as a word. For example, 4123*6666*7777*8888 would not return a valid match against a credit card number field. Symantec Brightmail Gateway interprets this content as the word: 4123*6666*7777*8888.
Use a delimiter other than commas to separate adjacent number patterns	Use a tab or pipe (|) instead of a comma as your field delimiter to separate two adjacent fields that are number patterns. Otherwise, the Record validator may interpret the two numbers from adjacent fields as a single number. For example, assume that you have adjacent fields: Age and Weight. And assume that you have separated the fields for Age and Weight by a comma; for example: 25,150. Symantec Brightmail Gateway might interpret 25,150 as belonging to the Age field instead of 25 belonging to the Age field and 150 belonging to the Weight field.

Table 15-13 Data source file requirements and considerations *(continued)*

Requirement or consideration	Description
Data source file must have the minimum number of columns required by the Structured Data template	Ensure that your data source file contains the columns that you want to use to define a view. For example, assume that you use a Structured Data policy that calls for a minimum of three fields to trigger a violation. Those three fields must be mapped in the record so that Symantec Brightmail Gateway can reference them. For example, assume that you use the EU Data Protection Directives policy template. Any view that accesses the EU Data Protection Directives policy should be configured to match entries in at least four of five fields: Last name, email, phone, account number, and user name. See "Structured Data policy templates" on page 344. See "Creating views" on page 397.
Pattern fields should match the data types that are used in the policy template	The pattern fields must correspond to the data types that your Structured Data policy template uses. For example, assume that you want to create a policy with the Customer Data Protection template. You should use the pattern fields that correspond to Social Security number, credit card number, phone, and email columns. See "Structured Data policy templates" on page 344.
Mappings must match header row columns	If the mappings in your record do not match the columns in the header row, Symantec Brightmail Gateway counts the actual header row as invalid. The header row is considered invalid because it returns values other than those expected.
Credit card numbers must pass the Luhn checksum test	All credit card numbers must pass the Luhn checksum test, where total modulus 10 is congruent to 0, to produce a match. The Luhn test is used to distinguish valid numbers from random collections of digits.

Table 15-13 Data source file requirements and considerations *(continued)*

Requirement or consideration	Description
CRLF line breaks that precede rows in a data source file are included in row counts	If the data source file does not contain CRLFs, Symantec Brightmail Gateway skips the header row in the row count. If a data source file contains CRLFs, Symantec Brightmail Gateway treats the first CRLF as the header row. So it returns the values from subsequent rows, including those for the actual header row. For example, assume that one column is mapped to recognize the US ZIP code pattern and one or more CRLFs begin the data set. Symantec Brightmail Gateway counts the actual header row as a normal row. It expects to return a 5-digit number in that column. When the actual header row returns a word value instead of a 5-digit number, Symantec Brightmail Gateway counts it as an invalid row. Symantec Brightmail Gateway ignores the CRLFs that occur within a data set or at the end of a data set. Such CRLFs are not counted as rows.
Rows that occur more than 99 times are not matched	Because of implementation limitations, Symantec Brightmail Gateway cannot match any rows that occur more than 99 times.

See "Structured Data policy templates" on page 344.

See "About mapping your data source file columns" on page 383.

See "About preventing data loss with structured data" on page 375.

Defining records

After you obtain a data source file, you can create a record.

See "About your data source files" on page 378.

When you create a record, you define the following criteria:

Name and description

Specify a name and description for your record. Select a name that identifies the data that the record contains. A description for your record is optional.

If you change the name of a record, all conditions that reference views of that record reflect the name change.

See "Modifying records" on page 391.

Data source attributes	Specify the delimiter that your data source file uses. The supported delimiter characters are tab, comma (","), or pipe ("|"). Also indicate whether your data source file contains a header row. The header row is neither processed nor included in the record. You do not need to check the **Data source file contains a header row** box if you define all fields as system patterns. See "Mapping data source file columns to fields in Symantec Brightmail Gateway" on page 388.
Error threshold	Set **Maximum Allowable Errors** to a percentage of the total rows that can safely return errors when you upload the data source file and still continue processing. Setting a percentage that is too low may make it difficult to complete processing an otherwise useful data source file. Setting a percentage that is too high may hide the fact that the record file is partially corrupted.

To create a record, you must have Full Administration rights or rights to modify settings.

To define records

1. In the Control Center, click **Content > Resources > Records**.
2. Click **Add**.
3. In the **Record Resource Name** field, type a name for the record.
4. In the **Optional description** field, type a description for the record.
5. Under **Data Source Attributes**, click the **Delimiter** drop-down list and select the appropriate delimiter character for your data source file.
6. Check the **Data source file contains a header row** checkbox if your data source file contains a header row.
7. Under **Error Threshold**, type the maximum allowable percentage of errors that can occur before processing is halted.

   After you define your record, map your data source file columns to fields in Symantec Brightmail Gateway.

   See "Mapping data source file columns to fields in Symantec Brightmail Gateway" on page 388.

About mapping your data source file columns

Patterns are the named regular expressions or system patterns that describe a commonly known data object, such as the pattern for credit card numbers. You must map the pattern of each column in your data source file to corresponding fields in Symantec Brightmail Gateway.

Table 15-14 describes the field types that you can use to map your data source file.

Table 15-14 Field types

Field type	Description
System patterns	Symantec Brightmail Gateway matches data in each column of your data source file with a system pattern. A set of regular expressions defines a system pattern. Data in your data source file must conform to one of the regular expressions for a data file to be successfully uploaded and indexed. The system patterns are as follows: ■ Credit card number pattern ■ Email pattern ■ IP address pattern ■ Number pattern ■ Percent pattern ■ US phone pattern ■ US ZIP code pattern ■ SSN/ITIN pattern See "System patterns" on page 384.
Customized	If data in your data source file does not fall into one of the system pattern categories, you can create a customized field. Symantec Brightmail only recognizes a customized field as a WORD pattern. Customized fields are not indexed nor are they validated. Symantec recommends that you make the customized field names similar to the field names in your data source file. For example, if the first field in your data source file is first_name, then you could name the corresponding field in Symantec Brightmail Gateway first_name. Custom field names cannot be the same as any of the predefined list of field names. Nor can they be the same as any other custom field name in the record.

Table 15-15 provides an example of how a data source file is mapped.

Table 15-15 Data source file mapping example

Data source file column	Data that is contained in data source file column	Symantec Brightmail Gateway column	Field Name
1	Credit card numbers	1	**Credit card number pattern**
2	Email addresses	2	**Email pattern**
3	First names	3	**Customize** Customized field name: first_name
4	Last names	4	**Customize** Customized field name: last_name

See "About your data source files" on page 378.

See "Mapping data source file columns to fields in Symantec Brightmail Gateway" on page 388.

See "About preventing data loss with structured data" on page 375.

System patterns

Table 15-16 describes the system patterns that you can use to map your data source file columns to Symantec Brightmail Gateway.

Table 15-16 System patterns

System pattern	Examples	Description
Credit Card number	5369 7777 8888 9999	MasterCard: Any 16-digit number that begins with 5 and whose second digit is a number from 1 to 5, separated into four groups of four by spaces or hyphens.
	5369-7777-8888-9999	
	4567 1234 5678 9123	VISA: Any 16-digit number that begins with 4 and separated into four groups of four digits that are separated by a space or hyphen, or any 12-digit number that begins with 4.
	4123-6666-7777-8888	
	4123456789012	
	3442 456789 12345 3758 456789 12345	American Express: Any 15-digit number that begins with 34 or 37 and separated into three groups of 4, 6, and 5 digits, respectively, by spaces or hyphens.
	3056 123456 7890 3667 123456 7890 3878 123456 7890 3056-123456-7890 3667-123456-7890 3842-123456-7890	Diners Club card: Any 15-digit number that begins with 30, 36, or 38 and separated into three groups of 4, 6, and 5 digits, respectively, by spaces or hyphens.
	6011 1234 5678 9012	Discover card: Any 16-digit number that begins with 6011 and separated into groups of 4 by spaces or hyphens.
	6011-1234-5678-9012	
	2014 123456 78901	Enroute card: Any 15-digit number that begins with 2014 or 2149 separated into groups of 4, 6, and 5 by spaces or hyphens.
	2149-123456-78901	
	3123 4567 8901 1234	JCB: Any 16-digit number that is separated into four groups of four by a space or hyphen and begins with 3; or any 15-digit number that begins with 2131 or 1800 and is followed by 11 digits.
	3123-4567-8901-1234	
	213112345678901 1800123245678901	

Table 15-16 System patterns *(continued)*

System pattern	Examples	Description
Email	jabberwocky @symantec.com mister_smith @senate.gov tom.swift @gadgets.arpa t-rex9@nature. museum harry@hogwarts.edu.uk	Any alphanumeric string, that is divided by an underscore (_), hyphen (-), or period, followed by the @ sign and an alphanumeric string, a period, and one of the domain-name extensions listed. Symantec Brightmail Gateway cannot validate top-level domains of two letters, where one or both letters are uppercase. It does, however, validate uppercase three-letter domains. For example, it does not validate harry@hogwarts.edu.UK or bilbo@canterbury.ac.Nz. However, Symantec Brightmail Gateway validates mister_smith @senate.GOV.
IP address	1.2.3.4 10.0.0.0 18.255.30.41 10.0.10.0/24 10.0.10.0/1 10.0.10.0/0	Any grouping of four-digit numbers that start with three 1-, 2-, or 3-digit numbers less than 256 separated by periods. A CIDR address range can be indicated by a 1- or a 2-digit number from 0 to 32 inclusive and separated from the initial IP address by a forward slash. Symantec Brightmail Gateway does not parse any terminal characters other than a 1- or 2-digit numeral that is preceded by a forward slash. Thus, 10.113.14.10a is not interpreted as a valid IP address.

Table 15-16 System patterns *(continued)*

System pattern	Examples	Description
Number	10 10.99 0.33 9999 10,000 6.999,66 99.999.999 -9,999 -10.99	Symantec Brightmail Gateway recognizes European-style numbers, where the comma serves as decimal point and periods separate groups of three digits. Fractions must be preceded by a numeral, including zero (0) if necessary, and expressed as a decimal. Although numbers 8 digits or smaller with commas are supported, Symantec recommends that you use tab- or pipe-delimited data-source text files that contain numbers using commas. The use of the tab or pipe delimiters avoids the possibility that commas in numbers are mistaken as field delimiters. Numbers that are larger than eight digits are interpreted as of type WORD.
Percent	76% 23.4% 56.78% -1.089,01% -0.32%	Symantec Brightmail Gateway recognizes European-style numbers, where the comma serves as decimal point and periods separate groups of three digits. Fractions of a percent must be preceded by a numeral, including zero (0) if necessary, and expressed as a decimal. Only the numbers that are adjacent to the percent sign (%) (no space) are regarded as valid percentages. The following patterns do not produce a match: .32 32 percent 32per 5 3/4%

Table 15-16 System patterns *(continued)*

System pattern	Examples	Description
US phone	(238) 832 5555 (238) 832-5555 238-832-5555 238 8325555 1-238-832-5555 238.832.5555 1 - (238) 8325555 1 - (238) 832-5555 1 - (238) 832 5555	Any 10-digit phone number beginning with 2-9 and/or preceded by 1 followed by a hyphen or a space-hyphen-space. The 3-digit area code can be enclosed in parentheses or not, followed by a space, hyphen, or period and the 7-digit number grouped into 3 and 4 digits that are separated by a space, hyphen or period or not separated at all.
US ZIP code	90210 89412-4321	Any 5-digit ZIP code or combination of 5-digit code plus 4-digit extension that are separated by a hyphen.
SSN/ITIN	777-77-7777 777 77 7777 123456789	Any 9-digit number, either continuous or separated into 3 groups of 3, 2, and 4 digits that are separated by a hyphen or space. The first group of three digits cannot be 000. The second number group must be greater than or equal to one, and the last number group must be greater than zero.

See "About mapping your data source file columns" on page 383.

See "Mapping data source file columns to fields in Symantec Brightmail Gateway" on page 388.

Mapping data source file columns to fields in Symantec Brightmail Gateway

After you define a record, map the columns in your data source file to fields in Symantec Brightmail Gateway. You can use predefined system pattern fields, or you can create custom fields.

See "Defining records" on page 381.

See "About mapping your data source file columns" on page 383.

See "About preventing data loss with structured data" on page 375.

To map a data source file, you must have Full Administration rights or rights to modify settings.

To map data source file columns to fields in Symantec Brightmail Gateway

1 On the **Add Record Resource** page, under **Mapping**, click the **Field Names** drop-down list and do one of the following tasks:

To use a system pattern field	Select the system pattern that you want to associate with the corresponding column in your data source file. For example, if column 1 in your data source file consists of phone numbers, in column 1, select **US phone pattern**. See “System patterns” on page 384.
To create your own custom field	Select **Customize**, and in the adjacent text field, type a unique custom field name. Custom field names cannot be the same as any of the predefined list of field names. Nor can they be the same as any other custom field name in the record. For example, assume column 3 of your data source file consists of first names. In column 3, select **Customize** and in the adjacent box, type **first_name**.

2 To change the order of the fields, check the box beside the field name that you want to move, and then click **Move Up** or **Move Down**.

3 To add additional fields, click **Add** and repeat the tasks in step 1 to use a system pattern or to create your own custom field.

4 Click **Next**.

After you map your data source file columns, upload the data source file.

See “Uploading data source files” on page 389.

Uploading data source files

After you map your data source file columns, you must upload your data source file. If the upload is unsuccessful, an error message appears that indicates the reason for failure. If upload is successful, Symantec Brightmail Gateway indexes the record. Indexing records helps Symantec Brightmail Gateway find rows in your data source file faster than if the record were not indexed.

See “Mapping data source file columns to fields in Symantec Brightmail Gateway” on page 388.

See “Modifying records” on page 391.

The **Record Resource Status** page shows the status of the upload process. It displays the time that the upload process starts and finishes. It also indicates the number of rows that were successfully uploaded and unsuccessfully uploaded.

Uploading a large data source file can be processing-intensive and time consuming. Consider uploading data source files during off-peak hours.

Note: Symantec Brightmail Gateway does not support uploading a data source file that is larger than 1.5 GB.

To upload a record, you must have Full Administration rights or rights to modify settings.

To upload data source files

1. On the **Edit Record Resource** page, under **Record Resource Data Source**, click **Browse** to locate the data source file that you want to upload.
2. Click **Upload**.

   After you successfully upload a record, you must then replicate it to your Scanners.

   See “Replicating records” on page 390.

Replicating records

After you successfully upload your data source file, you must replicate it to all of your Scanners. Transferring a large record to one or more Scanners can take a significant amount of time and processing resources. So Symantec Brightmail Gateway lets you replicate records when you create a new record or at a later time.

See “Uploading data source files” on page 389.

The **Record Resource Status** page indicates the status, the time the replication started and finished, and the size of the record that was replicated.

To replicate records, you must have Full Administration rights or rights to modify settings.

To replicate records when you create a new record

1. On the **Record Resource Status** page, under **Processing Status**, ensure that the status is **Completed**.
2. Under **Replication Status**, click **Replicate Now**.
3. When replication is successfully completed, click **OK**.

To replicate records at a later time

1. In the Control Center, click **Content > Resources > Records**.
2. Check the box beside the record that you want to replicate, and then click **Status Details**.
3. If the **Record Resource Status** page shows that the data source file is uploaded, under **Replication Status**, click **Replicate Now**.

   If the status is **Upload Failed**, you must try to upload your data source file again before you can replicate it.

   See "Uploading data source files" on page 389.
4. Click **OK**.

   To cancel the replication of a Record resource to all enabled Scanners, click **Cancel Replication**.

   After you replicate the record, create a view.

   See "About views" on page 393.

   See "Creating views" on page 397.

Modifying records

You can modify a record, provided that no upload of the record is in progress or pending. When you modify a record, you must upload and replicate the record again.

Policies that reference a record's views continue to use that record's existing data until Symantec Brightmail Gateway uploads, indexes, and replicates the new data. If you change the name of a record, all conditions that reference views of that record reflect the name change.

In addition, if you modify any mappings, any existing views that a policy condition references remains valid during the update. After the updated record is uploaded and replicated, all previously defined policies remain valid. The new column fields are available to create views.

You cannot delete a named field from a record if it is part of any existing view. This restriction applies whether or not a policy condition references that view. If you attempt to delete such a field, Symantec Brightmail Gateway displays an error message at the top of the **Edit Record Resource** page.

See "About preventing data loss with structured data" on page 375.

See "Deleting records" on page 392.

To modify records

1. In the Control Center, click **Content > Resources > Records**.
2. Check the box beside the record that you want to modify and click **Edit**.

   You can also click on the name of the record to edit it.
3. On the **Edit Record Resource** page, make the desired changes and then click **Next**.

   See “Defining records” on page 381.
4. On the **Edit Record Resource** page, under **Record Resource Data Source**, click **Browse** to locate the data source file that you want to upload.
5. Click **Upload**.

   See “Uploading data source files” on page 389.
6. On the **Record Resource Status page** under **Processing Status**, ensure that the status is **Completed**.
7. Under **Replication Status**, click **Replicate Now**.
8. When replication is successfully completed, click **OK**.

   See “Replicating records” on page 390.

Deleting records

To delete the record, you must first delete any of the record's views from all content filtering policy conditions. The **Record Resource Views** page lists all of the views that you have created. This page indicates the number of fields that you have selected for the view. It also shows the number of policies that use that view as a condition. However, this page does not indicate which policies use the views. Symantec Brightmail Gateway lets you know which policies use the views for that record when you attempt to delete it.

See “Deleting views” on page 399.

To delete records, you must have Full Administration rights or rights to modify settings.

See “About preventing data loss with structured data” on page 375.

To delete records

1. In the Control Center, click **Content > Resources > Records**.
2. Check the box beside the record that you want to remove and click **Delete**.

   A message appears at the top of the page indicating which content filtering policies use the views that are created from this record. Delete or modify this condition in these policies first, then repeat this procedure.

   See "Editing content filtering policies" on page 372.

About views

You must create record views to use in a content filtering policy condition. You use views to specify which fields or combinations of fields in a record you want Symantec Brightmail Gateway to match in a message.

You can create one or more views from a record. The benefit of multiple views is that you can create several conditions for a policy, and each condition can use a different view. Multiple views also let you use the same record for several different policies, and each policy can use a different view.

The **Record Resource Views** page lists all of the views that you have created. This page indicates the number of fields that you have selected for the view. It also shows the number of policies that use that view as a condition.

Table 15-17 describes the options that you use to create a view. It also provides an example. For the sake of our example, assume that Record A exists. Also assume that Record A and View A are used in a content filtering policy that is enabled.

Table 15-17 How to configure a view

Option	Description	Example
View Name	You should specify a unique name that describes the purpose of the view.	In our example, the view is named View A.

Table 15-17 How to configure a view *(continued)*

Option	Description	Example
Field Selection	When you create a view, all of the fields that you mapped for that record appear in the **Fields** list on the **Add Record Resource page**. You specify which fields you want to use in the view.	Assume that Record A consists of the following fields: ■ Credit card number pattern ■ Email pattern ■ IP address pattern ■ Number pattern ■ Percent pattern ■ US phone pattern ■ US ZIP code pattern ■ SSN/ITIN pattern Now assume that you select the following fields for this view: ■ Credit card number pattern ■ Email pattern ■ IP address pattern ■ US phone pattern ■ US ZIP code pattern ■ SSN/ITIN pattern
Minimum number of matched selected fields required for an occurrence	This option specifies how many of the fields in your view must appear in a message to constitute an "occurrence." For example, assume that you select six fields for your view, and the **Minimum number of matched selected fields required for an occurrence** is four. If any of the four (or more) fields that you selected are found in a message, Symantec Brightmail Gateway considers it an occurrence. An occurrence is not a violation. When you create a policy that uses this view, you can specify how many occurrences can occur in a message before the policy is violated. See "About preventing data loss with structured data" on page 375.	Assume that you specify the **Minimum number of matched selected fields required for an occurrence** as three. A message that meets or exceeds three matches is considered an occurrence. So if the only match in a message is **Email address pattern**, there is no occurrence. If the only match in a message is **Email address pattern** and **US phone pattern**, there is no occurrence. However, if a message contains matches for **Email address pattern** and **US phone pattern**, and **SSN/ITIN** pattern, the result is an occurrence.

Table 15-17 How to configure a view *(continued)*

Option	Description	Example
Exception Combinations	Specific combinations of fields may be acceptable in a message. So Symantec Brightmail Gateway lets you exclude the combination of fields from being considered a match. You can create multiple exceptions.	

Table 15-17 How to configure a view *(continued)*

Option	Description	Example
		Assume that you create an exception that consists of **US phone pattern** and **Email address pattern**. If Symantec Brightmail Gateway detects this combination of fields in a message, it does not consider it an occurrence. *Example 1* A message contains the following matches: ■ **Email address pattern** ■ **US phone pattern** ■ **SSN/ITIN pattern** No occurrence is detected because the combination of **US phone pattern** and **Email address pattern** is permitted. Therefore, the only match is **SSN/ITIN pattern**. *Example 2* A message contains the following matches: ■ **Email address pattern** ■ **US phone pattern** ■ **SSN/ITIN pattern** ■ **US ZIP code pattern** No occurrence is detected because the combination of **US phone pattern** and **Email address pattern** is permitted. Therefore, the only two matches are **SSN/ITIN pattern** and **US ZIP code pattern**. *Example 3* A message contains the following matches: ■ **Email address pattern** ■ **US phone pattern** ■ **SSN/ITIN pattern** ■ **US ZIP code pattern** ■ **Credit card number pattern** An occurrence is detected. Even though the combination of **US phone pattern** and **Email address pattern** is permitted, the remaining three matches met the threshold for an occurrence. *Example 4*

Table 15-17 How to configure a view *(continued)*

Option	Description	Example
		A message contains the following matches: - **US phone pattern** - **SSN/ITIN pattern** - **US ZIP code pattern** - **IP address pattern** An occurrence is detected. The message contains four matches, which exceed the threshold for an occurrence and none of the matches are an exception.

See "About preventing data loss with structured data" on page 375.

See "About mapping your data source file columns" on page 383.

See "Creating views" on page 397.

Creating views

After the record is replicated to the Scanner, you can define a view. A view is used to identify which fields from the associated record to use in a content filtering policy conditions. Symantec Brightmail Gateway uses views (not records) to create policy conditions, so you must create at least one view per record.

See "Replicating records" on page 390.

See "About views" on page 393.

See "Editing views" on page 398.

To create views

1. In the Control Center, click **Content > Resources > Records**.
2. On the **Records** page, check the box beside the record for which you want to create a view.
3. Click **Views**.
4. On the **Record Resource Views** page, click **Add**.
5. In the **View Name** field, type a unique name for the view.
6. Under **Field Selection**, check the fields that you want to include in this view.
7. In the **Minimum number of matched selected fields required for an occurrence** box, specify the number of fields that must be matched in a message to create an occurrence.

8 To create an exception, under **Exception Combinations**, click **Add.**

9 Select the combination of fields that you want omitted from the match evaluation.

10 Click **Save**.

11 On the **Add Record Resource** page, click **Save**.

After you create a view, you can use it in a content filtering policy condition.

See "Creating content filtering policies" on page 334.

Editing views

You can modify a view at any time. You can also add, edit, or delete **Exception Combinations**. When you modify a view, the changes are automatically propagated to any content filtering policies that you have that use that view.

See "About views" on page 393.

See "Creating views" on page 397.

See "Deleting views" on page 399.

To edit views

1 In the Control Center, click **Content > Resources > Records**.

2 On the **Records** page, check the box beside the record that contains the view that you want to edit.

3 Click **Views**.

4 On the **Record Resource Views** page, check the box beside the view that you want to edit.

5 Click **Edit**.

6 Make the desired changes to the name, fields, and the minimum number of matched selected fields required for an occurrence.

7 Do any of the following tasks:

To add a new Exception Combination	Do all of the following tasks: ■ Under **Exception Combinations**, click **Add.** ■ Select the combination of fields that you want omitted from the match evaluation. ■ Click **Save**.

To modify an existing Exception Combination	Do all of the following tasks: ■ Under **Exception Combinations**, check the box beside the Exception Combination that you want to edit. ■ Click **Edit.** ■ Select the combination of fields that you want omitted from the match evaluation. ■ Click **Save**.
To delete an existing Exception Combination	Do all of the following tasks: ■ Under **Exception Combinations**, check the box beside the Exception Combination that you want to delete. ■ Click **Delete.**

8 On the **Record Resource Views** page, click **Save**.

Deleting views

To delete a view, you must first delete that view from any content filtering policy conditions in which it is used. The **Record Resource Views** page lists all of the views that you have created. This page indicates the number of fields that you have selected for the view. It also shows the number of policies that use that view as a condition. However, this page does not indicate which policies use the views. Symantec Brightmail Gateway lets you know which policies use the view when you attempt to delete it.

See “About views” on page 393.

See “Deleting records” on page 392.

See “Editing views” on page 398.

To delete views

1. In the Control Center, click **Content > Resources > Records**.
2. On the **Records** page, check the box beside the record that contains the view that you want to delete.
3. Click **Views**.

4 On the **Record Resource Views** page, check the box beside the view that you want to delete.

5 Click **Delete**.

A message appears at the top of the page to indicate the content filtering policies in which a condition to use this view occurs. Delete or modify this condition in these policies first, then repeat this procedure.

See "Editing content filtering policies" on page 372.

About policy violation notifications

Notifications are the messages that are sent when a policy is violated, and the policy action is to **Send notification**. Notifications are different from annotations in that annotations are added to the message, while a notification is a separate message. When you specify to send a notification, the original message is delivered to the original recipient, unless you specify an additional action to do otherwise. Notifications can be used for any type of policy. You can select the notation that you want to use when you specify the action for a policy.

See "Specifying content filtering policy actions" on page 365.

See "Content filtering policy actions" on page 365.

Symantec Brightmail Gateway provides several predefined notifications that you can use. You can also create your own notification message. Predefined notifications can be modified and can only be deleted if they are not used in a policy.

Notifications can be sent to any of the following people:

- Message sender
- Message recipient
- Third party (such as an administrator)

You can customize the notification message subject and body text with variables. For example:

```
The message concerning $subject$ sent by $sender$
to $recipients$ was stripped of $attachments$.
```

See "Creating policy violation notifications" on page 401.

Creating policy violation notifications

Notifications are the predefined messages that are sent when a message meets the conditions that are specified in a policy and the action is to send a notification. The name that you specify for the notification appears on the **Notifications** page. This name is also the name that appears in the **Notification** list when you choose the **Send notification** action when you configure a policy.

You can create notifications for violations to content filtering policies and IM policies.

See "Specifying content filtering policy actions" on page 365.

Notification recipients can reply to the email address that sends the notification. Replies go to the address that is the envelope address, not the From: header address. As a best practice, either provide an email address for an account that is monitored or include a statement in the notification that responses are not monitored.

Email notification variables let you customize your notification subject text and body text. You can use one or more notification variables in your text, or none at all.

Table 15-18 lists the notification variables that you can use to customize your message subject text and body text and what the tags are replaced with.

Table 15-18 Email notification variable attributes

Attribute name	Variable tag	Result
Sender	$sender$	Sender's email address
Recipients	$recipients$	Comma-separated list of recipient email addresses
Subject	$subject$	Subject line of the original message
Attachment names	$attachments$	Comma-separated list of top-level attachments from the original message

To create a notification, you must have Full Administration rights or rights to modify policies.

See "About policy violation notifications" on page 400.

To create policy violation notifications

1 In the Control Center, click **Content > Resources > Notifications**.

2 Click **Add**.

3 In the **Notification description** box, type a name for the notification.

4 Select the message protocol that you want to use.

5 In the **Send from** box, type an email address to appear in the `From` header of the notification message.

Specify the full email address including the domain name, such as admin@symantecexample.com.

6 Under **Send to**, select one or more of the following options:

Sender	Check this box to send the notification to the sender that is listed in the message envelope (not the sender that is listed in the `From:` header).
Recipients	Check this box to send the notification to the recipients that are listed in the message envelope (not the recipients that are listed in the `To:` header).
Others	Check this box to send the notification to one or more complete email addresses that you specify. Then type the email address. Separate multiple email addresses with a comma, semicolon, or space.

7 Under **Subject**, click the **Encoding** drop-down list to change the character encoding for the notification subject line.

8 Optionally, click the **Include message attribute** drop-down list to select the type of variable you want to include in the subject line, and then click **Add**.

The variable appears in the **Subject** text box.

9 In the **Subject** text box, type the text for the subject header of the notification message.

10 Under **Body**, click the **Encoding** drop-down list to change the character encoding for the notification body.

11 Optionally, click the **Include message attribute** drop-down list to select the type of variable you want to include in the body, and then click **Add**.

The variable appears in the **Message body** text box.

12 In the **Message body** text box, type the text for the body of the notification message.

13 Optionally, check **Attach the original message** to attach the original message to the notification message.

14 Click **Save**.

Editing policy violation notifications

You can modify predefined and custom policy violation notifications notifications as needed. However, the protocol for the predefined policy violation notifications cannot be changed.

To edit a policy violation notification, you must have Full Administration rights or rights to modify policies.

To edit policy violation notifications

1 In the Control Center, click **Content > Resources > Notifications**.

2 Check the box beside the notification that you want to edit.

3 Click **Edit**.

You can also click on the notification that you want to edit.

4 Make the desired changes.

5 Click **Save**.

Deleting policy violation notifications

You can delete any custom policy violation notifications that you create when they are no longer needed. To delete a policy violation notification, you must first delete that notification from any content filtering policy actions in which it is used. Symantec Brightmail Gateway lets you know which policies use that notification when you attempt to delete it.

To delete a policy violation notification, you must have Full Administration rights or rights to modify policies.

See “About policy violation notifications” on page 400.

To delete policy violation notifications

1 In the Control Center, click **Content > Resources > Notifications**.

2 Check the box beside the notification that you want to delete.

3 Click **Delete**.

A message appears at the top of the page to indicate the content filtering policies in which an action to use this notification occurs. Delete or modify this action in these policies first, then repeat this procedure.

See “Editing content filtering policies” on page 372.

About content filtering dictionaries

A dictionary is a list of the words, phrases, file names, and the file name extensions that pertain to a specific topic. When a dictionary is part of a content filtering policy condition, Symantec Brightmail Gateway searches your email messages for the contents of the dictionary. If you configure Symantec Brightmail Gateway to scan non-plain text file attachments, they too are scanned for contents of the dictionary.

See “Scanning non-plain text file attachments for content filtering violations” on page 407.

Content filtering policies evaluate matches to a referenced dictionary with substring text analysis, not regular expression analysis. When a substring in a message matches a dictionary entry, Symantec Brightmail Gateway takes the actions that you specify in the policy.

The types of dictionaries that you can use are as follows:

Premium dictionaries and custom dictionaries	Most of the premium dictionaries and custom dictionaries that Symantec Brightmail Gateway provides already contain predefined words, phrases, and characters. You can view the contents of the premium dictionaries and custom dictionaries, enable and disable existing keywords, and add words, phrases, or characters. Examine the contents of the premium dictionaries and custom dictionaries before you use them to determine if they meet your needs. If they do not, you can add the new user-defined words and phrases that you need. You can also disable the predefined words and phrases that you do not want to use. Some premium dictionaries and custom dictionaries are empty. You must populate them with user-defined words or phrases before you can use them in a content filtering policy. **Note:** The dictionaries that are marked as ambiguous (such as **Profanity (ambiguous)**) contain the terms that can be legitimate when used in certain contexts. See "Premium and custom content filtering dictionaries" on page 413.
User-defined dictionaries	In addition to premium dictionaries and custom dictionaries, you can create user-defined dictionaries. You populate user-defined dictionaries to suit a site-specific or business need. These dictionaries must contain at least one word or phrase.

Table 15-19 describes the tasks that you can perform with the different types of dictionaries.

Table 15-19 Tasks that you can perform with dictionaries based on type

Task	Premium or custom	User- defined
Create a new dictionary See "Creating user-defined dictionaries" on page 409.		X
Edit the dictionary See "Editing dictionaries" on page 411.	X	X

Table 15-19 Tasks that you can perform with dictionaries based on type *(continued)*

Task	Premium or custom	User- defined
Delete the dictionary See "Deleting dictionaries" on page 412.		X
Import words or phrases into the dictionary See "Importing words or phrases into dictionaries" on page 410.	X	X
Export words or phrases from the dictionary See "Exporting words or phrases from dictionaries" on page 407.	X	
Enable and disable words and phrases in the dictionary See "Disabling and enabling predefined words or phrases in dictionaries" on page 408.		X
Delete words or phrases in the dictionary See "Deleting user-defined words or phrases from dictionaries" on page 411.	X	
Search for words or phrases in the dictionary See "Finding words or phrases in dictionaries" on page 412.	X	X

See "Considerations when you use dictionaries" on page 406.

Considerations when you use dictionaries

Note the following additional information about dictionaries:

- Words in dictionaries are detected only if they match exactly with words as they appear in the dictionary. Words in dictionaries are not detected if they are variations, such as verb tenses. For example, "selling" is not a match with the word "sell."
- Wildcards are not supported in dictionaries.
- Up to 100 dictionaries are supported, and each dictionary can contain up to 10,000 words.
- Words and phrases in dictionaries cannot be prioritized.

- A dictionary can be used in multiple content filtering policies.
- When you add words to a dictionary, keep in mind that some words can be considered both profane and legitimate, depending on the context. For example, the word "breast" might be inappropriate until it is coupled with the word "cancer."
- Symantec Brightmail Gateway does not search for dictionary matches in the HTML headers or tags of HTML messages or HTML attachments.
- Words are not case sensitive.
- Punctuation characters (such as a question mark (?)) cannot be used as the first character of a word.

See “About content filtering dictionaries” on page 404.

See “Creating user-defined dictionaries” on page 409.

Scanning non-plain text file attachments for content filtering violations

Symantec Brightmail Gateway can check any file attachments that are not plain text files against dictionaries. Scanning non-plain text files maximizes the effect of content filtering, but it can affect the system load and slow performance.

To scan non-plain text file attachments for content filtering violations

1. Click **Protocols > SMTP > Settings**.
2. Under **Content Filtering Settings**, check **Enable scanning of non-plain text attachments for words in dictionaries**.
3. Click **Save**.

Exporting words or phrases from dictionaries

Symantec Brightmail Gateway lets you export words and phrases from premium dictionaries and custom dictionaries. By default, the exported words and phrases are saved in a text file named `DictionaryWords.txt`. You cannot export specific words or phrases to export. You must export all of the words and phrases in the dictionary.

Note: You cannot export words or phrases from user-defined dictionaries.

This feature is helpful when you set up a new Control Center and Scanner deployment. Assume that your existing Control Center and Scanner deployment contain some dictionary words and phrases that you want to use in the new deployment. You can export the words and phrases that you want to use in the new deployment. Then you can import those words and phrases into the new deployment from the Control Center computer.

See "Importing words or phrases into dictionaries" on page 410.

You may also want to export a premium dictionary or custom dictionary before you delete it to maintain for your records. You could also restore the exported file, if needed.

See "Deleting dictionaries" on page 412.

To export words or phrases from premium or custom dictionaries, you must have Full Administration rights or rights to modify settings.

To export words or phrases from dictionaries

1. In the Control Center, click **Content > Resources > Dictionaries**.
2. Check the box beside the dictionary that contains the words or phrases that you want to export, and click **Edit**.
3. Click **Export**.
4. In the **File Download** dialog box, click **Save**.
5. In the **Save As** dialog box, type the file name and select the location where you want to save the file.
6. Click **Save**.
7. In the **Download Complete** dialog box, click **Open** to view the text file or **Close** to close the dialog box.

Disabling and enabling predefined words or phrases in dictionaries

Predefined words and phrases are words or phrases that Symantec Brightmail Gateway provides. You cannot delete predefined dictionary words or phrases. However, you can disable the predefined words or phrases that you do not want to use in content filtering policies.

Note: You cannot disable user-defined words or phrases in any of the dictionaries (premium, custom, or user-defined). If you do not want to use user-defined words, you must delete them.

See "Deleting user-defined words or phrases from dictionaries" on page 411.

To enable or disable predefined words or phrases, you must have Full Administration rights or rights to modify policies.

See "About content filtering dictionaries" on page 404.

See “Creating user-defined dictionaries” on page 409.

To disable or enable predefined words or phrases in dictionaries

1 In the Control Center, click **Content > Resources > Dictionaries**.

2 Check the box beside the dictionary that contains the words or phrases that you want to disable or enable, and click **Edit**.

  See “Finding words or phrases in dictionaries” on page 412.

3 Check the box beside the predefined word or phrase that you want to enable or disable.

  Check **Word or Phrase** to select all of the words and phrases in the list.

4 Do one of the following tasks:

Click **Disable**.	If the dictionary is a condition of a policy, Symantec Brightmail Gateway does not search for this word or phrase when it evaluates messages.
Click **Enable**.	Symantec Brightmail Gateway searches for the word or phrase when it evaluates messages.

5 Click **Save**.

Creating user-defined dictionaries

Symantec Brightmail Gateway provides the premium dictionaries and custom dictionaries that you can use in your content filtering policies. However, you can create your own user-defined dictionary to suit your specific needs.

To add a user-defined dictionary, you must have Full Administration rights or rights to modify policies.

See “About content filtering dictionaries” on page 404.

See “Importing words or phrases into dictionaries” on page 410.

See “Editing dictionaries” on page 411.

See “Disabling and enabling predefined words or phrases in dictionaries” on page 408.

To create user-defined dictionaries

1 In the Control Center, click **Content > Resources > Dictionaries**.

2 Click **Add**.

3 In the **Dictionary name** field, type a name for the dictionary.

4 In the **Optional description** field, type a description of the dictionary.

5 In the **Enter a word or phrase** field, type the word or phrase that you want to add to your dictionary.

6 Click **Add** to add the word or phrase to the **Word or Phrase** list.

The word or phrase is automatically enabled when you add it.

7 Repeat steps 5 and 6 to add more words and phrases.

8 Click **Save**.

Importing words or phrases into dictionaries

You can import words or phrases into all types of dictionaries (premium, custom, and user-defined). The file that you import must be a newline-delimited text file in UTF-8 format, and each word or phrase must be on a separate line. The file must be accessible to the Control Center computer.

Note: Files with extended ASCII are not supported.

See "About content filtering dictionaries" on page 404.

See "Creating user-defined dictionaries" on page 409.

See "Editing dictionaries" on page 411.

See "Disabling and enabling predefined words or phrases in dictionaries" on page 408.

To import words or phrases into dictionaries

1 In the Control Center, click **Content > Resources > Dictionaries**.

2 Check the box beside the dictionary that you want to import words or phrases into, and then click **Edit**.

3 On the **Add Dictionary** page, click **Import**.

4 Under **Import**, do one of the following tasks:

- In the **File name** box, type the name of the text file that you want to import.
- Click **Browse** to locate the text file that you want to import.

The newly imported words or phrases are enabled by default.

5 Click **Save**.

Editing dictionaries

You can make modifications to any type of dictionary (premium, custom, user-defined) as needed.

See "Creating user-defined dictionaries" on page 409.

Note: Once you create a user-defined dictionary and add at least one word, you must always have at least one word in that dictionary. If you no longer need the user-defined dictionary, you can delete it.

See "Deleting dictionaries" on page 412.

To edit a dictionary, you must have Full Administration rights or rights to modify policies.

To edit dictionaries

1. In the Control Center, click **Content > Resources > Dictionaries**.
2. Check the box beside the dictionary that you want to edit, and click **Edit**.
3. Make the desired modifications.

   See "Creating user-defined dictionaries" on page 409.

   See "Importing words or phrases into dictionaries" on page 410.

   See "Disabling and enabling predefined words or phrases in dictionaries" on page 408.

   See "Deleting user-defined words or phrases from dictionaries" on page 411.
4. Click **Save**.

Deleting user-defined words or phrases from dictionaries

User-defined words or phrases are the words or phrases that you add to any type of dictionary (premium, custom, and user-defined). You can delete any of the user-defined words or phrases when they are no longer needed.

Note: You cannot delete predefined words or phrases. However, you can disable the predefined word or phrase that you do not want to use in content filtering policies.

See "Disabling and enabling predefined words or phrases in dictionaries" on page 408.

To delete user-defined words or phrases from dictionaries, you must have Full Administration rights or rights to modify policies.

To delete user-defined words or phrases from dictionaries

1. In the Control Center, click **Content > Resources > Dictionaries**.
2. Check the box beside the dictionary that contains the user-defined word or phrase that you want to delete, and click **Edit**.

   See "Finding words or phrases in dictionaries" on page 412.
3. On the **Edit Dictionary** page, check the box beside the user-defined word or phrase that you want to remove.

   Check **Word or Phrase** to select all of the words and phrases in the list.
4. Click **Delete**.
5. Click **Save**.

Finding words or phrases in dictionaries

You may have an instance in which you want to enable or disable, modify, or delete a word or phrase in multiple dictionaries. Symantec Brightmail Gateway lets you search for the words and phrases that are used in all types of dictionaries (premium, custom, and user-defined).

When you search for a word or phrase, Symantec Brightmail Gateway displays a message at the top of the **Dictionaries** page. The message lists all of the dictionaries that contain the word or phrase. If no match is found, the following message appears: No dictionaries match the word or phrase.

To search for words or phrases in dictionaries, you must have Full Administration rights or rights to modify policies.

To find words or phrases in dictionaries

1. In the Control Center, click **Content > Resources > Dictionaries**.
2. On the **Dictionaries** page, in the **Find word or phrase** field, type the word or phrase that you want to find.
3. Click **Find**.

Deleting dictionaries

You can delete any user-defined dictionary that you create unless it is used in a custom filtering policy. Before you delete a user-defined dictionary, you may want to export the words and phrases to a text file to maintain for your records. The text file can be reimported, if needed.

Note: You cannot delete premium or custom dictionaries.

See "Exporting words or phrases from dictionaries" on page 407.

See "Importing words or phrases into dictionaries" on page 410.

To delete a user-defined dictionary, you must first delete that dictionary from any content filtering policy conditions in which it is used. Symantec Brightmail Gateway lets you know which policies use that dictionary when you attempt to delete it.

To delete a user-defined dictionary, you must have Full Administration rights or rights to modify policies.

See "About content filtering dictionaries" on page 404.

To delete dictionaries

1 In the Control Center, click **Content > Resources > Dictionaries**.

2 On the **Dictionaries** page, check the box beside the dictionary that you want to delete.

3 Click **Delete**.

A message appears at the top of the page to indicate the content filtering policies in which a condition to use this dictionary occurs. Delete or modify this condition in these policies first, then repeat this procedure.

See "Editing content filtering policies" on page 372.

Premium and custom content filtering dictionaries

Symantec Brightmail Gateway provides premium dictionaries and custom content filtering dictionaries that you can use when you create content filtering policy conditions.

See "About content filtering dictionaries" on page 404.

Table 15-20 lists the premium dictionaries and custom dictionaries and the recommended policy templates with which to use them.

Table 15-20 Premium dictionaries, custom dictionaries, and related templates

Dictionary name	Associated templates
ABA Routing Number Keywords	PIPEDA; Employee Data Protection; Customer Data Protection; Swift Codes; Employee Data Protection; State Data Privacy; Gramm-Leach-Bliley

Table 15-20 Premium dictionaries, custom dictionaries, and related templates *(continued)*

Dictionary name	Associated templates
Affiliate Domains	State Data Privacy
Analysts' Email Addresses	NASD Rule 2711 and NYSE Rules 351 and 472
California Keywords Illinois Keywords Letter/12 Num. DLN State Words New Jersey Keywords New York Keywords	State Data Privacy
Canadian Social Ins. No. Words	PIPEDA; Canadian Social Insurance Number
Company Name Keywords (user-defined)	SEC Fair Disclosure Regulation; Sarbanes-Oxley
Competitor Domains	Competitor Communications
Confidential Keywords (user-defined) Internal Use Only Keywords (user-defined) Proprietary Keywords (user-defined) Not for Distribution Words (user-defined)	Confidential Documents
Confidential/Proprietary Words	Sarbanes-Oxley
Credit Card Number Keywords	Credit Card Numbers; Payment Card Industry Data Security Standard; Customer Data Protection; Employee Data Protection; PIPEDA
Design Documents Extensions	Design Documents
Disease Names	HIPAA (including PHI); Caldicott Report
EAR CCL Keywords; EAR Country Codes	Export Administration Regulations (EAR)
EU Country Codes	EU Data Protection Directives
Financial Keywords	Sarbanes-Oxley Financial Information
Gambling Keywords, Confirmed; Gambling Keywords, Suspect	Gambling

Table 15-20 Premium dictionaries, custom dictionaries, and related templates *(continued)*

Dictionary name	Associated templates
GPG Encryption Keywords PGP file extensions PGP8 Keywords	Encrypted Data
Hacker Keywords; Keylogger Keywords	Network Security
ITAR Country Codes ITAR Munition Names	International Traffic in Arms Regulations (ITAR)
Job Search Keywords, Education Job Search Keywords, General Job Search Keywords, Work	Resumes
M & A Project Code Names (user-defined)	Mergers and Acquisitions Data
Manufd. Controlled Substances Street Drug Names	Illegal Drugs
Media Files Extensions	Media Files
Medical Treatment Keywords	HIPAA (including PHI)
NASD 2711 Keywords (user-defined)	NASD Rule 2711 and NYSE Rules 351 and 472
NASD 3010 General Keywords NASD 3010 Stock Keywords NASD 3010 Buy/Sell Keywords	NASD Rule 3010 and NYSE Rule 342
OFAC Country Codes OFAC SDN Country Codes SDN List	Office of Foreign Assets Control (OFAC)
Offensive Language, Explicit Offensive Language, General	Offensive Language
Other Sensitive Information	US Intelligence Control Markings (CAPCO) & DCID 1/7
Password Filenames	Password Files

Table 15-20 Premium dictionaries, custom dictionaries, and related templates *(continued)*

Dictionary name	Associated templates
Prescription Drug Names	HIPAA (including PHI) Caldicott Report
Publishing Document Extensions	Publishing Documents
Racist Language	Racist Language
Restricted Recipients	Restricted Recipients
SEC Fair Disclosure Keywords	SEC Fair Disclosure Regulation
Secret Top Secret Classified or Restricted	Defense Message System (DMS) GENSER Classification US Intelligence Control Markings (CAPCO) & DCID 1/7
Sensitive Keywords	NERC Security Guidelines for Electric Utilities
Sensitive Project Code Names	Project Data
Sex. Explicit Words, Confirmed Sex. Explicit Words, Possible Sex. Explicit Words, Suspect	Sexually Explicit Language
Source Code Extensions	Source Code
SWIFT Code Keywords	SWIFT Codes
TPO Email Addresses	HIPAA (including PHI)
UK Electoral Roll Number Words	UK Electoral Roll Numbers Data Protection Act 1998 Human Rights Act 1998
UK Tax ID Number Keywords	UK Tax ID Numbers Data Protection Act 1998
UK NIN Keywords	UK National Insurance Number Data Protection Act 1998 Caldicott Report

Table 15-20 Premium dictionaries, custom dictionaries, and related templates *(continued)*

Dictionary name	Associated templates
UK Keywords	UK Drivers License Numbers UK Electoral Roll Numbers Data Protection Act 1998 Human Rights Act 1998
UK Passport Keywords	UK Passport Numbers Data Protection Act 1998
UK Personal Data Keywords	Human Rights Act 1998
US ITIN Keywords	Individual Taxpayer Identification Numbers
US SSN Keywords	US Social Security Numbers Employee Data Protection
Violence Keywords Weapons Keywords	Violence and Weapons
Vulnerability Keywords	NERC Security Guidelines for Electric Utilities

About patterns

Symantec Brightmail Gateway provides a feature that lets you use patterns to detect policy violations. Patterns are regular expressions or system patterns that depict commonly known forms of content, such as the pattern for credit card numbers. Symantec Brightmail Gateway provides predefined basic patterns and premium patterns. You can also create your own custom patterns.

All of these patterns are available in the **Matches pattern/does not match pattern** drop-down list on the **Content Filtering Policy Conditions** page when you define a policy condition.

Table 15-21 describes the types of patterns that you can use.

Table 15-21 Patterns for content filtering policy conditions

<table>
<tr><th>Pattern</th><th>Description</th></tr>
<tr><td>Basic</td><td>Basic patterns are known and performance-tested regular expressions that you can use in a policy or as examples when you create a custom pattern. You can view basic patterns, but these patterns cannot be edited or deleted.
The predefined basic patterns and their syntax are as follows:
■ Credit Card
\b(?<!-)((4\d{3})|(5[1-5]\d{2})|(6011))(-?)\d{4}\5\d{4}\5\d{4}(?!-)\b|(\b)(?<!-)3[4,7]\d{13}(?!-)\b
■ Email Address
\b([0-9a-zA-Z]([-.\w]*[0-9a-zA-Z])*@([0-9a-zA-Z][-\w]*[0-9a-zA-Z]\.)+[a-zA-Z]{2,9})\b
■ Social Security Number
\b(?<!-)(\d{3})([-]?)(\d{2})\2(\d{4})(?!-)\b
■ US Phone Number
\b(?<!-)((1([\s.-]?)[2-9]\d{2}\3\d{3}\3\d{4})|([2-9]\d(2)([.-]?)\d{3}\2\d{4})|(\d{3}([.-]?)\d{4}))(?!-)\b
■ US Zipcode
\b(?<!-)((\d{5})|(\d{5}\-\d{4}))(?!-)\b</td></tr>
<tr><td>Premium</td><td>Premium patterns perform additional checking and validation (such as Luhn checking) beyond regular expression definitions to reduce false positives. These patterns cannot be viewed, edited, or deleted.
The predefined premium patterns and the policy templates that you can use them with are as follows:
■ Valid Credit Card
Credit Card Numbers; Customer Data Protection; Employee Data Protection; PIPEDA; State Data Privacy; Payment Card Industry Data Security Standard; Gramm-Leach-Bliley
■ Valid IP Address
Network Diagrams
■ Valid Social Security Number
HIPAA; Gramm-Leach-Bliley; US Social Security Numbers; Individual Taxpayer Identification Numbers (ITIN); State Data Privacy; Customer Data Protection; Employee Data Protection
Note: The Valid Credit Card and Valid Social Security Number patterns differ from the Credit Card and Social Security Number patterns that are provided in the basic patterns. The valid patterns contain a rules checking feature to help insure their completeness and accuracy.</td></tr>
</table>

Table 15-21 Patterns for content filtering policy conditions *(continued)*

Pattern	Description
Custom	You can create your own custom pattern with the syntax that you copy from a basic pattern or regular expressions. You can view, edit, and delete custom patterns. An example of a custom pattern is a pattern for non-US phone numbers. You should consult an advanced resource on regular expressions to construct your own patterns.

See “Creating content filtering policies” on page 334.

See “Creating your own custom patterns” on page 419.

See “Editing patterns” on page 420.

See “Deleting patterns” on page 421.

Creating your own custom patterns

Symantec Brightmail Gateway provides the basic patterns and premium patterns that you can reference in your policy conditions. Symantec Brightmail Gateway uses these patterns to search messages for the character patterns that indicate the data that you may want to restrict.

You can copy all or any portion of a basic pattern to use in your custom pattern, or you can use regular expressions.

See “About patterns” on page 417.

See “Putting predefined regular expressions into content filtering policy conditions” on page 357.

See “Predefined regular expressions” on page 358.

An example of a custom pattern that you can create is for non-US telephone numbers. For example, Australia has a variety of formats for phone numbers. One such format follows the pattern: nn nnnn nnnn

To detect Australian telephone numbers that follow this format, you could create the following pattern:

```
\b(([(]0\d[)]|0\d)\s)?\d{4}\s\d{4}\b
```

Your pattern might also include formats with country codes for international dialing.

Note: You should consult an advanced resource on regular expressions to construct your own patterns.

To create a custom pattern, you must have Full Administration rights or rights to modify policies.

See “Editing patterns” on page 420.

See “Deleting patterns” on page 421.

To create your own custom pattern

1. In the Control Center, click **Content > Resources > Patterns**.
2. Click the **Custom** tab.
3. Click **Add**.
4. In the **Pattern name** field, type a name for your custom pattern.
5. In the **Regular expression** field, type the regular expression.
6. Click **Save**.

Editing patterns

You cannot edit basic patterns and premium patterns. However, you can edit the custom patterns that you create. You can only edit one pattern at a time. To edit a custom pattern, you must have Full Administration rights or rights to modify policies.

See “About patterns” on page 417.

See “Putting predefined regular expressions into content filtering policy conditions” on page 357.

See “Creating your own custom patterns” on page 419.

See “Deleting patterns” on page 421.

To edit patterns

1. In the Control Center, click **Content > Resources > Patterns**.
2. Click the **Custom** tab.
3. Check the box beside the pattern that you want to edit, and then click **Edit**.

   You can also click the custom pattern name to edit the pattern.
4. In the **Pattern name** field, edit the pattern name as needed.

   If a pattern's name is changed, that name change is propagated to any policy that references the custom pattern.

5 In the **Regular expression** field, edit the regular expression as needed.

6 Click **Save**.

Deleting patterns

You cannot edit the basic patterns or premium patterns. However, you can delete the custom patterns that you create when you no longer need them.

To delete a custom pattern, you must first delete that pattern as a condition from any content filtering policies in which it is used. Symantec Brightmail Gateway lets you know which policies use that pattern when you attempt to delete it.

See “Editing patterns” on page 420.

To delete a custom pattern, you must have Full Administration rights or rights to modify policies.

See “About patterns” on page 417.

See “Creating your own custom patterns” on page 419.

To delete patterns

1 In the Control Center, click **Content > Resources > Patterns**.

2 Click the **Custom** tab.

3 Check the box beside the pattern that you want to delete.

Check the box beside **Patterns** to select all of the patterns for deletion.

4 Click **Delete**.

A message appears at the top of the page to indicate the content filtering policies in which a condition to use this pattern occurs. Delete or modify this condition in these policies first, then repeat this procedure.

See “Editing content filtering policies” on page 372.

Annotating messages that contain violations

Symantec Brightmail Gateway lets you add annotations to the content filtering messages and IM messages that contain violations. To do so, when you create a policy you must select the action to add an annotation. For email, you can prepend the annotation to the beginning of the message body or append it to the end of the message. If you prepend, you may want to end your annotation text with a blank line or a line of dashes. Some type of separator between the annotation and message body provides the user a visual cue as to where the message body begins. For IM, the annotation is automatically prepended to the message.

See “Selecting virus policies for a policy group” on page 322.

See “Selecting spam and spim policies for a policy group” on page 324.

Annotations are different from policy violation notifications. Annotations are added to the message. A policy violation notification is a separate message that is sent when a policy is violated and the policy action is to **Send notification**.

See “About policy violation notifications” on page 400.

Symantec Brightmail Gateway supports the following types of annotations:

Plain text	Required for both content filtering messages and IM messages. Text for content filtering messages should not exceed 65,000 characters. Text for IM messages should not exceed 128 characters.
HTML	Optional for content filtering messages. Unsupported for IM messages. You can use HTML formatting tags, such as <b>bold text here</b>. HTML structure tags are unsupported. Text should not exceed 128 characters.

All IM message and content filtering message violations for which the action is to **Add annotation** receive a plain text annotation.

If a content filtering message's body part is in both text and HTML, one of the following events occurs:

You specify a plain text annotation and an HTML annotation	The plain text annotation is added to the plain text part, and the HTML annotation is added to the HTML part. Only the HTML version appears in the resulting email message.
You specify only a plain text annotation	Only the plain text annotation is added.

Note: As a best practice, ensure that both the plain text message and the HTML message are the same.

To add an annotation, you must have Full Administration rights or rights to modify policies.

See “Editing annotations” on page 423.

To annotate messages that contain violations

1. In the Control Center, click **Content > Resources > Annotations**.
2. On the **Annotations** page, click **Add**.
3. In the **Annotation description** field, type a name for the annotation.
4. Beside **Protocol**, click **Email** or **Instant Messaging**.
5. Choose a character encoding for the plain text annotation.
6. In the **Plain text** box, type the annotation text.

   If this annotation is for instant messaging, go to step 9.
7. Choose a character encoding for the HTML annotation.
8. In the **HTML** box, type the annotation text.
9. Click **Save**.

Editing annotations

You can edit an existing annotation to change the wording.

Symantec Brightmail Gateway supports the following types of annotations:

Plain text	Required for both content filtering messages and IM messages. For content filtering messages, text should not exceed 1,000 characters. For IM messages, text should not exceed 128 characters.
HTML	Optional for content filtering messages. Unsupported for IM messages. You can use HTML formatting tags, such as <b>bold text here</b>. HTML structure tags are unsupported. Text should not exceed 128 characters.

To edit an annotation, you must have Full Administration rights or rights to modify policies.

See "Annotating messages that contain violations" on page 421.

To edit annotations

1. In the Control Center, click **Content > Resources > Annotations**.
2. On the **Annotations** page, check the box beside the annotation that you want to edit, and click **Edit**.

   You can also click on the name of the annotation to edit it.
3. Change the annotation text as desired.
4. Click **Save**.

Deleting annotations

When you no longer need an annotation, you can delete it. To delete an annotation, you must first delete that annotation from any content filtering policy actions in which it is used. Symantec Brightmail Gateway lets you know which policies use that annotation when you attempt to delete it.

To delete an annotation, you must have Full Administration rights or rights to modify policies.

See “Annotating messages that contain violations” on page 421.

To delete an annotation

1. In the Control Center, click **Content > Resources > Annotations**.
2. On the **Annotations** page, check the box beside the annotation that you want to delete, and click **Delete**.

   A message appears at the top of the page to indicate the content filtering policies in which an action to use this annotation occurs. Delete or modify this action in these policies first, then repeat this procedure.

   See “Editing content filtering policies” on page 372.

About attachment lists

An attachment list contains the file extensions and the file application types that you want Symantec Brightmail Gateway to detect. You use an attachment list as a condition of a content filtering policy. When Symantec Brightmail Gateway detects an attachment that has an extension that is on the file extension list, it applies the action that you specify. Symantec Brightmail Gateway can determine the email attachment type based on the application that created it, regardless of its file extension.

You can create an attachment list with any combination of the following:

True file classes	True file classes are categories in which similar file types and applications are grouped. For example, **Desktop publishing** and **Movie file** are true file classes. Within a true file class are the true file types. For example, FrameMaker is a true file type within the Desktop publishing true file class. You cannot modify the list of true file types within a true file class. You must use the list as is.
True file types	True file types are contained within a true file class. True file types consist of file extensions or applications. For example, a true file type within the category **Desktop publishing** is FrameMaker. A true file type within the category **Movie file** is MPEG movie. You can add and delete items in a true file types list. **Note:** The **Archive Files** attachment list does not support true file type MIME. The Document Files attachment list does not support true file type HTML.
Common file extensions and MIME-types	You can specify common file extensions and MIME-types. For example, you can filter out all .jpeg file attachments, regardless of the application that created it. For a technical description of MIME, see the following RFC: http://www.ietf.org/rfc/rfc2045.txt

See “Creating lists to detect file attachment types” on page 425.

Symantec Brightmail Gateway provides predefined attachment lists. These lists contain the pre-populated file classes and file types that you can use to create your attachment list.

See “Predefined attachment lists” on page 429.

Creating lists to detect file attachment types

Use attachment lists in content filtering policy conditions to detect and act on email attachment file types. For example, you can assign an attachment list that detects messages with .exe file attachments as a condition in a content filtering policy. You can then specify what action you want Symantec Brightmail Gateway to take if messages with an .exe file attachment are detected.

Symantec Brightmail Gateway provides predefined attachment lists that you can use to create your attachment list. All of the attachments lists (predefined and custom) appear in the **Conditions** section on the **Email Content Filtering Policy** page. So you can conveniently select the attachment list that you want to use as a condition when you configure a content filtering policy.

When you create an attachment list, you can use any combination of true file classes, true file types, and common file extensions and MIME-types.

Note: MIME-type and file extension-based attachment lists are ineffective for stripping the attachments that are not described in the MIME header (for example, in the case of some forwarded messages). To ensure that you strip all instances of an attachment file type, configure your attachment lists to use true file type whenever possible.

To create an attachment list, you must have Full Administration rights or rights to modify policies.

See "About attachment lists" on page 424.

See "Editing attachment lists" on page 428.

To create lists to detect file attachment types

1. In the Control Center, click **Content > Resources > Attachment Lists**.
2. Click **Add**.
3. In the **Attachment list name** box, type a name for the attachment list.
4. Under **Add Attachment Types**, do any of the following tasks:

To add a true file class to your attachment list	Do all of the following tasks: ■ Above the **File classes** list, click the radio button beside **If the**. ■ Click the attachment type drop-down list and select **True file class**. ■ In the **File classes** list, select the class or classes that you want to add to your attachment list. You can press and hold **Ctrl** to select multiple, non-consecutive items. You can also press and hold **Shift** to select consecutive blocks of items. When you select a true file class, you add every true file type that is within that class to your attachment list. You cannot add or remove any file types.

To add a true file type to your attachment list	Do all of the following tasks: ■ Above the **File classes** list, click the radio button beside **If the**. ■ Click the attachment type drop-down list and select **True file type**. **True file type** is the default option. ■ In the **File classes** list, select the true file class that contains the true file types that you want to add to your attachment list. You can press and hold **Ctrl** to select multiple, non-consecutive items. You can also press and hold **Shift** to select consecutive blocks of items. See "Predefined attachment lists" on page 429. ■ In the **File types** box, select the true file types that you want to add to your attachment list.
To add a file extension based on file names, extensions, or MIME types	Do all of the following tasks: ■ To add your own extension, under the **File classes** list, click the radio button beside **If the**. ■ Click the **Extension** drop-down list and select one of the following options: ■ File name ■ Extension This value is the default. ■ MIME-type ■ Click the **is** drop-down list and select one of the following options: ■ contains ■ begins with ■ ends with ■ is This value is the default. ■ In the adjacent text box, type only one file name, extension, or MIME type. Type the MIME type completely, such as image or image/gif.

5 Click **Add** to add the attachment type to the attachment list.

6 Repeat steps 4 and 5 to add more attachment types to the attachment list.

7 Click **Save**.

Editing attachment lists

You can edit the attachment lists that you create.

You can only edit the following predefined attachment lists:

- Archive Files
- True Type Executable Files
- Document Files
- Executable Files
- Image Files
- Multimedia Files

See "Predefined attachment lists" on page 429.

You can modify the list name, add file extensions to the list, or remove files extensions from the list.

To edit an attachment list, you must have Full Administration rights and rights to modify policies.

See "About attachment lists" on page 424.

To edit attachment lists

1. In the Control Center, click **Content > Resources > Attachment Lists**.
2. On the **Attachment Lists** page, check the box beside the attachment list that you want to modify, and then click **Edit**.

   You can also click on the name of the attachment list to edit it.
3. Do any of the following tasks:
   - Modify the attachment list name.
   - Add file classes, file extensions, or file types to the list.
     See "Creating lists to detect file attachment types" on page 425.
   - In the **Attachment Types** list, select an attachment type that you want to remove, and then click **Delete**.
4. Click **Save**.

Deleting attachment lists

You can delete the attachment lists that you create when they are no longer needed. Predefined attachment lists cannot be deleted.

To delete an attachment list, you must first delete the condition to use that attachment list from all of the content filtering policies in which it is used. Symantec Brightmail Gateway lets you know which policies use that attachment list when you attempt to delete it.

To delete an attachment list, you must have Full Administration rights or rights to modify policies.

See "About attachment lists" on page 424.

To delete attachment lists

1 In the Control Center, click **Content > Resources > Attachment Lists**.

2 On the **Attachment Lists** page, check the box beside the attachment list that you want to delete.

3 Click **Delete**.

A message appears at the top of the page to indicate the content filtering policies in which a condition to use this attachment list occurs. Delete or modify this condition in these policies first, then repeat this procedure.

See "Editing content filtering policies" on page 372.

Predefined attachment lists

Table 15-22 describes the predefined attachment lists.

Table 15-22 Predefined Attachment Lists

Attachment list	True file class	Predefined true file type
Archive Files	Encapsulation format	Apple Double, Apple Single, ASCII-armored PGP encoded, ASCII-armored PGP Public Keyring, ASCII-armored PGP signed, BinHex, Compactor / Compact Pro, cpio archive (CHR Header), cpio archive (CRC Header), Disk Doubler, GZ Compress, LHA Archive, IBM Lotus Notes Database NSF/NTF, MacBinary, Microsoft Outlook, Microsoft Outlook PST, MIME, OLE Compound Document, PAK/ARC Archive, OpenPGP Message Format (with new packet format), PGP Compressed Data, PGP Encrypted Data, PGP Public Keyring, PGP Secret Keyring, PGP Signature Certificate, PGP Signed and Encrypted Data, PGP Signed Data, RAR, Serialized Object Format (SOF), SHAR, SMTP, StuffIt (MAC), SUN PEX Binary Archive, TAR, Unix Compress, UU encoded , WANG Office GDL Header, ZIP Archive, Mac Disk Copy Disk Image File, Microsoft Entourage Database Format, 7 Zip Format(7z), Bzip 2 Compressed File, ISO-9660 CD Disc Image Format, Microsoft Cabinet File (CAB), Nero Encrypted File(.nef), Legato EMailXtender Archives Format (EMX)
Design Documents	Vector graphic Document F97, Lotus Freelance for DOS, Lotus Freelance for OS/2, Lotus Freelance for Windows; Persuasion, Microsoft PPT 2007 XML, Microsoft PPT Macro 2007 XML, Microsoft Visio, PowerPoint 95, PowerPoint 97, PowerPoint MAC, PowerPoint PC, Extensible Forms Description Language (.xfdl, .xfd), Apple iWork Keynote format	
	Scheduling/Planning Format	Microsoft Project 2000, Microsoft Project 4, Microsoft Project 4.1, Microsoft Project 98, Microsoft Project Activity, Microsoft Project Calculation, Microsoft Project Resource, PlanPerfect

Table 15-22 Predefined Attachment Lists *(continued)*

Attachment list	True file class	Predefined true file type
Executable Files	Executable File	ELF Executable MS-DOS Batch File MSDOS Device Driver MSDOS/Windows Program PC (.COM) Unix Executable (3B20) Unix Executable (Basic-16) Unix Executable (Bell 5.0) Unix Executable (iAPX 286) Unix Executable (MC680x0) Unix Executable (PDP-11/pre-System V VAX) Unix Executable (VAX) Unix Executable (WE32000) Unix Executable (x86)
Financial Information	Spreadsheet formats	Applix Spreadsheets, CSV (Comma Separated Values), Lotus 1-2-3 , Microsoft Excel, Microsoft Excel Chart, Microsoft Excel Macro, Microsoft Excel XML, Microsoft Works for DOS Spreadsheet, Microsoft Works for MAC Spreadsheet, Microsoft Works for Windows Spreadsheet, Multiplan (Mac), Multiplan (PC), Quattro Pro for DOS, Quattro Pro for Windows, SYLK
Image Files	Vector graphic Documents	Ability Word Processor AutoCAD DXF Binary AutoCAD DXF Text AutoDesk Drawing (DWG) AutoDesk WHIP AutoShade Rendering, CCITT G3 1D, Curses Screen Image, Encapsulated PostScript, Encapsulated PostScript with Preview, FPX Format, GEM Bit Image, Computer Graphics Metafile (CGM) Binary, Computer Graphics Metafile (CGM) Character, Computer Graphics Metafile (CGM) Clear Text, Corel CMX, Corel Draw, DeVice Independent file (DVI), Enhanced Metafile, Freehand MAC, GEM VDI, Graphics Interchange Format (GIF87a), Graphics Interchange Format (GIF89a), JPEG Interchange Format, Micrografx Designer Microsoft Office Drawing MicroStation V8 DGN (OLE) OS/2 PM Metafile, PostScript QuickDraw 3D Metafile Simple Vector Format (SVF) VRML Windows Draw (Micrografx) Windows Metafile (no header), CATIA Formats (CAT*), ODF Drawing, Omni Graffle (.graffle) XML File, SolidWorks Format (.sldasm, .sldprt, .slddrw)
	Raster image Documents	MacPaint, MS Windows Device Independent Bitmap, OLE DIB object, OS/2 PM Metafile, PC Paintbrush Graphics (PCX), PCD Format, Portable Bitmap Utilities ASCII Format, Portable Bitmap Utilities Binary Format, Portable Greymap Utilities ASCII Format, Portable Greymap Utilities Binary Format, Microsoft Document Imaging Format, Photoshop Document (.psd)

Table 15-22 Predefined Attachment Lists *(continued)*

Attachment list	True file class	Predefined true file type
Confidential Documents Documents Not For Distribution Internal Use Only Documents Proprietary Documents	Spreadsheet formats	Applix Spreadsheets, CSV (Comma Separated Values), Lotus 1-2-3, Lotus 1-2-3 97, Lotus 1-2-3 Release 9, Microsoft Excel, Microsoft Excel 2000, Microsoft Excel 2007, Microsoft Excel 95, Microsoft Excel 97, Microsoft Excel Chart, Microsoft Excel Macro, Microsoft Excel XML, Microsoft Works for DOS Spreadsheet, Microsoft Works for MAC Spreadsheet, Microsoft Works for Windows Spreadsheet, Multiplan (Mac), Multiplan (PC), Portable Document Format, Quattro Pro for DOS, Quattro Pro for Windows, SYLK
	Word Processing	Microsoft Word 2000, Microsoft Word 2007, Microsoft Word 95, Microsoft Word 97, Microsoft Word for Macintosh, Microsoft Word for PC, Microsoft Word for Windows, Microsoft Word UNIX, Microsoft Word XML
	Presentation	Microsoft PowerPoint 2000, PowerPoint 95, PowerPoint 97, PowerPoint MAC, PowerPoint PC
Media Files	Animation File	Macromedia Director, Macromedia Flash
	Vector graphic Document	Quickdraw 3D Metafile, VRML
	Movie File	QuickTime Movie, RIFF Multimedia Movie, Video for Windows (AVI), RealMedia Streaming Media (.rm, .ra)
	Sound File Format	Microsoft Wave, MIDI, MP3 (MPEG Audio), RealAudio, RIFF MIDI, AC3 Audio File Format (.ac3)
MSAccess files and Executables	Library	DOS/Windows Object Library
	Database Documents	Microsoft Access, Microsoft Access 2000, Microsoft Access 95 Microsoft, Access 97
	Executables	MSDOS Device Driver, PC (.COM), MSDOS/Windows Program, Unix Executable (3B20), Unix Executable (Basic-16), Unix Executable (Bell 5.0), Unix Executable (iAPX 286), Unix Executable (MC680x0), Unix Executable (PDP-11/pre-System V VAX), Unix Executable (VAX), Unix Executable (WE32000), Unix Executable (x86)
	Object Module Format	Unix Object Module (old MS 8086), Unix Object Module (VAX Demand), Unix Object Module (Z8000)

Table 15-22 Predefined Attachment Lists *(continued)*

Attachment list	True file class	Predefined true file type
Multimedia Files	Movie File	MPEG Movie, QuickTime Movie, RIFF Multimedia Movie, Video for Windows (AVI)
	Sound File Format	Amiga IFF (8SVX) Sound, Amiga MOD, Audio Interchange File Format (AIFF), Creative Voice (VOC), Microsoft Wave, MIDI, MPEG Audio, NeXT/Sun Audio Data, Real Audio, RIFF MIDI, AC3 Audio File Format(.ac3), RealMedia Streaming Media (.rm, .ra)
Network Diagrams with IP Address Keyword Network Diagrams with IP Addresses	Presentation Document	Microsoft Visio
	Word Processor Document	Microsoft Visio XML
Password Protected Files	Spreadsheet formats (password protected)	Microsoft Excel, Microsoft Excel 2000, Microsoft Excel 2007, Microsoft Excel 95, Microsoft Excel 97, Microsoft Excel Chart, Microsoft Excel XML
	Presentation Document (password protected)	Microsoft PowerPoint2000, PowerPoint 95, PowerPoint 97, PowerPoint MAC, PowerPoint PC
	Word Processor Document (password protected)	Microsoft Word 2000, Microsoft Word 2007, Microsoft Word 95, Microsoft Word 97, Microsoft Word for Macintosh, Microsoft Word for PC, Microsoft Word for Windows, Microsoft Word UNIX, Microsoft Word XML
	Encapsulation Format (password protected)	ZIP Archive
PGP Files	Encapsulation Format	ASCII-armored PGP encoded,ASCII-armored PGP Public Keyring, ASCII-armored PGP signed, PGP Compressed Data, PGP Encrypted Data, PGP Public Keyring, PGP Secret Keyring, PGP Signature Certificate, PGP Signed and Encrypted Data, PGP Signed Data
Publishing Documents	Desktop Publishing	FrameMaker, FrameMaker Book, Microsoft Publisher, PageMaker for Macintosh, PageMaker for Windows, Quark Xpress MAC
Resumes, All Resumes, Employee	Word Processor Document	Microsoft Word 2000, Microsoft Word 2007, Microsoft Word 95, Microsoft Word 97, Microsoft Word for Macintosh, Microsoft Word for PC, Microsoft Word for Windows, Microsoft Word UNIX, Microsoft Word XML

Table 15-22 Predefined Attachment Lists *(continued)*

Attachment list	True file class	Predefined true file type
SEC Fair Disclosure Regulation	Spreadsheet format	Applix Spreadsheets, CSV (Comma Separated Values), Lotus 1-2-3, Lotus 1-2-3 97, Lotus 1-2-3 Release 9, Microsoft Excel, Microsoft Excel 2000, Microsoft Excel 2007, Microsoft Excel 95, Microsoft Excel 97, Microsoft Excel Chart, Microsoft Excel Macro, Microsoft Excel XML, Microsoft Works for DOS Spreadsheet, Microsoft Works for MAC Spreadsheet, Microsoft Works for Windows Spreadsheet, Multiplan (Mac), Multiplan (PC), Quattro Pro for DOS, Quattro Pro for Windows, SYLK
	Word Processor Document	Microsoft Word 2000, Microsoft Word 2007, Microsoft Word 95, Microsoft Word 97, Microsoft Word for Macintosh, Microsoft Word for PC, Microsoft Word for Windows, Microsoft Word UNIX, Microsoft Word XML, Portable Document Format, WordPerfect, WordPerfect MAC, WordPerfect VAX

There are several true file classes that are not referenced by any Attachment List. These file classes and their associated true file types, however, can be accessed from within the File classes and File types lists the same way that other file classes and types can be selected. You can add any of these file types to an existing Attachment List or create a new Attachment List.

Table 15-23 Additional file classes and types

True file class	True file types
Communications Format	Ability Communications, FTP Session Data, Microsoft Works for MAC Communications, SmartWare II Communications
FAX Format	DCX FAX Format(PCX images)
Font Type Document	NeWS bitmap font, SUN vfont, Definition TrueType Font
General Purpose Document	Microsoft Office 2007, Program Information File (PIF), Windows Group, Windows Help File, WordPerfect auxiliary file
Mixed Type Document	Framework, Framework II, Windows C++ Object Storage
Planning/Outline Format	MORE Database MAC

Perl-compatible regular expressions

This section describes the Perl-compatible regular expressions that can be used in content filtering policies.

See "Expressions for content filtering policy conditions" on page 362.

See "Perl-compatible regular expression examples" on page 436.

Table 15-24 Regular expression metacharacters

Metacharacter/ construct	Description
.	Period: Matches any single character of the input sequence.
^	Circumflex: Represents the beginning of the input line. For example, ^A is a regular expression that matches the letter A at the beginning of a line. The ^ character is the only special character allowed at the beginning of a regular expression or after the (or | characters.
$	Dollar sign: Represents the end of the input line. For example, A$ is a regular expression that matches the letter A at the end of a line. The $character is the only special character allowed at the end of a regular expression or before the) or | characters. Note that this metacharacter cannot be used for line-by-line matching in a message body.
*	Asterisk: Matches zero or more instances of the string to the immediate left of the asterisk. For example, A* matches A, AA, AAA, and so on. It also matches the null string (zero occurrences of A).
?	Question mark: Matches zero or one instance of the string to the immediate left of the question mark.
+	Plus sign: Matches one or more instances of the string to the immediate left of the plus sign.
\	Escape: Turns on or off the special meaning of metacharacters. For example, \. only matches a dot character. \$ matches a literal dollar sign character. Note that \\ matches a literal \ character.
|	Pipe: Matches either expression on either side of the pipe. For example, exe|com|zip matches exe, com, or zip.

Table 15-24 Regular expression metacharacters *(continued)*

Metacharacter/ construct	Description
[string]	Brackets: Inside the brackets, matches a single character or collating element, as in a list. Characters within brackets are not case sensitive. The string inside the brackets is evaluated literally, as if an escape character (\) were placed before each character in the string. If the initial character in the bracket is a circumflex (^), then the expression matches any character or collating element except those inside the bracket expression. If the first character after any potential circumflex (^) is a dash (-) or a closing bracket (]), then that character matches only a literal dash or closing bracket.
(string) \(string\)	Parentheses: Groups parts of regular expressions, which gives the string inside the parentheses precedence over the rest.

Perl-compatible regular expression examples

Table 15-25 gives some examples of Perl-compatible expressions.

See "Perl-compatible regular expressions" on page 435.

See "Expressions for content filtering policy conditions" on page 362.

Table 15-25 Sample Perl-compatible regular expressions

Character	Description	Example	Sample matches
.	Match any one character	j.n	jen, jon, j2n, j$n
..	Match any two characters	jo..	john, josh, jo4#
.*	Match zero or more characters	sara.*	sara, sarah, sarahjane, saraabc%123
.*		s.*m.*	sm, sam, simone, s321m$xyz
.+	Match one or more characters	sara.+	sarah, sarahjane, saraabc%123
.+		s.+m.+	simone, s321m$xyz
\.	Match a period	stop\.	stop.

Table 15-25 Sample Perl-compatible regular expressions *(continued)*

Character	Description	Example	Sample matches
*	Match an asterisk	b**	b**
\+	Match a plus character	18\+	18+
\/	Match a forward slash	18\/	18/
[0-9]{n}	Match any numeral n times, for example, match a social security number	[0-9]{3}-[0-9]{2}-[0-9]{4}	123-45-6789

Specifying where to save archived messages

Your organization can maintain copies of the messages that violate specific policies. To use the archive feature, you must specify the action **Archive the message** when you create a content filtering policy.

You can also configure your archive email destination on a per-policy basis.

See "Content filtering policy actions" on page 365.

This feature can be useful to your organization in any of the following scenarios, just to name a few:

- You want to prepare for a potential lawsuit against an employee or other organization
- You must maintain records for regulatory compliance purposes
- You want to retain records for certain groups of employees

When messages trigger a policy violation with the action to **Archive the message**, Symantec Brightmail Gateway copies the messages. It sends the copy to one or both an email address that you specify or an archive server.

To archive messages, you must have Full Administration rights or rights to modify policies.

To specify where to save archived messages

1. In the Control Center, click **Content > Settings > Archive**.
2. In the **Archive email address** box, type a complete email address, such as kyi@symantecexample.com.
3. In the **Archive server host** field, type the name of the archive server host.

   This server host is the host name or IP address for the archive email address that you provided in step 2.

4 If you provided an **Archive server host**, in the **Archive server port**, type the server host's port number.

5 If you want to route archive messages with MX Lookup to locate the information that corresponds to the archive server host, check **Enable MX Lookup**.

6 To make the archive server information available to all of your existing policies, click **Apply to all current policies**.

7 Click **Save**.

About content incident folders

Content incident folders help you organize, monitor, and manage the incidents of content filtering policy violations. As a best practice, you should create a content incident folder for each type of content filtering policy that you use. For example, if you use the HIPAA template to create a policy, create a HIPAA content incident folder. You can use that folder to monitor HIPAA policy incidents.

See "About content filtering policy templates" on page 336.

Create the content incident folder before you configure a content filtering policy. That way, when you create a content filtering policy, the folder is an available option when you select the action to create an incident. If you do not specify a folder in which to create an incident, messages that violate that policy are filed in the **Informational Incidents** folder.

See "Creating content filtering policies" on page 334.

When you create a content filtering folder, you must choose the type of folder you want to use.

Symantec Brightmail Gateway provides the following types of content incident folders:

Hold for Review (Content Quarantine)	Use this type of folder for the incidents that you want to review. This folder lets you retain the messages that trigger content filtering violations so that you can review them and determine how to act on them. Any additional actions for that policy are deferred until the incident is reviewed. **Note:** Messages in the **Hold for Review (Content Quarantine)** folder are expunged based on the settings that you specify, even if they have not been reviewed. See "About managing the size of content incident folders" on page 439.

Informational Incidents	Use this type of folder to track the incidents that are at a lower priority than the ones that you want to hold for review.

You can also configure the following settings for each folder:

- Archive tag and encoding for the tag
  This setting is optional.
- Expunger settings
  The maximum content incident folder size is required. The setting for the number of days to hold an incident before it is expunged is optional.
- Notification message that indicates an incident has been logged to that folder and the people that you want to receive the notification
  This setting is optional.

Symantec Brightmail Gateway supports up to 1,000 content incident folders, including the default Informational Incidents and Quarantine Incidents folders. You must have Full Administration rights or rights to modify settings to create content incident folders.

See "Creating content incident folders" on page 441.

About managing the size of content incident folders

Your content incident folders can fill up quickly and consume a lot of disk space when one or both of the following conditions exist:

- You have a large number of content filtering policies whose actions are to create incidents in content incident folders.
- You have a large number of messages that violate policies.

Symantec Brightmail Gateway provides an Expunger that can manage the size of your content incident folders. The Expunger automatically runs at the frequency that you specify.

When the Expunger runs, it expunges incidents from your content incident folders based on the following criteria:

By the maximum content incident folder size	The maximum content incident folder size is the sum of the block size (the size on disk) of each message file in the folder. The actual disk usage may be higher. When a folder reaches its maximum size, a message is sent to the BrightmailLog.log. When the folder reaches the maximum size that you specify for alerts (for example, 120% of the maximum size), then an alert is sent. Only one alert is sent if multiple folders reach their maximum size at the same time. However, you continue to receive alert emails to notify you of any subsequent folders that exceed their maximum size. The oldest incidents are deleted until the folder returns to its maximum size. See "Configuring alerts" on page 614. Symantec recommends that you expunge folders based on size while you fine-tune content filtering policies. Then you can modify the Expunger settings as desired. You must specify a maximum content incident folder size. The default maximum size is 5 GB, but you can modify this setting.
By the number of days that you want to store incidents	Optionally, you can specify the number of days to store an incident before the Expunger performs one of the following actions: ■ Approve This option is only available for Hold for Review (Content Quarantine) folders. After the Expunger approves the incident, it is deleted from the folder. ■ Reject This option is only available for Hold for Review (Content Quarantine) folders. After the Expunger rejects the incident, it is deleted from the folder. ■ Delete This option is available for both Hold for Review (Content Quarantine) and Informational Incident folders. **Note:** When this threshold is met, incidents in Held for Review (Content Quarantine) folders are expunged, even if they have not been reviewed. Symantec Brightmail Gateway takes the action that you specify in the content filtering policy, starting with the oldest incidents. Symantec Brightmail Gateway determines the age of an email based on the time that the message is received.

You configure your Expunger criteria when you create the content incident folder. You can specify one or both the folder size thresholds and days to store thresholds. You can also modify the settings at any time thereafter. When a threshold is met, no further incidents are created in the folder.

See "Creating content incident folders" on page 441.

The Expunger cycle applies to all content incident folders.

See "Scheduling the content incident folder Expunger" on page 444.

See "About content incident folders" on page 438.

Creating content incident folders

Create content incident folders to help you organize, monitor, and manage the incidents that content policy violations generate. You can have as many as 1,000 folders including the default Informational Incidents and Quarantine Incidents folders.

To create a content incident folder, you must have Full Administration rights or rights to modify policies.

See "About content incident folders" on page 438.

See "Editing content incident folders" on page 442.

See "Scheduling the content incident folder Expunger" on page 444.

See "Deleting content incident folders" on page 443.

To create content incident folders

1. In the Control Center, click **Content > Settings > Content Incident Folders**.
2. Click **Add**.
3. In the **Content Incident folder name** field, type a name for the content incident folder.

   As a best practice, you should use a name that reflects the type of content that you intend to collect in that folder. For example, if you want to place the messages that violate a HIPPA policy in this folder, you might name the folder HIPPA.
4. Click the **Content incident folder** type drop-down list to select the type of folder you want to use.
5. Optionally, in the **Optional archive tag** field, type the text that you want to identify this folder with for archival purposes.

   When you specify an archive action on an incident in this content incident folder, that text accompanies the incident.

6 If you specified an incident archive tag, in the **Encoding** drop-down list, choose the character encoding set to use for the tag text.

7 Under **Expunger Settings** do the following tasks:

To specify the number of days to retain incidents before they are expunged (optional)	Do all of the following tasks: ■ Check **Days to store before default action occurs**, and in the field beside it, type the number of days. ■ Under **Default expunger action**, specify the action that you want to take when an incident is expunged. If the folder is an **Informational Incidents** folder, the only option that is available is **Delete**.
To specify a maximum folder size (required)	Under **Thresholds**, specify the size in kilobytes, megabytes, or gigabytes.

See "About managing the size of content incident folders" on page 439.

8 Optionally, click the **Notification format** drop-down list and select the format for incident notifications.

This notification is sent to whomever you specify in step 10 and step 11 that indicates that an incident is added to this folder.

See "Creating incident notifications" on page 445.

9 Click **Edit** to view or edit the incident notification template.

10 In the **Notification recipient addresses** field, type the email addresses of the content filtering officers who should be notified of incidents in this folder. If you do not want to have notifications sent, leave this field blank.

Separate multiple addresses with commas, semicolons, or spaces.

11 In the **Administrator Notifications** list, check the names of administrators whom you want to be notified of incidents in this folder. If you do not want to have notifications sent, leave this field blank.

You can specify administrators to notify of incidents in this folder when you add administrators or modify administrator settings.

12 Click **Save**.

Editing content incident folders

You can edit any content incident folder as needed. To edit a content filtering folder, you must have Full Administration rights or rights to modify policies.

See "About content incident folders" on page 438.

See “Deleting content incident folders” on page 443.

See “Creating content incident folders” on page 441.

To edit content incident folders

1. In the Control Center, click **Content > Settings > Content Incident Folders**.
2. Check the box beside the content incident folder that you want to edit, and then click **Edit**.

   You can also click on the content incident folder name to edit it.
3. Edit the content incident folder settings as necessary.
4. Click **Save**.

Deleting content incident folders

You can delete a content incident folder that you create when you no longer need it. However, you cannot delete the default Informational Incidents folder or the Quarantine Incidents folder.

Before you delete a folder, you may want to ensure that all of the items in the folder have been reviewed and addressed as needed. You must also delete the action to log incidents to the folder from any content filtering policies in which it is used. Symantec Brightmail Gateway lets you know which policies use that folder in an action when you attempt to delete it.

To delete a content incident folder, you must have Full Administration rights or rights to modify policies.

See “About content incident folders” on page 438.

See “Editing content incident folders” on page 442.

See “Scheduling the content incident folder Expunger” on page 444.

To delete content incident folders

1. In the Control Center, click **Content > Settings > Content Incident Folders**.
2. Check the box beside the content incident folder that you want to delete.

   Check the box beside **Content Incident Folders** to select all of the folders.
3. Click **Delete**.
4. In the confirmation dialog, click **Delete**.

   A message appears at the top of the page to indicate the content filtering policies in which an action to log incidents to this folder occurs. Delete or modify this action in these policies first, then repeat this procedure.

Scheduling the content incident folder Expunger

Symantec Brightmail Gateway provides an Expunger that manages the size of your content incident folders. You can specify the frequency at which you want the Expunger to run. The default frequency is every day at 4:00 A.M. When the Expunger runs, it takes the default actions that you specify in your content filtering policies. Then it deletes the incidents based on the threshold criteria: folder size or days to store in the folder.

See "Creating content incident folders" on page 441.

See "About managing the size of content incident folders" on page 439.

See "About content incident folders" on page 438.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See "About scheduled tasks" on page 621.

You must have Full Administration or rights to modify settings to schedule the content incident folder Expunger.

See "Administrator rights" on page 684.

To schedule the content incident folder Expunger

1. In the Control Center, click **Content > Settings > Content Incident Folders**.
2. Under **Content Filtering Expunging Cycle**, in the **Incident expunger frequency** drop-down list, select how often you want the Expunger to run.
3. In the **Incident expunger start time** drop-down lists, specify the hour and minute in which you want the Expunger to run.
4. Click **Save**.

About monitoring and acting on incidents

Content incident folders contain incidents of the messages that violate a policy's conditions. Information about incidents can help you understand, prevent, respond to, and audit potential violations. For example, you can use an incident folder to monitor content filtering policy violations at your company before you adopt permanent policies.

See "About content incident folders" on page 438.

Content incident folders store the incidents that let you do the following:

Monitor incidents	You can monitor the incidents that occur, but you do not have to act on them. You must create a content filtering policy in which at least one action is to create an incident in an informational incidents folder. If a policy violation occurs, an incident is created in that informational incidents folder for you to review at your convenience.
Hold incidents for review	You can hold messages in a Content Quarantine for you to review and determine the action that you want to take. Messages in the content quarantine are not acted upon until the content filtering officer approves, rejects, or deletes them or they are expunged. You must create a policy in which at least one action is to create an incident in the content quarantine incidents folder that you specify. If a policy violation occurs, an incident is created in that folder for your review and action.

See “Viewing incidents in content incident folders” on page 447.

See “Specifying content filtering policy actions” on page 365.

See “Content incident actions” on page 449.

Messages that are held in content incident folders remain there until they are acted upon (for example, deleted or forwarded) or until they are expunged.

See “About managing the size of content incident folders” on page 439.

Creating incident notifications

Symantec Brightmail Gateway can notify the persons that you specify that an incident has been added to a content incident folder. Symantec Brightmail Gateway provides a default notification message that you can modify as necessary.

In the notification message text, Symantec Brightmail Gateway replaces the variable *%NEW_COMPLIANCE_MESSAGES%* with a list of incident numbers and the policies that triggered those incidents. You can type text before or after this variable, or you can delete the variable.

For example, you can type the following text to precede the variable and replace *Name* with the folder name:

```
A new incident has been created. Please access the Name Content
Incident Folder for incident details.
```

You can have one notification message for each incident folder. That message can be unique or it can be the same message that you use for a different incident folder.

Note: This notification differs from the Notifications resource. The Notifications resource lets you send notification messages to the message sender, message recipient, or a third party to indicate that a policy is violated. The Notifications resource is the notification that is sent when the policy action is to **Send notification**.

See "About policy violation notifications" on page 400.

You must have Full Administration or rights to modify settings to create incident notifications.

See "Administrator rights" on page 684.

To create incident notifications

1. In the Control Center, click **Content > Settings > Content Incident Folders**.
2. In the **Content Incident Folder** list, click on the name of a content incident folder for which you want to create an incident notification message.

   Alternatively, check the box beside the name of the content incident folder, and click **Edit**.
3. On the **Content Incident Folder Settings** page under **Notification Settings**, click the **Notification format** drop-down list and select the format that you want to use the notification message.

   You can select multi-part (HTML and text), HTML only, or text only.
4. If you want to modify the notification message, click **Edit** beside **Notification template**.
5. On the **Notification template** page, under **Encoding Type**, click the drop-down list and select the character encoding for the notification message.
6. In the **Send from** box, type the address of the `From` header to appear in incident notification messages.
7. In the **Subject** box, type the `Subject` header that you want to appear on the incident notification messages.
8. In the **Notification** box, modify the text as desired.
9. Click **Save**.

10 To specify who should receive this notification, do any of the following tasks:

To notify an administrator	Under **Administrator Notifications**, check the box beside the administrators that you want to receive this notification.
To notify people who are not administrators	In the **Notification recipient addresses** box, type the email addresses of anyone that you want to receive the notification message. Separate multiple entries with commas.

11 Click **Save**.

Viewing incidents in content incident folders

You can do any of the following tasks on the **Incident Management** page to view your incidents:

- Use a filter to narrow or expand your search of incidents.
- View details about a particular incident from this page.
- Sort columns by ascending order or descending order.

Use the **Folder Overview** folder for a high-level summary all of your content incident folders or a single folder at a time.

Symantec Brightmail Gateway provides the following tools on the **Incident Management** page to help you customize your view and navigate through your incidents:

Entries per page	This setting lets you specify how many entries you want to display per page. You can specify 10, 25, 50, 100, 200, 500, or all.
Display	This setting lets you specify which range of entries you want to display.
Forward and back arrows	Use the back <| and forward |> arrows to move to the very first page or the very last page. Use the back < arrows and forward > arrows to move forward or back one page at a time.

Incidents remain in content incident folders until you act on them (for example, approve, delete, or reject) or the incidents are expunged.

See "About monitoring and acting on incidents" on page 444.

See “About managing the size of content incident folders” on page 439.

See “Content incident actions” on page 449.

To view the contents of incident folders, you must have Full Administration rights or rights to access each content incident folder separately.

To view a summary of what is in your content incident folders

1. In the Control Center, click **Content > Folder Overview**.

   By default, all of your folders appear in the **Content Incident** table.

2. In the **Content incident folder** drop-down list, select the folder for which you want to see an overview.

3. Click **Display**.

   The status appears in the **Content Incident Folders** table.

To view incidents in content incident folders

1. In the Control Center, click **Content**.

2. In the **Incident Management** task pane, select the folder that contains the incidents that you want to view.

Acting on multiple incidents at a time

You can perform the same action on multiple incidents at a time. For example, you might want to reject all of the incidents that are from a specific sender.

See “Content incident actions” on page 449.

To act on incidents in incident folders, you must have Full Administration rights or rights to access each content incident folder separately.

To act on multiple incidents at a time

1. In the Control Center, click **Content**.

2. In the **Incident Management** task pane, select the folder that contains the incidents that you want to act on.

3. Check the box beside the incident ID for all of the incidents that you want to act on.

   Check the box beside the Incident ID column heading to select all of the incidents that appear on the page.

4 Click the option for the task that you want to take on all of the incidents that you selected.

See "About monitoring and acting on incidents" on page 444.

5 Take the necessary steps to complete the action.

For example, if you want to exports incidents, you must specify the location to where you want the files exported.

Content incident actions

Table 15-26 lists the actions that you can take on incidents in content incident folders.

Table 15-26 Content incident actions

Action	Description
Archive	Archives incidents. You can archive a single incident at a time or multiple incidents at once. See "Archiving incidents" on page 451.
Export Incident History	Exports an incident's history to the location that you specify. You can export a single incident's history at a time or multiple incident histories at once. See "Exporting an incident's history" on page 452.
Forward an incident	Forwards an incident to the email address that you specify. You can add a subject to your email and a comment. You can also select the message encoding and specify whether you want to forward your message with the original message. You can forward a single incident at a time or multiple incidents at once. See "Forwarding incidents" on page 452.
Delete an incident	Deletes a single incident at a time or multiple incidents at once. See "Deleting incidents" on page 454.

Table 15-26 Content incident actions *(continued)*

Action	Description
Approve, reject, or hold an incident	Specifies how to act on a quarantined incident. When you approve or reject an incident in a Quarantine Incidents folder, Symantec Brightmail Gateway takes the action that you specified in the content filtering policy actions **Message Review Approved Actions** and **Message Review Rejected Actions**. You can approve or reject a single incident at a time or multiple incidents at once. You can specify to hold only a single incident at a time. See "Approving, rejecting, or holding quarantined incidents" on page 455.
Update incident status	Specifies the status of an incident. The status options are as follows: ■ New ■ Active ■ Confirmed ■ False positive You can only change the status of a single incident at a time. See "Updating an incident's status" on page 456.
Modify an incident's severity	Specifies the severity of an incident. The severity levels are as follows: ■ Unknown ■ High ■ Medium ■ Low You can only change the severity of a single incident at a time. See "Changing an incident's severity level" on page 457.

Table 15-26 Content incident actions *(continued)*

Action	Description
Review an incident's history	Shows the entire history of an incident. The history can include when the incident is created, when a status or severity is changed, or when an action is taken. The incident history contains the date and time that the event occurred and any comments that you create in connection with the event. You can only view the history of one incident at a time. See "Viewing an incident's history" on page 457.

See "About monitoring and acting on incidents" on page 444.

See "About content incident folders" on page 438.

Archiving incidents

After you review incidents and take the action that you want, you can archive the message. You can archive one or more messages from the Incident Manager page or one at a time from an incident's detail page.

See "About monitoring and acting on incidents" on page 444.

See "Specifying where to save archived messages" on page 437.

To archive incidents, you must have Full Administration rights or rights to view the specific content incident folder.

To archive incidents from the Incident Management page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident that you want to review and archive.
3. Check the box beside one or more incidents that you want to archive.
4. Click **Archive**.

To archive incidents from the Incident Management details page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident that you want to review and archive.

3. In the **Incident ID** column, click on the incident number that you want to review and archive.
4. On the **Incident Management** details page, in the right pane under **Incident Actions**, click **Archive**.

Exporting an incident's history

You can export all of the information about an incident's history. For example, when the incident is created, when a status changed, message review actions, and comments. You can select the encoding that you want to use for the exportation and the delimiter.

You can export one or more incident histories at a time.

To export an incident's history, you must have Full Administration rights or rights to view the specific content incident folder.

See "About monitoring and acting on incidents" on page 444.

See "Viewing an incident's history" on page 457.

To export an incident's history

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident history that you want to export.
3. Check the box beside one or more incidents whose histories you want to export.
4. Click the **Encoding** drop-down list to select the encoding.
5. Click the **Delimiter** drop-down list to select the delimiter that you want to use.
6. Click **Export Incident History** to export an incident history.
7. Select whether you want to save the file or open it. If you save the file, browse and select the desired location.

Forwarding incidents

You can forward incidents in incident folders for further review. When you forward an incident, you can create a comment for the person to whom you forward the message to provide instructions or deadlines. You can also specify whether you want to attach the original message.

You can forward one or more messages from the **Incident Management** page or one incident at a time from an incident's detail page.

To forward incidents, you must have Full Administration rights or rights to view the specific content incident folder.

See “About monitoring and acting on incidents” on page 444.

To forward incidents from the Incident Management page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident that you want to review and forward.
3. Check the box beside one or more incidents that you want to forward.
4. Click **Forward Incident**.
5. On the **Forwarding Message** page, under **Encoding**, click the drop-down list to select the encoding type.
6. Under **Message Content**, in the **Forward to** text box, type the email address.
7. In the **Subject** text box, type a subject name.
8. Check **Forward with original message** if you want to forward the message without any modifications.

   This action forwards the message in its original state with no policy-based actions that can affect the message (for example, removing or cleaning attachments).
9. In the **Comments** text box, type a message that you want to accompany the incident email.
10. Click **Send**.

To forward incidents from the Incident Management details page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident that you want to review and forward.
3. In the **Incident ID** column, click on the incident number that you want to review and forward.
4. On the **Incident Management** details page, in the right pane under **Incident Actions**, click **Forward**.
5. On the **Forwarding Message** page, under **Encoding**, click the drop-down list to select the encoding type.
6. Under **Message Content**, in the **Forward to** text box, type the email address.
7. In the **Subject** text box, type a subject name.

8 Check **Forward with original message** if you want to forward the message without any modifications.

This action forwards the message in its original state with no policy-based actions that can affect the message (for example, removing or cleaning attachments).

9 In the **Comments** text box, type the message that you want to accompany the email.

10 Click **Send**.

Deleting incidents

Symantec recommends that after you review and act on incidents, you delete them to reduce the volume of incidents in your incidents folders. When an incident is deleted, it no longer appears in the incident folder, and the incident cannot be retrieved.

Incidents that meet expunger thresholds are automatically deleted.

See "About managing the size of content incident folders" on page 439.

You can delete one or more messages from the Incident Manager page or one incident at a time from an incident's detail page.

To delete an incident, you must have Full Administration rights or rights to view the specific content incident folder.

See "About monitoring and acting on incidents" on page 444.

To delete incidents from the Incident Management page

1 In the Control Center, click **Content**.

2 Under **Incident Management**, select the folder that contains the incident that you want to review and delete.

3 Check the box beside one or more incidents that you want to delete.

4 Click **Delete**.

Optionally, you can click **Delete All** to delete all of the incidents on the page, and in the confirmation dialog box, click **OK**.

To delete incidents from the Incident Management details page

1 In the Control Center, click **Content**.

2 Under **Incident Management**, select the folder that contains the incident that you want to review and delete.

3. In the **Incident ID** column, click on the incident number that you want to review and delete.
4. On the **Incident Management** details page, in the right pane under **Incident Actions**, click **Delete**.

Approving, rejecting, or holding quarantined incidents

After you review an incident in a Quarantine Incidents folder, you can approve it, reject it, or continue to hold it.

When you approve or reject an incident in a Quarantine Incidents folder, Symantec Brightmail Gateway takes the action that you specified in the content filtering policy actions for **Message Review Approved Actions** and **Message Review Rejected Actions**.

See "Specifying content filtering policy actions" on page 365.

For example, assume that you create a content filtering policy whose actions is to create an incident in the **Quarantine Incidents** folder. The **Message Review Approved Action** is deliver the message. The **Message Review Rejected Action** is to delete the message. Assume that this policy is violated, and an incident is created in the **Quarantine Incidents** folder. If you approve the incident, the message is delivered. If you reject the incident, the message is deleted.

You can approve or reject one or more messages from the **Incident Management** page. You can approve, reject, and continue to hold one at a time from an incident's detail page.

See "About monitoring and acting on incidents" on page 444.

To approve or reject incidents, you must have Full Administration rights or rights to view the specific content incident folder.

To approve or reject incidents from the Incident Management page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the quarantine incident folder that contains the incident that you want to review and approve or reject.
3. Check the box beside one or more incidents that you want to approve or reject.
4. Click **Approve** or **Reject**.

To approve, reject, or hold incidents from the Incident Management details page

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the quarantine incident folder that contains the incident that you want to review, approve, reject, or continue to hold.

3. In the **Incident ID** column, click on the incident number that you want to review, approve, reject, or continue to hold.
4. On the **Incident Management** details page, in the right pane, click the **Message review action** drop-down list and select the action that you want to take.
5. Optionally, in the **Comment** box, type a comment about the action.
6. Click **Update**.

Updating an incident's status

You can update an incident's status to reflect changes to the incident. The incident's status lets reviewers know what stage of the review process the incident is currently in and what the results of the review are.

The status options are as follows:

New	The incident occurred, but the status has not yet been updated. This setting is the default status that is assigned to the incident when it is initially created.
Active	The incident is under review.
Confirmed	The incident is valid.
False positive	The incident is not valid.

You can only change the status of a single incident at a time.

See "About monitoring and acting on incidents" on page 444.

To update an incident's status, you must have Full Administration rights or rights to view the specific content incident folder.

To update an incident's status

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the quarantine incident folder that contains the incident that you want to review, approve, reject, or continue to hold.
3. In the **Incident ID** column, click on the incident number that you want to review, approve, reject, or continue to hold.
4. On the **Incident Management** details page, in the right pane, click the **Update status to** drop-down list and select the new status.

5 Optionally, in the **Comment** box, type a comment.

6 Click **Update**.

Changing an incident's severity level

As you review an incident, you can modify its severity level.

The severity levels are as follows:

- Unknown
  This setting is the default severity when an incident is created.
- High
- Medium
- Low

You can only change the severity of a single incident at a time.

See "About monitoring and acting on incidents" on page 444.

To update an incident's severity, you must have Full Administration rights or rights to view the specific content incident folder.

To update an incident's severity

1 In the Control Center, click **Content**.

2 Under **Incident Management**, select the folder that contains the incident whose severity you want to change.

3 In the **Incident ID** column, click on the incident number whose severity you want to change.

4 On the **Incident Management** details page, in the right pane, click the **Update severity to** drop-down list and select the new status.

5 Optionally, in the **Comment** box, type a comment.

6 Click **Update**.

Viewing an incident's history

You can view the history of an incident to see how the incident has progressed.

You can only view the history of a single incident at a time.

See "About monitoring and acting on incidents" on page 444.

To view an incident's history, you must have Full Administration rights or rights to view the specific content incident folder.

See "Exporting an incident's history" on page 452.

To view an incident's history

1. In the Control Center, click **Content**.
2. Under **Incident Management**, select the folder that contains the incident history that you want to review.
3. In the **Incident ID** column, click on the incident number whose history you want to review.

   The incident history appears on the **Incident Management** details page under **Incident History**.

About encrypting messages with Symantec Content Encryption

Symantec content encryption uses Symantec Hosted Services, powered by MessageLabs, to provide you the ability to encrypt outbound messages for greater security and to track statistics for those messages via the Control Center.

To encrypt messages, you must purchase the Symantec Content Encryption license, configure your system for encryption, and provision an encryption account. You then create and assign policies that encrypt outbound messages.

To encrypt outgoing messages, you must complete the following tasks:

- Obtain a Symantec Content Encryption license.
  You can learn more about the license by clicking **Administration > Hosts > Licenses** and then clicking on the **Click here** link.
  See “Licensing your product” on page 678.
- Configure your system for content encryption functionality by procuring and installing an SSL certificate and configuring your hosts for content encryption.
  See “Preparing your system for content encryption” on page 459.
- Provision your encryption account by providing system information to Symantec.
  See “ Provisioning a content encryption account” on page 461.
- Create a content policy with the action Deliver message with content encryption and select the groups to which the encryption policy should be applied. See “Creating content filtering policies” on page 334.

Once you begin processing encrypted mail messages, you can track message statistics in the Status dashboard of the Control Center and view message logs in the Message Audit Log reports.

See “Searching for a message in the Message Audit Log” on page 652.

See “Viewing the Dashboard” on page 605.

Preparing your system for content encryption

Once you have obtained a content encryption license, you must provision your account through Symantec and then also configure your system for content encryption. To configure your system for content encryption, you must complete the tasks described in this topic before or after your begin the provisioning process with Symantec. These tasks must be completed before you use your new encryption account.

See “About encrypting messages with Symantec Content Encryption” on page 458.

Preparing your certificates

1. Secure a domain, if necessary, and purchase a certificate for your domain

   You must use a Certificate Authority (CA) that is listed in your provisioning form provided by Symantec. If you already have a certificate, but it is not on the list of approved CAs, you will need to replace your certificate with one from a CA in this list.

2. In Symantec Brightmail Gateway, click **Administration > Settings > Certificates** and click the **TLS & HTTP Certificates** tab to request and create a new certificate.

   See “Requesting a Certificate Authority-signed certificate” on page 194.

3. Import the certificate issued by your CA. The certificate's Subject Common Name must match the fully-qualified domain name of the desired Symantec Brightmail Gateway host. For example, if “update5.brightmailtest.info” is the name of your Symantec Brightmail Gateway host used to accept mail, then the Subject Common Name “update5.brightmailtest.info” must be used by the certificate and by SMTP clients connecting to that host name.

   See “Importing a Certificate Authority-signed certificate” on page 197.

4. Add the CA certificate or the CA certificate bundle to your system.

   See “Adding a CA certificate” on page 193.

Configuring your hosts for content encryption

1. Click **Administration > Hosts > Configuration** to configure your hosts for encryption.

2. Select the first host and click **Edit** then click the **SMTP** tab.

3. Under **Mail Filtering**, make sure that you have selected either **Inbound mail filtering only** or **Inbound and outbound mail filtering**.

4 Under **Inbound Mail Settings**, check **Accept TLS encryption** and select the new certificate.

5 Optionally, you can check **Request client certificate**.

This option adds information about the client certificate verification transaction into the Received header. This information is potentially useful for building policies.

6 If your host uses only one IP address, save your changes and repeat steps 2 through 5 to enable TLS for all of your hosts. If your host uses more than one IP address, go to step 7.

7 Click **Advanced Settings** at the bottom of the page to display the SMTP configuration page.

8 On the **Delivery** tab, under **SMTP Delivery Bindings**, set **Dynamically routed messages** to **Auto**.

9 Click **Continue** to return to the **Edit Host Configuration** page.

10 Click **Save** to save your changes for this host.

Use this procedure to configure all of your hosts for TLS encryption.

Managing the host and port information for content encryption

The host information for your provisioned encryption account is determined by your Symantec provisioning representative, who will determine the host and port that is most appropriate for your system.

This section describes how to manage host and port information for an encryption account. To provision an encryption account you must obtain a license, configure your system for encryption, work with your provisioning representative to provision the account and then create policy groups that direct the system to encrypt outbound messages.

See "About encrypting messages with Symantec Content Encryption" on page 458.

Editing host and port information for content encryption

1 To view your Content Encryption host and port information, click **Content > Settings > Content Encryption**, then click **Show Advanced**.

2 To change the host information, provide the new MessageLabs encryption host name in the **Host** field.

3 To change the port information, provide the new port in the **Port** field.

4 Click **Save** to save your new host and port information.

Provisioning a content encryption account

To apply content encryption to your messages, you must work with Symantec to provision an encryption account.

This section describes how to provision an existing license. If you have not acquired a Symantec Content Encryption license, you can learn more about the service by clicking **Administration > Hosts > Licenses**. In addition to provisioning your account, you must also configure your system for encryption and create policy groups that direct the system to encrypt outbound messages.

See "About encrypting messages with Symantec Content Encryption" on page 458.

To provision your content encryption license

1 Select **Content > Settings > Content Encryption**.

- If you have not yet purchased a Symantec Content Encryption license, you can learn more about the service via the Content Encryption page. Click **click here.** Once you have purchased a license, go to step 2.
- If you have not already provisioned your encryption account, the **Content Encryption** page provides a link to facilitate the provisioning process. Go to step 2.
- If you have already provisioned your encryption account, the **Content Encryption** page displays your licensing information and provides your host and port setting defaults. If you want to edit your connection information at this time, click **Show Advanced**.

  See "Managing the host and port information for content encryption" on page 460.

2 In the Content Encryption page, click **click here**.

A page provides a downloadable provisioning form along with instructions for provisioning your content encryption account.

3 Complete the Content Encryption Provisioning form as directed. For encryption to be fully effective, you must include information for all domains and Scanners in the provisioning form.

If you add a new domain or Scanner at a later date, you must inform your Symantec provisioning representative. For more information, see the Symantec Content Encryption Provisioning page by clicking **Content > Settings > Content Encryption** and then clicking the **click here** link.

New hosts must also be configured for encryption.

See "Preparing your system for content encryption" on page 459.

4 When you are finished, email the provisioning form to the provided email address.

Once you have mailed your completed form, a Symantec employee on the MessageLabs team will contact you to facilitate the provisioning of your account and will notify you when your account is active.

See "Managing the host and port information for content encryption" on page 460.

About Symantec Network Prevent

Symantec Brightmail Gateway integrates with Symantec Network Prevent to deliver, route, hold, or block email traffic. Symantec Network Prevent is a component of Symantec Data Loss Prevention, which discovers, monitors, and protects confidential data wherever it is stored or used. You install Symantec Network Prevent on a separate server. You must have a Scanner configured for outbound mail filtering to route email to Symantec Network Prevent.

Note: You cannot route inbound mail through Symantec Network Prevent.

You can configure Symantec Network Prevent policies to perform the following actions on messages, depending on the type of data that is detected:

Block	Block messages and return a customized bounce message back to senders.
Redirect	Route the messages to different recipients.
Tag	Modify the subject line or add a new header to messages.

See "Common Symantec Network Prevent actions" on page 465.

Based on message modification by Symantec Network Prevent, you can also configure policies on Symantec Brightmail Gateway to perform actions such as the following:

Archive	Send the messages to a specific email address for archiving.
Create an incident	Route messages to a content filtering folder and review them before they are delivered. You can optionally configure notification for the messages that are routed to content filtering folders.
Encrypt	Government regulations or your own policies may require that you encrypt sensitive messages.

See "How Symantec Brightmail Gateway and Symantec Network Prevent interact" on page 463.

Required ports for Symantec Network Prevent integration

Table 15-27 describes the default ports to use to route email to one or more Symantec Network Prevent servers.

Table 15-27 Ports used for Symantec Network Prevent integration

Port	Protocol	Origin	Destination	Description	Notes
25	SMTP	Internal mail servers	Scanners	Outbound email	–
10025	SMTP	Scanner	Symantec Network Prevent	Outbound email for processing by Symantec Network Prevent	–
25	SMTP	Symantec Network Prevent	Scanners	Outbound email that was processed by Symantec Network Prevent	By default, Symantec Network Prevent returns email to port 10026. You may need to change the **Remote SMTP Listener Port** setting on Symantec Network Prevent to match the port that Symantec Brightmail Gateway expects, such as port 25.

How Symantec Brightmail Gateway and Symantec Network Prevent interact

If you configure Symantec Brightmail Gateway to route email to Symantec Network Prevent, email is typically routed in the following order:

- Symantec Brightmail Gateway accepts outbound messages at the gateway on port 25, by default.
- Symantec Brightmail Gateway passes outbound messages to Symantec Network Prevent on port 10025, by default.
- Symantec Network Prevent scans messages and blocks, redirects, or tags messages for further action by the MTA.
- See "Common Symantec Network Prevent actions" on page 465.
- Symantec Network Prevent passes messages back to Symantec Brightmail Gateway on port 25 (default) unless the Symantec Network Prevent rejects the message. In that case, Symantec Brightmail Gateway returns the message back to the sender with an SMTP 5xx failure response code. The message includes the text that you specify.
- Symantec Brightmail Gateway processes messages as configured. Symantec Brightmail Gateway can process messages based on subject or header markup of messages by Symantec Network Prevent. Redirected messages are delivered to the alternate recipient or recipients.

The port numbers that are listed are suggested. Actual port numbers may differ at your site.

See "Required ports for Symantec Network Prevent integration" on page 463.

If you have multiple Scanners, the **Symantec Data Loss Prevention Setup** settings in the Control Center apply to all Scanners.

Each Scanner routes email to all configured Symantec Network Prevent servers according to the preference order as follows:

Reflecting mode	If Symantec Network Prevent is configured in reflecting mode, then each Symantec Network Prevent server returns each message to the Scanner from which it received the message.
Forwarding mode	If Symantec Network Prevent is configured in forwarding mode, then Symantec Network Prevent servers pass messages to the next destination.

See "About Symantec Network Prevent preference order" on page 471.

Symantec Brightmail Gateway integrates with Symantec Network Prevent through SMTP Client IP Address-based Routing. Refer to the *MTA Integration Guide for Network Prevent* for more information.

Common Symantec Network Prevent actions

Symantec Brightmail Gateway and Symantec Network Prevent interoperate by exchanging SMTP messages. Through response rules and policies, Symantec Network Prevent can modify, reroute, and reject messages. Some actions that you configure Symantec Network Prevent to take on email require no further action by Symantec Brightmail Gateway. Other actions require you to configure Symantec Brightmail Gateway.

Table 15-28 Common Symantec Network Prevent actions

Action	Description
Block	Bounces the message back to the sender
Redirect	Sends the message to different recipients You can configure Symantec Network Prevent to redirect messages to a new recipient, such as an administrator email address.
Tagging - Header markup	Adds a custom email header to messages You can add a header to an email message, such as X-Sensitive-Data: SSN. You can configure Symantec Brightmail Gateway to search for the custom header and act on the message. Symantec Brightmail Gateway can delete, archive, create an incident, and hold the message for review. You can also run reports on the actions that Symantec Brightmail Gateway takes on matching messages. See "Creating an incident based on Symantec Network Prevent header markup" on page 475.
Tagging - Subject modification	Changes the subject line Like header markup, you can configure Symantec Brightmail Gateway to search for the specific text in the subject line and act on matching messages.

Supported Symantec Network Prevent delivery modes

You can integrate Symantec Network Prevent into your network architecture. Symantec Network Prevent does not make its own on-disk copies of messages. Messages are retained in memory only. The incoming message transaction is not committed until the outbound message transaction succeeds.

You can integrate Symantec Network Prevent into your network architecture through the following methods:

Reflecting	After it processes messages, Symantec Network Prevent returns the message to the Scanner from which it came.
Forwarding	After it processes messages, Symantec Network Prevent passes messages to the MTA that you specify.

The method that you choose depends on the particular requirements at your site.

About failure behavior with Symantec Network Prevent

If Symantec Network Prevent is unreachable, email may bypass Symantec Network Prevent or wait in a mail queue. The behavior depends on how many Symantec Network Prevent servers are unreachable and whether bypass is enabled. By default, bypass is enabled.

See "Enabling or disabling bypass for Symantec Network Prevent" on page 469.

Symantec Network Prevent servers may be unreachable because of the following reasons:

- Failures in the network
- Failures on the hardware on which Symantec Network Prevent is running
- The network bandwidth, hardware speed, or number of Symantec Network Prevent servers are not adequate for the mail flow

Table 15-29 Failure behavior with Symantec Network Prevent

Bypass status	Symantec Network Prevent server status	System behavior
Bypass disabled or enabled	One unreachable of two or more	The unavailable Symantec Network Prevent server is bypassed. Email is routed to the next Symantec Network Prevent server according to the preference list, MX record, or both.

Table 15-29 Failure behavior with Symantec Network Prevent *(continued)*

Bypass status	Symantec Network Prevent server status	System behavior
Bypass disabled	All unreachable	Email is stored in Symantec Brightmail Gateway 's delivery queue. No outbound email is delivered. If the Symantec Network Prevent servers continue to be unavailable, the delivery queue grows larger as time passes. Messages can get stuck in Symantec Brightmail Gateway if the following conditions are met: ■ The Maximum number of messages in the delivery queue limit is reached. ■ The option **Defer new connections when delivery queue is full** is enabled. New inbound connections and outbound connections are deferred when the delivery queue becomes full. See "How to resolve a delivery queue back up to Symantec Network Prevent" on page 467. See "Configuring SMTP advanced settings" on page 95.
Bypass enabled (default)	All unreachable	Email is not routed to Symantec Network Prevent servers, but Symantec Brightmail Gateway does process it. Sensitive data can leave your site unscanned unless you configure Symantec Brightmail Gateway appropriately. See "Creating a policy to detect unscanned email if Symantec Network Prevent bypass is enabled" on page 469.

How to resolve a delivery queue back up to Symantec Network Prevent

If you disable bypass to the Symantec Network Prevent servers and the servers are unreachable, messages back up in the delivery queue. Normally when the Symantec Network Prevent servers become reachable again, the delivery queue automatically drains as the servers process the messages.

The delivery queue may not automatically drain for any of the following reasons:

- The amount of free hard disk space on Symantec Brightmail Gateway is too low.

- The amount of free memory on Symantec Brightmail Gateway is too low.
- The following conditions on Symantec Brightmail Gateway are met:
  - The maximum number of messages in the delivery queue limit is reached.
  - The **Defer new connections when delivery queue is full** option is enabled.

See "About failure behavior with Symantec Network Prevent" on page 466.

If Symantec Network Prevent servers continue to be unavailable, the queue may grow to a large size and consume large amounts of disk space. This lowered disk space can impact Symantec Brightmail Gateway's ability to deliver messages.

Table 15-30 Process to resolve delivery queue backups

Step	Task
1	Correct the issue with the Symantec Network Prevent servers. Confirm that the Symantec Network Prevent servers are running, have sufficient capacity for processing the given email volumes, and network paths are reachable.
2	If disk space is low on the Symantec Brightmail Gateway Scanners, free disk space. Examples of ways to free disk space are to delete report data or logs. See delete on page 753.
3	Temporarily or permanently enable bypass. Sensitive data can leave your site unscanned unless you configure Symantec Brightmail Gateway appropriately. See "Creating a policy to detect unscanned email if Symantec Network Prevent bypass is enabled" on page 469. See "Enabling or disabling bypass for Symantec Network Prevent" on page 469.
4	Temporarily divert incoming message flow to a different Scanner.
5	In the Control Center, access **Administration > Hosts > Configuration**. Select a Scanner, access the **SMTP** tab, and then select **Advanced Settings**. Temporarily change the following settings: ■ **Maximum number of messages in delivery queue**: Increase. ■ **Defer new connections when delivery queue is full**: Uncheck. See "Configuring SMTP advanced settings" on page 95.
6	Optionally, in the Control Center, click **Status > SMTP > Messages Queues**, and flush all of the queues. The MTA flushes queues automatically, but this manual action expedites queue processing.

Enabling or disabling bypass for Symantec Network Prevent

By default, outbound email bypasses Symantec Network Prevent if all Symantec Network Prevent servers are unavailable. Bypass is triggered only if a connection cannot be established with the Symantec Network Prevent servers.

Bypass is not triggered in the following cases:

- The connection to Symantec Network Prevent server is established but the connection is deferred.
- The email results in an SMTP 4xx temporary failure. The Symantec Brightmail Gateway MTA attempts to redeliver the message later.
- The email results in an SMTP 5xx permanent failure. The Symantec Brightmail Gateway MTA sends a bounce message to the sender.
- The Symantec Network Prevent server is slow in processing the SMTP connection. However, if the SMTP connection times out, bypass is triggered if no other Symantec Network Prevent servers are available.

Sensitive data can leave your site unscanned if Symantec Network Prevent servers are unreachable. If you disable bypass and Symantec Network Prevent is unavailable, all outbound email waits in the delivery queue which prevents timely delivery.

See “About failure behavior with Symantec Network Prevent” on page 466.

See “Creating a policy to detect unscanned email if Symantec Network Prevent bypass is enabled” on page 469.

To enable or disable bypass for Symantec Network Prevent

1. In the Control Center, click **Content > Settings > DLP Connect**.
2. To enable bypass, check **Enable bypass when all DLP servers are unreachable**.

   To disable bypass, uncheck **Enable bypass when all DLP servers are unreachable**.
3. Click **Save**.

Creating a policy to detect unscanned email if Symantec Network Prevent bypass is enabled

By default, outbound email bypasses Symantec Network Prevent if Symantec Network Prevent is unavailable. Sensitive data can leave your site unscanned if Symantec Network Prevent servers are unreachable. However, you can configure a policy on Symantec Brightmail Gateway to prevent unscanned messages from leaving your site.

You can prevent unscanned messages from leaving your site when you create a content filtering policy for unscanned messages. In this policy configuration, the administrator of the content filtering folder receives notification if an unscanned message is detected. Unscanned messages are held in the content incident folder for review.

Table 15-31 describes the process that you can follow to prevent bypassed email from leaving your site.

Table 15-31 Process to create a content filtering policy to prevent bypassed email from leaving your site

Step	Task
1	Ensure that the Symantec Network Prevent server adds a header to messages it processes. By default, Symantec Network Prevent adds the header `X-CFilter-Loop:` to messages it processes.
2	In Symantec Brightmail Gateway, create a content filtering folder for unscanned messages, and configure email notification for the content filtering folder.
3	In Symantec Brightmail Gateway, create an email content filtering policy with the following characteristics: ■ Policy template: Blank ■ Apply to: Outbound messages ■ Condition - Text in this specific part of the message: Message header ■ Condition - Header name: `X-Cfilter-Loop` ■ Condition - The message header: does not exist ■ Perform the following action: Create an incident. Add additional actions for approved and rejected. ■ In content incident folder: The folder that you created ■ Hold message for review: checked See "Creating content filtering policies" on page 334.

Alternatively, you may not need to store unscanned messages but want to be notified if messages are unscanned. You can create an email content filtering policy on Symantec Brightmail Gateway to send an email notification if unscanned messages are detected.

See "About policy violation notifications" on page 400.

About Symantec Network Prevent reports

Symantec Brightmail Gateway does not provide reports specifically for Symantec Network Prevent activity. Use Symantec Network Prevent for reports on its activity. Some activity that is related to Symantec Network Prevent may be part of reports on Symantec Brightmail Gateway. For example, you can create an incident for the messages that are marked up by Symantec Network Prevent. That incident activity is reflected in some content filtering reports.

See "About working with reports" on page 568.

About performance implications for Symantec Network Prevent integration

Due to the additional processing that is involved, the integration with Symantec Network Prevent may add latency to the outbound email delivery speed.

The amount of latency depends on the following factors:

- The volume of outbound email at your site
- The number and complexity of content filtering policies in Symantec Network Prevent
- Message content and size
- The number of Symantec Network Prevent servers compared to the volume of mail

You may be able to decrease latency and increase throughput by doing any of the following tasks:

- Add additional Symantec Network Prevent servers
- Tune the **Maximum number of connections** and **Maximum number of connections from a single IP address** setting for outbound SMTP and SMTP delivery. This setting is a Symantec Brightmail Gateway Scanner setting.

  See "Configuring SMTP advanced settings" on page 95.
- Tune the **NumThreads** setting on the Symantec Network Prevent server to optimize throughput depending on the volume of outbound email at your site.

About Symantec Network Prevent preference order

If you have multiple Symantec Network Prevent servers, you can configure the order in which Symantec Brightmail Gateway employs Symantec Network Prevent servers. You can use the preference order for server priority, server load balancing, or failover. You set the preference either in the Control Center or in MX records.

Table 15-32 Available preference types

Preference type	Description
Preference on Symantec Data Loss Prevention Setup page	You configure this type of preference on the **Symantec Data Loss Prevention Setup** page in the Control Center. You must specify a preference value for every Symantec Network Prevent server that you specify. The valid preference range is 1 - 100.
MX record preference	You configure MX record preference in the DNS records for Symantec Network Prevent with your DNS software. You must configure MX record preference if the host name that you specify routes to more than one Symantec Network Prevent server. The valid MX preference range is 0 - 65535. Typical values are every 10 digits between 10 and 100.

For both types of preference, lower numbers are attempted before higher numbers. For example, a Symantec Network Prevent server with a preference of 1 is tried before a Symantec Network Prevent server with a preference of 2.

You can use both types of preference in combination. The preference settings on the **Symantec Data Loss Prevention Setup** page are compared before the preference is checked in the MX record for the chosen host name.

Table 15-33 Uses for preference

Preference use	Description
Server priority and server failover	Choose a lower preference number for Symantec Network Prevent servers that run on high network bandwidth computers with ample CPU and hard disk resources. Choose a higher preference number for Symantec Network Prevent servers that run on lower network bandwidth computers with fewer CPU and hard disk resources. In case a Symantec Network Prevent server is unreachable, the server with the next higher preference number is tried.
Load balancing	If Symantec Network Prevent servers have the same preference number, Symantec Brightmail Gateway randomly chooses a Symantec Network Prevent server for each outbound message. This random selection of Symantec Network Prevent servers creates load balancing among the Symantec Network Prevent servers.

Troubleshooting Symantec Network Prevent integration: messages bounce

If you have recently enabled routing to Symantec Network Prevent and email messages bounce, check the following:

- **Accept Scanned Mail from DLP Servers** IP address is configured correctly in Symantec Brightmail Gateway.
- The Prevent Server is correctly configured with the reflected port setting that points to the Scanner port.

Troubleshooting Symantec Network Prevent integration: deferred messages

If outbound email returns with the following error:

```
421 Forwarding agent unavailable. Closing connection.
```

the **Remote SMTP Listener Port** for Symantec Network Prevent does not match the inbound email port on Symantec Brightmail Gateway.

See "Configuring Symantec Network Prevent to return email to Symantec Brightmail Gateway" on page 473.

This message describes the error that is seen in Symantec Data Loss Prevention from Symantec Version 8.1. Future versions may produce different errors.

Configuring Symantec Network Prevent to return email to Symantec Brightmail Gateway

You must configure Symantec Network Prevent server to return email to Symantec Brightmail Gateway on the expected port. The default outbound email port for Symantec Brightmail Gateway is 25. The default Remote SMTP Listener Port in Symantec Network Prevent is 10026. Change the Symantec Network Prevent port to match the Symantec Brightmail Gateway outbound email port.

Table 15-34 Process to configure Symantec Network Prevent to return mail to Symantec Brightmail Gateway

Step	Action	Description
1	Check the default outbound email port on Symantec Brightmail Gateway.	The outbound mail IP address port appears on the **SMTP** tab of the **Host Configuration** page in the Control Center.

Table 15-34 Process to configure Symantec Network Prevent to return mail to Symantec Brightmail Gateway *(continued)*

Step	Action	Description
2	Set the **Remote SMTP Listener Port** for Symantec Network Prevent.	Set the **Remote SMTP Listener Port** for Symantec Network Prevent to the outbound email port that Symantec Brightmail Gateway uses. See the Symantec Network Prevent documentation for configuration details.

Configuring email connections to and from Symantec Network Prevent

If you have one or more Symantec Network Prevent servers, you can route email to them from Symantec Brightmail Gateway. You also must configure Symantec Network Prevent to route email back to Symantec Brightmail Gateway.

See "Configuring Symantec Network Prevent to return email to Symantec Brightmail Gateway" on page 473.

To configure email connections to and from Symantec Network Prevent

1 In the Control Center, click **Content > Settings > DLP Connect**.

2 Check **Enable** to route email to Symantec Network Prevent.

3 Under **Route Outbound Mail to DLP servers**, click **Add** to add a blank row.

4 Under **Host or IP Address**, specify the domain name or IP address of a Symantec Network Prevent server.

The domain name can be of the form server1.symantecexample.com or symantecexample.com. Specify a domain name (not IP address) if MX records are configured for the Symantec Network Prevent server.

5 Under **Port**, specify the port number on the Symantec Network Prevent server to which the outbound email should be routed.

The default port is 10025. Ensure that the Local SMTP Listener Port on the Symantec Network Prevent server is set to the same port number.

6 Check **MX Lookup** to enable MX lookup for the Symantec Network Prevent server.

If you check **MX Lookup**, ensure that you have specified a domain name (not IP address) in the form server1.symantecexample.com or symantecexample.com.

7 Under **Preference (1 - 100)**, specify the preference of this Symantec Network Prevent server as compared to all the defined Symantec Network Prevent servers.

See "About Symantec Network Prevent preference order" on page 471.

Enable bypass when all DLP servers are unreachable is described in another section.

See "Enabling or disabling bypass for Symantec Network Prevent" on page 469.

8 Under **Accept Scanned Mail from DLP servers**, click **Add** to add a blank row.

9 Specify an IP address from which Symantec Brightmail Gateway should expect email from Symantec Network Prevent.

Add additional rows to specify the IP addresses of all Symantec Network Prevent servers from which Symantec Brightmail Gateway should expect email. For example, you may have only one host name specified under **Route Outbound Mail to DLP servers**. But if that host name resolves to multiple Symantec Network Prevent servers, add the IP addresses of all of those servers.

10 Click **Save**.

Creating an incident based on Symantec Network Prevent header markup

Below is a sample method to handle any sensitive data that Symantec Network Prevent detects. This method requires that you configure both Symantec Network Prevent and Symantec Brightmail Gateway.

In this example, Symantec Network Prevent adds a custom header to matching messages. Symantec Brightmail Gateway creates an incident for messages with the custom header and holds the messages for review. So messages are not delivered to the original recipient but are instead routed to a content filtering folder on Symantec Brightmail Gateway. An administrator can approve, reject, forward, archive, delete, and manage the messages in the content incident folder.

Table 15-35 describes the process to create an incident based on Symantec Network Prevent header markup.

Table 15-35 How to create an incident based on Symantec Network Prevent header markup

Step	Task
1	Configure Symantec Network Prevent to add a custom header to the messages that it detects with sensitive data on the **Add/Edit Response Rule** screen under **Create a Modify SMTP Message response** rule. You can add up to three RFC 2822 header lines. Symantec recommends that you use the header `X-Cfilter:` with different values depending upon the wanted action on Symantec Brightmail Gateway or scan verdict. For example, you can specify `X-Cfilter: Symantec Incident` to mark messages for a content incident folder or `X-Cfilter: SSN` for any messages that contain social security numbers.
2	In Symantec Brightmail Gateway, create a content filtering folder, such as "Symantec Incidents". You may want to enable email notification for the content filtering folder. See "About content incident folders" on page 438.
3	In Symantec Brightmail Gateway, create an email content filtering policy with the following characteristics: ■ Policy template: Blank ■ Apply to: Outbound messages ■ Condition - Text in this specific part of the message: Message header ■ Condition - Header name: The header you configured, such as `X-Cfilter: Symantec Incident` ■ Condition - The message header: exists ■ Perform the following action: Create an incident. Add additional actions for approved and rejected. ■ In content incident folder: The folder that you created ■ Hold message for review: checked

About taking Symantec Network Prevent servers offline for maintenance

Occasionally you may need to take one or more Symantec Network Prevent servers offline to perform maintenance, such as to install a new software release.

Table 15-36 describes the effect on mail flow for the options available.

Table 15-36 Options for taking Symantec Network Prevent servers offline

Option	Description
If you have multiple Symantec Network Prevent servers, perform maintenance on one server at a time	You can perform maintenance on one server at a time if either of the following conditions are true: ■ You have configured multiple Symantec Network Prevent servers in Symantec Brightmail Gateway ■ The address that you specify for Symantec Network Prevent in Symantec Brightmail Gateway resolves to multiple Symantec Network Prevent servers using MX records The Symantec Network Prevent server that is not available is ignored and the next available Symantec Network Prevent server in the preference list is used.
Disable routing to Symantec Network Prevent	Outbound email bypasses Symantec Network Prevent servers. You must manually re-enable routing to Symantec Network Prevent servers when they become available again. Sensitive data can leave your site unscanned by Symantec Network Prevent servers. See "Creating a policy to detect unscanned email if Symantec Network Prevent bypass is enabled" on page 469.

Table 15-36 Options for taking Symantec Network Prevent servers offline *(continued)*

Option	Description
Enable bypass	Outbound email bypasses Symantec Network Prevent servers if none are available. Sensitive data can leave your site unscanned by Symantec Network Prevent servers. When Symantec Network Prevent servers are available again, Symantec Brightmail Gateway automatically routes outbound email to them. Bypass is enabled by default. If you disable bypass, outbound email remains in the delivery queue. When you later activate Symantec Network Prevent servers, the delayed email is scanned for sensitive data. However, outbound email may be bounced back to the senders if the connection is not reestablished in time. Messages are queued up for three days by default before they are bounced. If bypass is disabled, you can perform maintenance on Symantec Network Prevent servers when you do the following tasks: ■ Enable bypass in the Control Center of Symantec Brightmail Gateway. ■ Perform maintenance on the Symantec Network Prevent servers. ■ When work on the servers is complete, bring the Symantec Network Prevent servers online again. ■ Disable bypass in the Control Center of Symantec Brightmail Gateway. Ensure that your organization's policies let you enable bypass temporarily.
Stop outbound message queue	No outbound messages are delivered while the queue is stopped. Symantec Network Prevent servers and Symantec Brightmail Gateway eventually scan all outbound messages. You must manually reenable the outbound message queue when Symantec Network Prevent servers become available again. See "MTA and message queue behavior" on page 613.
Enable routing to Symantec Network Prevent but disable bypass	Outbound messages back up in the delivery queue.

Chapter 16

Configuring directory data integration

This chapter includes the following topics:

About using the directory data service

The directory data service lets you use the information that is stored in your Lightweight Directory Access Protocol (LDAP) directories for features in the Symantec Brightmail Gateway.

Symantec Brightmail Gateway provides four functions that you can enable for your data source: authentication, address resolution, routing, and recipient validation. You enable one or a combination of these functions for each data source.

See "About data sources and functions" on page 482.

To use your directory data for these features, you configure data sources for desired functionality, and also configure additional features in Symantec Brightmail Gateway.

See "Creating a data source" on page 491.

You can use these functions to support feature configuration within Symantec Brightmail Gateway. Table 16-1describes the functionality in Symantec Brightmail Gateway that relies on directory data service configurations.

Table 16-1 Features in Symantec Brightmail Gateway that use configured data sources

Feature	Data source function to enable	Additional setup tasks
Reject or drop invalid recipients	Recipient validation	Configure invalid recipient handling and enable recipient validation for at least one domain See "Setting up invalid recipient handling" on page 130. See "Adding or editing domains" on page 117.
Use directory harvest attack recognition	Recipient validation	Configure directory harvest attack and domains See "Adding or editing domains" on page 117. See "Configuring directory harvest attack recognition" on page 172.
Participate in the Symantec Probe Network.	Recipient validation	Enable recipient validation for at least one domain. See "To add or edit a domain" on page 119. See "Creating probe accounts from invalid recipient email addresses" on page 253.

Table 16-1 Features in Symantec Brightmail Gateway that use configured data sources *(continued)*

Feature	Data source function to enable	Additional setup tasks
Implement end-user quarantine and allow end-users access to spam quarantine messages	Authentication	Configure at least one policy group to quarantine spam messages, configure quarantine settings, and configure spam filtering. See "About assigning filter policies to policy groups" on page 321. See "About quarantining spam" on page 258. See "About filtering spam" on page 239. See "About quarantining spam" on page 258.
Let authenticated end users remotely send email messages using the SMTP authentication protocol.	Authentication	Set up SMTP authentication See "Using SMTP authentication" on page 145.
Use LDAP groups and distribution lists to apply policies.	Address resolution	Add an LDAP group or distribution list as a member of a policy group See "About assigning filter policies to policy groups" on page 321.
Allow end users to configure email language preferences and Good and Bad Senders lists	Address resolution and authentication User preference functionality requires that a data source is enabled for both authentication (to let the user log in and set preferences) and for address resolution (to replicate the preferences to the scanner.)	Enable at least one policy group for end-user settings. See "About assigning filter policies to policy groups" on page 321.

Table 16-1 Features in Symantec Brightmail Gateway that use configured data sources *(continued)*

Feature	Data source function to enable	Additional setup tasks
Re-route user email to an alternate address or alternate mail host based on directory information	Routing	Configure domains See "Adding or editing domains" on page 117.

About data sources and functions

To integrate your LDAP directories with Symantec Brightmail Gateway you must first configure data sources. A data source is a set of configuration data that allows your system to easily connect to and use your LDAP server.

When you configure a data source, you configure the parameters for connecting to your LDAP server and the functions that each data source provides.

You can configure multiple data sources for one LDAP server.

See "Creating a data source" on page 491.

Once you have configured your directory data sources, you can configure other features across Symantec Brightmail Gateway.

For example, if you want to use your new data source to apply a policy using an LDAP group, you must enable the address resolution function for that data source and configure an address resolution query.

See "About using the directory data service" on page 479.

You can configure one or more of the following functions for each of your data sources:

- Authentication
  Use the authentication function to securely authenticate end-user login to the Control Center and to authenticate users wishing to send email. The authentication function includes quarantine address resolution, which lets you better manage spam quarantine through improved handling of aliases, distribution lists, and invalid addresses.
  See "About using the authentication function with your data source" on page 487.
- Recipient validation
  Use the recipient validation function to validate email addresses against directory data and drop messages or reject connections for invalid recipients.

See "About using the recipient validation function with your data source" on page 488.

- Routing
  Use the routing function to route messages to alternate addresses or mail hosts on a per-user basis using directory data.
  See "About using the routing function with your data source" on page 489.
- Address resolution
  Use the address resolution function to apply policies to users and groups consistently.
  See "About using the address resolution function with your data source" on page 490.

When the directory data service cannot properly communicate with an LDAP server (for example, if the network link to the LDAP server is down or when a data integrity problem is encountered) message processing and user authentication can be affected. The resulting behavior varies based on the directory data service function that is affected.

You may need to contact your system administrator for assistance, or you can attempt the following strategies through your directory data service configuration.

- Use a network load balancer between the Symantec Brightmail Gateway host(s) and the LDAP directory to distribute requests between replicas of the directory data.
  This can improve performance and provide for failover if one of the servers becomes unavailable. If a load balancer cannot be employed, the directory data service rotates connections among multiple IP addresses assigned to an LDAP server hostname.
- Ensure that your data source cache size and time-to-live values are appropriate for your deployment. If the directory data service cannot contact an LDAP server, that cached result can be used, even if the time-to-live (TTL) of that cached entry has expired.
  See "Editing advanced settings for a data source" on page 558.

Note: To avoid message delivery impacts due to long network timeouts on requests to a faltering LDAP server host, if 10 data source access errors are encountered over a period of 60 seconds, then the data source is marked as unavailable to the directory data service for a period of 300 seconds, and all requests are served from the cache during that time.

About data source queries

Symantec Brightmail Gateway provides a default query for each function that is suitable for most common mail system configurations. Use the **Customize Query** option on the function pages to view and confirm those settings. If your LDAP schema requires unique queries you can create queries for any enabled function to change scope and support custom directory schemas.

First, you can use the **Base DN** field to set the scope of directory searches.

You provide the distinguished name (DN) for the subset of the directory you want to search. The smaller the subset of your directory, the more quickly the system can return data.

For example, the default query for your global company might be: dc=company, dc=com. If you need to limit searches to your New York office, you can set your base DN to limit the search to that sub-container in the directory. In this case your custom Base DN might be: ou=NewYork, dc=company, dc=com.

Note: The **Base DN** cannot be provided in this field for all directory types but is not applied to Domino or Active Directory Global Catalog queries.

In the **Query filter** field, you can define the rules the directory data service uses to search for entries in the directory. Use custom query filters that direct the directory data service to return data by using attributes and tokens to describe a specific LDAP query syntax.

The query filter is specified in standard LDAP query syntax. For example, email=%s describes a filter that finds all records with an attribute named "email" that matches the provided email address in it's entirety (as indicated by the %s query filter token).

At least one token is required for the **Query filter**. Table 16-2 describes the tokens that you can use to construct your query filters. The tokens you use are also dependent upon your directory type. Consult the documentation for your directory type for information about which of the following tokens your directory supports:

Table 16-2 Query filter tokens supported by the Symantec Brightmail Gateway

Syntax	Instructs the directory data service to look for the following information when it retrieves an entry:	Returns entries with the following types of information
%s	Email address to find, including user ID and fully-qualified domain name.	joe_smith@example.com
%u	Local portion of the email address to find.	joe_smith

Table 16-2 Query filter tokens supported by the Symantec Brightmail Gateway *(continued)*

Syntax	Instructs the directory data service to look for the following information when it retrieves an entry:	Returns entries with the following types of information
%d	Domain portion of the email address to find.	example.com
%n	Full name to find. Periods and underscores are replaced with space characters. Consecutive periods or underscores are replaced with a single space. This benefits quarantine address resolution by consolidating multiple versions of the same addresses into a single entry in the cache This token is particularly useful if using Domino, since Domino can be configured to consider the FullName (cn) attribute as a deliverable email address.	joe smith

You can combine attributes with and without tokens to refine returned data. For example, if you use the query filter mail=%s with a recipient validation function and want to only return individual users (as opposed to distribution lists) you can use the following filter: (&(objectClass=person)(mail=%s)).

Table 16-3describes the operators that you can use to construct you query filters.

Table 16-3 Query filter operators supported by Symantec Brightmail Gateway

Operator	Indicates the value is	Example	Data returned
=	Equal to	o=symantec	Entries for which the organization name is Symantec.

Table 16-3 Query filter operators supported by Symantec Brightmail Gateway *(continued)*

Operator	Indicates the value is	Example	Data returned
~=	Approximately equal to	(sn~=Johnson)	Entries with a name that is similar to "Johnson". Looks for entries that resemble the name provided. For example, the directory server might return entries with the surnames "Jonson" and "Johnsen". Rules for approximate matches vary based on directory type, and not all directory types support this operator. Consult your directory documentation for more information.
<=	Less than or equal to	(st<=ca)	Entries that occur before or at the location in the alphabet for the information provided. In this case the name "ca" and names appearing before "ca" in an alphabetical state attribute list, for example, "ak" can be returned.
>=	Greater than or equal to	(st>=ca)	Entries that occur after or at the location in the alphabet for the information provided. In this case the name "ca" and names appearing after "ca" in an alphabetical state attribute list, for example, "va" can be returned.
&	AND	(&(o=symantec)(co=france))	Entries in which the organization name is Symantec and the country is France.
=*	Exists	(&(mail=%s)(mailboxFile=*))	Entries that have a mail attribute that exactly matches the full recipient address, and have any mailboxFile attribute value to indicate that the user has an active email mailbox.

Table 16-3 Query filter operators supported by Symantec Brightmail Gateway *(continued)*

Operator	Indicates the value is	Example	Data returned
|	OR	(|(o=symantec)(sn=simpson) (cn=simpson))	Entries for the organization Symantec, the surname is Simpson, or the common name is Simpson.
!	NOT	(!(STREET=*))	Entries that do not contain a StreetAddress attribute. Note that when using the ! operator, you must use it outside of the parenthesis as follows: (!(<filter>)).
*	Wildcard	(o=*yma*)	Organization names that contain the string "yma", such as "Symantec" as well as the exact string match "yma" if present.

Finally, you can customize attributes such as the **Primary email attribute** and object classes such as **Distribution list object classes** as necessary to meet your unique needs.

Be sure that the administrator credentials used to connect to the LDAP server has access to the attribute you provide. If the account does not have proper access, or if the provided attribute does not exist or is misspelled, the query results indicate that no data is associated with the attribute.

See “Creating a data source” on page 491.

See “Creating and testing a custom authentication and quarantine address resolution query” on page 505.

See “Creating and testing a custom recipient validation query” on page 514.

See “Creating and testing a custom routing query” on page 517.

See “Creating and testing a custom address resolution query” on page 523.

About using the authentication function with your data source

An authentication data source provides end-user functions such as end-user preferences, end-user access to spam quarantine, and SMTP authentication.

See “About using the directory data service” on page 479.

When a data source is enabled for authentication, the system can do the following tasks:

- Enable and authenticate end-user access to Spam Quarantine.
  You must also configure at least one policy group to quarantine spam messages and configure quarantine settings and spam filtering.
  See "About assigning filter policies to policy groups" on page 321.
  See "About quarantining spam" on page 258.
  See "About filtering spam" on page 239.
- Let end users configure User Preferences, including personal good and bad sender lists and personal language settings (address resolution must also be enabled).
  See "About using the address resolution function with your data source" on page 490.
  See "Enabling and disabling end user settings for policy groups" on page 327.
- Let authenticated end users remotely send email messages using the SMTP authentication protocol.
  You must also set up SMTP authentication.
  See "Using SMTP authentication" on page 145.

When authentication is enabled, quarantine address resolution is automatically enabled as well. This query obtains information from inbound messages and uses that information to resolve email aliases to a primary email address for recipients of quarantined messages. As a result, the user can access all of their quarantined messages regardless of which alias the message was originally sent to.

See "About quarantining spam" on page 258.

See "Creating an authentication data source" on page 500.

About using the recipient validation function with your data source

Recipient validation determines whether your recipient addresses are valid. A valid address is an address that is found in at least one data source that is enabled for recipient validation. If a lookup address is valid and caching is enabled, the corresponding recipient entry is cached for later use (the cache is enabled by default). If a lookup address is not valid, it is cached in the invalid recipient cache.

Recipient validation works with other features within Symantec Brightmail Gateway to help you identify and manage invalid recipients and messages that are sent to those recipients.

See "About using the directory data service" on page 479.

When a data source is enabled for recipient validation, the system can the following:

- Validate email recipients and reject connections or drop messages for invalid recipients. You must also properly configure invalid recipient handling to enable this behavior.
  See "Setting up invalid recipient handling" on page 130.
- Perform directory harvest attack recognition. You must also configure directory harvest attack recognition to enable this behavior.
- Capture invalid recipient email addresses for use as probe accounts by the Symantec Probe Network .
  See "About probe accounts" on page 249.

See "Creating a recipient validation data source" on page 512.

Once invalid recipient handling has been enabled to reject or drop messages to invalid recipients, that source cannot be deleted or disabled. To delete or disable such a data source, you must reconfigure invalid recipient handling to accept all messages.

See "Setting up invalid recipient handling" on page 130.

About using the routing function with your data source

The routing function lets you configure per-user routing based on directory data. For example, you can use this function to route email messages addressed to users who have specified (as attributes associated with their own accounts) an alternate email address or mail host.

Each domain can be configured to route messages based on a directory data source, destination hosts, or MX lookup. This section talks about routing that uses a directory data source. For the other types of routing, consult the domain documentation.

See "Adding or editing domains" on page 117.

When a data source is enabled for directory-based routing, you can configure the directory data service to route email messages to alternate email addresses and mail hosts.

See "Creating a routing data source" on page 516.

If you configure both an **Alternate address attribute** and an **Alternate mailhost attribute**, the directory data service resolves all alternate email addresses first and then resolves the alternate mailhost. If an alternate address attribute is found, the data service attempts to assess the new address. If the address is local, the data service then attempts to find the new alternate address.

Once all local alternate addresses are resolved, the rules for that final address are then applied for purposes of alternate mail host routing.

The following is an example of what might happen if you send mail to a routing user who has both an alternate mailhost and an alternate address value:

1 Mail sent to this user resolves to the alternate address and is forwarded to the new address.

2 If the new address is local and that user also has an associated alternate address established, then mail is forwarded to that next address.

3 If the user associated with the last address only has an alternate mailhost value established, then the mail is delivered through the defined mailhost to the last email address.

Note: If an LDAP entry has multiple alternate mailhost attribute values, the MTA randomly chooses one of the values to use each time a message is delivered to that recipient. When an alternate email address attribute is supplied, mail that is bound for the original user is rerouted to the new email address. If there is more than one alternate address value provided, the message is rerouted to all of them.

To route messages to alternate addresses or mail hosts you must associate the desired domains with the routing data source.

See "Adding or editing domains" on page 117.

See "About email domains" on page 115.

See "Creating a routing data source" on page 516.

About using the address resolution function with your data source

The address resolution function resolves alias and distribution list data to let you apply policies on a per-user basis. You accomplish this using LDAP-based group memberships to define policies.

When a data source is enabled for address resolution, the system can do the following:

- Resolve directory group and distribution-list membership to enable scanners to apply filter policies to policy groups. You must configure policy groups. Distribution list expansion (which is enabled by default) is required to apply policies to users in distribution lists.
  See "About assigning filter policies to policy groups" on page 321.
  See "About expanding distribution lists and preserving recipient addresses" on page 531.
- Rewrite alias email addresses to primary email addresses to ensure that policies are applied consistently.

- Let end users configure User Preferences, such as personal language settings and good and bad sender lists based on email address (authentication must also be enabled).

  See "About using the authentication function with your data source" on page 487.

  See "Enabling and disabling end user settings for policy groups" on page 327.

The (optional) group listing query is enabled automatically when address resolution is enabled. This function lets you more easily define policy group membership by selecting from a list of LDAP users and groups. When you configure the group listing query, the groups list in the **Add Policy Groups** page is populated with a list of all LDAP groups (including LDAP distribution lists) retrieved by the query.

You may select groups from this list to add to policies or add groups by entering the LDAP group distinguished name yourself. If you have a large number of groups in your organization, you may want to disable this query by leaving the group query filter blank.

See "Creating an address resolution data source" on page 520.

Creating a data source

Symantec Brightmail Gateway provides a wizard that walks you through the process of creating a new data source. This basic initial configuration consists of the following steps:

- Provide details for the LDAP server hosting your directory data, such as host settings, bind DN, credentials, and SSL status.

  See "Adding a data source " on page 492.
- If desired, adjust your connection and cache settings. You can choose to use the default values, or you can edit cache and connection based settings to suit your needs.

  See "Configuring data source advanced settings" on page 494.
- Enable and configure the functions that you want your data source to provide: (authentication, recipient validation, routing or address resolution). You can configure one or more functions for a single data source.

  See "Enabling functions on a new data source" on page 498.

  See "About data source queries" on page 484.

Once you have saved your new data source, you can modify it as needed.

See "Editing a data source" on page 537.

Adding a data source

The **Add Directory Data Source** wizard guides you through the tasks necessary to create and configure new data sources. The first step in this process is to configure the LDAP server that provides data for your directory source.

See "Creating a data source" on page 491.

For authentication, address resolution, or routing data sources, the results that are returned must be unique. Therefore you cannot have identical or overlapping (where two data sources can potentially return the same result) data sources. If the **Directory data integrity errors** alert is configured on the **Alerts** page, that alert is triggered when this condition is detected.

See "Configuring alerts" on page 614.

When the directory data service cannot properly communicate with an LDAP server (for example, if the network link to the LDAP server is down or when a data integrity problem is encountered) message processing and user authentication can be affected.

See "About data sources and functions" on page 482.

To add a data source

1. Click **Administration > Settings > Directory Integration**.
2. On the **Directory Integration Settings - Directory Data Sources** page, click **Add**.

   The **LDAP Server Configuration** page appears.
3. In the **Data source name** field, provide a unique name for the data source.
4. Select the **Directory type** that your LDAP source uses. Available choices are as follows:
   - Active Directory
   - Active Directory Global Catalog
   - iPlanet/Sun ONE/Java Directory Server
   - Domino
   - Other (can be any LDAPv3 compliant directory)

   Though Symantec Brightmail Gateway is compatible with any LDAPv3 directory, it is specifically designed to support configurations with the four specific directory types listed. If you select "other" for your directory type, you might need to consult your directory documentation to ensure proper functionality.
5. In the **Host name** field, type the host name or IP address of the LDAP server.

6 The **Port** is used to access the LDAP server. The port is automatically populated based on your directory type and SSL settings but can be modified by typing a new value into this field.

7 Check **Enable SSL** if you want to enable SSL on all connections to the LDAP server host. Encryption is provided regardless of the certificate authority that is used to sign the LDAP server x.509 certificate. If you change this checkbox, the port is automatically updated to the default ports for your directory type. Confirm the port if modifying this setting.

8 Check **Anonymous bind** if you want to let the directory data service connect to the LDAP server without providing specific user ID and password information. Or check **Use the Following** to provide the directory data service with specific authentication credentials.

9 If you checked **Use the following**, provide the bind credentials in the following fields:

- **Name (Bind DN)**
  The distinguished name (DN) that is used for authenticating to the LDAP server.
  For Active Directory or Global Catalog server, you can optionally use the full DN, the NetBIOS and logon name (NetBIOS\SAM Account Name), or the User Principal Name.

- **Password**
  Password to be used to authenticate to the LDAP server.

10 Click **Test login** to validate your authentication to the LDAP server.

This test only verifies that the LDAP server can be reached and that the account has read access to the root of the directory data tree on a directory server. Therefore, a successful test result does not guarantee that the credentials can succeed elsewhere in the directory. This is particularly true for anonymous access. You should verify access before deployment by using the Test Query function when configuring individual functions for your data source.

11 Click **Show Advanced Settings** if you want to configure optional LDAP server and cache settings.

See “Configuring data source advanced settings” on page 494.

12 When you are finished, click **Next** to configure the functions for the data source.

See “Enabling functions on a new data source” on page 498.

Configuring data source advanced settings

In addition to setting up your LDAP server in the Configure LDAP Server page, you can further adjust your cache size and network settings in the Advanced Settings fields.

See “About the directory data cache” on page 526.

To configure advanced settings for a data source

1. Add your data source on the **Directory Integration Settings - Directory Data Sources** page then click **Next**.

   See “Adding a data source ” on page 492.

2. On the **LDAP Server Configuration** page, configure your server integration settings then click **Show Advanced Settings**.

3 You can configure any of the following settings for your new data source. All fields are optional and only need to be configured if your system has special requirements.

Item	Description
Maximum connections	Specify the maximum number of client connections (bind operations) that can be created at one time. If this field is set to zero, connection pooling is turned off and a new connection is created for each request.
Minimum connections	Specify the number of connections that are added to the connection pool when the source services the first request. A value of zero indicates that connections are created only to service actual pending requests. Created connections are released if they are idle for longer than the specified idle timeout.
Connection timeout	Specify the amount of time that should elapse before an attempt to connect to the LDAP server host is timed out and automatically ended. A value of zero indicates that the connection should never be timed out (though the LDAP server can also close the connection).
Idle timeout	Specify how long a client connection can remain idle before the connection is automatically closed. A value of zero indicates that the connection can remain idle indefinitely without being closed (though the LDAP server can also close the connection).
Search timeout	Specify how long (in seconds) a request /search operation should run before the directory data service ends the operation and displays the partial results. Please note that the LDAP server can impose a search timeout lower than this value.

Item	Description
Page size	Determine the maximum number of initial entries to return when a query is successful. Setting this value too low can impede performance. A page size higher than the limit set by the LDAP server can cause the operation to fail when the page size exceeds that limit. The default server-side limit for this value varies according to directory type. To change this limit on your directory server, see your directory server documentation. To disable paged searching entirely, set this value to 0. **Note:** Make sure that the Administration Credentials you provide have sufficient access rights to bypass search limits.
Chase referrals	When a request is processed, a server may return a query response that suggests that the directory data service query another LDAP server. When this field is enabled, the directory data service follows such referrals when it executes queries. The directory data service uses the same bind credentials to connect to the referred to server. If the referred-to LDAP server does not recognize the same bind credentials, a query can return an error.
Enable cache	The cache stores address entry data from previous requests. The cache lets the directory data service process requests faster by using this cached data instead of consulting the LDAP server directly for each lookup. See "About the directory data cache" on page 526. The cache is enabled by default.

Item	Description
Cache size	Specify the maximum number of entries that can be stored in the cache. When the cache size limit is exceeded, the least recently used entry is deleted to make room for a new entry. You should set the cache size based on your system's needs and memory availability. Symantec recommends that you set this value equal to or greater than the number of users and groups (including distribution lists, contacts, and public folders) in your environment. If you set the cache to zero entries, the cache is effectively disabled.
Cache index size multiplier	Specify the size of the email address index in relation to the cache. This setting allows the index to store multiple aliases for each entry in the cache. For example, a multiplier of 2 would be twice the size of the cache and allow for an average of two aliases per cache entry. Symantec recommends that you set this value equal to or greater than the average number of primary and alias addresses that are associated with your users.
Minimum cache TTL	Set the minimum Time to Live (TTL) for entries in the cache. When a cache TTL is reached, the entries in that cache expire and are refreshed upon query. A minimum value and maximum value for cache expiration creates a period of time over which these entries can be expired and refreshed. This helps pace your system's workload by ensuring that all entries do not expire at the same time. For address resolution sources, this setting also determines how often the membership information is rebuilt for LDAP groups used in policies. See “About the directory data cache” on page 526.

Item	Description
Maximum cache TTL	Set the maximum Time to Live (TTL) for entries in the cache. When a cache TTL is reached, the entries in that cache expire and are refreshed upon query. A minimum value and maximum value for cache expiration creates a period of time over which these entries can be expired and refreshed. This helps pace your system's workload by ensuring that all entries do not expire at the same time. See "About the directory data cache" on page 526.
Invalid recipient cache size	Provide the maximum number of entries that can be stored in the invalid recipient cache. When the directory data service cannot find an entry in the directory, the entry is considered invalid and is added to the invalid recipient cache. When the entry is queried in the future, load time is reduced because the directory data service checks the cache before it queries the directories. You should set the cache size based on your system's needs and memory availability. If you set the cache to 0 the cache is effectively disabled.
Invalid recipient cache TTL	Provide the period of time an email address entry should be kept in the invalid recipient cache.

4 To undo your changes, click **Restore Defaults**.

The original default values are restored to all Advanced Settings fields.

5 To hide these fields on the LDAP Server Configuration page, click **Hide Advanced Settings**.

6 When you are finished, click **Next** to configure the functions for the data source.

See "Enabling functions on a new data source" on page 498.

Enabling functions on a new data source

The second step in the **Add Directory Data Source** wizard is to enable and configure functions for your data source. Before you can enable functions for your data source, you must complete the first step in the wizard and confirm or configure the LDAP connection settings for that source on the **LDAP Server Configuration** page.

See “Adding a data source ” on page 492.

Functions enable certain behaviors for your data sources and let you take advantage of many Symantec Brightmail Gateway features. You can configure your data source for one or multiple functions.

See “About data sources and functions” on page 482.

To enable functions on a new data source

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See “Adding a data source ” on page 492.

2. In the **Add Directory Data Source - Directory Data Source Functions** page do the following tasks:

   - To enable and configure your data source for end-user authentication, quarantine address resolution, or SMTP authentication, check **Authentication**.
     See “Creating an authentication data source” on page 500.
     See “Using SMTP authentication” on page 145.
     To customize your authentication, quarantine address resolution, or SMTP authentication query, check **Customize Query**.
     See “Creating and testing a custom authentication and quarantine address resolution query” on page 505.

   - To enable and configure your data source for address resolution, check **Address resolution**.
     See “Creating an address resolution data source” on page 520.
     To customize your address resolution configuration, click **Customize Query**.
     See “Creating and testing a custom address resolution query” on page 523.

   - To enable and configure your data source for routing, check **Routing**.
     See “Creating a routing data source” on page 516.
     To customize your routing configuration, click **Customize Query**.
     See “Creating and testing a custom routing query” on page 517.

   - To enable and configure your data source for recipient validation, check **Recipient validation**.
     See “Creating a recipient validation data source” on page 512.

To customize your recipient validation configuration, click **Customize Query**.

See "Creating and testing a custom recipient validation query" on page 514.

3 When you have configured all of the functions for your new data source, click **Next** and verify your configuration.

4 When you are satisfied with your configuration, click **Save** to save and deploy your data source.

Creating an authentication data source

Authentication lets end users authenticate to the LDAP server and configure user preferences. It also lets you configure SMTP authentication.

See "About using the authentication function with your data source" on page 487.

See "Using SMTP authentication" on page 145.

Enabling your functions is the second step in the **Add Directory Data Source** wizard. Before you can enable functions for your data source, you must confirm or configure the LDAP connection settings for that source on the **LDAP Server Configuration** page.

See "Adding a data source " on page 492.

To prevent directory data service errors, be sure that your data sources do not produce overlapping results. Table 16-4describes how directory data service processes authentication errors for the most common conditions.

See "About data sources and functions" on page 482.

Table 16-4 Common directory data service authentication errors

Authentication error	System behavior
The directory data service cannot properly communicate with the LDAP directory server when it attempts to authenticate a user to either Control Center or SMTP server. This can happen, for example, if the network link to the LDAP server is down.	The error "service is temporarily unavailable" appears. If the proper alerts are enabled on the **Alerts** page, a **Directory data access errors** alert is triggered. See "Configuring alerts" on page 614.
The directory data service cannot uniquely determine the identity of the user from the user name provided.	The error "service is temporarily unavailable" appears. If the proper alerts are enabled on the **Alerts** page, a **Directory data access errors** alert is triggered. See "Configuring alerts" on page 614.

Table 16-4 Common directory data service authentication errors *(continued)*

Authentication error	System behavior
The directory data service cannot properly communicate with the LDAP directory server when it attempts to resolve the recipient of a quarantined message to the recipient's primary email address or uniquely determine the primary email address of the recipient.	The message is kept in the delivery queue. If the alerts are enabled on the **Alerts** page, the appropriate alert is triggered. See "Configuring alerts" on page 614.

To configure authentication for a new data source

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See "Adding a data source " on page 492.

2. On the **Add Directory Data Source - Directory Data Source Functions** page, check **Authentication**.

3. In the **Authentication type** pull-down menu, choose one of the following:

Control Center authentication only	To authenticate end users for accessing quarantine or setting end-user preferences only.
SMTP authentication only	To enable remote users to use Symantec Brightmail Gateway to send email using SMTP authentication.
Control Center and SMTP authentication	To enable both Control Center and SMTP authentication.

4 If you selected **Control Center authentication only** or **Control Center and SMTP authentication** as the authentication type, confirm or configure the following test data or click **Customize Query** to view or customize your query.

See "Creating and testing a custom authentication and quarantine address resolution query" on page 505.

Test user name

Type a user name that can be used to test and validate your authentication configuration.

Symantec recommends that you verify your settings using a variety of input designed to produce both successful authentication and unsuccessful authentication. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

Test password

Type the password for the test user name.

Test domain (optional)

Type a NetBIOS domain for the provided test credentials.

You only need supply the test domain if you use Active Directory or the Active Directory Global Catalog and the provided **Test user name** is domain-specific.

Test Query

Click to validate your query using the test information provided.

This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source being tested. The results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Customize Query

Click **Customize Query** if you want to examine or modify the default settings for your queries. You can create custom queries to more accurately reflect your system's configuration.

See "Creating and testing a custom authentication and quarantine address resolution query" on page 505.

Test email address

Type an email address that can be used to test and validate your quarantine address resolution configuration. Symantec recommends that you test your data source using a combination of addresses, including a user address, distribution list address, alias address, and an invalid address.

Test Query

Click to validate the defined quarantine address resolution query using the provided test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source being tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Customize Query

Click **Customize Query** if you want to examine or modify the default settings for your queries.

See “Creating and testing a custom authentication and quarantine address resolution query” on page 505.

5 If you selected **SMTP authentication only** or **Control Center and SMTP authentication** as the authentication type, configure the following fields or click **Customize Query** to view or customize your query.

See "Creating and testing a custom authentication and quarantine address resolution query" on page 505.

If you selected Control Center and SMTP authentication, you can also check **Share Control Center and SMTP Authentication query details** to populate your SMTP details with those provided for the Control Center and skip this step.

Test user name	Type a user name that can be used to test and validate your SMTP authentication configuration. Symantec recommends that you verify your settings using a variety of input designed to produce both successful authentication and unsuccessful authentication. If the test produces unexpected results, click **Customize Query** to verify your settings and directory data.
Test password	Type the password for the test user name.
Test Query	Click to validate the defined SMTP Authentication query using the provided test user name and password. This test is conducted against the directory data service instance that is running on the Control Center host. It cannot be used to verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source being tested and. The results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.
Customize Query	Click **Customize Query** to create a customer query, or to examine the default settings for your query. You can create custom queries to more accurately reflect your system's configuration. See "Creating and testing a custom authentication and quarantine address resolution query" on page 505.

6 You can configure your data source for multiple functions or click **Next** to review your changes and save your deployment.

See "Enabling functions on a new data source" on page 498.

Creating and testing a custom authentication and quarantine address resolution query

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See “About data source queries” on page 484.

To customize the authentication query for a new data source

1 Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

See “Adding a data source ” on page 492.

2 On the **Add Directory Data Source - Directory Data Source Functions** page, check **Authentication**, select the type of authentication you want to use, and then click **Customize Query**.

You can also click **Customize Query** to examine the default settings for your query.

See “Enabling functions on a new data source” on page 498.

3 In the **Authentication Type** drop-down list, select the type of authentication you want to configure for your data source.

Control Center authentication only	To authenticate end users for accessing quarantine or setting end-user preferences only. You must configure Control Center authentication and quarantine address resolution queries.
SMTP authentication only	To enable remote users to use Symantec Brightmail Gateway to send email using SMTP authentication. You must configure the SMTP authentication query.
Control Center and SMTP authentication	To allow all of the capabilities of both Control Center and SMTP authentication. You must configure Control Center authentication, quarantine address resolution, and SMTP authentication queries. You can check **Share Control Center and SMTP Authentication query details** if you want to share your query details.

4. If you have selected Control Center and SMTP authentication you can check **Share Control Center and SMTP Authentication query details** if you want to use the Control Center authentication information as your SMTP authentication information.

5. Provide the following information for the Control Center authentication query if you selected **Control Center authentication only** or **Control Center and SMTP authentication** as the authentication type:

NetBIOS domain name (optional)

Provide a list of NetBIOS domain names that the system can use to resolve login ambiguities.

This field is optional but recommended if all of the following are true:

- Your directory type is Active Directory or Active Directory Global Catalog
- Your directory uses more than one NetBIOS domain
- User name uniqueness is not guaranteed.

Base DN

Provide a **Base DN** for the routing query.

A default value is provided, if customize is selected, a custom query start can be provided in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

Query filter

Modify the default **Query filter** for the authentication query, if desired.

The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Primary email attribute

This field is not relevant to SMTP authentication but is required for Control Center authentication.

Type the attribute in your LDAP schema that is used to store the primary email address information for the authentication query.

If you specify a primary email attribute of "proxyAddresses," the directory data service automatically identifies the attribute value that is prepended with "SMTP" as the address. You do not need to specify this prefix in the Primary email attribute field.

Authentication Method

Choose your **Authentication Method**.

Click **Simple bind** to attempt login using the password provided by the user.

Or, you can use password fetching, which fetches the password from the data source and compares it to the password provided by the user. Click **Password attribute** to choose password fetching.

You can optionally specify a **Default hash type** for password fetching. If you choose **none**, your directory data must prepend the password attribute with a prefix indicating the scheme, such as **{PLAIN}**. If you choose **plaintext**, a scheme prefix is not needed.

Symantec Brightmail Gateway supports SMTP authentication via LDAP using simple bind for all supported LDAP directory types. For SMTP authentication via LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported.

Test user name

Type a user name that can be used to test and validate your authentication configuration.

Symantec recommends that you verify your settings using a variety of input intended to produce both successful authentication and unsuccessful authentication. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

Test password

Type the password for the test user name.

Restore Defaults

Click **Restore Defaults** to restore the settings to the default settings.

Test domain (optional)

Type a NetBIOS domain for the provided test credentials.

You only need supply the test domain if you use Active Directory and the provided **Test user name** is domain-specific.

Test Query	Click to validate the defined quarantine address resolution query using the provided test email address. This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source tested. Results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

6 Provide the following information for the quarantine address resolution query if you selected **Control Center authentication only** or **Control Center and SMTP authentication** as the authentication type:

Query filter

In the **Query filter** field under Quarantine Address Resolution Query, modify the default filter if desired. The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Distribution list object classes

In the **Distribution list object classes** field, type the object classes in your LDAP schema that should be used to identify distribution list entries.

Restore Defaults

Click **Restore Defaults** to restore the settings to the default settings.

Test email address

In the **Test email address** field, type an address that can be used to test and validate your quarantine address resolution configuration. Symantec recommends that you test your data source using a combination addresses including a user address, distribution list address, alias address, and an invalid address.

Test Query

Click **Test Query** to validate the defined quarantine address resolution query using the provided test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source being tested. Results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

7 Provide the following information for the SMTP authentication query if you selected **SMTP authentication only** or **Control Center and SMTP authentication**.

If configuring both Control Center and SMTP authentication, you can check **Share Control Center and SMTP Authentication query details** to populate your SMTP details with those provided for the Control Center and skip this step.

Base DN

Provide a **Query start** for the routing query.

A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

You can customize the base DN to refocus the search to a specific part of the directory tree. This lets you configure the query to fit your particular needs and return data more quickly.

Query filter

Provide a **Login query filter** for the SMTP authentication query, if desired.

The login query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Primary email attribute

This field is not relevant to SMTP authentication. It is required for Control Center authentication.

In the **Primary email attribute** field, type the attribute in your LDAP schema that is used to store the primary email address information for the SMTP authentication query.

Authentication Method

Choose your **Authentication Method**.

Click **Simple bind** to attempt login using the password provided by the user.

Or, you can use password fetching, which fetches the password from the data source and compares it to the password provided by the user. Click **Password attribute** to choose password fetching. You can use the default password attribute or change it.

You can optionally specify a **Default hash type** for password fetching. If you choose **none**, your directory data must prepend the password attribute with a prefix indicating the scheme, such as **{PLAIN}**. If you choose **plaintext**, a scheme prefix is not needed.

Symantec Brightmail Gateway supports SMTP authentication via LDAP using simple bind for all supported LDAP directory types. For SMTP authentication via LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported.

Restore Defaults

Click **Restore Defaults** to restore the settings to the default settings.

Test user name

Type a user name that can be used to test and validate your authentication configuration.

Symantec recommends that you verify your settings using a variety of input intended to produce both successful authentication and unsuccessful authentication. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

Test password

Type the password for the test user name.

Test Query	Click to validate the defined SMTP authentication query using the provided test user name and password. This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source tested. Results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

8 Click **Save** to return to the **Add Directory Data Source - Directory Data Source Functions** page.

See “Enabling functions on a new data source” on page 498.

Creating a recipient validation data source

Recipient validation works with other features within Symantec Brightmail Gateway to help you identify and manage invalid recipients and the messages sent to those recipients.

See “About using the recipient validation function with your data source” on page 488.

Enabling your functions is the second step in the **Add Directory Data Source** wizard. Before you perform the task described in this topic, you must confirm or configure the LDAP connection settings for that source on the **LDAP Server Configuration** page.

See “Adding a data source ” on page 492.

See “Enabling functions on a new data source” on page 498.

If you want to use invalid recipient handling functionality you must also configure invalid recipient handling and enable recipient validation for at least one domain

See “Setting up invalid recipient handling” on page 130.

See “Adding or editing domains” on page 117.

Note: Be sure that your data sources do not produce overlapping results to prevent directory data service errors. If the directory data service cannot properly communicate with the LDAP directory server (for example, if the network link to the LDAP server is down) when it attempts to determine the validity of a message recipient, the MTA returns an error indicating that the delivery attempt should be retried later. See "About data sources and functions" on page 482.

To configure recipient validation for a new data source

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See "Adding a data source " on page 492.

2. On the **Add Directory Data Source - Directory Data Source Functions** page, check **Recipient validation**.

3. In the **Test email address** field, type an email address that can be used to test and validate your recipient validation configuration. Symantec recommends that you test at least one valid and one invalid address. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

4. If you want to examine or modify the default settings for your queries, click **Customize Query**. Create custom queries to more accurately reflect your system's configuration.

   See "Creating and testing a custom recipient validation query" on page 514.

5. Click **Test Query** to validate your query using the test email address provided.

   This test is conducted against the directory data service instance that is running on the Control Center host. It cannot be used to verify connectivity from attached scanners to your LDAP server.

   Test results reflect only the data source being tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

6. You can configure your data source for multiple functions or click **Next** to review your changes and save your deployment.

   See "Enabling functions on a new data source" on page 498.

Creating and testing a custom recipient validation query

You can view your default query information or customize your queries by configuring custom attributes and filters to change query scope and support unique directory schemas.

See "About data source queries" on page 484.

Symantec recommends that you test all queries before deploying a new data source.

To configure a custom recipient validation query for a new data source:

1 Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

  See "Adding a data source " on page 492.

2 On the **Add Directory Data Source - Directory Data Source Functions** page, check **Recipient validation**, and then click **Customize Query**.

  See "Enabling functions on a new data source" on page 498.

3 Provide the following information:

Base DN

Provide a **Base DN** for the query.

A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.

Query filter

Provide a **Query filter** for the query . The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Restore Defaults

Click **Restore Defaults** if you want to remove your edits to the recipient validation query configuration fields and replace them with the default values.

Test email address

Provide a **Test email address** that can be used to test and validate your recipient validation configuration. Symantec recommends that you test at least one valid and one invalid address.

Test Query

Click **Test Query** to validate the defined query using the provided test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. This test cannot be used to verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source being tested. Test results and do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

4 Click **Save** to return to the **Add Directory Data Source - Directory Data Source Functions** page.

See “Enabling functions on a new data source” on page 498.

Creating a routing data source

The routing function uses the information provided by your data source to route messages to addresses and domains.

See "About using the routing function with your data source" on page 489.

Each domain can be configured to route messages based on directory data source, destination hosts, or MX lookup. This section talks about routing that uses a directory data source. For the other types of routing, you should consult the domain documentation.

See "Adding or editing domains" on page 117.

Enabling your functions is the second step in the **Add Directory Data Source** wizard. Before you can enable functions for your data source you must configure the LDAP connection settings for that source on the **LDAP Server Configuration** page.

See "Creating a data source" on page 491.

Be sure that your data sources do not produce overlapping results to prevent directory data service errors. Table 16-5describes how directory data service processes routing errors for the most common conditions.

See "About data sources and functions" on page 482.

Table 16-5 Directory data service routing errors

Directory data service routing process error	System behavior
The directory data service cannot properly communicate with the LDAP directory server when it attempts to determine the routing information for a message recipient. This can happen, for example, if the network link to the LDAP server is down.	The MTA queues the message and periodically retries to perform the operation. A **Directory data access errors** alert is triggered. See "Configuring alerts" on page 614.
The directory data service cannot uniquely determine the LDAP entry that corresponds to the recipient.	The message is delivered without any attempt to apply routing information. If the proper alerts are enabled on the **Alerts** page, a **Directory data integrity errors** alert is triggered. See "Configuring alerts" on page 614.

.

To configure routing for a new data source:

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See “Adding a data source ” on page 492.

2. In the **Add Directory Data Source - Directory Data Source Functions** page, check **Routing**.

3. In the **Test email address** field, type an email address that can be used to test and validate your routing configuration. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

   The wizard provides a set of default attribute values for an LDAP configuration that uses the SunOne directory type. If your configuration uses a different directory type, or if you want to route mail using a different routing attribute, click **Customize Query** to configure your query.

   See “Creating and testing a custom routing query” on page 517.

4. Click **Test Query** to validate your query data using the provided test email address.

   This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot be used to verify connectivity from attached scanners to your LDAP server.

   Test results reflect only the data source tested. The test does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

5. You can configure your data source for multiple functions or click **Next** to review your changes and save your deployment.

   See “Enabling functions on a new data source” on page 498.

Creating and testing a custom routing query

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query page to customize your routing query to change scope and support custom directory schemas. Depending on the directory type selected, your routing attribute names may not be prepopulated with defaults. If your attribute names are not populated you cannot save your routing query until you provide at least one attribute for your LDAP schema.

See “About data source queries” on page 484.

This section covers directory based routing. For information about the other routing methods of destination host routing or MX lookup, consult the domain documentation.

See “Adding or editing domains” on page 117.

Symantec recommends that you test all queries before deploying a new data source.

To configure a custom routing query for a new data source:

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See “Adding a data source ” on page 492.

2. In the **Add Directory Data Source - Directory Data Source Functions** page, check **Routing**, and then click **Customize Query**.

   See “Enabling functions on a new data source” on page 498.

3. Provide the following information:

   Base DN

   Provide a **Base DN** for the query.

   A default value is provided, or you can select "Customize" to provide your own values in the **Custom query start** field.

   Custom query start

   If you select "Customize" for the **Base DN** field, you must provide a **Custom query start**.

   You can customize the base DN to refocus the search to a specific part of the directory tree. The **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.

   Query filter

   Provide a **Query filter** for the routing query, if desired. The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Alternate address attribute

If you want to enable per-user, directory-based, alternate-address routing, provide an **Alternate address attribute**. When an alternate email address attribute is supplied, mail bound for the original user is rerouted to the new email address. If there is more than one alternate address, the message is rerouted to all of them.

Note: To save your routing source, you provide an **Alternate address attribute**, an **Alternate mailhost attribute**, or both. If you configure both an **Alternate address attribute** and an **Alternate mailhost attribute**, the directory data service resolves all alternate email addresses before it resolves the alternate mailhost.

See “About using the routing function with your data source” on page 489.

Alternate mailhost attribute

If you want to enable directory-based alternate mail host routing, provide an **Alternate mailhost attribute** that corresponds to the attribute in the directory schema that you have selected to store this information.

The mail host should be specified as either an IP address or a fully qualified host name. A port can be specified by using a colon followed by the port, for example ":25". By default, the SMTP protocol uses port 25 for email transmission.

You can use the following valid attributes:

- 10.32.100.64
- 10.32.100.64:25
- myothermailserver.mycompany.com
- myothermailserver.mycompany.com:25

If an LDAP entry has multiple alternate mailhost attribute values, the MTA randomly selects one of the values to use each time a message is delivered to that recipient.

Note: To save your routing source, you provide an **Alternate address attribute**, an **Alternate mailhost attribute**, or both. If you configure both an **Alternate address attribute** and an **Alternate mailhost attribute**, the directory data service resolves all alternate email addresses before it resolves the alternate mailhost.

See “About using the routing function with your data source” on page 489.

Perform MX Lookup on transport value

Check **Perform MX Lookup on transport value** to allow MX lookup for the specified transport value. MX Lookup performs a directory lookup for the domain and returns the hosts responsible for accepting mail destined to that domain.

Restore Defaults

Click **Restore Defaults** if you want to remove your edits to the routing query configuration fields and replace them with the default values.

Test email address

Provide a **Test email address** that can be used to test and validate your routing configuration.

Test Query

Click **Test Query** to validate the defined routing query using the provided test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source being tested. It does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

4 Click **Save** to return to the **Add Directory Data Source - Directory Data Source Functions** page.

See “Enabling functions on a new data source” on page 498.

Creating an address resolution data source

The address resolution function resolves alias and distribution list data. This lets you apply policies on a per-user basis by using LDAP-based group memberships to define policies.

See “About using the address resolution function with your data source” on page 490.

Enabling your functions is the second step in the **Add Directory Data Source** wizard. Before you perform the task described in this topic you must confirm or configure the LDAP connection settings for that source on the **LDAP Server Configuration** page.

See “Adding a data source ” on page 492.

See “Enabling functions on a new data source” on page 498.

Be sure that your data sources do not produce overlapping results to prevent directory data service errors. Table 16-6describes how the directory data service processes address resolution errors for the more common error conditions. See “About data sources and functions” on page 482.

Table 16-6 directory data service processes address resolution errors

Address resolution error condition	System behavior
The directory data service cannot properly communicate with the LDAP directory server or uniquely determine the LDAP entry that corresponds to the recipient, when it attempts to resolve the recipient of a message for BATV purposes. An example of this condition is if the network link to the LDAP server is down.	The MTA returns a "451 Requested action aborted: error in processing" error. Assuming that the directory data service alerts are enabled on the **Alerts** page, the appropriate directory data service alert is triggered. See "Configuring alerts" on page 614.
The directory data service cannot properly communicate with the LDAP directory server when it attempts to resolve the recipient of a message for scanning purposes.,	The message is placed in the inbound deferred queue and periodically retried. If the alert is enabled, a **Directory data access errors** alert is triggered. See "Configuring alerts" on page 614.
The directory data service cannot uniquely determine the LDAP entry that corresponds to the recipient.	The message is delivered with no attempt to apply address resolution information. Policy groups are applied according to the original recipient address only. If the alert is enabled, a **Directory data integrity errors** alert is also triggered. See "Configuring alerts" on page 614.

To configure address resolution for a new data source

1 Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

  See "Adding a data source " on page 492.

2 On the **Add Directory Data Source - Directory Data Source Functions** page, check **Address resolution**.

3 In the **Test email address** field, type an email address that can be used to test and validate your configuration.

Symantec recommends that you test a variety of addresses that include a primary address, an alias address, a distribution list address, and an invalid address. If the test produces unexpected results, click **Customize Query** to verify your settings and directory data.

Since test lookups are conducted as real-time queries (as opposed to cached queries) Symantec recommends that you test with one of your smaller distribution lists.

4 Click **Test Query** to validate your address resolution query data using the test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot verify connectivity from attached scanners to your LDAP server.

If your query is successful, you can click the icon next to the **Test Query** option to display all query results. This test reports all email addresses and user preferences that are associated with the test email address. If the recipient is a distribution list, this information is provided for all users belonging to that distribution list.

Test results reflect only the data source being tested. The test does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

5 If you want to examine or modify the default settings for your queries, click **Customize Query**. Create custom queries to more accurately reflect your system's configuration.

See “Creating and testing a custom address resolution query” on page 523.

6 Click **Test Query** to validate your group listing query data using the test email address.

Test results reflect only the data source being tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

7 If you want to examine or modify the default settings for your group listing query, click **Customize Query**.

See “Creating and testing a custom address resolution query” on page 523.

8 You can configure your data source for multiple functions or click **Next** to review your changes and save your deployment.

See “Enabling functions on a new data source” on page 498.

Creating and testing a custom address resolution query

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See “About data source queries” on page 484.

Symantec recommends that you test all queries before deploying a new data source.

To configure a custom address resolution query for a new data source:

1. Add your data source and configure your server integration settings in the **Directory Integration Settings - LDAP Server Configuration** page then click **Next**.

   See “Adding a data source ” on page 492.

2. In the **Add Directory Data Source - Directory Data Source Functions** page, check **Address resolution**, and then click **Customize Query**.

   See “Enabling functions on a new data source” on page 498.

3 Provide the following information:

Base DN

Provide a **Base DN** for the custom query.

A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.

Primary email attribute

Provide a **Primary email attribute** for the address resolution query, if desired.

The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Primary email attribute

In the **Primary email attribute** field, provide the attribute in your LDAP schema that is used to store the primary email address information for the query.

If you provide multiple primary email attributes, the system selects the first attribute (based on alphabetical order) to use as the primary attribute for query purposes. The subsequent values appear as aliases, but only if the primary attribute and alias attribute names defined for the data source function are the same.

If you specify a primary email attribute of "proxyAddresses", the directory data service automatically identifies the attribute value that is prepended with "SMTP:" as the address. You do not need to specify this prefix in the field.

Email alias attribute

In the **Email alias attribute** field, the attribute in your LDAP schema that is used to store the email alias address information.

Distribution list object classes

In the **Distribution list object classes** field, list the object classes in your LDAP schema to be used to identify distribution list entries.

Child membership attributes

In the **Child membership attributes** field, provide the names of the attributes, separated by semicolons, that are in your schema used to define members of a group.

If you do not provide a child membership attribute, distribution lists, and groups are not expanded. Choosing not to expand groups does create a performance benefit. It also means, however, that policies can only be applied to the email address of the recipient since LDAP group membership are not evaluated. Indirect policy groups through email aliases are still honored.

Restore Defaults

Click **Restore Defaults** to remove your edits to the address resolution query fields and replace them with the default values.

Test email address

Provide a **Test email address** that can be used to test and validate your query.

Test Query

To validate the defined address resolution query against the data source click **Test Query**.

This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot verify connectivity from attached scanners to your LDAP server.

If your query is successful, you can click the icon next to the **Test Query** option to display all query results. This test reports all email addresses and user preferences that are associated with the test email address. If the recipient is a distribution list, this information is provided for all users belonging to that distribution list.

Test results reflect only the data source being tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Query filter (optional)

Provide a **Query filter (optional)** for the group listing query, if desired.

The attribute describes the email address or attribute element to be searched and the token describes the parameters that are used to return data.

For example, for a SunONE data source, you might use the following query filter to identify all groups within the directory:

```
(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames))
```

If this field is left blank, then groups are not listed on the Add Policy Groups page.

Group name attribute (optional)	In the **Group name attribute (optional)** field, provide the attribute from your schema that is used as the display name for the groups that are returned. A representative sample of groups is returned. If this field is left blank, then the groups created for association with policy groups are listed on the **Add Policy Groups** page by the DN name only. See "About policy groups" on page 315.
Restore Defaults	Click **Restore Defaults** to remove your edits to the group listing query configuration fields and replace them with the default values.
Test Query	Click **Test Query** to validate the defined group listing query against the data source instance that is running on the Control Center host. Test results reflect only the data source tested. The query returns a representative sample of groups found in the directory and is conducted against the directory data service instance that is running on the Control Center host. It cannot be used to verify connectivity from attached scanners to your LDAP server.

4 Click **Save** to return to the **Add Directory Data Source - Directory Data Source Functions** page.

See "Enabling functions on a new data source" on page 498.

About the directory data cache

Symantec Brightmail Gateway provides caching functionality that stores directory data from previous requests. This lets the directory data service process requests faster by using this cached data instead of the more time-consuming LDAP directory query.

The cache is enabled by default and entries are added to the cache as they are received and validated. When the directory data service receives a request, it searches all caches until the recipient is found or there are no more caches to search. If the recipient is not found in a cache, the directory data service then searches the data source directories. If the recipient is found in a data source, it adds the recipient to the associated cache.

If the recipient is not found in a data source then the directory data service adds a cache entry to the invalid recipient cache for that data source.

Configure an appropriate cache size based on your system's needs and memory availability. Symantec recommends that you set this value equal to or greater than the number of users and groups in your environment. This list should include distribution lists, contacts, and public folders.

You can control aspects of each of your caches as follows:

- Set the cache size and entry life span using the Advanced Settings page.
  See "Editing advanced settings for a data source" on page 558.
  Use the **Minimum TTL** (Time to Live) and **Maximum TTL** settings on the Advanced Settings page to find the right balance of refresh frequency versus data freshness.
- Set alerts to notify you when your cache size has been exceeded.
  See "Configuring alerts" on page 614.
- If your directory is very large or if you have a slow LDAP connection, you can preload your cache to avoid a backup in your mail queues caused by the cache building process.
  See "About preloading your directory data cache" on page 528.
- Disable your cache. Note that if the cache is disabled, then the LDAP host is queried for each request. Disabling the cache can increase processing time.
  See "Editing advanced settings for a data source" on page 558.

Enabling cache persistence

Cache persistence saves the existing cache to a file any time the directory data service is stopped as a result of the Symantec Brightmail Gateway being shut down or rebooted. This cache can then be read at reboot to prevent the loss of cache data. The cache is enabled by default.

See "About the directory data cache" on page 526.

To enable or disable cache persistence for a data source

1. Click **Administration > Settings > Directory integration**.
2. On the **Directory Integration Settings - Directory Data Sources** page, check **Enable cache persistence** to enable cache persistence, or uncheck the box to disable cache persistence.

   You are prompted to **Cancel** or **Change** your distribution list expansion setting.
3. Click **Change**.

   Cache persistence is enabled or disabled across all data sources.

   See "About modifying existing data sources" on page 533.

About preloading your directory data cache

If you have a large number of LDAP entries or a slow LDAP connection, you can preload your directory caches to prevent mail from backing up in the system during the initial directory cache building process. Under most circumstances, cache preload is not necessary because the caches can be built gradually during the normal course of operation with satisfactory system performance.

Cache preload is performed on a per-function basis and is supported for the following functions:

- Quarantine address resolution
- Recipient validation
- Routing
- Address resolution

For all functions except routing, a single input file populates all data sources that are enabled for that function. For routing, a separate input file is required to populate each data source.

If you disable a data source, the cache for all enabled functions for that data source is cleared. If you create a new data source or enable a disabled data source you must preload the caches again.

See “Preloading your data cache” on page 528.

See “About the directory data cache” on page 526.

Preloading your data cache

Caches are built gradually during the normal course of operation with satisfactory system performance. If your system configuration is complex or your data sources are large, you can preload your directory caches to prevent mail from backing up in the system during the initial directory cache building process.

You must upload all your input files before you begin the preloading process. Each import file must be comprised of newline-separated email addresses and use ASCII or UTF8 encoding.

To preload a data source cache

1. Select **Administration > Settings > Directory Integration**.
2. Click **Preload Cache**.

3 On the **Preload Configuration and Status** page, browse to the Input file you want to use to preload your cache.

This file must be comprised of newline-separated email addresses using either ASCII or UTF format.

4 In the **Function** field, select the function for which the uploaded file is used.

Cache preload is performed in a per-function basis and is supported for the following enabled and configured functions: quarantine address resolution, recipient validation, routing, and address resolution. For all functions except routing, a single input file is shared across all data sources that are enabled for that function. For routing, you must associate your data source input file or files with each routing function.

5 Click **Upload**.

6 Continue to upload your files until you have uploaded all of the files you want to preload. You must upload the data source file for each function. If the list of email addresses is shared across functions, you must upload the same input file for each function.

7 When you have uploaded all the files, click **Start Preload**.

You can return to this page at anytime to check your status until the Control Center restarts. The Preload % Complete column shows updates for each function as the page is refreshed automatically. You can also manually refresh the status by clicking **Refresh**.

The preload status appears as follows:

- In progress
  The preloading process has begun but is not yet completed.
- Successfully completed
  All files have been successfully loaded.
- Cancel initiated
  Cancel was initiated but not completed.
- Canceled
  The file loading process has been stopped.
- Failed
  The file load process was unsuccessful.

You can return at anytime to check your status until the Control Center restarts. You can also consult the Control Center logs for additional information.

See “About logs” on page 625.

Chapter 17

Managing directory data integration

This chapter includes the following topics:

- About expanding distribution lists and preserving recipient addresses
- About modifying existing data sources
- Best practices for security
- Best practices for scalability

About expanding distribution lists and preserving recipient addresses

Distribution list expansion lets you resolve addresses and apply filter policies to policy group members within distribution lists. Preserve recipient address stops the rewrite from alias to primary that occurs with address resolution and quarantine address resolution.

If **Preserve recipient addresses** and **Enable distribution list** expansion are both checked, the **Preserve recipient addresses** functionality overrides the **Enable distribution list** expansion functionality, and thus distribution lists are not expanded but instead retain the recipient email address.

See "Enabling distribution list expansion for your data sources" on page 532.

See "Preserving recipient addresses for your data sources" on page 533.

Table 17-1describes resolution output based on the **Preserve recipient addresses** and **Expand distribution list** configuration combinations when address resolution is enabled.

Table 17-1 Address resolution based on recipient address preservation and distribution list expansion input

Input type	Preserve recipient addresses (disabled by default)	Expand distribution list (enabled by default)	The address resolves to
user's alias address	disabled	enabled	user's primary address
distribution list alias address	disabled	enabled	full list of primary user addresses within the distribution list
user's alias address	disabled	disabled	user's primary address
distribution list alias address	disabled	disabled	distribution list primary address
user's alias address	enabled	enabled	user's alias address
distribution list alias address	enabled	enabled	distribution list alias
user's alias address	enabled	disabled	user's alias address
distribution list alias address	enabled	disabled	distribution list alias
distribution list primary address	enabled	enabled	distribution list primary

Enabling distribution list expansion for your data sources

Distribution list expansion lets the system use your LDAP directories to resolve addresses within distribution lists and apply filter policies to policy groups more consistently.

This field is enabled by default. If you disable this feature, you cannot apply policies to individuals within distribution lists when distribution lists are used as members of policy groups.

If **Preserve recipient addresses** and **Enable distribution list expansion** are both checked, the features may affect each other and produce unintended consequences.

See "About expanding distribution lists and preserving recipient addresses" on page 531.

See “Editing the LDAP server configuration for a data source” on page 538.

To enable or disable distribution list expansion

1. Click **Administration > Settings > Directory Integration**.
2. On the **Directory Integration Settings - Directory Data Sources** page, check **Enable distribution list expansion** to enable distribution list expansion, or uncheck the box to disable distribution list expansion.

   You are prompted to **Cancel** or **Change** your distribution list expansion setting.
3. Click **Change**.

   Distribution list expansion is enabled or disabled across all data sources.

   See “About modifying existing data sources” on page 533.

Preserving recipient addresses for your data sources

Recipient address preservation retains addresses by stopping the system from rewriting addresses from alias to primary, as normally occurs when address resolution is enabled. Enabling this field affects your address resolution functionality by preventing this behavior.

If **Preserve recipient addresses** and **Enable distribution list expansion** are both checked, the features may affect each other and produce unintended consequences.

See “About expanding distribution lists and preserving recipient addresses” on page 531.

See “Enabling distribution list expansion for your data sources” on page 532.

To preserve recipient addresses for a data source

1. Click **Administration > Settings > Directory Integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page, check **Preserve recipient addresses**. You are prompted to **Cancel** or **Change** your recipient address preservation setting.
3. Click **Change**.

   See “About modifying existing data sources” on page 533.

About modifying existing data sources

Once you create a data source, you can modify certain aspects of the source configuration, functions, and queries. Before you modify a saved data source, you should carefully assess any unintended effects that your changes may have on your mail configuration.

Editing a data source can affect your system configuration as follows:

- When a data source is disabled or deleted, the cache for that data source is cleared. The cache is then discarded and repopulated if the data source is enabled or added again in the future.
- If you change a query or any setting that alters query results while mail is being processed, the directory data service discards the existing cache, and a new one is built as the messages are received.
- If you reconfigure your LDAP settings to use a different directory type, you query data is not updated and therefore may become outdated. You must reconfigure all queries. You can alternately use the **Restore Defaults** option for each enabled function to populate the correct defaults for each query based on the new directory type.

Modifying or attempting to modify a data source so that functions are no longer available can affect your system as follows:

Table 17-2

Task	Description
Authentication	■ Disabling an authentication data source automatically disables the authentication query. If this data source is configured for end-user quarantine, users are no longer able to log into the Control Center and access the end-user quarantine. The **Administrator-only Quarantine** field in the **Quarantine Settings** page is checked and cannot be edited unless a data source is made available. See “Configuring Spam Quarantine for administrator-only access” on page 260. ■ When an authentication source is deleted or disabled, quarantine address resolution no longer occurs for that source. Emails are no longer quarantined under a single account that is tied to a single address. See “Setting up invalid recipient handling” on page 130. ■ When an SMTP authentication source is disabled or deleted, SMTP authentication no longer occurs for that source. Users who attempt to send out mail through a queue that requires SMTP authentication are no longer able to do so.

Table 17-2 *(continued)*

Task	Description
Recipient validation	■ Deleting or disabling a recipient validation source prevents the data directory from applying invalid recipient handling. You must have at least one data source configured for recipient validation to drop or reject invalid recipients. If your system is configured to drop or reject, you are not able to delete the last recipient validation data source. See "Setting up invalid recipient handling" on page 130. ■ Deleting or disabling all recipient validation sources affects the behavior of your directory harvest attack configuration. If there is not at least one recipient validation data source, directory harvest attack only acts upon non-local recipients, which are counted as invalid recipients for purposes of statistics. Directory harvest attack does not evaluate local messages since there is no directory data available for validation, and the system attempts to deliver all local messages normally.
Address resolution	■ Deleting or disabling an address resolution source prevents the resolution of alias addresses to the primary address, as well as the expansion of distribution lists for entries within that source. ■ Groups associated with a deleted or disabled data source are no longer available for selection in the **Add Policy Groups** page. See "Creating a policy group" on page 316.
Routing	A routing data source cannot be deleted or disabled if an enabled domain is associated with the source. To modify or remove that data source you must first remove the domain associations using the **Domains** page. See "Adding or editing domains" on page 117.

You can modify data sources as follows:

- You can change the LDAP information for a data source.
  See "Editing the LDAP server configuration for a data source" on page 538.
- You can delete a data source.
  See "Deleting a data source" on page 537.
- You can disable or enable a data source.
  See "Disabling or enabling a data source" on page 536.
- You can modify your advanced settings data, such as enabling or disabling cache functionality, setting TTL or TTL spread, and configuring client settings.
  See "Editing advanced settings for a data source" on page 558.

- You can disable, enable, or modify the Control Center authentication, quarantine address resolution function or SMTP authentication and related queries for a data source.
  See "Enabling or editing the authentication function" on page 540.
- You can disable, enable, or modify the address resolution function and query for a data source.
  See "Enabling or editing the address resolution function" on page 547.
- You can disable, enable, or modify the routing function and queries for a data source.
  See "Enabling or editing the routing function" on page 552.
- You can disable, enable, or modify the recipient validation function and queries (including the group listing query) for a data source.
  See "Enabling or editing the recipient validation function" on page 556.
- You can enable or disable distribution list expansion for a data source.
  See "Enabling distribution list expansion for your data sources" on page 532.
- You can enable or disable recipient address preservation for a data source.
  See "Preserving recipient addresses for your data sources" on page 533.

Disabling or enabling a data source

You can disable a saved data source to make a data source unavailable to the system without losing your configuration. Disabling the data source saves configuration time if you need to enable the data source again at a later date. The disable feature is also useful if you need to temporarily take the source offline for troubleshooting purposes. Although settings are saved, the cache for a disabled data source is cleared and is rebuilt if enabled in the future.

Depending on your system configuration, disabling a data source can affect the functionality of other features. Affected functionality can include invalid recipient handling, directory harvest attack, and end-user preferences. In the case of certain feature dependencies, the system does not let you disable a saved source.

See "About modifying existing data sources" on page 533.

To disable or enable a data source

1. Click **Administration > Settings > Directory Integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page, check the boxes for the data sources you want to enable or disable.
3. Click **Disable** to disable the selected data sources, or **Enable** to enable the selected data sources.

Deleting a data source

You can delete an existing data source if you need to make the data source or directory information unavailable to Symantec Brightmail Gateway.

Depending on your existing configuration, modifying a saved data source can affect or disable existing functionality. You also may be prevented from deleting some data sources based on such configuration. Be aware of any unintended consequences your actions may have on your mail configuration before you attempt to delete your data source.

See "About modifying existing data sources" on page 533.

To delete a data source

1. Click **Administration > Settings > Directory Integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page and check the boxes for the data sources that you want to delete.
3. Click **Delete**.

Editing a data source

Depending on your system configuration, editing a data source can affect the functionality of other features.

In the case of certain feature dependencies, the system does not let you disable a saved source or some of the source's configured functions.

Note: If you disable a function for a data source, the settings and query configuration information for that function are not saved. Query settings must be reconfigured if the function is enabled at a later date.

See "About modifying existing data sources" on page 533.

To edit a data source

1. Click **Administration > Settings > Directory Integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page, check that the data source you want to modify.
3. Click **Edit**.
4. Click one the following tabs to edit your data source settings and configuration:
   - LDAP server
     See "Editing the LDAP server configuration for a data source" on page 538.

- Authentication
  See "Enabling or editing the authentication function" on page 540.
- Recipient validation
  See "Enabling or editing the recipient validation function" on page 556.
- Routing
  See "Enabling or editing the routing function" on page 552.
- Address resolution
  See "Enabling or editing the address resolution function" on page 547.
- Advanced
  See "Editing advanced settings for a data source" on page 558.

Editing the LDAP server configuration for a data source

If the LDAP server configuration changes for an existing data source, you must update that information in Symantec Brightmail Gateway.

Before you modify a deployed data source, be sure to first assess the impact your changes may have on any related configuration or processes.

Any modification that would affect query results cause the cache to be cleared and subsequently rebuilt, which can slow down mail delivery. Use the test functionality to help diagnose issues before the data source begins production service.

See "About modifying existing data sources" on page 533.

To edit LDAP server configuration for a data source

1. Click **Administration > Settings > Directory integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page, check the box for the data source you want to edit.
3. Click **Edit** and then click the **LDAP Server** tab.
4. You can edit any of the following fields:

Field	Description
Data source name	You can edit the unique name of the data source.

Field	Description
Directory type	You can change the directory type the LDAP source uses. If you change the directory type for a saved data source, your query default values are not updated accordingly. You must individually reconfigure your queries or use the **Restore Defaults** option for each function. The available choices are as follows: ■ Active Directory ■ Active Directory Global Catalog ■ iPlanet/Sun ONE/Java Directory Server ■ Domino ■ Other (can be any LDAPv3 compliant directory type) Though Symantec Brightmail Gateway is compatible with any LDAPv3 directory, it is specifically designed to support configurations using the four specific directory types listed. If you select "other" for your directory type, consult your directory documentation for more information. If you use the Active Directory Global catalog, you may experience authentication issues depending on your query configuration. These issues can be resolved by replicating the Global Catalog.
Host name	Modify the host name or IP address of the LDAP server.
Port	Modify the TCP/IP port that is used to access the LDAP server.
Enable SSL	Enable or disable SSL on all connections to the LDAP server host. Verify your port setting when you modify this field.
Anonymous bind	Check to let the directory data service connect to the LDAP server without providing specific user ID and password information.
Use the following	Check if you want to configure specific login credentials for authentication to the LDAP server.
Name (Bind DN)	The distinguished name (DN) that is used for authenticating to the LDAP server. **Note:** For an Active Directory or Global Catalog server, you can optionally use the full DN, the NetBIOS and logon name (NetBIOS\SAM Account Name), or the User Principal Name.
Password	The password to be used for authenticating to the LDAP server. A password is required if you checked **Use the following**.

5. When you are finished editing the desired fields, click **Test Login** to connect to the client and ensure that your LDAP connection settings are valid.

   This test only verifies that the LDAP server can be reached and that the provided account has read access to the root of the directory data tree on a directory server. A successful test result does not guarantee that the credentials can succeed elsewhere in the directory. This is particularly true for anonymous access. You should verify access before deployment.

6. Click **Save**.

Enabling or editing the authentication function

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See "About data source queries" on page 484.

Depending on your system configuration, editing an authentication data source can affect the functionality of other features. You may not be able to disable an authentication source if other features are dependent upon it.

See "About modifying existing data sources" on page 533.

Changing the authentication type for a data source can have the same effects as disabling it.

For example, assume that you configure end-user quarantine for a data source that is enabled for Control Center authentication. If you later enable the data source for "SMTP authentication only," end-user quarantine is automatically disabled because that functionality is dependent upon Control Center authentication. If, however, you enable the data source for "Control Center and SMTP authentication," you enable SMTP functionality and preserve all configuration that is based on Control Center authentication.

When a data source is enabled for authentication, the system can:

- Authenticate end-user login to Spam Quarantine.
- Let end users configure personal Good and Bad Sender lists based on email address (address resolution also required).
- Validate email recipients for messages that are stored in Spam Quarantine.
- Let end users configure personal language settings (address resolution also required).
- Authenticate users who use SMTP authentication to send messages.

See "About using the authentication function with your data source" on page 487.

See “About modifying existing data sources” on page 533.

See “Using SMTP authentication” on page 145.

To edit authentication for a data source

1 Click **Administration > Settings > Directory integration**.

2 In the **Directory Integration Settings - Directory Data Sources** page, check the data source that you want to edit.

3 Click **Edit**.

4 Click the **Authentication** tab.

5 Check **Enable Authentication** to enable authentication, or uncheck it to disable the data source for authentication.

When you remove a function from a data source by disabling the function for that source, the settings and query configuration information for that source are not saved and must be reconfigured.

6 In the drop-down list, select the type of authentication you want to configure for your data source.

Note: Editing from one exclusive authentication type to another authentication type can disable certain dependent features in your messaging system. If you are unsure as to how a change might affect your system, Symantec recommends that you select "Control Center and SMTP authentication" to preserve existing functionality.

Control Center authentication only	To authenticate end users for accessing quarantine or setting end-user preferences only (authentication is also required). You must configure Control Center authentication and quarantine address resolution details.
SMTP authentication only	To enable remote users to send email without a VPN connection by using the SMTP authentication protocol. You must configure SMTP authentication details.
Control Center and SMTP authentication	To enable features for both Control Center and SMTP authentication. You must configure Control Center authentication, quarantine address resolution, and SMTP authentication details. You can check **Share Control Center and SMTP Authentication query details** if you want to share your query details.

7. If you have selected Control Center and SMTP authentication you can check **Share Control Center and SMTP Authentication query details** if you want to use the same query details for Control Center and SMTP authentication.

8. Provide the following information for the Control Center authentication query if you selected **Control Center authentication only** or **Control Center and SMTP authentication** (but did not check **Share Control Center and SMTP Authentication query details**) as the authentication type:

   Test user name

   Type a user name that can be used to test and validate your authentication configuration.

   Symantec recommends that you verify your settings using a variety of input designed to produce both successful results and unsuccessful results. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

   Test password

   Type the password for the test user name.

   Test domain (optional)

   Type a NetBIOS domain for the provided test credentials.

   You can optionally specify a NetBIOS domain for the authentication query test. This field is required if your **Test user name** is not unique across multiple domains within the scope of your data source. This field is only available if you use Active Directory or Active Directory Global Catalog.

   Test Query

   Click to validate the quarantine address resolution query using the test information provided.

   This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server.

   Test results reflect only the data source tested. Results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.

   Customize Query

   Click to display the custom query fields.

   You can examine the default settings for your query or create a custom query to change scope and support custom directory schemas.

   See "About data source queries" on page 484.

NetBIOS domain name (optional)

Provide a list of NetBIOS domain names that the system can use to resolve login ambiguities.

This field is optional but recommended if all of the following are true:

- Your directory type is Active Directory or Active Directory Global Catalog
- Your directory uses more than one NetBIOS domain
- User name uniqueness is not guaranteed.

Base DN

Provide a **Base DN** for the routing query.

A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.

Query filter

Provide a **Query filter** for the authentication query, if desired.

The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Primary email attribute

This field is not relevant to SMTP authentication. It is required for Control Center authentication.

In the **Primary email attribute** field, type the attribute in your LDAP schema that is used to store the primary email address information for the authentication query.

If you specify a primary email attribute of "proxyAddresses," the directory data service automatically identifies the attribute value that is prepended with "SMTP" as the address. There is no need to specify this prefix in the Primary email attribute field.

Authentication Method	Choose your **Authentication Method**. Click **Simple bind** to attempt login using the user-provided password. Or, you can use password fetching, which fetches the password from the data source and compares it to the user-provided password. Click **Password attribute** to choose password fetching. You can use the default password attribute or change it. You can optionally specify a **Default hash type** for password fetching. If you choose **none**, your directory data must prepend the password attribute with a prefix indicating the scheme, such as **{PLAIN}**. If you choose **plaintext**, a scheme prefix is not needed. Symantec Brightmail Gateway supports SMTP authentication through LDAP using simple bind for all supported LDAP directory types. For SMTP authentication through LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported.
Hide query	Click to hide the custom query fields.
Restore Defaults	Click to remove your edits to the query configuration fields and replace them with the default values.

9 Provide the following information for the quarantine address resolution query if you selected **Control Center authentication only** or **Control Center and SMTP authentication** as the authentication type. You can check **Share Control Center and SMTP Authentication Query details** to use your Control Center information:

Test email address

In the **Test email address** field, type an address that can be used to test and validate your quarantine address resolution configuration. Symantec recommends that you test your data source using a combination addresses including a user address, distribution list address, alias address, and an invalid address.

Test Query

Click **Test Query** to validate the quarantine address resolution query using the provided test email address.

This test is conducted against the directory data service instance that is running on the Control Center host. It cannot be used to verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source tested. This test does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Customize Query

Click to display the custom query fields. You can view the default query or customize the query to change scope and support custom directory schemas.

See “About data source queries” on page 484.

Query filter

In the **Query filter** field under Quarantine Address Resolution Query, type a custom filter for the quarantine address resolution query, if desired. The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Distribution list object classes

In the **Distribution list object classes** field, type the object classes in your LDAP schema that should be used to identify distribution list entries.

Hide Query

Click to hide the custom query fields.

Restore Defaults

Click to remove your edits to the query configuration fields and replace them with the default values.

10 Provide the following information for the SMTP authentication query if you selected **SMTP authentication only** or **Control Center and SMTP authentication** as the authentication type:

Test user name

Type a user name that can be used to test and validate your authentication configuration.

Symantec recommends that you verify your settings using a variety of input designed to produce both successful results and unsuccessful results. If the test produces unexpected results, use **Customize Query** to verify your settings and directory data.

Test password

Type the password for the test user name.

Test Query

Click to validate the quarantine address resolution query using the test information provided.

This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server.

Test results reflect only the data source tested. It does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Customize Query

Click to display the custom query fields. Create a custom query to change scope and support custom directory schemas.

See “About data source queries” on page 484.

Base DN

Provide a **Base DN** for the routing query.

A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Custom query start

If you select "Customize" for the **Base DN** field, provide a **Custom query start**.

You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.

Query filter

Provide a **Query filter** for the authentication query, if desired.

The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.

Primary email attribute (optional)

This field is not relevant to SMTP authentication. It is required for Control Center authentication.

In the **Primary email attribute** field, type the attribute in your LDAP schema that is used to store the primary email address information for the authentication query.

Authentication Method

Choose your **Authentication Method**.

Click **Simple bind** to attempt login using the user-provided password.

Or, you can use password fetching, which fetches the password from the data source and compares it to the user-provided password. Click **Password attribute** to choose password fetching. You can use the default password attribute or change it.

You can optionally specify a **Default hash type** for password fetching. If you choose **none**, your directory data must prepend the password attribute with a prefix indicating the scheme, such as **{PLAIN}**. If you choose **plaintext**, a scheme prefix is not needed.

Symantec Brightmail Gateway supports SMTP authentication through LDAP using simple bind for all supported LDAP directory types. For SMTP authentication through LDAP using password fetching, all supported directory types except Active Directory, Active Directory Global Catalog, and Domino are supported.

Hide Query

Click to hide the custom query fields.

Restore Defaults

Click to remove your edits to the query configuration fields and replace them with the default values.

11 Click **Save**.

Enabling or editing the address resolution function

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See “About data source queries” on page 484.

Depending on your system configuration, editing or disabling the address resolution function for a saved data source can affect the functionality of other features.

See “About modifying existing data sources” on page 533.

To edit address resolution configuration for a data source

1 Click **Administration > Settings > Directory integration**.

2 In the **Directory Integration Settings - Directory Data Sources** page, check the boxes for the data source that you want to edit.

3 Click **Edit**.

4 Click the **Address resolution** tab.

5 You can edit any of the following fields:

Item	Description
Data Source Name	You can edit the unique name of the data source.
Enable Address Resolution	The address resolution function resolves alias and distribution list data to let you apply policies on a per-user basis by using LDAP based group memberships to define policies. See "About using the address resolution function with your data source" on page 490. If you remove a function from a data source by disabling the function for that source, settings and query configuration information for that source are not saved. This information must be reconfigured if the function is enabled in the future. See "About modifying existing data sources" on page 533.
Test email address	Provide an email address that can be used to test and validate your address resolution configuration.
Test Query	Click to test and validate the defined address resolution query using the provided test email address. This test reports all email addresses, distribution list memberships, and directory group memberships that are associated with the test email address. This test is conducted against the directory data service instance that is running on the Control Center host. It cannot be used to verify connectivity from attached scanners to your LDAP server. If your query is successful, you can click the icon next to the **Test Query** option to display all query results. This test reports all email addresses and user preferences that are associated with the test email address. If the recipient is a distribution list, this information is provided for all users associated with that distribution list. Test results reflect only the data source being tested. It does not provide information about the effects of other data sources or system settings such aliasing and masquerading.

Item	Description
Customize Query	Click to view or customize the default address resolution query. Create a custom query to change scope and support custom directory schemas. See "About data source queries" on page 484.
Base DN	Specify the name for the place in the directory from which to start searching for entries to authenticate. A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.
Custom query start	If you select "Customize" for the **Base DN** field, provide a **Custom query start**. You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.
Query filter	The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.
Primary email attribute	Specify the attribute in your LDAP schema that is used to store the primary email address information. If you specify a primary email attribute of "proxyAddresses," the directory data service automatically identifies the attribute value that is prepended with "SMTP" as the address. You do not need to specify this prefix in the **Primary email attribute** field.
Email alias attribute (Optional)	Specify the attribute in your LDAP schema that is used to store the email alias address information.
Distribution list object classes	List the object classes in your LDAP schema to be used to identify distribution list entries.

Item	Description
Child membership attributes	Provide the names of the attribute in your schema, separated by semicolons, that are used to define members of a group. If you do not provide a child membership attribute, distribution lists and groups are not expanded. This can create a performance benefit. It also means, however, that policies can only be applied to the email address of the recipient since LDAP group membership are not evaluated. Indirect policy groups through email aliases are still honored
Hide Query	Click to hide the custom query fields.
Restore Defaults	Click to remove your edits to the query configuration fields and replace them with the default values.
Test Query	Click to test and validate the group listing query using the provided test information. This test is conducted against the directory data service instance that is running on the Control Center host. It does not verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source tested.
Customize Query	Click to create a custom group listing query. Create a custom query to change scope and support custom directory schemas. See "About data source queries" on page 484.
Query filter (optional)	The query filter defines the rules the directory data service uses to search for groups in the directory. It is specified in standard LDAP query syntax. For example, for a SunONE data source, you might use the following query filter to identify all groups within the directory: `(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames))` If this field is left blank, then groups are not listed on the Add Policy Groups page. See "About policy groups" on page 315.

Item	Description
Group name attribute (optional)	Specify the attribute in your schema that is used as the display name for the groups that are returned. See "About policy groups" on page 315.
Hide Query	Click to hide the custom query fields.
Restore Defaults	Click to remove your edits to the query configuration fields and replace them with the default values.

6 Click **Save**.

Enabling or editing the routing function

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See "About data source queries" on page 484.

Depending on your system configuration, editing the routing function for a saved data source can affect the functionality of other features.

See "About modifying existing data sources" on page 533.

To edit the routing function for a data source

1 Click **Administration > Settings > Directory integration**.

2 In the **Directory Integration Settings - Directory Data Sources** page, check the boxes for the data source you want to edit.

3 Click **Edit**.

4 Click the **Routing** tab.

5 You can edit any of the following fields:

Item	Description
Data Source Name	You can edit the unique name of the data source.

Item	Description
Enable Routing	The routing function uses the information provided by your data source to route messages to addresses and domains. See "About using the routing function with your data source" on page 489. You cannot delete a routing source that has been associated with a domain. If you want to disable such a routing source, you must remove the association on the **Domains** page. When you remove a function from a data source by disabling the function for that source, the settings and query configuration information for that source are not saved. This information will need to be reconfigured if the function is enabled in the future. See "About modifying existing data sources" on page 533.
Test email address	Provide an email address that can be used to test and validate your routing configuration.
Test Query	Click to test and validate the defined routing query using the provided test email address. This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot be used to verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.
Customize Query	Click to expose the custom query fields. Create a custom query to change scope and support custom directory schemas. See "About data source queries" on page 484.
Base DN	Specify the name for the place in the directory from which to start searching for entries to authenticate. A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.

Item	Description
Custom query start	If you select "Customize" for the **Base DN** field, provide a **Custom query start**. You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.
Query filter	The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.
Alternate address attribute	If you want to enable per-user, directory-based alternate address routing, provide an **Alternate address attribute**. When an alternate email address attribute is supplied, mail bound for the original user is rerouted to the new email address. **Note:** To save your routing source, you provide an **Alternate address attribute**, an **Alternate mailhost attribute**, or both. If you configure both an **Alternate address attribute** and an **Alternate mailhost attribute**, the directory data service resolves all alternate email addresses before it resolves the alternate mailhost values. See "About using the routing function with your data source" on page 489.

Item	Description
Alternate mailhost attribute	If you want to enable directory-based alternate mail host routing, provide an **Alternate mailhost attribute** that corresponds to the attribute in the directory schema that you have selected to store this information. The mail host should be specified as either an IP address or a fully qualified host name. A port can be specified by using a colon followed by the port, for example ":25". By default, the SMTP protocol uses port 25 for email transmission. You can use the following valid attributes: ■ 10.32.100.64 ■ 10.32.100.64:25 ■ myothermailserver.mycompany.com ■ myothermailserver.mycompany.com:25 If an LDAP entry has multiple alternate mailhost attribute values, the MTA randomly selects one of the values to use each time a message is delivered to that recipient. **Note:** To save your routing source, you provide an **Alternate address attribute**, an **Alternate mailhost attribute**, or both. If you configure both an **Alternate address attribute** and an **Alternate mailhost attribute**, the directory data service resolves all alternate email addresses before it resolves the alternate mailhost values. See "About using the routing function with your data source" on page 489. For information about the other routing methods of destination host routing or MX lookup, consult the domain documentation. See "Adding or editing domains" on page 117.
Perform MX Lookup on transport value	Allows MX lookup for the specified transport.
Hide Query	Click to hide the custom query fields.
Restore Defaults	Click to remove your edits to the query configuration fields and replace them with the default values.

6 Click **Save**.

Enabling or editing the recipient validation function

You can customize your queries by configuring custom attributes and filters to suit your system's configuration. Use the custom query to change scope and support unique directory schemas.

See "About data source queries" on page 484.

Depending on your system configuration, editing the recipient validation function for a saved data source can affect the functionality of other features.

See "About modifying existing data sources" on page 533.

To edit the recipient validation function for a data source

1. In the Control Center, click **Administration > Settings > Directory Integration**.
2. In the **Directory Integration Settings - Directory Data Sources** page, check the box for the data source you want to edit.
3. Click **Edit**.
4. Click the **Recipient Validation** tab.
5. To enable recipient validation on your data source, check **Enable Recipient Validation**.

6 You can edit any of the following fields:

Item	Description
Data Source Name	You can edit the unique name of the data source.
Enable Recipient Validation	Check **Enable Recipient Validation**. Recipient validation works with other features within Symantec Brightmail Gateway to help you identify and manage invalid recipients and messages sent to those recipients. See "About using the recipient validation function with your data source" on page 488. You cannot disable or delete the last recipient validation source if you have configured invalid recipient handling to drop messages or reject invalid recipients. You must set invalid recipient handling to accept all messages to delete or disable the last recipient validation source. When you remove a function from a data source by disabling the function for that source, the settings and query configuration information for that source are not saved. This information must be reconfigured if the function is enabled in the future. See "About modifying existing data sources" on page 533.
Test email address	Enter an email address that can be used to test and validate your recipient validation configuration.
Test Query	Click to test and validate the defined recipient validation query using the provided test email address. This test is conducted against the directory data service instance that is running on the Control Center host. The test cannot be used to verify connectivity from attached scanners to your LDAP server. Test results reflect only the data source tested. Test results do not provide information about the effects of other data sources or system settings such aliasing and masquerading.
Customize Query	Click to expose the custom query fields. You can review the default query information or create a custom recipient validation query to change scope and support custom directory schemas. See "About data source queries" on page 484.

Item	Description
Base DN	Specify the name for the place in the directory from which to start searching for entries to authenticate. A default value is provided, select "Customize" to provide your own values in the **Custom query start** field.
Custom query start	If you select "Customize" for the **Base DN** field, provide a **Custom query start**. You can customize the base DN to refocus the search to a specific part of the directory tree. **Custom query start** lets you configure the query to fit your particular needs and return data more quickly.
Query filter	The query filter instructs the directory data service to return data using attributes and tokens that describe a specific LDAP query syntax.
Hide Query	Click to hide the custom query fields.
Restore Defaults	Click to remove your edits to the query configuration fields and replace them with the default values.

7 Click **Save**.

Editing advanced settings for a data source

Use caution when editing your settings on a deployed data source, as changes may affect your current configuration and processes. For example, modifying cache settings may clear the cache, while modifying referral chasing may stop referrals from working.

See “About the directory data cache” on page 526.

See “About modifying existing data sources” on page 533.

To edit advanced settings for a data source

1 Click **Administration > Settings > Directory Integration**.

2 In the **Directory Integration Settings - Directory Data Sources** page, check the boxes for the data source that you want to edit.

3 Click **Edit**.

4 Click the **Advanced settings** tab.

5 You can edit any of the following fields:

Item	Description
Maximum connections	Specify the maximum number of client connections (bind operations) that can be created at one time. If this field is set to zero, connection pooling is turned off and a new connection is created for each request.
Minimum connections	Specify the number of connections that are added to the connection pool when the source services the first request. A value of zero indicates that connections are created only to service actual pending requests. Connections are released if they are idle for longer than the specified idle timeout.
Connection timeout	Specify the amount of time that should elapse before an attempt to connect to the LDAP server host is timed out and automatically ended. A value of zero indicates that the connection can never time out on the client side (the LDAP server can still close the connection).
Idle timeout	Specify how long a client connection can remain idle before the connection is automatically closed. A value of zero indicates that the connection can remain idle indefinitely without being closed (though the LDAP server can also close the connection).
Search timeout	Specify how long (in seconds) a request /search operation should run before the directory data service ends the operation and displays the partial results. The directory server can impose a search timeout lower than this value.

Item	Description
Page size	Determine the maximum number of initial entries to return when a query is successful. Setting this value too low can impede performance. A page size higher than the limit set by the LDAP server can cause the operation to fail when the page size exceeds that limit. The default server-side limit for this value varies according to directory type. To change this limit on your directory server, see your directory server documentation. To disable paged searching entirely, set this value to 0. **Note:** Make sure that the Name (bind DN) configured for the Administration Credentials used by the LDAP source has sufficient access rights to bypass search limits.
Chase referrals	A server may return a query response that suggests that the directory data service query another LDAP server. When this field is enabled, the directory data service follows such referrals when it executes queries. The directory data service uses the same bind credentials to connect to the referred server. If the referred-to LDAP server does not recognize the same bind credentials, a query can return an error.
Enable cache	The cache stores address entry data from previous requests. This allows the directory data service to process requests faster by using this cached data instead of consulting the LDAP server. See "About the directory data cache" on page 526. The cache is enabled by default.

Item	Description
Cache size	Specify the maximum number of entries that can be stored in the cache. When the cache size limit is exceeded, the least recently used entry is deleted to make room for a new entry. You should set the cache size based on your system's needs and memory availability. Symantec recommends that you set this value equal to or greater than the number of users and groups (including distribution lists, contacts, and public folders) in your environment. If you set the cache to zero entries, each new invalid recipient is trimmed as it is added to the cache. This results in a total size of 0 at any given time. See "About the directory data cache" on page 526.
Cache index size multiplier	Specify the size of the email address index in relation to the cache. This setting allows the index to store multiple aliases for each entry in the cache. For example, a multiplier of two would be twice the size of the cache and allow for an average of two aliases per cache entry.
Minimum cache TTL	Set the minimum Time to Live (TTL) for entries in the cache. When a cache TTL is reached, the entries in that cache expire and are refreshed upon query. A minimum value and maximum value for cache expiration creates a period of time over which these entries can be expired and refreshed. This helps pace your system's workload. For address resolution sources, this setting also determines how often the membership information is rebuilt for LDAP groups used in policies. See "About the directory data cache" on page 526.

Item	Description
Maximum cache TTL	Set the maximum Time to Live (TTL) for entries in the cache. When a cache TTL is reached, the entries in that cache expire and are refreshed upon query. A minimum value and maximum value for cache expiration creates a period of time over which these entries can be expired and refreshed. This helps pace your system's workload. See “About the directory data cache” on page 526.
Invalid recipient cache size	Provide the maximum number of entries that can be stored in the invalid recipient cache. When the directory data service cannot find an entry in the directory, the entry is considered invalid and is added to the invalid recipient cache. When the entry is queried in the future, load time is reduced because the directory data service checks the cache before it queries the directories. You should set the cache size based on your system's needs and memory availability. If you set the cache to 0, each new invalid recipient is trimmed as it is added to the cache, resulting in a total of 0 at any given time.
Invalid recipient cache TTL	Period of time an email address entry should be kept in the invalid recipient cache.
Clear Cache	Click to clear the cache of all entries.
Restore Defaults	Click to restore the default values to all Advanced Settings fields.

6 Click **Save**.

Best practices for security

When integrating the directory data service, it is very important to take appropriate precautions to safeguard your directory systems and data.

Symantec recommends that you implement some or all of the following to increase security for your deployment:

- Enable secure socket layer functionality by checking the SSL checkbox when setting up your LDAP server configuration.

  SSL can be used alone or can be used to augment password hashing to provide an additional layer of protection. Symantec recommends that you use SSL for your data sources whenever possible, and very strongly recommends it for authentication sources. You enable SSL when you set up or edit your data source's server configuration.

  See "Adding a data source " on page 492.

  Additionally, you must make sure that your directory servers are configured for SSL. See your directory administrator or directory documentation for more information.

- Use password hash algorithms to protect the passwords that are stored in your directory server.

  Although hash algorithms provide some additional protection, used alone they do not provide adequate security for user password data and should not be used as a substitute for SSL. SSL is essential for authentication sources to protect user name and password data in transit.

- Use an account that is restricted to read-only access for your directory data source administrator credentials.

  If possible, avoid anonymous bind for such accounts.

  See "Adding a data source " on page 492.

  See "Editing the LDAP server configuration for a data source" on page 538.

- Limit access to your server by using minimal access rights.

  Restrict access to namespaces, object classes, or entry attributes of interest. For example, in addition to configuring the Bind DN for a data source to have read-only access, the access control for the directory should only grant read access to the objectClass and the attributes used for email addresses, mailhosts, and group membership. You can find these attribute names by looking at the custom query information for your existing source.

  See "Enabling or editing the authentication function" on page 540.

  For directory-specific recommendations, consult your directory documentation.

- If your appliance has two NICs, you can place one on an internal subnet and configure the system only that one for your LDAP traffic.

  This isolates your LDAP servers and directory data from potential exposure to the Internet.

  See "Modifying Scanner configurations" on page 107.

Best practices for scalability

If using the directory data service in a large or distributed environment, consider the following best practices to improve system performance and scalability:

- Use directory data service caching functionality to improve throughput and reduce the load on your directory servers.

  You should set the cache size based on your system's needs and memory availability. Symantec recommends that you set this value equal to or greater than the number of users and groups in your environment. This number should include distribution lists, contacts, public folders, and any other LDAP entry that lists a deliverable email address or a username.

  See "Configuring data source advanced settings" on page 494.

  See "Editing advanced settings for a data source" on page 558.

- Use alert settings to manage your processes and cache.

  The **Monitor swap space utilization** alert triggers when swapping exceeds the specified utilization. Use the swap alert to make sure that your systems have adequate RAM for all Symantec Brightmail Gateway processes, including the directory data service cache.

  The default value to trigger this alert is 50% and can be modified to suit your needs. For high performance deployments that are adequately provisioned with memory, there should be little or no swap space utilization. Symantec recommends setting the swap space alert threshold to only a few percent for such deployments

  The **Undersized data source cache** alert lets you know immediately when you need to increase the cache size to support your data.

  See "Types of alerts" on page 615.

- Use the Symantec Brightmail Gateway directory cache preloader to complete the cache building process before directing your production mailflow to the system.

  For most deployments, caches can be built gradually through normal system activity with adequate system performance and preloading the cache is not necessary. For some deployments with very large directories or slow LDAP connections, however, the preloader can be used to avoid temporary performance problems that may occur while a very large cache is built.

  Perform this task offline (for example, during a maintenance window) as the preloaded caches are not available to your Scanners until the process is complete.

  See "About preloading your directory data cache" on page 528.

- Use the **Minimum TTL** (Time to Live) and **Maximum TTL** settings on the **Advanced Settings** page to find the right balance of cache refresh frequency versus data freshness.

  Adequate spread smooths out the load on your LDAP servers by randomizing the expiration of cache entries. Refreshes that occur too frequently can increase processing time, but failing to refresh often enough results in stale data. Work with your directory administrator to determine the right refresh rate for your system.

  See "About the directory data cache" on page 526.

  See "Configuring data source advanced settings" on page 494.

  See "Editing advanced settings for a data source" on page 558.

- Improve Symantec Brightmail Gateway system performance by turning off distribution list expansion.

  Disabling **Distribution list expansion** can significantly increase mail delivery throughput. If **Distribution list expansion** is disabled, however, distribution lists are not resolved into their individual members for policy evaluation.

  This means that mail sent to a distribution list is subject only to the policies associated with the distribution list itself (either through an email address or a distinguished name). The policies associated with its individual members are not applied, even if they have higher precedence.

  See "Enabling distribution list expansion for your data sources" on page 532.

- Limit the number and size of LDAP groups and distribution lists associated with your policy groups.

  If you clear a data source cache or make a configuration change to your policy groups or a directory data source, the Symantec Brightmail Gateway must reload group information from your directory. This can result in the growth of your inbound or outbound message queues.

  For most deployments this process takes only a few seconds and results in an insignificant queue backup if any at all. However, in cases where LDAP access is slow, or the policy groups references many thousands of LDAP users, a noticeable backup can occur. For best performance, Symantec recommends that you use the default group to implement the most common behavior and then assign specific policies to smaller groups as necessary.

  See "Creating a policy group" on page 316.

- To improve performance for queries, restrict the Base DN for your LDAP queries to cover only the data that is needed for your data source.

  The larger the scope of the query, the longer the searches take. Poor query performance for quarantine address resolution can lead to a backup in your delivery queues. Poor query performance for address resolution can cause inbound or outbound queues to back up.

  See "About data source queries" on page 484.

If your data source uses the Active Directory Global Catalog, be sure to configure the directory data service to use the global catalog port (default 3268) instead of the domain controller port (default 389).

See "Adding a data source " on page 492.

- Create read-only copies of firewalled LDAP servers and place them outside of the firewall to improve connection time.

  In an environment where Scanners hosts are located outside the firewall and LDAP servers reside inside the firewall, you can speed up connection and query times by setting up replicas of those LDAP servers outside the firewall and near the Scanners in the network.

  Figure 17-1 provides an example of a firewalled server configuration and how you might use an LDAP server replica to improve processing time.

  Your directory administrator can determine the best path for this action based on your system configuration.

Figure 17-1 Mail configuration example for a firewalled server

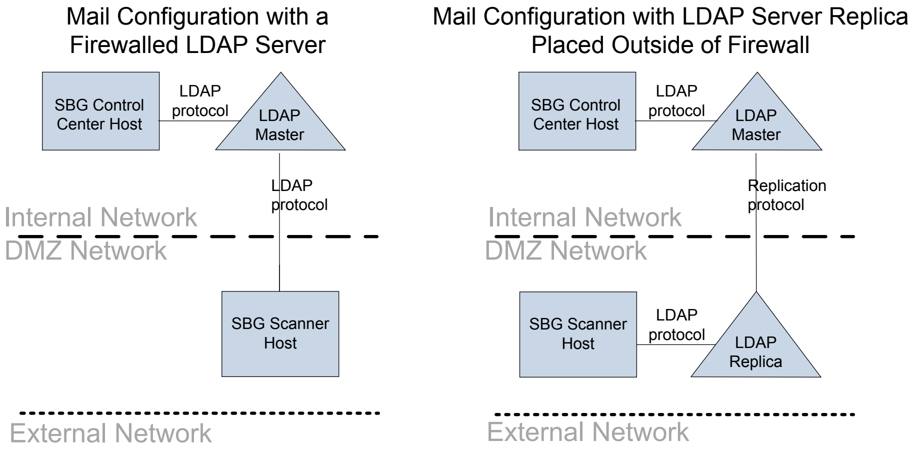

Chapter 18

Working with reports

This chapter includes the following topics:

About working with reports

Symantec Brightmail Gateway reporting capabilities provide you with information about filtering activity at your site. You can analyze consolidated filtering performance for all Scanners and investigate the spam attacks and virus attacks that target your organization.

See "Report types" on page 571.

You can create the following reports:

One-time reports	Create reports to run one-time on an as-needed basis. You can print, save, or email one-time the reports that you generate on demand.
Favorite reports	When you save reports to your Favorite Reports page, you can generate them on demand or schedule them to run automatically. You can print, save, or email the favorite reports that you generate on demand. However, automatically generated reports can only be emailed.

See "Creating and configuring reports" on page 569.

See "Saving favorite reports" on page 583.

See "Running reports on demand" on page 590.

See "Generating reports automatically" on page 590.

See "Saving generated reports" on page 595.

See "Printing generated reports" on page 595.

See "Emailing generated reports" on page 592.

Symantec Brightmail Gateway tracks some of the data that is included in reports automatically. You must specify the additional data that you want Symantec Brightmail Gateway to track and maintain.

See "Report types" on page 571.

See "Selecting the data to track for reports" on page 569.

Symantec Brightmail Gateway contains a utility that runs periodically to purge old report data and Dashboard data. You can specify how often and when data is automatically purged.

See "About purging report data" on page 596.

Selecting the data to track for reports

By default, Symantec Brightmail Gateway tracks data for several basic reports. Before you can generate other reports, you must tell Symantec Brightmail Gateway to track and store data appropriate for the report. Extended statistics such as Top Senders, Top Sender Domains, Instant Messaging, and Invalid Recipient data are only collected when the appropriate options are selected.

Note: Because the data storage requirements for some reports can be high, choose an appropriate length of time to store report data. In particular, the sender statistics usually consume a large amount of disk space.

See "Specifying how long report data is retained" on page 596.

See "About generated reports layout and data" on page 585.

Once you select the data to track, you can view the reports for that data after 24 hours.

See "Creating and configuring reports" on page 569.

To select the data to track for reports

1. In the Control Center, click **Administration > Settings > Reports**.
2. Under **Email Reports Data**, check the box beside the report data that you want to track.

   Note: To track extended statistics, ensure that all of the **Sender**-related check boxes are selected. To see the **Top Probe Accounts** report, check **Invalid Recipients**.

3. Click **Save**.

Creating and configuring reports

You can create and configure a report from the available report types. Then you can customize the report configuration to filter the data to include in the report. For example, you can specify time ranges and message flow direction. You must have at least full administration rights or rights to view or modify reports to create reports.

See "Report types" on page 571.

Before you create a report, ensure that you configure Symantec Brightmail Gateway to track the appropriate data for the report.

See "Selecting the data to track for reports" on page 569.

After you create the report, you can save it or run it.

See "Saving favorite reports" on page 583.

See "Running reports on demand" on page 590.

See "About generated reports layout and data" on page 585.

To create and configure reports

1. In the Control Center, click **Reports > View > Create a Report**.
2. In the **Report type** drop-down list, select a report category.
3. In a drop-down list beside **Report type**, select a specific report.

   This step does not apply to the **Executive Summary** report.

   Note: If you participate in the Symantec Probe Network, and want to see a **Top Probe Accounts** report, select **Invalid Recipients > Top Probe Accounts.**
4. If applicable, for the reports that filter on specific elements, click the drop-down menu to select the criteria.

   For certain reports, you can filter by **Sender name**. You can use the null sender address <> to filter for messages that do not contain Sender names.
5. If applicable, in the **Direction** drop-down list, select the message directions to include in the report.
6. For IM reports, in the **Network** drop-down list, select the network.
7. In the **Time range** drop-down list, do one of the following:

Select a preset range.	Select one of the following: ■ Past Hour ■ Past 24 hours ■ Past 7 days ■ Past 30 days

Specify a custom time range.	Do all of the following: ■ Click **Customize**. ■ Click in the Start Date field, then click the pop-up calendar and select the start date. ■ Click in the End Date field, then click the pop-up calendar and select the end date. You must enable JavaScript in your browser to use the pop-up calendar.

8 In the **Group By** drop-down list, select one of the following:

- Hour
- Day
- Week
- Month

9 Select one or more of the following:

- Graph
- Table

See “About generated reports layout and data” on page 585.

10 For the reports that rank results, in the **Entries** box, type the maximum number for each time range that is specified in the **Group by** drop-down list.

11 Select the columns that you want to display in the report table.

This option is only available for certain reports.

Report types

Symantec Brightmail Gateway contains predefined report types. When you configure a report, you specify the report type that you want to use. The tables in this section show the report types that you can choose. The last column in each table lists the type of data you want the system to track. You must specify one or more options before you can generate that report.

See “Selecting the data to track for reports” on page 569.

The categories of reports that you can generate are as follows:

- Summary
  See Table 18-1 on page 572.

- Content filtering
  See Table 18-2 on page 573.
- Email message
  See Table 18-3 on page 575.
- Instant messaging
  See Table 18-4 on page 576.
- Invalid Recipients
  See Table 18-5 on page 577.
- IP connection
  See Table 18-6 on page 578.
- Spam
  See Table 18-7 on page 578.
- Virus
  See Table 18-8 on page 581.

See "About generated reports layout and data" on page 585.

Table 18-1 describes the summary reports.

Table 18-1 Summary reports

Report	Description	Required data storage options
Executive	Overview of your security profile, which includes total messages and threats processed, and virus and content filtering summaries.	None
Content filtering	Overview of the content filtering violations and trends affecting your organization. Includes number of policies triggered, and percentage of policies triggered versus total processed messages.	None
Email Messages	Overview of email message threat counts and types of threats.	None
Instant Messages	There is no Summary report for Instant Messages.	None
Invalid Recipients	Overview of invalid recipient data.	None
IP Connections	Overview of the IP connections of email entering your system.	None

Table 18-1 Summary reports *(continued)*

Report	Description	Required data storage options
Spam	Overview of the email message spam.	None
Virus	Overview of the current viral threats to your organization. Includes a message summary, virus summary, suspect virus outcomes, and separate tables showing stats for known and potential viral threats.	None

Table 18-2 describes the available content filtering reports.

Table 18-2 Content filtering reports

Report	Description	Required data storage options
Summary	A summary of total detected content filtering violations.	None
Top Sender Domains	Domains from which the most content filtering matches have been detected. For each domain, the total messages processed and number and percentage of content filtering policies triggered are listed.	Sender domains
Top Senders	Email addresses from which the most content filtering matches have been detected. For each email address, the total messages processed and number and percentage of content filtering policies triggered are listed.	Senders, Sender domains
Top Sender HELO Domains	SMTP HELO domain names from which the most content filtering matches have been detected. For each HELO domain, the total messages processed and number and percentage of content filtering policies triggered are listed. Specify the maximum number of HELO domains to list for the specified time range.	Sender HELO domains
Top Sender IP Connections	IP addresses from which the most content filtering matches have been detected. For each IP address, the total messages processed and number and percentage of content filtering policies triggered are listed. Specify the maximum number of IP addresses to list for the specified time range.	Sender IP connections

Table 18-2 Content filtering reports *(continued)*

Report	Description	Required data storage options
Top Recipient Domains	Recipient domains for which the most content filtering matches have been detected. For each recipient domain, the total messages processed and number and percentage of content filtering policies triggered are listed. Specify the maximum number of recipient domains to list for the specified time range.	Recipient domains
Top Recipients	Email addresses for which the most content filtering matches have been detected. For each email address, the total messages processed and number and percentage of content filtering policies triggered are listed. Specify the maximum number of email addresses to list for the specified time range.	Recipients, Recipient domains
Specific Senders	Number of content filtering policies triggered from a sender email address that you specify. For each grouping, the total messages processed and number and percentage of content filtering policies triggered are listed.	Senders, Sender domains
Specific Recipients	Number of content filtering policies triggered for a recipient email address that you specify. For each grouping, the total messages processed and number and percentage of content filtering policies triggered are listed.	Recipients, Recipient domains
Top Policies	Names of the most common content filtering matches, number of policies triggered, and percentage of policies triggered versus total processed messages. Optionally, you can limit the report to a particular content incident folder.	None

Table 18-3 describes the available email messages reports.

Table 18-3 Email messages reports

Report	Description	Required data storage options
Summary	Summary of total messages and messages that matched filters for spam, suspected spam, attacks, bad sender groups, good sender groups, viruses, suspicious attachments, worms, unscannable messages, malware (spyware/adware), encrypted attachments, and content filtering.	None
Custom	Lets you select the columns that you want to appear in the report.	None
Average Message Size	The average size of messages in KB.	None
Total Message Size	Total size in KB of all messages in the report, and total size of each grouping.	None
Number of Messages	Number of all messages in the report, and number for each grouping.	None
Number of Recipients	Number of recipients in the report, and number of recipients in each grouping. Every recipient in a message (`To:`, `Cc:`, and `Bcc`)counts as one.	None
Top Sender Domains	Domains from which the most messages have been processed. For each domain, the total processed and number of virus and spam messages are listed. Specify the maximum number of domains to list for the specified time range.	Sender domains
Top Senders	Email addresses from which the most messages have been processed. For each email address, the total processed and number of virus and spam messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Senders, Sender domains
Top Sender HELO Domains	SMTP HELO domain names from which the most messages have been processed. For each HELO domain, the total processed and number of virus and spam messages are listed. Specify the maximum number of HELO domains to list for the specified time range.	Sender HELO domains

Table 18-3 Email messages reports *(continued)*

Report	Description	Required data storage options
Top Sender IP Connections	IP addresses from which the most messages have been processed. For each IP address, the total processed and number of virus and spam messages are listed. Specify the maximum number of IP addresses to list for the specified time range.	Sender IP connections
Top Recipient Domains	Recipient domains for which the most messages have been processed. For each recipient domain, the total processed and number of virus and spam messages are listed. Specify the maximum number of recipient domains to list for the specified time range.	Recipient domains
Top Recipients	Email addresses for which the most messages have been processed. For each email address, the total processed and number of virus and spam messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Recipients, Recipient domains
Specific Senders	Number of messages that were processed for a sender email address that you specify. For each grouping, the total processed and number of virus and spam messages are listed.	Senders, Sender domains
Specific Recipients	Number of messages that were processed for a recipient email address that you specify. For each grouping, the total processed and number of virus and spam messages are listed.	Recipients, Recipient domains

Table 18-4 describes the available instant messages reports.

Table 18-4 Instant messages reports

Report	Description	Required data storage options
Overview	Total number of IM messages, messages sent, messages received, number and percentage of messages that were spim, number and percentage of messages blocked due to spim.	None

Table 18-4 Instant messages reports *(continued)*

Report	Description	Required data storage options
Top Spim Senders	Screen names from which the most IM messages have been sent. For each screen name, the report lists the total IM messages sent, the number of spim messages sent, the number of spim messages blocked, and the percentages of each.	Senders
Top Spim Recipients	Screen names for which the most IM messages have been received. For each screen name, the report lists the total IM messages received, the number of spim messages received, the number of spim messages blocked, and the percentages of each.	Recipients
File Transfers	The total number of files transferred, and number and percentage of files that contained a virus, were blocked, could not be scanned, had a scanner error, contained malware, or were encrypted.	None

Table 18-5 describes the available types of invalid recipient reports.

Table 18-5 Invalid recipients reports

Report	Description	Required data storage options
Summary	Summary of all invalid recipient email that enters your network.	None
Top Invalid Recipients	Invalid recipient email addresses with the most incoming messages.	Recipients, Recipient domains, Invalid Recipients
Top Probe Accounts	Probe accounts with the most incoming messages.	Recipients, Recipient domains, Invalid Recipients
Specific Invalid Recepients	Specified invalid email address with the most incoming messages.	Recipients, Recipient domains, Invalid Recipients

Table 18-6 describes the available IP connection-related reports.

Table 18-6 IP Connections reports

Report	Description	Required data storage options
Reputation Summary	Summary of all good reputation and bad reputation verdicts on messages that enter your network. Includes Fastpass and Connection Classification.	Sender IP connections
Connection Summary	Number of connections that are attempted, accepted, rejected, and deferred at connection time.	Sender IP connections
Connection Classification Summary	Summary of connections based on their connection class. See "About managing connection load at the gateway" on page 166.	Sender IP connections
Top Accepted Connections	IP addresses from which the most successful SMTP connections were detected.	Sender IP connections
Top Deferred Connections	IP addresses from which the most failed SMTP connections were detected.	Sender IP connections
Top Rejected Connections	IP addresses from which the most rejected SMTP connections were detected.	Sender IP connections
Top Virus Attacks	IP addresses from which the most virus attacks have been detected. For each IP address, the total messages that were processed and number and percentage of virus attacks are listed.	Sender IP connections
Top Directory Harvest Attacks	IP addresses from which the most directory harvest attacks have been detected. For each IP address, the total messages that were processed and number and percentage of directory harvest attacks are listed.	Sender IP connections

Table 18-7 describes the available spam reports.

Table 18-7 Available spam reports

Report	Description	Required data storage options
Summary	Summary of total detected spam messages (spam, bad senders, and suspected spam messages).	None

Table 18-7 Available spam reports *(continued)*

Report	Description	Required data storage options
Top Sender Domains	Domains from which the most spam messages have been detected. For each domain, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of senders to list for the specified time range.	Sender domains
Top Senders	Email addresses from which the most spam messages have been detected. For each email address, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Senders, Sender domains
Top Sender HELO Domains	SMTP HELO domain names from which the most spam messages have been detected. For each HELO domain, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of HELO domains to list for the specified time range.	Sender HELO domains
Top Sender IP Connections	IP addresses from which the most spam messages have been detected. For each IP address, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of IP addresses to list for the specified time range.	Sender IP connections
Top Recipient Domains	Recipient domains for which the most spam messages have been detected. For each recipient domain, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of recipient domains to list for the specified time range.	Recipient Domains

Table 18-7 Available spam reports *(continued)*

Report	Description	Required data storage options
Top Recipients	Email addresses for which the most spam messages have been detected. For each email address, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Recipients, Recipient domains
Specific Senders	Number of spam messages that were detected from a sender email address that you specify. For each grouping, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed.	Senders, Sender domains
Specific Recipients	Number of spam messages that were detected for a recipient email address that you specify. For each grouping, the spam-to-total-processed percentage, total processed, and the number of spam, suspected spam, and bad sender messages are listed.	Recipients, Recipient domains
Sender Authentication Overview	Total messages that were processed and number and percentage of the sender authentication sessions that were attempted, not attempted, successful, or failed.	None
Top Attempted Senders	Email addresses from which the most sender authentication attempts have been detected. For each email address, the total messages that were processed and number and percentage of sender authentication attempts are listed.	Senders
Top Not Attempted Senders	Email addresses from which the fewest sender authentication attempts have been detected. For each email address, the total messages that were processed and number and percentage of not attempted sender authentication sessions are listed.	Senders
Top Succeeded Senders	Email addresses from which the most successful sender authentication attempts have been detected. For each email address, the total messages that were processed and number and percentage of successful sender authentication attempts are listed.	Senders

Table 18-7 Available spam reports *(continued)*

Report	Description	Required data storage options
Top Failed Senders	Email addresses from which the most failed sender authentication attempts have been detected. For each email address, the total messages that were processed and number and percentage of failed sender authentication attempts are listed.	Senders
Quarantine	Total number of quarantined messages and quarantine releases.	None

Table 18-8 describes the available virus reports.

Table 18-8 Available virus reports

Report	Description	Required data storage options
Summary	Summary of the total number of viruses detected.	None
Top Sender Domains	Domains from which the most virus messages have been detected. For each domain, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of senders to list for the specified time range.	Sender domains
Top Senders	Email addresses from which the most virus messages have been detected. For each email address, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Senders, Sender domains
Top Sender HELO Domains	SMTP HELO domain names from which the most virus messages have been detected. For each HELO domain, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of HELO domains to list for the specified time range.	Sender HELO domains

Table 18-8 Available virus reports *(continued)*

Report	Description	Required data storage options
Top Sender IP Connections	IP addresses from which the most virus messages have been detected. For each IP address, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of IP addresses to list for the specified time range.	Sender IP connections
Top Recipient Domains	Recipient domains for which the most virus messages have been detected. For each recipient domain, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of recipient domains to list for the specified time range.	Recipient Domains
Top Recipients	Email addresses for which the most virus messages have been detected. For each email address, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed. Specify the maximum number of email addresses to list for the specified time range.	Recipients, Recipient domains
Top Virus and Worms	Names of the most common viruses detected. For each grouping, the virus-to-total-processed percentage, virus to total virus and worm percentage, and last occurrence of the virus are listed.	None
Specific Senders	Number of virus messages that were detected from a sender email address that you specify. For each grouping, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed.	Senders, Sender domains
Specific Recipients	Number of virus messages that were detected for a recipient email address that you specify. For each grouping, the virus-to-total-processed percentage, total processed, and the number of viruses, worms, and unscannable messages are listed.	Recipients, Recipient domains

Saving favorite reports

When you create a report, you can specify a name for the report and save it on your Favorite Reports page in the Control Center. You can run saved reports on demand. You can also schedule generating the report automatically.

You must have at least full administration rights or rights to view or modify reports to save reports.

See "Generating reports automatically" on page 590.

See "Deleting favorite reports" on page 585.

You must create and configure a report before you can save it.

See "Creating and configuring reports" on page 569.

To save a favorite report

1 On the **Report Filter** page, under **Report Options**, in the Report name box, type the name of the favorite report.

2 Click **Save to Favorites**.

Editing a favorite report's filter options

You may want to modify a favorite report to expand or condense the information that it provides. You can modify the schedule in which reports are automatically generated and the report's filter options.

You must have full administration rights or rights to view or modify reports to edit reports.

See "Editing a favorite report's schedule" on page 584.

See "Creating and configuring reports" on page 569.

To edit a favorite report's filter options

1 In the Control Center, click **Reports > View > Favorite Reports**.

2 Click on the underlined name of a report in the list.

3 Change the values in the report as necessary.

4 Click **Save**.

Editing a favorite report's schedule

You may want to modify a favorite report's schedule to expand or condense the information that it provides. You must have full administration rights or rights to view or modify reports to edit reports.

See "Editing a favorite report's filter options" on page 583.

To edit a favorite report's schedule

1. In the Control Center, click **Reports > View > Favorite Reports**.
2. Check the box beside the report that you want to edit, and then click **View Schedule**.
3. Make any changes to the settings.
4. Click **Save**.

Copying favorite reports

You may have instances in which you create a favorite report and want to create a similar report with only a few variances. Symantec Brightmail Gateway lets you copy favorite reports.

When you copy a report, the name of the new report is: Copy of <original report name>. Specify a unique name for the new report to help you remember the scope of the report.

See "Saving favorite reports" on page 583.

You must have full administration rights or rights to view or modify reports to copy favorite reports.

To copy favorite reports

1. In the Control Center, click **Reports > View > Favorite Reports**.
2. In the **Report Name** column, check the box beside the report that you want to copy.
3. Click **Copy**.
4. Click on the underlined name of the new, copied report.
5. Change the values in the report as necessary.
6. Click **Save**.

Deleting favorite reports

You can delete the favorite reports that you no longer need. However, when you delete a report, the report configuration cannot be retrieved.

You must have full administration rights or rights to view or modify reports to delete a favorite report.

See “Editing a favorite report's filter options” on page 583.

To delete favorite reports

1. In the Control Center, click **Reports > View > Favorite Reports**.
2. Check the box beside the report that you want to delete.
3. Click **Delete**.

About generated reports layout and data

Use the following information to help you understand the layout and data that appears in the reports that you generate.

See “Report types” on page 571.

Table 18-9 provides information about how reports are displayed.

Table 18-9 Report layout

Element	Description
Graphs and tables	You can specify whether you want the report data to appear in a graph, table, or both. Graph and table options are not available for the Executive Summary report. The options for displaying report data for graphs and tables are as follows: ■ Graph—overview Graphs each category of report data. This graph does not contain the summary information (sums and averages for the entire time period) listed in the overview table. ■ Graph—all others (non-overview) Displays bar graph(s) for each item in the report type chosen. For the reports other than the summary reports, a maximum of 20 items can be displayed in a bar graph. ■ Table Creates numeric a representation of the report data. For all reports, a table report can list more than 20 items. The method to save graphs and tables to files depends on the report, its format, and whether you save or email the report. See "Saving generated reports" on page 595. See "Emailing generated reports" on page 592.
Number of rows	The maximum size for any report (including a scheduled report) is 1,000 rows. If you encounter this limitation, shorten the time range, group by a longer time interval, or decrease the top entries field (applicable to some reports). **Note:** This limitation is not configurable.

Table 18-9 Report layout *(continued)*

Element	Description
Extra bars in report graphs	The current fractional hour is included in report graphs in its own bar. This information ensures that the entirety of the selected time range is displayed. This extra bar usually portrays noticeably less data than the rest of the bars. Consider the following examples: ■ You run a report for the past hour at 2:22 P.M. Tuesday: ■ The resulting data set is from 1:00 P.M. until 2:22 P.M. ■ The data appears by hour, spread across two bars. ■ You run a report for the past 24 hours at 2:22 P.M. Tuesday: ■ The resulting data is from 2:00 P.M. Monday until 2:22 P.M. Tuesday. ■ The data appears by hour, spread across 25 bars.
Time ranges	Report statistics are stored in units from 0 minutes, 0 seconds to 59 minutes, 59 seconds of every hour. For example, from 1:00 A.M. to 1:59 A.M. is one unit and from 2:00 A.M. to 2:59 A.M. is another unit. Because of this scheme, reports cannot be displayed with a time range less than an hour or grouped by a period less than an hour.

Table 18-10 provides the information to help you interpret the information in reports.

Table 18-10 Report data details

Issue	Description
What constitutes a threat	The summary reports and the Dashboard contain threat summary graphs and tables. A threat is a harmful attribute or potentially harmful attribute of an email message. For example, threats include spam, viruses, and content filtering policy violations. Similar message verdicts are grouped into threat categories.
Single threat, multiple threat, and clean messages	The summary reports and the Dashboard categorize messages into single threat, multiple threat, and clean messages. Multiple threat messages contain more than one type of threat. For example, a message that contains spam and a virus is a multiple threat message. Clean messages contain no known threats.

Table 18-10 Report data details *(continued)*

Issue	Description
Message and connection counts	The appliance uses many technologies to track email and filter email. Some of these technologies function at the email connection level before an actual email message can be generated and sent. When a connection is rejected or deferred because it triggered a bad reputation filter, that connection is counted as one message.
Verdicts of suspect viruses messages	If a message is routed to the Suspect Virus Quarantine, the outcome of rescanning the message is not counted toward total threat counts. However, the outcome of rescanning the message is displayed in the Suspect Virus Outcomes graph. The graph indicates whether quarantined suspect viruses were deleted, determined to be viruses or not, or are still in the Suspect Virus Quarantine.
Sender HELO domain or IP connection shows gateway information	If any Scanners accept relayed messages from a gateway computer, the SMTP HELO name or IP connection address is the name or connection of the gateway computer. Affected reports are as follows: ■ Top Sender HELO Domains All Top Sender HELO Domain reports are affected ■ Top Sender IP Connections All Top Sender IP Connections reports are affected ■ Top Succeeded Connections SMTP report ■ Top Failed Connections SMTP report ■ Top Rejected Connections SMTP report
Processed message count	For the reports that list the number of processed messages, the number of processed messages is counted per message, not per recipient. For example, if a single message lists 12 recipients, the processed message count increases by 1, not 12.

Table 18-10 Report data details *(continued)*

Issue	Description
How duplicate verdicts per messages are reported	Each email message can have multiple recipients and multiple threats. Different recipients in the same email message may have different threats triggered. This situation occurs because the different recipients may belong to different policy groups. For example, recipients in group A may have content filtering enabled for employee data protection terms, while recipients in group B may not. Some verdicts have names associated with them to describe unique instances of that verdict type. For example, a known virus may be called W32.Zoltan or VBS.Throckmorton. Each named verdict is counted separately. If both W32.Zoltan and VBS.Throckmorton are found one or more times in a message, the virus count increases by two. The message is considered a multiple threat message. The following verdicts have unique names: ■ Content filtering policies ■ Malware ■ Viruses ■ Worms Verdicts that are not included in this list are counted once per message regardless of the number of occurrences of the verdict in the message. For example, a single message is sent to three recipients. The message to recipient A has two matches for encrypted content. The same message that is sent to recipient B has two matches for encrypted content. That same message that is sent to recipient C has no matches. The total count of encrypted content for the message is one. The virus threat count for the message is one (encrypted content counts as a virus without a unique name). If no other threats are detected in the message, it is considered a single threat message. See "Threat category components" on page 606.
IM message count	IM messages between users within your site are counted twice in the reports.

See "Creating and configuring reports" on page 569.

Running reports on demand

You can run reports as needed. You can run the on-demand reports that you create and the reports that are saved on your Favorite Reports page. You can even run scheduled reports whenever you want.

After you run the report, if there is data available, the report appears in a new browser window. Based on how much data is available for that report, this process may take several minutes.

You must have full administration rights or rights to view or modify reports to run reports.

See "Troubleshooting report generation" on page 598.

To run reports on demand

1. Create a report.

   See "Creating and configuring reports" on page 569.

2. Click **Run**.
3. If a "Pop-up blocked" message appears in the Control Center, click the message and permit pop-ups from Symantec Brightmail Gateway Control Center.

   The report appears in a separate browser window.

To run scheduled reports on demand

1. In the Control Center, click **Reports > View > Favorite Reports**.
2. In the **Report Name** column, check the box besides the report that you want to run.
3. Click **Run**.
4. If a "Pop-up blocked" message appears in the Control Center, click the message and permit pop-ups from Symantec Brightmail Gateway Control Center.

   The report appears in a separate browser window.

Generating reports automatically

You can schedule a favorite report to run automatically at specified intervals. Scheduled reports cannot be automatically saved to the host computer. They must be emailed to at least one recipient. Recipients of the email report can manually save the report on their local computer.

See "Saving generated reports" on page 595.

See "Printing generated reports" on page 595.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See "About scheduled tasks" on page 621.

You must have full administration rights or rights to view or modify reports to create automatically generated reports.

To generate reports automatically

1. In the Control Center, click **Reports > View > Favorite Reports**.
2. Check the box beside the report that you want to schedule and click **View Schedule**.
3. On the **Schedule** tab under **Report Schedule**, in the **Generate report at** drop-down lists, set the time of day to generate the report.
4. Specify when you want the report to be generated as follows:

Daily	Click **Daily** and specify whether you want the report every day or only weekdays.
Weekly	Click **Weekly** and check the boxes for the days of the week that you want to generate the report. You can select multiple days.
Monthly	Click **Monthly** and specify whether you want the report to be created on the same day of each month or the last day of every month. If you specify 29, 30, or 31 in the **Day of every month** box, and a month does not have one of those days, the report is not sent. Instead, click **Last day of every month** to avoid this problem.

5. Click the **Export** tab.
6. Under **Report Format**, select one of the following to specify the format:
   - HTML
   - PDF
   - CSV (this file format is not available for the Executive Summary report)
     In the **CVS Delimiter** drop-down list, select a delimiter.
     In the **File Encoding** drop-down list, select an encoding for the CSV file.

7 Under **Report Sender and Destination Addresses**, in the **Send from the following email address** box, type the email address.

For example, r1b3s@symantecexample.com.

Separate multiple email addresses with a space, comma, or semi-colon.

8 In the **Send to the following email addresses** box, type at least one email address.

For example, r1b3s@symantecexample.com.

Separate multiple email addresses with a space, comma, or semi-colon.

9 In the **Character Set** drop-down list, select a character set appropriate for the recipient of the email message.

10 Click **Save**.

Canceling scheduled reports

You can cancel a scheduled report from being automatically generated without deleting the report. You can still manually generate the report at any time or modify the report later to set up a new schedule.

You must have full administration rights or rights to view or modify reports to cancel automatically generated reports.

See "Generating reports automatically" on page 590.

To cancel scheduled reports

1 In the Control Center, click **Reports > View > Favorite Reports**.

2 Check the box beside the report that you no longer want scheduled, and then click **Clear Schedule**.

Emailing generated reports

You can email the reports that you generate on demand to one or more recipients. Reports that you generate on demand are attached to an email message in a compressed file. The compressed file contains an HTML file and graphic files. When recipients open the HTML file, the report appears as it did in the Web browser when you generated it.

See "Running reports on demand" on page 590.

Scheduled reports must be emailed to at least one recipient, and you must specify the recipient when you create the scheduled report. For more information about how to create scheduled reports, see See “Generating reports automatically” on page 590.

Ensure that you have configured from whom the email report is sent. You can also customize the subject line.

See “Specifying the report email notification sender and subject line” on page 593.

You must create a report before you can email it. You must have full administration rights or rights to view or modify reports to create and email reports.

See “Running reports on demand” on page 590.

To email generated reports

1. In the Control Center, click **Reports > View > Create a Report**.
2. Specify the report type and other choices you want for your report.

   See “Creating and configuring reports” on page 569.
3. In the **Recipient addresses** box type the address of the person to whom you want the report emailed.

   For example, r1b3s@symantecexample.com.

   Separate multiple email addresses with a comma, semi-colon, or space.
4. Select an appropriate encoding for the message from the **Character Set** drop-down list.

   The following options are available:

   - Western European (ISO-8859-1)
   - Unicode (UTF-8)
   - Japanese (Shift_JIS, ISO-2022-JP, or EUC-JP)
   - Simplified Chinese (GB2312 or GB18030)
   - Traditional Chinese (Big5)
   - Korean (KS_C_5601-1987)
5. Click **Email**.

Specifying the report email notification sender and subject line

Symantec Brightmail Gateway can email reports to the recipients that you specify. However, you must specify from whom the email notification is sent. You can also customize the subject line of the email notice. The email address that you specify must be a valid email address for your domain.

When you install the product, the default email address is ReportAdmin@yourcompany.com. If you do not change the default email address and users reply to the notification email, the reply message is undeliverable.

See “Emailing generated reports” on page 592.

The default subject line for reports consists of the title of the report and the report range date. You can specify a custom subject line using a combination of static text and variables.

The default report subject lines are as follows:

Report by hour or day	%TITLE% ({MMMM d, yyyy hh:mm a} to {MMMM d, yyyy hh:mm a zz}) For example: DailyReport (December 28, 2008 12:20 PM to December 29, 2008 12:20 PM PST)
Report by week or month	%TITLE% ({MMMM d, yyyy} to {MMMM d, yyyy zz}) For example: MonthlyReport (December 1, 2008 to January 1, 2009 PST)

See “Date format and time format pattern syntax” on page 704.

You must have full administration rights or rights to view or modify reports to modify report settings.

To specify the report email notification sender and subject line

1. In the Control Center, click **Administration > Settings > Reports**.
2. Under **Report Export Settings**, in the **Email send from** box, type the email address that you want to appear on the report notification email as the sender.
3. If you want to apply a custom subject line for the report email, check **Apply custom subject and filename format**.

   If unchecked, default text is used for the subject line.
4. Modify the subject, if needed.
5. Click **Save**.

Printing generated reports

When you generate a report, if there is data available, the report appears in a new browser window. You can print the report from your browser. You must have full administration rights or rights to view reports.

You must generate a report before you can print it.

See "Creating and configuring reports" on page 569.

See "Running reports on demand" on page 590.

To print generated reports

1 In the **Report** browser window, click **Print**.

2 Choose the appropriate options in your print dialog box to print the report.

Saving generated reports

When you generate a report, if there is data available, the report appears in a new browser window. You can save the report from your browser.

Table 18-11 lists the report file formats that you can use.

Table 18-11 Report file formats

Report file format	Description
HTML	Save the report as a compressed file that contains an HTML file and graphic files. When you unzip the file and open it in a Web browser, the HTML file appears as it appeared originally, with the graphics inline.
CSV	You can save the report as a .csv text file with comma, semicolon, or tab-delimited data. After you save the report as a .csv file, you can import it into a spreadsheet. The report is saved as a delimited text file no matter which Table box or Graph box is checked. To view a CSV file that contains double-byte characters in Microsoft Excel, specify a comma-delimited, UTF-8 file in the MS Excel Text Import Wizard. Alternatively, you can open the CSV file in a text editor that can convert UTF-8 to Unicode and save the CSV file as Unicode. This file format is not available for the Executive Summary report.

Table 18-11 Report file formats *(continued)*

Report file format	Description
PDF	You can save the report as an Adobe Acrobat .pdf file. The PDF file looks similar to the HTML version of the file, though the sections are broken into pages. You can view the saved PDF file in a PDF reader such as Adobe Acrobat Reader.

You must generate a report before you can save it.

See "Running reports on demand" on page 590.

To save generated reports

1 On the browser window, select the report file format that you want to use.

See Table 18-11 on page 595.

Before you click **Save as CSV**, click the **Delimiter** drop-down list. Delimiters include the comma, semicolon, or tab character.

2 Choose the appropriate options in your browser dialog box to save the report.

About purging report data

Symantec Brightmail Gateway contains a utility (Expunger) that purges old report data and Dashboard data. You can configure the amount of time report data is kept before it is purged, the frequency data is purged, and the Expunger start time.

Symantec Brightmail Gateway performs better when you configure it to store less report data. Try to balance your reporting needs against the performance of Symantec Brightmail Gateway. You can accomplish this balance by adjusting the retention period and the types of report data to store.

See "Specifying how long report data is retained" on page 596.

See "Specifying when and how often report data is purged" on page 597.

Specifying how long report data is retained

By default, report and Dashboard data is retained for seven days. If Symantec Brightmail Gateway already has seven days of data, the oldest hour of data is deleted as each new hour of data is stored. Based on your organization's size and message volume, the disk storage requirements for report data can become quite large. You can modify how long Symantec Brightmail Gateway maintains report data.

You must have full administration rights or rights to view or modify reports to modify report settings.

See "Deleting all report data at one time" on page 598.

See "About purging report data" on page 596.

See "Specifying when and how often report data is purged" on page 597.

To specify how long report data is retained

1. In the Control Center, click **Administration > Settings > Reports**.
2. Under **Report and Dashboard Expunger Settings**, click the **Delete data older than** drop-down list to select how long Symantec Brightmail Gateway keeps your report data.
3. Click **Save**.

Specifying when and how often report data is purged

You can specify when the Expunger utility begins the purge process and how frequently the purge process occurs. The Expunger lets you keep the report data that is maintained at a manageable size. Because it is resource intensive, you may want to configure the Expunger to run at off hours. Report data that is purged cannot be retrieved.

See "About purging report data" on page 596.

See "Deleting all report data at one time" on page 598.

See "Specifying how long report data is retained" on page 596.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See "About scheduled tasks" on page 621.

Note: You must have **Full Administration** or **Modify** rights to schedule report purges. See "Administrator rights" on page 684.

To specify when and how often report data is purged

1. In the Control Center, click **Administration > Settings > Reports**.
2. Under **Report and Dashboard Expunger Settings**, in the **Run Expunger** drop-down list, select the Expunger frequency.

3 Click the **Start Expunger at** drop-down lists, and specify the Expunger start time.

The **hour** drop-down list uses a 24 hour clock.

The default setting is 3:00 A.M.

4 Click **Save**.

Deleting all report data at one time

Based on your organization's size and message volume, the disk storage requirements for reports data can become quite large. You can delete all of the report data that Symantec Brightmail Gateway has retained at one time. Once the data is deleted, it cannot be retrieved. If you are unsure about deleting all of the report data, you can delete data based on a specified time range.

See "Specifying how long report data is retained" on page 596.

See "About purging report data" on page 596.

You must have full administration rights or rights to view or modify reports to modify report settings.

To delete all report data at one time

1 In the Control Center, click **Administration > Settings > Reports**.

2 Beside **Delete all reporting and dashboard data now**, click **Delete Data Now** to remove all report data that is stored to date.

3 In the confirmation dialog box, click **OK**.

Troubleshooting report generation

Table 18-12 lists issues you might encounter when you generate reports.

Table 18-12 Report generation issues

Issue	Information
No data available for the report type specified	Instead of displaying the expected reports, Symantec Brightmail Gateway might display the following message: `No data is available for the report type and time range specified.` If you receive this message, verify the following: ■ Data exists for the filter you specified. For example, perhaps you specified a recipient address that received no mail during the specified period for a Specific Recipients report. ■ Symantec Brightmail Gateway is configured to keep data for that report type. Some reports require that you enable report data before those reports can be run. See "Selecting the data to track for reports" on page 569. Occasionally you can produce reports even if data collection is not currently enabled. This situation can happen if you enabled data collection in the past and then turned off data collection. The data that are collected are available for report generation until they are old enough to be automatically purged. After that period, report generation fails. The "Delete data older than" setting on the Report Settings page controls this retention period.

Table 18-12 Report generation issues *(continued)*

Issue	Information
Discrepancies in Suspect Virus Outcomes	The graph part of the Virus Summary report contains a section near the bottom called Suspect Virus Outcomes. The table part of the same report contains a Suspect Virus column. The total suspect virus outcomes may not match the suspect virus column. The reasons for this difference include the following: ■ The suspect virus outcomes are counted for messages only if the message matches a policy that contains the action “Strip and Delay in Suspect Virus Quarantine” or “Hold message in Suspect Virus Quarantine.” ■ Even if a matching policy might trigger one of those actions for a message, another policy may match the message, and take precedence. For example, if a message contains a virus and a suspect virus and the matching virus policy is "Delete message" and the matching suspect virus policy is "Hold message in Suspect Virus Quarantine," the message is deleted. The message is deleted because deletion takes precedence over "Hold message in Suspect Virus Quarantine."
Data in Content Filtering Summary report can look inconsistent	By default, the bottom of the Content Filtering Summary report contains a table of the top content filtering policies that were triggered. The table contains a **Policies Triggered** column and an **Incidents Created** column. Logically, the number of **Incidents Created** should never exceed the number of **Policies Triggered**. However, because of different data sources and timing issues, **Incidents Created** can sometimes exceed **Policies Triggered**.

See “About working with reports” on page 568.

Chapter 19

Monitoring the status of your product

This chapter includes the following topics:

- Managing the log database size
- Manually deleting log files
- About log disk space alerts
- Configuring low disk space alerts
- Clear disk space checklist
- Configuring remote logging to syslog
- Enabling the Message Audit Log
- Configuring log levels
- About message audit logging
- Searching for a message in the Message Audit Log
- Exporting Message Audit Log data
- About message queues
- Viewing IM users that are signed on
- Viewing the connection status of your IM networks

About monitoring the status of your product

You can monitor the status of Symantec Brightmail Gateway from the Control Center. The ability to monitor the status of your product lets you stay up-to-date on product performance and activity. In some cases, you can use filters to customize the status information, such as specifying time ranges.

Table 19-1 describes the items that you can monitor.

Table 19-1 Status items that you can monitor

Category	Description
System	You can monitor the following system statuses: ■ Dashboard View the Dashboard to obtain a dynamic view of product status and filtering activity for various timeframes. See "About the Dashboard" on page 604. See "Viewing the Dashboard" on page 605. ■ Hosts You can monitor the status of your hardware and the size and volume of your message queues. You can also view information about the hardware, software, and services that are installed. See "Viewing information about your hardware" on page 608. See "Viewing the status of software and services" on page 609. See "Viewing the status of your hardware" on page 607. See "Monitoring message queue size and volume" on page 612. ■ Logs Symantec Brightmail Gateway logs information about the Control Center, Spam Quarantine, directory data service, and logs on each Scanner. You can view these logs to monitor the status of your product and troubleshoot issues. See "About logs" on page 625.
SMTP	You can monitor the following SMTP statuses: ■ Message audit logs Symantec Brightmail Gateway provides a message auditing component that lets you search for messages to find out what has happened to them. You can view the message audit log to determine the trail of messages that Scanners accept and process. See "About message audit logging" on page 646. ■ Message queues A message queue is a temporary holding area for messages before they reach their destination. You can view the messages that are queued in any of the message queues. See "About message queues" on page 657.

Table 19-1 Status items that you can monitor *(continued)*

Category	Description
Instant messaging	You can monitor the following instant message statuses: ■ Active users You can view all of the registered and unregistered IM users that are currently signed on. See "Viewing IM users that are signed on" on page 662. ■ Network status You can view the connection status of each IM network that you support from each Scanner that is in your corporate network. See "Viewing the connection status of your IM networks" on page 665.

About the Dashboard

Table 19-2 describes the information that appears on the Dashboard.

Table 19-2 Dashboard contents

Section	Description
Messages graph	Shows the single threats, multiple threats, and clean messages in proportion to the total message volume for the specified timeframe.
Messages table	Lists the individual counts for single threat, multiple threat, and clean messages. Also lists each type's percentage of total message volume for the specified timeframe.
Threats graph	Shows all known threats for a specified timeframe.
Threats table	Lists the individual counts for common message verdicts as well as each type's percentage of total message volume for the specified timeframe. See "Threat category components" on page 606.
Top 5 Named Viruses	Lists the most prevalent viruses and worms for the specified timeframe, ranked by the number of times they were detected. The Details link takes you to the Virus Summary.
Top 5 Content Filtering Policies	Lists the most prevalent types of Content Filtering Policy violations for the specified timeframe, ranked by the number of times they were detected. The Details link takes you to the Content Filtering Summary.

Table 19-2 Dashboard contents *(continued)*

Section	Description
System Status	This pane displays the following information: ■ System Status of the hardware components and software components in your system. If available, the link that is adjacent the component provides a more detailed status. ■ Definitions Status of spam, spim, and virus definitions available through LiveUpdate. If available, the date that is adjacent to the Virus definitions field links to the LiveUpdate Settings page. If IM is not enabled, "IM filtering not enabled" appears as a link that is adjacent to the Spim definitions field. This link provides information about how to enable IM and configure IM-related policies. ■ Licenses Status of the licenses you purchased from Symantec. **Note:** The Definitions column and License column show the oldest definition or license across all Scanners that are licensed for that feature.
Symantec ThreatCon level	To the left of the Dashboard is Symantec ThreatCon. This rating is a measurement of the global threat exposure that is delivered as part of Symantec DeepSight Threat Management System.

See "Viewing the Dashboard" on page 605.

Viewing the Dashboard

The Dashboard provides a dynamic view of your appliance's status and filtering activity for various timeframes. Color-coded graphs show the total volume of processed email and detected threats, sub-divided by message category and verdict, respectively. Accompanying tables present the same data numerically. You can tailor the display to show data according to the direction of email through your system: inbound mail, outbound mail, or both. You can also drill down to more detailed data on viral threats and content filtering policy violations.

Note: Some statistics are relevant only to one direction of the mail stream. For example, the Invalid recipients verdict applies only to inbound mail.

You must have Full Administration rights or Manage Status and Log view or modify rights to view the Dashboard.

See "About the Dashboard" on page 604.

To view the Dashboard

1. Do one of the following actions:
   - Log onto Symantec Brightmail Gateway.
     When you log onto the product, the Dashboard is the first page that appears.
   - In the Control Center, click **Status > System > Dashboard**.
2. Select an email message processing direction from the drop-down list: inbound, outbound, or both.
3. Select a timeframe from the adjacent drop-down list.

   When the screen refreshes, the specified timeframe appears in the top-right corner.

Threat category components

The Dashboard contains data for several threat categories.

Table 19-3 lists the verdicts that make up each threat category so that you can better interpret and analyze the data.

Table 19-3 Threat category components

Threat category	Verdicts that make up the threat category
Content Filtering	■ Policies Triggered ■ Incidents Created ■ Held Messages ■ Approved ■ Rejected ■ Currently Held Content filtering policies count as threats if "Track violations of this policy in the dashboard and reports" is checked.

Table 19-3 Threat category components *(continued)*

Threat category	Verdicts that make up the threat category
Virus	Virus threat category components are as follows: ■ Viruses ■ Worms ■ Malware Potential Virus threat category components are as follows: ■ Suspect Virus ■ Unscannable Attachments ■ Encrypted Attachment
Invalid recipients	■ Invalid recipients
Bad reputation	Bad reputation threat category components are as follows: ■ Directory Harvest Attacks ■ Virus Attacks ■ Bad IPs ■ Connection Class ■ Symantec Global Bad Senders Good reputation threat category components are as follows: ■ Symantec Global Good Senders ■ Good IPs ■ Fastpass
Spam	■ Sender authentication failure ■ Spam ■ Suspected spam ■ Bounce attack

See "About the Dashboard" on page 604.

See "Viewing the Dashboard" on page 605.

Viewing the status of your hardware

You can monitor the hardware status for all of the Scanners that the Control Center administers. Some items appear in red if there is an error condition. The status that appears is based on stored data, so it may be a few minutes old. A dash (–) in a column indicates that the data is not available for that hardware.

You must have Full Administration rights or Manage Status and Logs view or modify rights to view information the status of your hardware.

To view the status of your hardware

1. In the Control Center, click **Status > System > Hosts**.
2. Click the **Hardware Status** tab.
3. To view additional status information, click the hostname.

   When you click the hostname, the data is read from the host to provide real-time status.

Viewing information about your hardware

You can view information about the Control Center host and every Scanner that the Control Center administers. Some hardware information does not apply to the virtual computers that are configured with VMware products. In some cases, a hardware information field may not display for Dell hardware.

The information that appears is as follows:

- Hostname
- Model
- Processor Type
- Processor Cores
- Dell Service Tag
- Total Memory
- Disk

Hardware information is updated dynamically. Some data may not be available while an appliance starts up.

You must have Full Administration rights or Manage Status and Logs view or modify rights to view information about your hardware.

To view information about your hardware

1. In the Control Center, click **Status > System > Hosts**.
2. Click the **Hardware Information** tab.

Viewing the status of software and services

You can view the version of software that is installed on the components that the Control Center administers. You can also see the status of the services that are running for each components. Available details vary depending on how each appliance is configured: as a Control Center, a Scanner, or both.

You must have Full Administration rights or Manage Status and Logs view or modify rights to view information about your software and services.

Table 19-4 describes the software and services status for the Control Center.

Table 19-4 Control Center software and services status

Item	Description
Control Center	Click the Control Center name link for additional details. The **Services** page appears in a tree mode. The information for that Host updates when you click **Back to Services Status**. A green check indicates that the current version of Symantec Brightmail Gateway is installed. A red X indicates that a newer version of Symantec Brightmail Gateway is available.
Version	The currently installed version of Symantec Brightmail Gateway. The version is displayed in red if a newer version is available.
Spam Quarantine	Spam Quarantine contains quarantined spam messages. You can search, delete, sort, and release quarantined messages.
Spam Quarantine Disk Usage	The amount of disk space that Spam Quarantine uses.
Suspect Virus Quarantine	Suspect Virus Quarantine contains messages that could potentially contain viruses. You can search, delete, sort, and release quarantined messages.
Suspect Virus Quarantine Disk Usage	The amount of disk space that Suspect Virus Quarantine uses.
Content Incident Folders	Content incident folders help you organize, monitor, and manage the incidents that trigger content filtering policies in which the action is to create an incident in a content incident folder.
Content Incident Folders Disk Usage	The amount of disk space that content filtering quarantine uses.
Directory Data Service	The status of LDAP, either **Running** or **Stopped**.

Table 19-5 describes the software and services status for Scanners.

Table 19-5 Scanner software and services status

Item	Description
Scanner	Click the Scanner name link for additional details. The **Services** page appears in a tree mode. The information for that Host updates when you click **Back to Services Status**. A green check indicates that the current version and virus definitions are current. A red X indicates an issue with the version or virus definitions.
Version	The currently installed version of Symantec Brightmail Gateway. The version is displayed in red if a newer version is available.
Virus Definitions	The date of the last virus definition update. The date is displayed in red if the virus definitions are out of date or the last LiveUpdate attempt failed.
Scanner	The status of the Scanner, either enabled or disabled.
Agent	The status of the Agent, either **Running** (in black) or **Stopped** (in red). If the Agent crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The Agent transfers configuration information between the Control Center and attached and enabled Scanners.
Conduit	The status of the Conduit, either **Running** (in black) or **Stopped** (in red). If the Conduit crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The Conduit retrieves new and updated filters from Symantec Security Response through secure HTTPS file transfer.
Directory Data Service	The status of the directory data service, either **Running** (in black) or **Stopped** (in red). If the directory data service crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The directory data service lets you use the information that is stored in your Lightweight Directory Access Protocol (LDAP) directories for features in the Symantec Brightmail Gateway.

Table 19-5 Scanner software and services status *(continued)*

Item	Description
LiveUpdate	The status of LiveUpdate, either **Running** (in black) or **Stopped** (in red). If LiveUpdate crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. LiveUpdate automatically downloads virus definitions from Symantec Security Response to the Scanner.
Brightmail Engine	The status of the Brightmail Engine, either **Running** (in black) or **Stopped** (in red). If the Brightmail Engine crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The Brightmail Engine scans the following categories for viruses, spam, and content filtering according to the filter polices that you have configured: ■ Email and attachments ■ Instant messages ■ IM file transfers
MTA	The status of the mail transfer agent, either **Running** (in black) or **Stopped** (in red). If the MTA crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The MTA routes inbound and outbound messages to the Brightmail Engine for processing and delivers filtered messages to their internal destinations or to the Internet.
IM Relay	The status of the IM Relay, either **Running** (in black) or **Stopped** (in red). If the IM Relay crashed in the last 24 hours, a red underline appears beneath the status. Hover your mouse over the red underline to view the number of crashes that occurred in the last 24 hours. The IM Relay retrieves new and updated virus definitions and Spim filters from Symantec Security Response.

If you want to make any modifications, you can do the following from the same page that you view the status:

- Modify a Scanner.

See “Modifying Scanner configurations” on page 107.

- Enable or disable a service.
  See “Working with Services” on page 112.

To view the status of software and services

1. In the Control Center, click **Status > System > Hosts**.
2. Click the **Software and Services** tab.
3. To view additional information about a host, click the host that you want to examine.

   Click the plus sign, where available, next to any component to view additional information on that component.

To modify a Scanner

◆ On the **Software and Services** tab, click any linked word to modify a Scanner.

The Edit Host Configuration page appears.

To enable or disable a service

1. On the **Software and Services** tab, click on a host.
2. On the **Services** page, click the linked word that follows **Status** next to the component.

   The linked word is either **Running** or **Stopped**. The **Services** tab of the **Edit Host Configuration** page appears.
3. On the **Edit Host Configuration** page on the **Services** tab, check the service that you want to enable or disable, and click **Start** or **Stop**.

Monitoring message queue size and volume

You can view the number of queued messages and the size of the queues for all of your message queues. Monitor this status to determine if the message queue is clogged. You can set the maximum size for each message queue and decide whether to defer messages when the queue is full on the SMTP Advanced Settings page. You can configure alerts for message queues on the Alerts page.

See “Configuring SMTP advanced settings” on page 95.

See “Types of alerts” on page 615.

You must have Full Administration rights or Manage Status and Logs view rights.

See “About message queues” on page 657.

See “Viewing queued messages” on page 658.

See "Flushing message queues" on page 661.

To monitor message queue size and volume

1. In the Control Center, click **Status > System > Hosts**.
2. Click the **Message Queues** tab.

MTA and message queue behavior

Each Scanner includes an MTA and corresponding message queues: inbound, outbound, and delivery. Each message queue is managed by a corresponding listener: inbound, outbound, and delivery. In the Control Center, you can perform a variety of actions on the MTA and on message queues.

Note: If you use SMTP authentication, Symantec Brightmail Gateway employs an additional listener for authentication. This listener is controlled by the outbound listener controls. If you stop the outbound listener, the authentication listener also stops.

Table 19-6 describes the expected behavior for new messages and messages in queues when you perform specific actions on the **Services** tab of the **Administration > Hosts > Configuration/Edit** page.

Table 19-6 Manage MTA and message queues on the Edit Host Configuration page, Services tab

Action performed	New messages are	Messages in queues are	Message delivery
Click **MTA**, then click **Stop**	Not accepted. There is no MTA running. External MTAs treat this as an SMTP 4xx error.	Not scanned, not delivered.	Stops.
Click **Pause message scanning and delivery**	Accepted.	Not scanned, not delivered. Accumulate in the inbound and outbound message queues.	Stops.
Click **Do not accept incoming messages** **Note:** This is equivalent to issuing the `mta-control pause-mode pause-accept` command, assuming that delivery is running.	Rejected, issuing SMTP `service not available` (450) error messages.	Scanned and delivered.	Continues.

Table 19-7 describes the expected behavior for new messages and messages in queues when you perform specific actions on the **Status > SMTP > Message Queues** page.

Note: Each action in Table 19-7 affects only one message queue. For example, stopping the inbound message queue has no effect on the outbound or delivery message queues.

Table 19-7 Manage MTA and message queues on the Message Queues page

Action performed	New messages are	Messages in queues are	Message delivery
Display a message queue, click **Flush All**.	Accepted.	Any messages deferred due to delivery problems are retried.	Continues.
Display a message queue, click **Delete All**.	Accepted.	Deleted.	Continues.
Display the inbound message queue, click **Stop**.	Not accepted. External MTAs treat this as an SMTP 4xx error.	Not scanned.	Continues.
Display the outbound message queue, click **Stop**.	Not accepted. External MTAs treat this as an SMTP 4xx error.	Not scanned.	Continues.
Display the delivery message queue, click **Stop**.	Accepted.	Scanned, not delivered. Accumulate in the delivery message queue.	Stops.

See “Managing services and MTA operations” on page 110.

See “Working with Services” on page 112.

See “Components of Symantec Brightmail Gateway” on page 31.

See “Turning off an appliance ” on page 690.

See “About message queues” on page 657.

Configuring alerts

You can configure alerts for Symantec Brightmail Gateway events such as email queue size, Spam Quarantine size, license expiration, and hardware issues. You can specify the email address that appears in the alert notification email and how frequently alerts are sent. Alerts are sent to the administrators that you specify when you create administrators.

See "Adding administrators" on page 682.

Except for the UPS status alert, alerts are not sent at the exact time that the alert condition occurs. Instead, alerts are sent at configurable intervals (the default is hourly).

See "Types of alerts" on page 615.

To configure alert notification criteria

1. In the Control Center, click **Administration > Settings > Alerts**.
2. Under **Notification Sender**, in the **Send from** box, type the email address.
3. In the **Notification Frequency** box, specify how frequently you want notifications sent.
4. Click through each tab and check the alert conditions for which alerts are to be sent.

   Specify duration, percentage, or size parameters where necessary, with the appropriate boxes and drop-down lists.
5. Click **Save**.

Types of alerts

Alerts are automatic email notifications sent to inform administrators of the conditions that potentially require attention. You can choose the types of alerts sent, the `From:` header that appears in alerts messages, and which administrators receive them.

See "Configuring alerts" on page 614.

Table 19-8 describes the available alert settings.

Table 19-8 Alerts page

Alert setting	Explanation
Notification Sender	These settings apply to all alerts.
Send from	The email address that appears in the alert notification's `From:` header.
Notification Frequency	The interval at which notifications are sent. The default is hourly.
Outbreaks	Alerts related to virus outbreaks

Table 19-8 Alerts page *(continued)*

Alert setting	Explanation
Outbreak detection	An alert is sent when a designated number of viruses have been detected over the specified number of hours, days, weeks, or months.
Filters	Alerts related to spam or virus filters
Spam filters are older than	An alert is sent because of the period of time between updates of spam filters. Spam filters update periodically, at different intervals for different types of filters. To avoid unnecessary alerts, a minimum setting of two hours is recommended.
Virus filters are older than	An alert is sent because of the period of time between the virus filter updates which typically occur several times a week. To avoid unnecessary alerts, a minimum setting of seven days is recommended. The default setting is 10 days.
New virus filters are available	An alert is sent because new virus rules are available for download from Symantec Security Response. New virus rules are updated daily. Rapid Response rules are updated hourly.
Queues	Alerts related to message queues
The combined message queue is larger than	An alert is sent when the total combined size of all three message queues exceeds the size specified next to the alert description. Message queues include Inbound, Outbound and Delivery. Queues can grow if the particular queue has stopped or paused, or if an undeliverable message is blocking a queue.
A queue reaches the message limit	An alert is sent when one of the three message queues (inbound, outbound, or delivery) exceeds the maximum number of messages set on the SMTP Advanced Settings page. See “Configuring SMTP advanced settings” on page 95.
Disk Space	Alerts related to available disk space
Available disk space is less than	An alert is sent when the amount of free disk space on the appliance is less than the size that you specify. Low disk space can cause performance and stability issues in Symantec Brightmail Gateway.
Usage of the maximum configured disk space for Spam Quarantine exceeds	An alert is sent when the disk space used by Spam Quarantine exceeds the percentage of the configured maximum size of Spam Quarantine. Set the maximum size of Spam Quarantine on the **Spam > Settings > Quarantine Settings** page. See “Modifying Spam Quarantine thresholds” on page 272.

Table 19-8 Alerts page *(continued)*

Alert setting	Explanation
Usage of the maximum configured disk space for a Content Incident Folder exceeds	An alert is sent when the disk space that is available for any content incident folder exceeds the percentage that is specified. See "About managing the size of content incident folders" on page 439.
Low disk space prompts reduced or halted logging	An alert is sent when the Scanner disk space nears or reaches capacity and the system reduces or halts logging. See "About log disk space alerts" on page 633.
SMTP	Alerts related to SMTP authentication
Login failures occur for a single user	An alert is sent when a user attempts to login using SMTP authentication and fails, if the failures for that user <equal or> exceed the number specified during the time interval specified. You can also specify maximum number of users to display in each notification.
Login failures occur from a single IP	An alert is sent when a user attempts to login using SMTP authentication and fails, if the failures for that IP address <equal or> exceed the number specified during the time interval specified. You can also specify maximum number of IP addresses to display in each notification.
DDS	Alerts related to directory data service
Directory Data access errors	An alert is sent because the directory data service failed to read data from an LDAP server.
Directory data integrity errors	Indicates a problem with the customer's directory data that prevents a directory data service operation from succeeding.
Undersized data source cache	An alert is sent because a directory data source cache or cache index is not large enough to hold all of the requested data in memory.
User preference replication errors	An alert is sent because of an error replicating user preferences to the Scanner.
License/Updates	Alerts related to licenses, certificates, and software updates
Symantec Premium Content Control license expired	An alert is sent when the PCC license approaches expiration. Another alert is sent when your license expires. Contact your Symantec sales representative for assistance.

Table 19-8 Alerts page *(continued)*

Alert setting	Explanation
Symantec Antivirus license expired	An alert is sent when your antivirus license approaches expiration. Another alert is sent when your license expires. Contact your Symantec sales representative for assistance.
Symantec Antispam license expired	An alert is sent when your Symantec Antispam license approaches expiration. Another alert is sent when your license expires. Contact your Symantec sales representative for assistance.
Symantec Content Encryption license expired	An alert is sent when your Symantec Content Encryption license approaches expiration. Another alert is sent when your license expires. Contact your Symantec sales representative for assistance.
Software Updates license expired	An alert is sent when your software update license approaches expiration. Another alert is sent when your license expires. Contact your Symantec sales representative for assistance.
SSL/TLS certificate expiration warning	An alert is sent when a certificate expires. You can check the status of your certificates by going to the Certificate Settings page and clicking View. The first expiration warning is sent seven days before the expiration date. A second warning is sent one hour later. No more than two warnings per certificate are sent.
New software release update available	An alert is sent to indicate that a new software update release is available.
Frequency of checking for updates:	Specify the frequency in minutes, hours, or days at which Symantec Brightmail Gateway checks for new software updates.
Events	Alerts related to system operations
Swap space utilization exceeds	An alert is sent when the available memory for swap exceeds the percentage you specify.
A service is not responding or working	An alert is sent because of a nonresponsive service. Services include the Conduit, LiveUpdate, Brightmail Engine, MTA, IM Relay, and Directory Data Service.
Hardware failures	An alert is sent due to a hardware problem such as a fan failure or disk failure.
Service start after improper shutdown	An alert is sent because a service restarted after an improper shutdown. Services include the Conduit, LiveUpdate, Brightmail Engine, MTA, IM Relay, and Directory Data Service.

Table 19-8 Alerts page *(continued)*

Alert setting	Explanation
Service shutdown	An alert is sent because a service was shut down normally. Services include the Conduit, LiveUpdate, Brightmail Engine, MTA, IM Relay, and Directory Data Service.
Service start	An alert is sent because a service was started. Services include the Conduit, LiveUpdate, Brightmail Engine, MTA, IM Relay, and Directory Data Service.
UPS status	An alert is sent because the uninterruptible power supply status has changed. This alert can be sent as frequently as every seven minutes.
Failed Scheduled Tasks	An alert is sent when a task you have scheduled fails to execute. Scheduled tasks are configured in their respective tabs but the status of all scheduled tasks can be viewed on the **Status > System > Scheduled Tasks** page.

Monitoring devices through SNMP

Simple Network Management Protocol (SNMP) lets administrators monitor network devices, such as the Control Center and Scanners. You can specify an SNMP community string and trap. You can also manage access privileges to the SNMP agent for up to four hosts in your environment.

Before you configure SNMP settings, you must first download the Management Information Base (MIB) database and import it to your SNMP client.

See "Downloading a Management Information Base for SNMP" on page 620.

To monitor devices through SNMP

1. In the Control Center, click **Administration > Settings > SNMP**.
2. Check **Enable SNMP**.
3. In the **SNMP community string** box, type the SNMP Agent's community string.
4. In the **SNMP listen port** box, type the port at which the SNMP agent listens for network traffic.
5. In the **SNMP trap host** box, type the IP address or host name of the device that receives SNMP trap alerts.
6. Under **SNMP Client Access**, specify which hosts can access the SNMP agent by doing one of the following tasks:

- Click **All hosts** to grant all hosts access to the SNMP client.
- Click **Only the following hosts** and then type the IP addresses, host names, or CIDR ranges of specific hosts that you want to have access to the SNMP client.

7 Click **Add**.

8 You can delete any currently SNMP-access-enabled hosts by checking the box next to their names and clicking **Delete**.

9 Click **Save**.

Downloading a Management Information Base for SNMP

Before you configure SNMP settings, you must first download the Management Information Base (MIB) database and import it to your SNMP client.

The following SNMP MIBs are provided for hardware related alerts:

- LSI-AdapterSAS.txt
- LSI-AdapterSASIR.txt
- PERC-MIB.txt
- afa-MIB.txt

The following SNMP MIB is provided for Symantec Brightmail Gateway applications:

- SYMANTEC-EMAIL-SECURITY.txt

See “Monitoring devices through SNMP” on page 619.

To download a MIB for SNMP

1 Log on to the Control Center.

2 Type `/snmp-mibs` in the URL path in your browser after **/brightmail** and press **Enter**.

For example: `https://your_hostname/brightmail/snmp-mibs`

The MIBs page appears.

3 Download the appropriate MIB for your appliance and save it to the computer that runs the SNMP monitoring program.

4 Import the MIB file into any SNMP v2c-compliant monitoring program.

Configuring UPS settings

Symantec Brightmail Gateway can monitor USB attached APC UPS devices. It can also perform a graceful shutdown due to loss of power when any one of the following conditions are met:

Battery level	If during a power failure, the remaining battery percentage (as reported by the UPS) is less than or equal to the specified value
Runtime minutes	If during a power failure, the remaining runtime in minutes (as calculated internally by the UPS) is less than or equal the specified value
Timeout minutes	If during a power failure, the UPS has run on batteries for the timeout minutes **Note:** If you have a Smart UPS, you can disable the Timeout minutes feature and use the other settings to control when a shutdown is initiated.

See shutdown on page 809.

To configure UPS settings

1. In the Control Center, click **Administration > Settings > UPS**.
2. Check **Enable UPS Monitoring** and select the conditions under which the appliance turns itself off.

   To disable a feature, type a value of 0.

About scheduled tasks

The Symantec Brightmail Gateway has a number of automated and scheduled tasks which you configure throughout the Control Center. You can view and track the status of your scheduled tasks under the **Status** tab. Scheduled tasks are grouped into subtabs by associated task types. Each subtab page contains the name and status of the scheduled task. If a task completes successfully, the table displays the start time, the finish time, and the time that the task is scheduled to run next. If a task fails to complete, the **Finished** column displays a status of **Failed**.

When a scheduled task fails, Symantec Brightmail Gateway provides an email alert option which contains details about the failure. If you enable the email alert option, the system notifies the designated administrator when any of the Control Center's scheduled tasks fail.

The system tracks the following types of scheduled tasks:

- Backups
- Expungers
- Notifications
- Quarantine
- Reports
- Preferences

See “Scheduled tasks types” on page 622.

See “Setting scheduled tasks failures alerts” on page 622.

Setting scheduled tasks failures alerts

Symantec Brightmail Gateway lets you set an email alert when a scheduled task fails to execute. If a task fails, an email notification is sent to the designated administrator, indicating which task failed and on which host. Control Center logs provide further details about the failure.

To set scheduled tasks failures alerts

1. In the Control Center, click **Administration > Settings > Alerts**.
2. Under the **Events** tab, check **Scheduled tasks failures**, and click **Save**.

See “About scheduled tasks” on page 621.

See “Scheduled tasks types” on page 622.

Scheduled tasks types

The Symantec Brightmail Gateway has a number of automated and scheduled tasks. You can view the status of all Control Center tasks on the **Status > System > Scheduled Tasks** page.

Tasks are grouped by type under the following tabs:

- Backup
- Expungers
- Notifications
- Quarantine
- Reports
- Preferences

Under each tab you see a page that contains a table of related tasks. Each scheduled task shows a start time, a finish time, and the time the task runs next. If a task fails for any reason, the term **Failed** appears in the **Finished** column.

You can set up an alert to be notified if a scheduled task fails.

See "Setting scheduled tasks failures alerts" on page 622.

See "About scheduled tasks" on page 621.

The **Backups** tab reports on tasks that create copies of the Control Center database or database configurations such as content filtering messages, incidents, logs, and reports.

Note: This table is blank until you add a backup configuration.

Table 19-9 Backups

Task	Description
Backups	See "Scheduling backups" on page 692.

The **Expungers** tab reports on tasks that delete data from the Control Center repository, such as messages in the spam quarantine, log files in the log database, reports, and data in your content incident folders.

Table 19-10 Expungers

Task	Description
Spam Expunger	See "Specifying when and how often Spam Quarantine is expunged" on page 284.
Log Expunger	See "Managing the log database size" on page 630.
Reports Expunger	See "Specifying when and how often report data is purged" on page 597.
Content Expunger	See "Scheduling the content incident folder Expunger" on page 444.

The **Notifications** tab reports on tasks that produce automatic emails to system administrators, such as alerts involving system changes or conditions that potentially require attention.

Table 19-11 Notifications

Task	Description
Alert Notifications	See "Types of alerts" on page 615.
Content Filtering notifications (Informational Incidents)	See "Creating incident notifications" on page 445.
Content Filtering notifications (Quarantine Incidents)	See "Creating incident notifications" on page 445.
Spam Notifications	See "Specifying when to notify users of spam messages in their quarantine" on page 277.

The **Quarantine** tab reports on the task that releases quarantined suspect viruses.

Table 19-12 Quarantine

Task	Description
Suspect Virus Message Quarantine Release	See "Specifying how long suspect virus messages are retained in quarantine" on page 237.

The **Reports** tab reports on the tasks that run **Favorite** reports at specified intervals which are then emailed to a specified recipient. Report types include, Summary, Content Filtering, Email Message, Instant Messaging, Invalid Recipients, IP Connection, Spam and Virus reports.

Table 19-13 Reports

Task	Description
Reports	See "Generating reports automatically" on page 590.

The **Preferences** tab reports on the task that propagates Directory Data Service data to Scanners.

Table 19-14 Preferences

Task	Description
User Preferences	See "Configuring the replication of end user preference data" on page 676.

About logs

Symantec Brightmail Gateway keeps log records of component activity on Symantec Brightmail Gateway. If any part of Symantec Brightmail Gateway does not work properly, you can view logs to investigate the problem. Even if you do not notice any problems, Symantec recommends that you view logs regularly to check the status of the components.

You can specify the level of information that you want to log for each of your Scanners. Some levels can yield high volumes of data. Set the log levels to the lowest level that meets your monitoring needs. The default log level of warning is usually appropriate.

You view log data in the Control Center. You can sort log data, create log reports, and clear log files from the database. You can also view logs using the command line interface with the `cat`, `grep`, `more`, and `tail` commands. The command line interface displays real-time log data.

See "Viewing log files" on page 627.

See "Saving log files" on page 629.

See more on page 791.

In addition to viewing logs using the Control Center, some Scanner logs can be sent to syslog on a remote server.

See "Configuring remote logging to syslog" on page 637.

Symantec Brightmail Gateway lets you configure the log database size to prevent the database from filling the hard disk. It also provides a log Expunger utility that lets you set automatic purges at specified intervals.

See "Managing the log database size" on page 630.

See "Log types" on page 626.

Scanner logs are made available in the Control Center in two steps: First, incoming data files are captured to a temporary location on the Scanner. Next, the Control Center accesses the temporary location and copies the data to the Control Center log database. Once the log data from the Scanner is transferred to the Control Center, the system deletes log data on the Scanner to conserve disk space.

In some circumstances the Scanner logs can fill up the disk faster than the Control Center can copy the data to the log database. If a Scanner disk nears or reaches capacity, the system reduces or halts logging and no logs are copied to the Control Center log database. You must free disk space for normal logging to resume.

See "About log disk space alerts" on page 633.

See "Clear disk space checklist" on page 634.

Log types

Symantec Brightmail Gateway logs let you monitor events for your Scanners, the Control Center, Directory Data Service, and Spam Quarantine.

Table 19-15 describes the log types that are available for Scanners.

Table 19-15 Scanner log types

Scanner log type	Description
Conduit	Records the status about downloading antispam rules and uploading statistics.
Brightmail Client	Records the status about message filtering.
Brightmail Engine	Records the status of the Brightmail Engine.
JLU Controller	Records the status about Java LiveUpdate virus definition downloads. This log is the primary log file that you should use for troubleshooting LiveUpdate issues.
JLU Client	An auxiliary log file to the JLU Controller log file that records the status about Java LiveUpdate virus definition downloads. Use this log file only when the JLU Controller log file does not contain enough information for troubleshooting an issue.
MTA	Records the status about sending and receiving email.
IM Relay	Records the status about instant messaging activities, such as scanning instant messages for viruses.
Content Filtering	Records the status about Content Filtering.

Table 19-16 describes the log types that are available for the Control Center console.

Table 19-16 Control Center log types—Console

Console log file	Description
BrightmailLog.log	Records the status about Control Center interactions.
catalina<date>.log and catalina.out	Records the status from the Tomcat Web server . The Control Center runs inside the Tomcat server. These files contain the messages that are generated from the Tomcat Server and also the applications that run within Tomcat.

Table 19-17 describes the log types that are available for the Control Center database.

Table 19-17 Control Center log types—Database

Database log file	Description
error.log and error.log.#.gz	Records any errors that occur while the Control Center accesses the MySQL database.
slow-queries.log and slow-queries.log.#.gz	Records the slow MySQL queries.

Table 19-18 describes the log type that is available for the Control Center events.

Table 19-18 Control Center log types—Events

Event log file	Description
Brightmail_Admin_Events.<yyy-mm-dd>.log	Records all changes made in the Control Center for the date that is indicated in the log file name.

Table 19-19 describes the log that is available for Spam Quarantine.

Table 19-19 Quarantine log

Quarantine log	Description
Release	Records the To address, From address, and Subject of each message that is released from Spam Quarantine. It also records the user who released each message and a timestamp.

Table 19-20 describes the log that is available for directory data service.

Table 19-20 Directory data service log

DDS log	Description
Directory Data Service	Records the status about directory data service.

See "About logs" on page 625.

See "Viewing log files" on page 627.

Viewing log files

You can view the events that are logged for the Control Center, Scanners, Directory Data Service, and Quarantine. If any part of Symantec Brightmail Gateway does not work properly, you can view logs to investigate the problem. Even if you do not notice any problems, Symantec recommends that you view logs regularly to check the status of the components.

When you specify a filter and click **Display**, Symantec Brightmail Gateway displays the log events that match the filter. You must have one of the following access levels to view logs:

- **Full Administration Rights**
- **Manage Status and Logs** with **View** or **Modify** rights

See "About logs" on page 625.

For Scanner logs, the amount of log data available depends on the log level that you set for the Scanner component. For example, if you set the Conduit log level to warnings, no log data is saved or available for notice, information, or debug level events. Set the Scanner component log levels on the **Administration > Settings > Logs** page.

See "Configuring log levels" on page 643.

The Control Center displays a question mark icon next to the description of error level Scanner logs. Click the icon to display a Web page that contains more information about the error. The Web page opens in a new Web browser window. The Web page may indicate that Symantec has not published information about the error. Symantec tracks the error information requests and adds new error information continually.

Table 19-21 lists the filter options that you can use to view specific log events.

Table 19-21 Log view options

Item	Description
Host (drop-down)	Select a host from the list. This option is only available for Scanner and directory data service logs.
Severity (drop-down)	Select a severity level from the list. This option is only available for Scanner and directory data service logs.
Time range (drop-down)	Select a time range from the list or create a custom time range. If you have recently changed time zones on the Control Center, this change is not reflected immediately, but requires that you restart the appliance.
Component (drop-down)	Select a component for which to view logs: **Scanner**, **Control Center**, **Quarantine**, or **Directory Data Service**.

Table 19-21 Log view options *(continued)*

Item	Description
Log type (drop-down)	Select a log type from the list. See "Log types" on page 626.
Log action (drop-down)	Select the type of actions to display: **System Events**, **Blocking Actions**, **Message Actions**, or **All**.

If a character in a Scanner log is not printable or is not ASCII, the sequence \xAB is printed instead of that character. AB is the hexadecimal value of the character. For example, a character with decimal value of 128 is displayed as \x80.

Since log information is dynamic, you can refresh the view at any time by clicking **Display**.

See "Saving log files" on page 629.

See "Managing the log database size" on page 630.

To view log files

1. In the Control Center, click **Status > System > Logs**.
2. Under **Filter**, specify selection criteria for the log events that you want to view.
3. Click **Display**.

   The results of the filter appear on the Logs page.

Saving log files

Symantec Brightmail Gateway lets you save log files. Log files are saved in .txt format.

When you save a log file, you can view and print the file with a text editor application. You can also email the file. You must have Full Administration rights or Manage Status and Logs view or modify rights to save a log file.

Scanner log files do not contain individual log file links like the Control Center and Quarantine logs. When you save a Scanner log file, the text file contains all of the detailed log items that appear on the Control Center page. The Control Center log and Quarantine log require that you save the Log files that you want individually.

See "About logs" on page 625.

To save log files

1. In the Control Center, click **Status > System > Logs**.
2. Create a log.

   See “Viewing log files” on page 627.
3. Do any of the following:

To save a Scanner log file	Click **Save Log**.
To save Control Center or Quarantine logs	Do all of the following: ■ In the **Log Files** column, click on the Log File that you want to save. ■ In the **File Download** dialog box, click **Save**.

4. In the **Save As** dialog box, specify the file name and the location where you want to save the file, and click **Save**.

   The default file name is LogDetails.txt.
5. In the **Download Complete** dialog box, click **Close**.

Managing the log database size

The Control Center log database acquires log data from Scanners at regularly scheduled intervals. Depending on your environment and the log level that you specify, the log database can quickly increase to an unmanageable size. To prevent expansion, Symantec Brightmail Gateway provides controls for specifying a maximum database size (the default is 50 MB). It also provides a log Expunger utility which purges older logs and prevents the database from exceeding the set limit (the default setting is daily). Users may also create custom configurations to meet specific or temporary needs.

Note: The settings you specify here should take into consideration the log levels you set. For example, if you have a log level that produces high volumes of log data (such as Debug level), you may want to temporarily increase the database size and log purge rate. When you resume to normal log levels, you can change these settings back.

See “Configuring log levels” on page 643.

See “About logs” on page 625.

See “About log disk space alerts” on page 633.

Note: If necessary, you can also manually purge files from the log database.

See “Manually deleting log files” on page 632.

Table 19-22 Log database size control options

Option	Description
Maximum log size	Specifies the maximum size of the database log. The maximum log size is enforced when the Log Expunger runs. The maximum log size may be exceeded temporarily until the next Log Expunger process runs. The default value is 50 MB.
Days to store log data before deleting	Specifies the number of days that log files are retained in the database log before they are purged. The default value is 7.
Log Expunger frequency	Specifies how frequently the Expunger utility runs. The Log Expunger deletes older log files to enforce the maximum log size. The default setting is Every day.
Log Expunger start time	Specifies the time in which the Expunger starts. The time is based on a 24-hour clock. For example, 23:00 is 11:00 P.M. The default setting is 02:00 (2:00 A.M.). You can check the status of your scheduled task from the **Status > Scheduled Tasks** page. See “About scheduled tasks” on page 621.

Note: When the Log Expunger runs, it deletes log entries in the log database. It does not compact the database. To compact the database and decrease its size, Symantec Brightmail Gateway runs an optimization process. This process occurs automatically based on disk usage, so you do not need to configure it. However, because the optimization process is processor-intensive, it normally runs during off-peak hours (the default setting is 2:00 A.M.). So the Expunger may delete rows from the log database, but the size of the database does not decrease until the optimizer runs.

Note: You must have **Full Administration** or **Modify** rights to change log database settings. See "Administrator rights" on page 684.

To manage the log database size

1. In the Control Center, click **Administration > Settings > Logs**.
2. Click the **Local** tab.
3. Under **Database Log Storage Limits**, check **Maximum log size**.
4. In the adjacent box, type the maximum size that you want to allocate for the log database.
5. In the **Days to store log data before deleting** box, type the number of days to store log data.
6. Under **Log Expunger**, choose a frequency and a start time when the Control Center runs the Log Expunger to delete log data.

   The hour drop-down list uses 24-hour format. For example, 23:00 is 11:00 P.M.
7. Click **Save**.

Manually deleting log files

The log Expunger utility lets you schedule purges of the log database at regularly recurring intervals. However, in some cases you may want to delete logs manually from the log database. You can clear Scanner log files and directory data service log files from the Control Center log database on the **Status > System > Logs** page. You can also manually delete log files using the `delete` command.

See delete on page 753.

Note: Deleting Scanner and directory data service log files from the Control Center log database does not remove the original log files from the Scanner appliance. To delete the original Scanner and directory data service log files, you must use the command line interface.

See "Clear disk space checklist" on page 634.

See "Managing the log database size" on page 630.

See "About log disk space alerts" on page 633.

See "Viewing log files" on page 627.

You must have full Administration, Manage Status, or Logs modify rights to manually delete logs.

To manually delete log files

1. In the Control Center, click **Status > System > Logs**.
2. In the Filter section, select **Scanner** or **Directory Data Service** from the component drop-down menu.
3. Click **Clear All Scanner Logs** for Scanner logs or **Clear All DDS Logs** for directory data service logs.

About log disk space alerts

Symantec Brightmail Gateway includes an enhanced logging feature that allows system administrators to enable alerts when logging disk space on Scanners nears or reaches capacity. When the default threshold levels are reached, the system shifts to a reduced or halted logging mode and sends an email notification to the specified administrator. You will need to enable the low disk space alerts to receive this notification.

See "Configuring low disk space alerts" on page 634.

A logging mode change is typically caused by a change in log levels in the Control Center. If an administrator sets the log level to a level that yields a high volume of data (such as the Debug level), the Control Center may not be able to copy data files from the Scanner to the log database at a rate that keeps pace with the incoming data feeds. Log files remain on the Scanner and disk space quickly fills up.

See "Managing the log database size" on page 630.

When the system shifts to reduced logging mode, only urgent log levels (e.g., Errors and Warnings) are recorded. Log levels that produce excess data of a non-urgent nature (e.g., Notices, Information, and Debug) are suspended until space is freed on the disk. When a halted situation occurs, all log levels are suspended. Normal logging resumes only after the log files have been transferred off the disk.

Note: Reduced and halted logging only applies to appliances operating as a Scanner or Directory Data Service (DDS).

See "Clear disk space checklist" on page 634.

See "About logs" on page 625.

See "Manually deleting log files" on page 632.

See "Configuring log levels" on page 643.

See "Types of alerts" on page 615.

Configuring low disk space alerts

To ensure that the system notifies you when Scanner disk space nears or reaches capacity, you must enable low disk space alerts.

Note: Disk space alerts are sent from the mail administrator that you specify on the **Administration > Settings > Alerts** page. See "Configuring alerts" on page 614.

To configure low disk space alerts

1. In the Control Center, click **Administration > Settings > Alerts**.
2. Click the **Disk Space** tab.
3. In the **Disk Space** section, ensure that one or more of the alerts are checked.
4. Click **Save**.

Clear disk space checklist

The following is a list of options for clearing disk space on the Control Center and Scanner appliances. Use these options if you have determined that low disk space availability currently affects system performance.

See "About log disk space alerts" on page 633.

See "About logs" on page 625.

See "Manually deleting log files" on page 632.

Warning: Deleting directory contents results in permanent data loss. The data is no longer available for reports or analysis. Using a subset of these actions may be sufficient to restore the disk to a working condition. Delete data conservatively.

Table 19-23 Control Center data

Control Center stored data	Description
Quarantined mail	Click **Delete All** on the **Spam > Quarantine > Email Spam** page to delete all spam messages that are stored in the Spam Quarantine.

Table 19-23 Control Center data *(continued)*

Control Center stored data	Description
Content filtering incidents	Old content filtering incidents can be deleted from **Content > Incident Management > [Content Incident management folder]**. Click **Delete** or **Delete All** to delete contents.
Reporting data	Click **Delete Data Now** on the **Administration > Settings > Reports** page to delete all reporting data and dashboard data.
Scanner data that is stored in the log database	On the **Status > System > Logs** page, select **Scanner** from the **Component** list and click **Clear All Scanner Logs** to delete Scanner logs from the log database.
Directory data service (DDS) data that is stored in the log database	On the **Status > System > Logs** page, select **Directory Data Service** from the **Component** list and click **Clear All DDS Logs** to delete directory data service logs.
Database backups that are stored locally	Backups are stored on the appliance in /data/backups/ or /data/backups/tmp/. Click **Delete** on the **Administration > Hosts > Version > Restore/Download** page to delete the old backups that are stored locally on the appliance.
Old software update data	Files in /data/apt/ directory and its subdirectories. To delete old software update data, type the following on the command line: `delete sudata`

You can delete data files using the command line.

See "Administering Symantec Brightmail Gateway through the command line" on page 711.

See delete on page 753.

See mta-control on page 792.

Table 19-24 Scanner data

Scanner stored data	Description
Core files	Files in /data/scanner/jobs/ directory and its subdirectories. To delete core files, type the following on the command line: `delete cores`

Table 19-24 Scanner data *(continued)*

Scanner stored data	Description
Log files	Files in /data/logs/ directory and its subdirectories that the Control Center has not retrieved. To delete all log files, type the following on the command line: `delete alllogs`
Stats files	Files in the /data/scanner/stats/ directory that the Control Center has not retrieved. To delete stats files, type the following on the command line: `delete alldata` This command deletes other data in addition to the stats files. Only use `delete alldata` if necessary.
Tmp files	Files in /tmp/ and /var/tmp/ directories and their subdirectories. To delete individual files, type the following on the command line: `delete file /tmp/`*`filename`* or `delete file /data/tmp/`*`filename`*
Audit logs	Files in the /data/logs/scanner/ directory whose names start with audit_. To delete audit logs, type the following on the command line: `delete mallogs`
Old messages that are stuck in inbound, outbound, and delivery mail queues	Files in /data/mta/queues/ directory and its subdirectories. To delete the messages that are stuck in the queues, type the following on the command line: `mta-control all delete-all-msgs`
Messages that are stored in bad-messages folder	Files in /data/mta/bad-messages. To delete bad messages, type the following on the command line: `mta-control all bad-msg-delete all`
Old software update data	Files in /data/apt/ directory and its subdirectories. To delete old software update data, type the following on the command line: `delete sudata`

Configuring remote logging to syslog

Some Scanner logs can be sent to syslog on a remote server. Ensure that the remote syslog is configured to match the settings in the Control Center. You must enable either local logging, remote logging, or both.

To configure remote logging to syslog

1. In the Control Center, click **Administration > Settings > Logs**.
2. Click the **Remote** tab.
3. Click **Enable Syslogs for the following host** and click a host to send log data from that host to a remote syslog.
4. In the **Host** field, specify the syslog server's IP address.
5. In the **Port** field, specify the port on the syslog server that handles log data.
6. In the **Protocol** field, specify the syslog protocol: UDP or TCP.
7. Under **Component Remote Log Levels**, specify the logging level and facility for each component.
8. Click **Enable message logs** to send message logs to the remote syslog.

   See “Format of audit logs” on page 647.
9. Click a **Message log facility**.
10. Click **Apply these Remote Logging settings to all hosts** if wanted.
11. Click **Save** to save your changes.

    Log components may need to be restarted.

Standard prefix for Scanner logs sent to remote syslog

You can configure Symantec Brightmail Gateway to send Scanner log data to a remote syslog. All Scanner log messages that are sent to the remote syslog take the following form:

Date and time	Facility Level	IP address	Original log message
`1-15-2009 15:42`	`Local3.Info`	`10.217.32.13`	`Jan 15 15:38:05 scanner1 jlu-controller: [Brightmail] (INFO:21145.3071248064): [54038] AV definitions are up-to-date.`

Date and time	Facility Level	IP address	Original log message
Date in the format month-date-year. Time in the format hour:minute. The time is in 24-hour clock notation. The date and time is the date and time that the log message was sent to the remote syslog.	The facility, a period, and the log level. The facility designates the facility on the remote syslog to which the log data is sent. The log level is the log level configured on the Scanner.	IP address of the Symantec Brightmail Gateway host sending the log message.	The original log message as it would appear on the Scanner. The format of this portion depends on the log component.

The first three columns make up a standard prefix that appears before all log messages send to a remote syslog. The following is a log message for one event as it would appear on the remote syslog.

```
1-15-2009 15:42 Local3.Info 10.217.32.13 Jan  15 15:38:05
scanner1 jlu-controller: [Brightmail] (INFO:21145.3071248064):
[54038] AV definitions are up-to-date.
```

Log format of boot.log, cron, message, and secure components for remote syslog

You can configure Symantec Brightmail Gateway to send Scanner log data to a remote syslog. If remote syslog is enabled, log data for the boot.log, cron, message, and secure components is sent to the remote syslog. These are standard UNIX log components. Log data for these components is sent to the standard syslog facility on the remote syslog. The log level for these components cannot be configured. All log messages that are sent to a remote syslog have the same prefix text.

See "Standard prefix for Scanner logs sent to remote syslog" on page 637.

The log messages that are sent to the remote syslog take the following form:

Standard prefix	Date and time	Scanner host name	Process[PID]:	Message
Date, time, facility, log level, and IP address	`Jan 15 11:51:33`	`scanner1`	`syslog-ng [25257]:`	`syslog-ng version 1.6.5 starting`

Standard prefix	Date and time	Scanner host name	Process[PID]:	Message
See "Standard prefix for Scanner logs sent to remote syslog" on page 637.	Date in the format month date. Time in the format hour:minute. The time is in 24-hour clock notation. The date and time is the date and time that the log message was recorded on the Scanner.	Name of the Scanner on which the log message was created.	Name and process ID of the process that generated the log message.	Log message.

The following is a log message for one event as it would appear on the remote syslog.

```
01-15-2009 11:52:44 Syslog.Notice 10.217.32.13 Jan 15 11:51:33
scanner1 syslog-ng[25257]:syslog-ng version 1.6.5 starting
```

Log format of Conduit, Brightmail Client, Brightmail Engine, JLU Controller, and IM Relay for remote syslog

You can configure Symantec Brightmail Gateway to send Scanner log data to a remote syslog. All log messages that are sent to a remote syslog have the same prefix text.

See "Standard prefix for Scanner logs sent to remote syslog" on page 637.

The following table contains sample log messages for the following components in the same order listed:

- Conduit
- Brightmail Client
- Brightmail Engine
- JLU Controller
- IM Relay

Standard prefix	Date and time	Scanner host name	Process: [Brightmail]	(LogLevel: PID: ThreadID):	[EC/Source: #:Function]	Message
Date, time, facility, log level, and IP address	Jan 15 11:34:51	scanner1	conduit: [Brightmail]	(DEBUG: 19713. 3071461056):	[src/rda_ controller.cc: 586:initial ize]	Appending HTTP header: 'Spamwall_ID: (null)'
	Jan 15 11:36:16	scanner1	ecelerity: [Brightmail]	(DEBUG: 20132. 3082545888):	[src/sms_ dpp.c:359: dpp_init]	DPP system initialized
	Jan 15 11:37:05	scanner1	bmserver: [Brightmail]	(DEBUG: 20516. 3066324672):	[src/rhk_ hint.c:497: rhk_hint_ parse]	rhk hint for rule 43731290 has been successfully parsed
	Jan 15 11:38:05	scanner1	jlu-controller: [Brightmail]	(INFO: 21145. 3071248064):	[54038]	AV definitions are up-to-date.
	Jan 15 11:41:39	scanner1	imrelay: [Brightmail]	(DEBUG: 1668. 2723883952):	[System Resource Monitor. cpp:199: Perform Resource Check]	Current imrelay virtual memory size is 483004416 bytes.

Standard prefix	Date and time	Scanner host name	Process: [Brightmail]	(LogLevel: PID: ThreadID):	[EC/Source: #:Function]	Message
See "Standard prefix for Scanner logs sent to remote syslog" on page 637.	Date in the format month date. Time in the format hour:minute. The time is in 24-hour clock notation. The date and time is the date and time that the log message was recorded on the Scanner.	Name of the Scanner on which the log message was created.	Process that generated the log message, a colon, a space, and the name Brightmail in square brackets.	In parenthesis, the log level, a colon, the process ID, a period, and the thread ID, followed by a colon.	In square brackets, the error code abbreviation (EC), source code file name, line number, and function call in the source code file. For the JLU controller this value is the internal error code number that corresponds to the log message.	Log message.

The following is a log message for one event as it would appear on the remote syslog.

```
01-15-2009 11:38:16 Local2.Debug 10.217.32.13 Jan  15 11:37:05
scanner1 bmserver: [Brightmail] (DEBUG:20516.3066324672):
[src/rhk_hint.c:497:rhk_hint_parse] rhk hint for rule 43731290
has been successfully parsed
```

Log format of mail transfer agent for remote syslog

You can configure Symantec Brightmail Gateway to send Scanner log data to a remote syslog. If remote syslog is enabled, MTA log data is sent to the remote syslog. All log messages that are sent to a remote syslog have the same prefix text.

See "Standard prefix for Scanner logs sent to remote syslog" on page 637.

The MTA log messages that are sent to the remote syslog take the following form:

Standard prefix	Date and time	Scanner host name	Process:[PID]	Message
Date, time, facility, log level, and IP address	`Jan 15 11:39`	`scanner1`	`ecelerity: [21911]`	`THPL-00150: Defer_queue_ suspect_bad_message thread -1696945232 starting`
The facility for MTA messages is always `mail`. See "Standard prefix for Scanner logs sent to remote syslog" on page 637.	Date in the format month date. Time in the format hour:minute. The time is in 24-hour clock notation. The date and time is the date and time that the log message was recorded on the Scanner.	Name of the Scanner on which the log message was created.	Name and process ID of the process that generated the log message. The name is always the MTA name: `ecelerity`.	Log message.

The following is a log message for one event as it would appear on the remote syslog.

```
01-15-2009 11:40:34 Mail.Debug 10.217.32.13 Jan  15 11:39:23
scanner1 ecelerity: [21911] THPL-00150: Defer_queue_suspect_bad_message
thread -1696945232 starting
```

Log format of message audit logs for remote syslog

You can configure Symantec Brightmail Gateway to send message audit log data to a remote syslog. All log messages sent to a remote syslog have the same prefix text.

See "Standard prefix for Scanner logs sent to remote syslog" on page 637.

The message audit log messages sent to the remote syslog take the following form:

Standard prefix	Date and time	Scanner host name	Process:	Message
Date, time, facility, log level, and IP address	`Jan 15 15:42`	`scanner1`	`ecelerity:`	`1230876822|0ad9200d-b7b61ae00000 5b81-00-495db08c9df2| DELIVER|10.217.32.13| rashu1@symantecs.org`

Standard prefix	Date and time	Scanner host name	Process:	Message
See "Standard prefix for Scanner logs sent to remote syslog" on page 637.	Date in the format month date. Time in the format hour:minute. The time is in 24-hour clock notation. This is the date and time that the log message was recorded on the Scanner.	Name of the Scanner on which the log message was created.	Name of the process that generated the log message.	Log message. See "Format of audit logs" on page 647.

The following is a log message for one event as it would appear on the remote syslog.

```
01-15-2009 11:44:53 Local3.Info 10.217.32.13 Jan  15 15:42:42
scanner1 ecelerity: 1230876822|0ad9200d-b7b61ae000005b81-00-
495db08c9df2|DELIVER|10.217.32.13|rashu1@symantecs.org
```

Enabling the Message Audit Log

By default, the Message Audit Log is disabled. You must enable this feature before any auditing information is available for viewing or searching. It is important to realize that storage for message auditing can become large, and searching the logs can create high demand for Scanner processing time.

See "About message audit logging" on page 646.

To enable the Message Audit Log

1. In the Control Center, click **Administration > Settings > Logs**.
2. On the Local tab, under Message Audit Logs, check **Enable message logs**.
3. Click **Save**.

Configuring log levels

You can specify the amount of log data to store for Scanners components. You can set Scanner log levels in five increments from error level to debug level. Each succeeding log level after errors includes log data for all of the previous log levels.

See “About logs” on page 625.

For example, if you set the log level for Conduit to the notice level, Conduit errors and warnings events are stored also. You can only view Conduit log data for the notice, errors, and warnings events. Conduit information and debug events are not stored and are not viewable in the Control Center or syslog in this example.

The level that you select should contain the type of information that you want to monitor. Avoid the information and debug log levels because they can consume large amounts of disk space and require extensive processing resources to create log reports.

You can modify logging levels at any time. For instance, if you need to troubleshoot a problem, you can increase the logging level. When you resolve the problem, you can return the log level to its previous setting.

Note: You may need to restart Scanner services after changing log levels.

Warning: High logging levels are those that produce high volumes of data but tend to be of a non-urgent nature. Low logging levels are those that produce lower volumes of data but tend to be of a more urgent nature.

High log levels (such as debug) can affect the logging process. If the Scanner disk cache nears or reaches capacity, the system switches to reduced logging mode or halted logging mode. In reduced mode, only urgent logs are copied to the Control Center log database. In halted mode, no log files are copied to the log database. You must clear disk cache, purge log database, and adjust log levels to resume normal logging.

See “About log disk space alerts” on page 633.

See “About maintaining adequate disk space” on page 689.

See “Managing the log database size” on page 630.

Table 19-25 lists the available log levels.

Table 19-25 Log levels

Level	Description
Errors	Provides the most important information. This level provides the least amount of log information.

Table 19-25 Log levels *(continued)*

Level	Description
Warnings	Provides warning and errors level data. This level is the default log level for all Scanner components (local and remote).
Notices	Provides notice information and warnings and errors level data.
Information	Provides informational messages and warnings, errors, and notices data.
Debug	Provides debugging information and warnings, errors, notices, and information data. This level provides the greatest amount of log information. **Warning:** Consult Symantec Technical Support before you use this log level.

See "Log types" on page 626.

You can specify log levels for local Scanners and remote Scanners. You can configure different logging levels for each local Scanner, or you can propagate settings to all of your local Scanners. You must have Full Administration rights and Manage Status and Logs modify rights to change log levels. You cannot change the log levels for the Control Center and Spam Quarantine.

See "Viewing log files" on page 627.

To configure log levels for specific local Scanner

1. In the Control Center, click **Administration > Settings > Logs**.
2. Click the **Local** tab.
3. Under **Local Logging**, check **Enable local logs for components of the following host** and click a host for which to store local logs.
4. Use the **Component Local Log Levels** drop-down lists to select the log level for each component.
5. Click **Save**.

To configure log levels for remote Scanners

1. In the Control Center, click **Administration > Settings > Logs**.
2. Click the **Remote** tab.

3 Under **Remote Logging**, check **Enable Syslogs for the following host** and click a host from which to send log data.

4 Under **Syslog Settings**, specify the host name, port, and protocol for the remote syslog.

5 Use the **Component Remote Log Levels** drop-down lists to select the log level and the facility for each component.

See "Configuring remote logging to syslog" on page 637.

6 Click **Save**.

About message audit logging

Symantec Brightmail Gateway provides a message auditing component that lets you search for messages and find out what has happened to them. When enabled, the Message Audit Log provides administrators with a trail of detailed information about every message that has been accepted and processed by a Scanner. Auditing information is used to track what decisions were made within a single Scanner framework. The Message Audit Log is not intended to replace debug or information level logging. Unlike standard Scanner logging, the Message Audit Log provides information specifically associated with a message.

Note: Log entries for messages are created after all policy actions applicable to a message have taken place. Because some actions, like Forward a copy of the message and Add BCC recipients, modify the envelope, it can be difficult to distinguish between the original and later email recipients.

Note: Messages that are rejected by the Spam Quarantine because they exceed the size limit appear in the Message Audit Log with no indication of the rejection. Instead, the rejection is recorded in the BrightmailLog.log file with the associated Audit ID that matches the entry in the Message Audit Log.

See "Checking the Control Center error log" on page 677.

For a description of the logged information, see the search instructions.

See "Searching for a message in the Message Audit Log" on page 652.

Note: The Message Audit Log provides information on each message received by each recipient. For example, if the same message is received by 10 recipients, you see 10 entries in the Message Audit Log. The number of messages that a query can return is limited to 1,000. However, to reach this limit Symantec Brightmail Gateway counts multiple entries for the different recipients of the same message as one message.

Enabling message audit logging results in approximately 800 bytes of audit logs per message. Message audit logging can cause performance and storage problems if your site receives more than 1,000,000 messages per day.

Format of audit logs

Each software component that writes message audit information writes to its own audit log. Each entry in the audit log consists of at least three fields.

See "About message audit logging" on page 646.

See "Searching for a message in the Message Audit Log" on page 652.

See malquery on page 780.

Table 19-26 describes the audit log parameters.

Table 19-26 Format of audit logs

Field	Description
UTC time stamp	The current time in UTC, encoded in UNIX epoch time format
UID	The unique audit ID for the message (not UNIX user identifier). The format of the audit ID is: <First IP address in hex>-<ThreadID & ProcessID in hex>-<Protected 2 digit counter>-<UTC Time in hex>-<Random number, seeded with millisecond timer, in 8 character padded hex>. The audit ID is always 41 bytes in length. The audit ID is stored in the message in a X-AuditID: header.
Event ID	The type of audit event being logged. Some event IDs are followed by parameters. See "Audit log events" on page 648.

Note: The order of events for a given UID in the audit log is not relevant. The data is intended to be gathered and interpreted independent of any implied order.

Audit log events

These audit log events can be queried using the Control Center or the `malquery` command.

See "Searching for a message in the Message Audit Log" on page 652.

See malquery on page 780.

ACCEPT

An inbound or outbound mailflow has accepted the message.

Logging module	Accepting MTA
Parameters	`<IP address>:<port>`
Count	One ACCEPT per audit ID
State	The message is at least in the inbound or outbound queue

Example:

```
UTC|UID|ACCEPT|10.240.190.2:25
```

ATTACH

The viruses or worms found in the message attachments.

Logging module	Brightmail engine
Parameters	The filenames of one or more attachments
Count	Potentially multiple ATTACH events per audit ID
State	The message has been filtered, and is at least in the delivery queue.

Example:

```
UTC|UID|ATTACH|virus.exe|internal.doc
```

DELIVER

Message was handed off to another MTA.

Note: There may be more than one DELIVER event for a single audited message, or recipient, as delivery can succeed or fail per recipient, and per target destination.

Logging module	MTA or Brightmail Engine
Parameters	■ The IP address and port of the MTA the message was to, in the form of `<IP address>:<port>` ■ One or more email addresses to which the message was successfully delivered
Count	Multiple DELIVER events per audit ID; potentially multiple DELIVER events per recipient
State	The message has been filtered and is at least in the delivery queue; some or all recipients have been delivered to the next hop

Example:

```
UTC|UID|DELIVER|10.240.190.45:25|recip@domain.tld|recip1@domain.tld
```

IRCPTACTION

The actions taken per recipient after distribution list expansion.

On the command line, specify `-e IRCPTACTION="search string"` to search for IRCPTACTION by itself or specify `-e RCPTS ="search string"` to search for both IRCPTACTION and ORCPTS

Logging module	Message transformation engine
Parameters	■ The recipient's email address ■ One or more action summaries for the recipient
Count	Multiple IRCPTACTION events per audit ID; potentially multiple IRCPTACTION events per recipient.
State	Message has been filtered, and is at least in the delivery queue

Example:

```
UTC|UID|IRCPTACTION|user2@domain.com|Subject Markup|Forward
```

MSGID

The contents of the Message-ID header, not the unique message ID added for the message audit logs.

Logging module	Brightmail engine
Parameters	The RFC822 Message-ID header
Count	One per audit ID
State	The message has been filtered, and is at least in the delivery queue

Example:

```
UTC|UID|MSGID|200607052345.LBB05394@symantecs.org
```

ORCPTS

The original recipients of a message before alias, distribution list, or masquerading changes as received by the accepting MTAs.

On the command line, specify `-e ORCPTS="search string"` to search for ORCPTS by itself or specify `-e RCPTS ="search string"` to search for both IRCPTACTION and ORCPTS

Logging module	Accepting MTA
Parameters	One or more RFC821 recipient addresses
Count	One per audit ID
State	The message is at least in the inbound or outbound queue

Example:

```
UTC|UID|ORCPTS|user1@example.com|user2@example.com
```

SENDER

The original sender of a message as received by the accepting MTAs.

Logging module	Accepting MTA
Parameters	One RFC821 sender address
Count	One per audit ID

State	Message is at least in the inbound or outbound queue

Example:

```
UTC|UID|SENDER|muir@symantecs.org
```

SOURCE

The origin, internal or external, of the accepting MTAs. Example:

Logging module	Accepting MTA
Parameters	The string "external" or the string "internal"
Count	One per audit ID
State	If "external," the message is at least in the inbound queue; if "internal," the message is at least in the outbound queue

Example:

```
UTC|UID|SOURCE|external
```

SUBJECT

The message subject line.

Logging module	Brightmail engine
Parameters	The subject line of the message
Count	One per audit ID
State	Message has been filtered, and is at least in the delivery queue

Example:

```
UTC|UID|SUBJECT|Make $$$ Fast!!!
```

VERDICT

The verdict, policy group, and filtering policy that was triggered on the message per recipient.

Logging module	Brightmail engine

Parameters	■ The intended recipient ■ The verdict for the intended recipient ■ The policy groupused to determine the filtering policy ■ The filtering policy used
Count	Multiple VERDICT events per message
State	Message has been filtered and is at least in the delivery queue

Example:

```
UTC|UID|VERDICT|user1@domain.tld|spam|default|Default Spam Policy
```

VIRUS

The viruses or worms found in the message.

Logging module	Brightmail engine
Parameters	One or more virus or worm identification strings
Count	One per audit ID
State	Message has been filtered and is at least in the delivery queue

Example:

```
UTC|UID|VIRUS|sobig|notsobig.win32
```

Searching for a message in the Message Audit Log

A query facility is provided to search the log to determine if one or more messages meet the criteria for the message you want to find. The **Status > SMTP > Message Audit Logs** page enables you to specify either one or two criteria and related supplementary information as follows:

Host	One or more Scanners running the Symantec Brightmail Gateway software. In order to find all details about a message, search on all attached Scanners.
Time range	Period of time for the search to query the audit log. While it is possible to search for longer periods, it is recommended that message searches not exceed one week.

Mandatory filter	Select the type of information for filtering messages. See Table 19-27.
Mandatory filter value	Enter a string that corresponds to the Mandatory filter type you selected. For example, if you chose to filter messages by sender, enter a valid email address here.
Optional filter	Select from the list of optional filtering criteria. See Table 19-28.
Optional filter value	If appropriate, enter a string or choose a value that corresponds to the Optional filter type you selected. For example, if you chose to filter messages by Connection IP, enter a valid IP address here. Or, if you choose to filter messages by Action taken, select the action for which you want to find messages.
Clear Filters	Clear the current filtering criteria from memory.
Display Filtered	Search for and display messages that fit your criteria.

Table 19-27 describes the items you can choose for your single required filter.

Table 19-27 Choices for the mandatory search criteria

Criteria	Description
Sender	Name of the message sender. Specify <> to filter for messages that do not contain Sender names.
Recipient	Name of the message recipient
Subject	Message subject
Audit ID	Unique identifier generated by Symantec Brightmail Gateway and included as a message header
Connection IP	IP address of the connecting server. In cases where Symantec Brightmail Gateway rejects an IP connection, this results in a row with the sender identified as none. Message details consist of the IP address and the reason for rejection.

Table 19-28 describes the items you can choose for your single optional filter.

Table 19-28 Choices for the optional search criteria

Criteria	Description
Sender	Name of the message sender. Specify <> to filter for messages that do not contain Sender names.

Table 19-28 Choices for the optional search criteria *(continued)*

Criteria	Description
Authenticated sender	Name of an authenticated sender
Recipient	Name of the message recipient
Subject	Message subject
Message ID	Unique identifier typically generated by the email software initiating the sending of the message and included as a message header. Because the Message ID is not generated by Symantec Brightmail Gateway, the uniqueness of the ID cannot be guaranteed. Spammers have used this header to mask the identity of a message originator.
Verdict	Verdict and/or other characteristics of a message such as Message has malformed mime. When this filter option is selected, a list of possible verdicts appears in the Option filter value drop-down list. Use these values to filter messages that resulted in a given verdict.
Untested verdict	An available verdict for which the Scanner did not test. A drop-down list of verdict choices is provided. These verdicts are available to inform you of any policies that are not currently associated with a recipient/sender group that would have been triggered for a given message.
Action taken	What happened to the message. When this filter option is selected, a list of possible actions appears in the Option filter value drop-down list. Use these values to filter messages that triggered policies that applied the given action. If you select Reject message from the Option filter value drop-down list, the reason for rejection appears in the message detail. ■ Rejected message for exceeding size limit ■ Rejected message due to lost connection ■ Rejected message for all recipients ■ Rejected message The Rejected message reason indicates that Symantec Brightmail Gateway was unable to tell why the message was rejected. In many cases this occurs when a connection is lost at an early stage, or when the sending mail server stops transmission at an incomplete state.
Connection IP	Connection IP used to receive the message.
Target IP	IP address of the message destination.

Table 19-28 Choices for the optional search criteria *(continued)*

Criteria	Description
Policy group	Name of the group (either the recipient's group or the sender's group) that determined which filter policy applied to the message.
Filter policy	Name of the filter policy applied to the message.
Virus	Name of the virus attached to the message.
Attachment	Name of a message attachment.
Suspect attachment	Name of a message attachment that triggered a content filtering policy.
Reason for unscannable verdict	Reason that the message matched the "If a message is unscannable for viruses" condition. A drop-down list of unscannable reasons is provided.
Source	Whether the message is internal or external.

While searching, the following rules are used:

- No more than 1,000 messages are allowed per search on each Scanner being searched.
- Freeform text fields are non-case-sensitive substring searches.

Note: The Message Audit Log provides information on each message received by each recipient. For example, if the same message is received by 10 recipients, you see 10 entries in the Message Audit Log. To reach the limit of 1,000 messages returned, Symantec Brightmail Gateway counts multiple entries for the different recipients of the same message as one message.

Email messages that fail delivery are tracked as delivery failures in the Message Audit Log. For example, messages to non-existent users that bounce are considered delivery failures. Delivery failures are indicated with a Delivery Failure heading on the Audit Logs page in the Delivery section. In addition to being indicated on the Audit Logs page, undelivered messages are logged with the new DELIVERY_FAILURE audit log event. DELIVERY_FAILURE events are logged in the following format: `utc|uid|DELIVERY_FAILURE|recipient|reason`

The **Actions** column indicates actions taken by the Scanner on messages, but does not indicate actions taken by administrators or users on messages. For example, if an administrator or user releases a message from Spam Quarantine, this activity is listed under **Spam Quarantine**, not **Actions**.

To search the message audit log and view message details

1. In the Control Center, click **Status > SMTP > Message Audit Logs**.
2. Select the Scanner whose logs you wish to search from the **Hosts** drop-down list, or select **All Scanners**.
3. Complete the desired search criteria.
4. Click **Display Filtered**.

   Use the **Entries per page** drop-down list to specify the number of records to show per page. Use the **Display _ of _** drop-down list to choose a range of data to display.
5. Click a message recipient in the To column to view processing details on that message.

To search the message audit log for content filtering incidents

1. In the Control Center, click **Status > SMTP > Message Audit Logs**.
2. Select the Scanner whose logs you want to search from the **Host** drop-down list, or select **All Scanners**.
3. Choose a selection from the **Mandatory filter** drop-down list and enter an appropriate value in the **Mandatory filter value** field.
4. Choose **Action taken** from the **Optional filter** drop-down list.
5. Choose either **Create an informational incident** or **Create a quarantine incident** from the **Optional filter value** drop-down list.
6. Click **Display Filtered**.

   Use the **Entries per page** drop-down list to specify the number of records to show per page. Use the **Display _ of _** drop-down list to choose a range of data to display.
7. Click a message recipient in the To column to view processing details on that message.

Exporting Message Audit Log data

After you select your filter criteria and click **Display Filtered**, you can export the log data to a CSV file.

To view a CSV file that contains double-byte characters in Microsoft Excel, specify a comma-delimited , UTF-8 file in the MS Excel Text Import Wizard. Alternatively, you can open the CSV file in a text editor that can convert UTF-8 to Unicode, such as Notepad, and save the CSV file as Unicode.

To export Message Audit Log data

1. In the Control Center, click **Status > SMTP > Message Audit Logs**.
2. Select the Scanner whose logs you want to search from the **Host** drop-down list, or select **All Scanners**.
3. Complete the search criteria.
4. Click **Display Filtered**.
5. Click **Export CSV**.
6. Click the **File Encoding** drop-down list to choose a character encoding for the CSV file.

   ISO-8859-1 and UTF-8 are appropriate for European languages. Windows-31j, EUC-JP, and ISO-2022-JP are appropriate for Japanese.
7. Click the **CSV Delimiter** drop-down list and choose a delimiter for the CSV file.

   Symantec Brightmail Gateway places the chosen delimiter between entries in the file.

About message queues

A message queue is a temporary holding area for messages before they reach their destination. The messages queues are: inbound, outbound, and delivery. The message queue size fluctuates based on mail flow.

If a queue continually grows without decreasing in size, there is a problem with message delivery. You can view the messages that are queued at any time to troubleshoot an issue. After you stop the message flow, you can attempt to redeliver the messages by flushing the queue. You can also delete the messages that block the queue. When you resolve the issue, you can restart the message flow.

See "Monitoring message queue size and volume" on page 612.

See "Viewing queued messages" on page 658.

See "Deleting queued messages" on page 660.

See "Stopping the mail flow" on page 660.

See "Flushing message queues" on page 661.

See "Troubleshooting the message queue" on page 662.

Viewing queued messages

You can view the messages that are in the message queues. At a minimum, you must specify the host and the queue that you want to view: inbound, outbound, or delivery. Symantec Brightmail Gateway also has filtering options that let you further customize the queued messages to view. Based on your filter criteria, Symantec Brightmail Gateway displays the messages that are in that queue at that moment. If there is an error associated with the message, it appears in the **Message** column.

You may have a situation in which you want to determine if a specific message is in the queue. You can filter for messages to or from specific email addresses.

For the delivery queue only, you can also search for messages from a specific route or set of routes. You can also reroute messages from one route to a different route.

See “Rerouting messages in the delivery queue” on page 658.

You can perform wild card searches for senders or recipients. For the delivery queue you can also perform wild card searches for custom routes. Search is not case sensitive.

If you do not stop mail flow, the messages in the queue can continually fluctuate as new messages enter the queue and older messages exit. You can refresh the view as needed.

You must have Full Administration rights or Manage Status and Logs view or modify rights to view message queues.

To view queued messages

1. In the Control Center, click **Status > SMTP > Message Queues**.
2. On the Message Queues page, select a host and queue.
3. Type search values for the fields that are provided.
4. Click **Display Filtered**.
5. To clear the **To** and **From** fields to begin a new search, click **Clear Filters**.

   All of the other drop-down menu selections that you made are retained.

To refresh the view

◆ On the Message Queues page, click **Refresh**.

Rerouting messages in the delivery queue

You can search for messages in the delivery queue from a specific route or set of routes. You can also reroute messages from one route to a different route.

See “Viewing queued messages” on page 658.

You can perform wild card searches for senders, recipients, or custom routes. Search is not case sensitive.

To reroute messages in the delivery queue

1 In the Control Center, click **Status > SMTP > Message Queues**.

2 On the Message Queues page, select a host and the **Delivery** queue.

3 In the **Route** field, select one of the following:

Default Local Route	The default route for local domains.
Default Non-local Route	The default route for non-local domains.
Control Center	The route this Scanner uses to communicate with the Control Center.
DLP	The route this Scanner uses to communicate with a Symantec Network Prevent server.
Custom	A route that you specify in the **Custom route** field. After you select **Custom**, the **Custom route** field appears.

4 If you selected **Custom**, enter the hostname or IP address:port for the route you want to search for in the **Custom route** field. If you made another selection, skip this step.

If you specify a CIDR block that returns messages from more than one route, you cannot reroute messages. If your filtered messages are all from the same route, the **Reroute All** button becomes active.

5 Optionally, complete additional search fields.

6 Click **Display Filtered**.

7 Under **Reroute all the messages originally destined for the filtered route**, type the hostname, IP address:port, CIDR block, or subnet for the new route in the **New Route** field.

8 If you typed a hostname, you can check **MX Lookup**, if desired.

9 Click **Reroute All**.

Deleting queued messages

When you view a message queue, you may find that a message has blocked the queue. You can delete that message so that the messages behind it can pass through the queue.

You can access the messages that are queued, and save them, via the `mta-control` command.

See mta-control on page 792.

Once you delete a queued message, you cannot retrieve it.

You can delete all of the messages in a queue. When you choose to **Delete All**, you delete the messages that are in that queue at that moment in time, including messages on additional pages. Before you delete all queued messages, you might want to temporarily stop the mail flow.

See "Stopping the mail flow" on page 660.

You must have Full Administration rights or Manage Status and Logs modify rights to delete queued messages.

See "Adding administrators" on page 682.

To delete queued messages

1. In the Control Center, click **Status > SMTP > Message Queues**.
2. Specify the queue that you want to view.

   See "Viewing queued messages" on page 658.
3. Do any of the following:

To delete a message	Select the message that you want to delete and click **Delete**.
To delete all of the messages in the queue, including messages on additional pages	Click **Delete All**.

Stopping the mail flow

You may need to stop the flow of mail (for example, before you flush a message queue).

See "Flushing message queues" on page 661.

See "Deleting queued messages" on page 660.

Consider the following implications before you stop the mail flow:

- If you stop the inbound mail flow, no inbound mail is accepted. The mail in the inbound message queue is not scanned, and mail delivery continues.
- If you stop the outbound mail flow, no outbound mail is accepted. The mail in the outbound message queue is not scanned, and mail delivery continues.
- If you stop the delivery mail flow, no mail is delivered to downstream local or remote mail servers. Mail in the inbound message queue and outbound message queue is scanned and accumulates in the delivery message queue.

See "MTA and message queue behavior" on page 613.

See "Turning off an appliance " on page 690.

See "Managing services and MTA operations" on page 110.

To stop the mail flow

1. In the Control Center, click **Status > SMTP > Message Queues**.
2. In the **Host** drop-down list, select a server.
3. From the **Queue** drop-down list, select **Inbound** and click **Display Filtered**.
4. If the queue is started, click **Stop**.
5. From the **Queues** drop-down list, select **Outbound** and click **Display Filtered**.
6. If the queue is started, click **Stop**.
7. From the **Queue** drop-down list, select **Delivery** and click **Display Filtered**.
8. If the queue is started, click **Stop**.

Flushing message queues

When you flush a message queue, you instruct the MTA to try to resend the messages that were deferred due to delivery problems. You may want to flush your inbound, outbound, and delivery email queues before you turn off an appliance.

Before you flush a message queue, temporarily stop the flow of inbound email.

Note: After you stop the mail flow or determine that a mail flow has stopped, Symantec recommends that you wait two minutes before flushing that message queue.

See "Stopping the mail flow" on page 660.

See "MTA and message queue behavior" on page 613.

See "Managing services and MTA operations" on page 110.

You must have Full Administration rights or Manage Status and Logs modify rights to flush message queues.

You can also flush email queues from the command line.

See mta-control on page 792.

To flush email queues from the Control Center

1. In the Control Center, click **Status > SMTP > Message Queues**.
2. Choose a host from the **Host** drop-down list.
3. In the **Queue** drop-down list, select the queue you want.

   See "Viewing queued messages" on page 658.
4. Click **Flush All**.

   All messages, including those on additional pages, will be flushed.
5. Wait until the message queue is empty (repeat the previous step as needed).

Troubleshooting the message queue

When a message queue becomes too large, Symantec Brightmail Gateway can become unresponsive or crash. To attempt to deter this issue, by default, Symantec Brightmail Gateway defers new messages when the queue is full. As a best practice, you should leave this setting enabled and keep queue limits below the recommended default thresholds.

See "Configuring SMTP advanced settings" on page 95.

If you experience issues with the message queue, try the following:

- Make sure that your downstream delivery host is functioning and accepting mail.
  See "Viewing the status of software and services" on page 609.
- Configure the MTA to reject incoming messages.
  See "Managing services and MTA operations" on page 110.
- Watch and monitor the queues until they reach acceptable limits.
  See "Viewing queued messages" on page 658.

Repeat these measures until the issue is resolved.

Viewing IM users that are signed on

You can view a list of both the registered and unregistered IM users that are currently signed on. IM users that are currently signed on are known as active IM users.

See “About registering IM users” on page 297.

You can view all active IM users, or you can create a filter to display only the active IM users that you want to view. The filter that you create is based on the values that you specify for one or more of the available user attributes. For example, you can create a filter to display the active IM users of a specific IM network, specific IP address, or both.

You can also use a wildcard in your filters to further specify the active IM users that you want to view. When creating a filter for Screen Name, Email Address, or IP Address, you can use the asterisk (*) to represent one or more characters in the filter value that you specify. In addition, you can use more than one wildcard in the same filter value. For example, if you specify jsmith*@*.com as the email address, IM users jsmith1@hotmail.com and jsmith10@gmail.com appear.

You can create a filter by using the following user attributes:

IM Account

The screen name of the IM user.

You can specify a complete screen name, or you can use a wildcard to represent one or more characters within a screen name. For example, if you specify jsmith*, users jsmith1 and jsmith10 appear.

Note: The screen name that you specify cannot begin with a wildcard. For example, you cannot specify *smith.

Email Address

The email address of the IM user.

You can specify a complete email address, or you can use a wildcard to represent one or more characters within the email address. For example, if you specify jsmith*@hotmail.com, users jsmith1@hotmail.com and jsmith10@hotmail.com appear.

Note: The email address that you specify cannot begin with a wildcard. For example, you cannot specify *smith@hotmail.com.

IM Network

The IM network of the IM user. These include:

- AOL
- Google Talk
- MSN Messenger
- Yahoo IM

You can select one of the IM networks from the drop-down list, or you can select All networks.

IP Address

The IP address of the IM network from which the IM user is signed on.

You can specify a single IP address, or you can use a wildcard to represent one or more characters within one of the IP address's bytes. For example, if you specify 192.255.255.*, all active IM users for that network appear.

The IP address that you specify cannot:

- Contain more than 3 dots
- Contain consecutive dots
- Begin or end with a dot
- Contain a number that is greater than 255

To view active IM users

1. In the Control Center, click **Status** > **Instant Messaging** > **Active Users**.
2. (Optional) Under Filter, type the value of one or more user attributes that you want to use to display the active IM users.

   The attributes include:

   - IM Account
   - Email Address
   - IM Network
   - IP Address

   See “Viewing IM users that are signed on” on page 662.
3. Do one of the following:
   - Click **Display Filtered** to display the active IM users that are based on the values of your current filter.
   - Click **Clear Filters** to clear the values of your current filter.

You can specify the number of active IM users that you want to appear on each page of your search results. This number is based on increments of 10, 25, 50 or 100. Based on the increment that you specify, you can navigate immediately to any page that contains additional search results. For example, if you specify that you want 10 users to appear on each page, and your search results yield 100 users, you can navigate to the page that contains users 21-30. Using the control buttons, you can also navigate to the first page, the previous page, the next page, or the last page at any time.

The search results contain the following information:

IM Account	The IM screen name of the user.
Email Address	The corporate email address of the user. The email address appears only if the user is registered. See “About registering IM users” on page 297.
IM Network	The IM network of the user.
Client Version	The version number of the user's IM client.
Status	The current status of the user's IM client, such as "online" or "away." The status is based on the various status types that are available for each IM client, including custom status types.
IP	The IP address from which the IM user is signed on.
Duration	The amount of time that the IM user has been signed on. The duration is based on days, hours, minutes, and seconds.

By default, the active IM users that appear in your search results are sorted in ascending order by the users' IM Account name. However, you can re-sort your search results in ascending or descending order by IM Account, Email Address, or IM Network.

After you display your search results, you can export them to a comma-separated values (CSV) file. The CSV file contains all the users that appear in your search results (not just the users that appear on the current page).

To work with the search results

1. From the Entries per page drop-down list, select the number of IM users that you want to appear on each page.

   The default setting is 25.

2. To view another page, do one of the following:

   - From the Display drop-down list, select the page that you want to view.
   - Click the **First Page**, **Previous Page**, **Next Page**, or **Last Page** button.

3. (Optional) To re-sort the search results, click the column heading by which you want to sort, and then click the ascending or descending arrow.

4. (Optional) To export the search results to a CSV file, click **Export CSV**.

Viewing the connection status of your IM networks

You can view the connection status of each IM network that you support from each Scanner that is in your corporate network. This feature may be helpful after you initially configure your DNS, after you configure a Scanner, or when IM connectivity problems are reported.

Each Scanner in your corporate network regularly attempts to connect to the public IM network servers that Symantec Brightmail Gateway supports.

Based on the success or failure of these connection attempts, the following results appear on the IM Network Status page:

- If a Scanner establishes a connection with an IM network, a green check mark appears in the column for that network. This includes IM networks to which you blocked access.
  See "Blocking access to an IM network" on page 307.
- If a Scanner is unable to establish a connection with an IM network, `Unknown` appears in the column for that network.
- If a Scanner is disabled, a dash appears in the column for each IM network that is associated with that Scanner.
- If an IM client is directed back to the Scanner instead of the Internet, `Loopback` appears in the column for the network that is associated with that client.

  This typically means that your DNS is configured incorrectly. See the *Symantec Brightmail Gateway Installation Guide*.

To view IM network status

- In the Control Center, click **Status** > **Instant Messaging** > **Network Status**.

Chapter 20

Administering your product through the Control Center

This chapter includes the following topics:

Configuring Control Center settings

About simultaneous Control Center access

Multiple end users can access the Control Center at the same time, for example to review messages in Spam Quarantine.

Multiple administrators can access the Control Center at the same time and perform administration tasks. However, each administrator may see errors in certain cases. In particular, errors can occur if each administrator attempts resource-intensive tasks in the Control Center at the same time. For example, querying message audit logs and IP reputation at the same time can cause errors.

About specifying host names for Control Center access

When you specify host names for Control Center access, you enable the Control Center to let clients connect based on the Control Center's DNS perspective. If the client's IP address resolves to a name that matches an allowed host name (a “reverse lookup”), then the Control Center permits access to the client.

The owner of a netblock (a range of ip addresses usually belonging to the same organization) controls the reverse lookup of an IP address. Typically, users don't control what name their IP addresses resolve to. Additionally, DNS servers may have unique mappings for the same netblock. Consider the following:

- A client's authoritative DNS server has a reverse lookup record of `m1.symantecexample.com` for the client's IP address.
- The DNS that is configured to be the Control Center's primary DNS server has a reverse mapping of `dhcp23.symantecexample.com` for the same IP address.

In this case, the Control Center sees the `dhcp23.symantecexample.com` name whenever the client connects. Therefore, the name the Control Center sees is the name you type in the host access control list in the Control Center. This situation happens more frequently on private networks than on the public Internet.

See “Specifying which hosts can access the Control Center” on page 669.

Specifying which hosts can access the Control Center

You access the Control Center through a Web browser. By default, anyone with the correct address and logon information has access from any host. But you can also choose to specify which hosts can access to the Control Center. Users that attempt to log in to the Control Center from unauthorized computers receive a 403 Forbidden page message in their Web browser.

Reverse Domain Name Server (DNS) lookup must be enabled in your DNS software for this feature to work with host names.

See “About specifying host names for Control Center access” on page 669.

Note: If you make an error when you type the host name, you block all access to the Control Center. If this situation occurs, use the command-line `clear bcchostacl` command to clear the list of computers that are permitted to access the Control Center. See delete on page 753.

To specify which hosts can access the Control Center

1. In the Control Center, click **Administration > Settings > Control Center.**
2. Click the **Access** tab.

3 Under **Control Center Access**, do one of the following tasks:

To permit any host access to the Control Center	Check **All hosts**.
To assign specific hosts to access the Control Center	Check **Only the following hosts**, and then type a host name, IP address, IP address with subnet mask, or Classless Inter-Domain Routing (CIDR) netblock. Specify additional computers or networks as needed. Hosts that are not in this list are not able to access the Control Center.

4 Click **Add**.

5 Click **Save**.

Setting the locale encoding and fallback encoding

Configure the Control Center for single- and double-byte character sets and for appropriate number, date, and time settings with the Locale setting. For example, the Locale setting affects the format of email messages that the Control Center sends, such as notifications, alerts, and reports.

You can also select a fallback encoding option for the Control Center to use in cases when the language identification feature is unable to correctly determine the language of a quarantined email message.

Language identification may fail if the message headers or body contain any of the following items:

- Raw 8-bit characters with missing or corrupted encoding information
- Too few characters with missing or corrupted encoding information
- Only a few characters with a mix of two or more types of encodings

Set the fallback encoding to the most common encoding that is used in your region or country.

See "Customizing the date format and time format" on page 703.

To set the locale encoding and fallback encoding

1 In the Control Center, click **Administration > Settings > Control Center.**

2 Click the **Locale** tab.

3 Under **System Locale and Fallback Encoding**, click the **System locale** drop-down list and select an encoding.

4 Click the **Quarantine fallback encoding** drop-down list and select an encoding.

5 Click **Save**.

Designating a Control Center certificate

You can designate a user interface HTTPS certificate through the Control Center. This certificate enhances the security for the Control Center and those logging into it.

See "Bypassing the security warning when you access the Control Center" on page 671.

To designate a Control Center certificate

1 In the Control Center, click **Administration > Settings > Control Center.**

2 Click th **Certificates** tab.

3 Under **Control Center Certificate**, click the **User interface HTTPS certificate** drop-down list to select the certificate that you want to use.

   See "About certificates" on page 191.

4 Click **Save**.

Bypassing the security warning when you access the Control Center

By default, the Control Center uses a demo certificate to authenticate access to the Control Center. The demo certificate causes a security warning in your browser when you access the Control Center. You can ignore the security warning and proceed to access the Control Center. However, you can install a certificate to enhance the security of the browser-to-Control Center communication and to prevent the security warning.

Determine if you want to use a self-signed certificate or a CA-signed certificate. A self-signed certificate does not provide the same level of security as a CA-signed certificate. To get a CA-signed certificate, you must submit a CSR to a Certificate Authority.

The following procedure assumes that you use a CA-signed certificate.

To bypass the security warning when you access the Control Center

1 Add a certificate in the Control Center and submit the CSR to a Certificate Authority to get a certificate.

Ensure that the hostname in the CSR matches the hostname of the Control Center.

See "Requesting a Certificate Authority-signed certificate" on page 194.

2 Import the certificate that you receive from a Certificate Authority.

See "Importing a Certificate Authority-signed certificate" on page 197.

3 Install an intermediate certificate, if needed.

See "Adding a CA certificate" on page 193.

4 Assign the certificate as the Control Center HTTPS certificate.

See "Assigning a user interface HTTPS certificate to the Control Center" on page 201.

5 Access the Control Center with the fully qualified domain name that you supplied on the CSR.

Configuring the Control Center listener port

By default, Spam Quarantine, Suspect Virus Quarantine, and content incident folders accept messages from the Scanner on port 41025. However, you can change this port if necessary. Only the administrators that have Full Administration rights or Manage Settings modify rights can modify these settings. You do not need to change any Scanner settings to match the change in the listener port.

You can also disable the listener port. Disable the listener port if your computer is not behind a firewall, and you are concerned about security risks. If you disable the listener port, disable any policies that quarantine messages. Otherwise, quarantined messages back up in the delivery mail flow queue until the expiration time elapses and then bounce back to the original sender.

See "About content incident folders" on page 438.

See "About quarantining spam" on page 258.

See "About quarantining suspected viruses" on page 229.

To modify or disable the listener port through which the Control Center accepts messages

1 In the Control Center, click **Administration > Settings > Control Center**.

2 Click the **Listeners** tab.

3 Under **Listener Port** in the **Port** box, do one of the following:

- Type the new port number.
- Type **0** to disable the listener port.

4 Click **Save**.

Configuring Control Center SMTP settings for alerts and reports

The Control Center sends the following information to designated email addresses and repositories at your site:

- Alert notifications
- Reports
- Spam Quarantine messages

You must supply the SMTP host IP address and port number to which you want the Control Center to send information.

Note: Symantec Brightmail Gateway verifies that the product version that runs on the Control Center and the Scanner are the same. If the product versions are not the same, Symantec Brightmail Gateway issues an error message. You should perform a software update on the Scanner and define the SMTP host on the Control Center Setting page again.

See "Updating your software" on page 700.

To configure Control Center SMTP settings for alerts and reports

1 In the Control Center, click **Administration > Settings > Control Center**.

2 Click the **SMTP** tab.

3 Under **Control Center SMTP Settings**do one of the following:

To specify that the email that the Control Center generates should use the non-local relay to send email	Click **Use existing non-local relay settings**. See “Configuring Scanner inbound email delivery settings” on page 91.
To specify the IP address or fully-qualified domain name of a computer that has a working MTA on it	Click **Define new host**. Change this setting from the default if a Scanner is not installed on the same appliance as the Control Center. Specify the port to use for SMTP. The default is 25.

4 Click **Save**.

Specifying a custom user Login help page

By default, when users click the **Need help logging in?** link on the Control Center **Login** page, the Symantec online help window opens. You can customize the Login help with a custom Login help page url. This change only affects the Login help page, not the rest of the online help.

Create a Web page that tells your users how to log in and make it available on your network. The Web page should be accessible from any computer where users log in.

See “Viewing spam and suspected messages in quarantine” on page 261.

See “Enabling users to bypass Control Center login credentials” on page 674.

To specify a custom user Login help page

1 In the Control Center, click **Administration > Settings > Control Center**.

2 Click the **Users** tab.

3 Under **User Help** in the **Login help URL** box, type the URL to the Web page that you want to use.

To disable your custom logon help page, delete the contents of the **Login help URL** box.

4 Click **Save**.

Enabling users to bypass Control Center login credentials

Symantec Brightmail Gateway contains a feature that, when enabled, remembers users' login credentials. A **Remember me** option appears on the Symantec

Brightmail Gateway Login page when the administrator enables the **Remember me** feature. When this feature is enabled, it is available to all users, including end users.

When users check the **Remember me** option as they log onto the Control Center, Symantec Brightmail Gateway places a cookie on their computer. (No personal information is stored on the client computer or in the browser.) Thereafter, when users access the Control Center URL, they bypass the Login screen and go directly to the Symantec Brightmail Gateway Dashboard page. If the user session times out (sessions timeout after 30 minutes of inactivity), the user is directed to the Dashboard page.

Warning: When users enable this feature, anyone who has access to their computer also has access to the Control Center. This feature is not recommended for administrator logins.

The **Remember me** feature is browser-specific. For example, assume the user logs onto the Control Center through Internet Explorer and enables the **Remember me** option. If the user subsequently logs onto the Control Center through Firefox, the user must enter login credentials.

The cookie expires after the number of days that you specify, up to 180 days. However, the expiration might change if you use the strong passwords feature.

The ability to use the **Remember me** feature with strong passwords depends on the order in which the features are enabled, as follows:

Strong passwords are enabled first, then the administrator subsequently enables the **Remember me** feature.	When the **Remember me** feature is enabled, if the bypass duration is greater than 60 days (the maximum number of days permitted for strong passwords), then the login bypass duration is set to the same expiration date as the strong password. This change is not reflected in the Control Center settings.
The **Remember me** feature is enabled first, then the administrator subsequently enables strong passwords.	The **Remember me** feature is disabled in the Control Center.

See “Enforcing strong passwords” on page 680.

See “Specifying a custom user Login help page” on page 674.

Enabling users to bypass Control Center login credentials

1. In the Control Center, click **Administration > Settings > Control Center**.
2. Click the **Users** tab.
3. Under **User Help**, check **Enable 'Remember me' feature**.
4. In the **Login bypass duration** box, type the number of days in which the users' login credentials are valid.

   You can enter a value from 1 to 180. The default value is 30 days.

   If you modify this value, the change applies to users the next time they select the **Remember me** option at login.
5. Click **Save**.

Configuring the replication of end user preference data

In the Control Center, replication refers to the process by which user preferences are propagated from the Control Center to attached and enabled Scanners. Global settings in the Control Center control the replication process.

You can check the status of automated user preferences from the **Status > Scheduled Tasks** page.

See “About scheduled tasks” on page 621.

To configure the replication of end user preference data

1. In the Control Center, click **Administration > Settings > Control Center.**
2. Click the **Users** tab.
3. In the **Replication frequency** box and drop-down list, you can set the replication frequency.

   You can choose **Never** to turn off replication.
4. In the **Replication start time** drop-down lists, you can specify the time at which replication starts.
5. You can click **Replicate Now** to have LDAP data replicated to all attached and enabled Scanners immediately.
6. You can click **Delete Now** to delete all end user preference data.
7. Click **Save**.

About backing up log data

In general, there is no reason to backup log files. For troubleshooting purposes, logs that are not set to Information or Debug (which provides the most detail) have limited use. The best practice is to view and save current logs as needed and set the appropriate retention period for logging data.

See "Viewing log files" on page 627.

See "Saving log files" on page 629.

See "Managing the log database size" on page 630.

Checking the Control Center error log

You might want to periodically view the Control Center error log to troubleshoot issues. All errors that are related to the Control Center are written to the `BrightmailLog.log` file.

Each issue results in a number of lines in the error log. For example, the following lines are the result of Spam Quarantine receiving a message that is too large to handle:

```
com.mysql.jdbc.PacketTooBigException:
Packet for query is too large (3595207 > 1048576)
at com.mysql.jdbc.MysqlIO.send(MysqlIO.java:1554)
at com.mysql.jdbc.MysqlIO.send(MysqlIO.java:1540)
at com.mysql.jdbc.MysqlIO.sendCommand(MysqlIO.java:1005)
at com.mysql.jdbc.MysqlIO.sqlQueryDirect(MysqlIO.java:1109)
at com.mysql.jdbc.Connection.execSQL(Connection.java:2030)
at com.mysql.jdbc.PreparedStatement.executeUpdate
(PreparedStatement.java:1750)
at com.mysql.jdbc.PreparedStatement.executeUpdate
(PreparedStatement.java:1596)
at org.apache.commons.dbcp.DelegatingPreparedStatement.executeUpdate
(DelegatingPreparedStatement.java:207)
at com.brightmail.dl.jdbc.impl.DatabaseSQLManager.handleUpdate
(Unknown Source)
at com.brightmail.dl.jdbc.impl.DatabaseSQLManager.handleUpdate
(Unknown Source)
at com.brightmail.dl.jdbc.impl.DatabaseSQLTransaction.create
(Unknown Source)
at com.brightmail.bl.bo.impl.SpamManager.create
(Unknown Source)
```

```
at com.brightmail.service.smtp.impl.SmtpConsumer.run
 (Unknown Source)
```

To check the Control Center error log

1. In the Control Center, click **Status > Logs**.
2. In the **Component** drop-down list, select **Control Center**.
3. In the Log Files table, click **BrightmailLog.log**.
4. Open the log file or save it to your local disk.

Running network utilities from the Control Center

You can run the following network utilities from the Control Center:

Nslookup	Query for DNS info about a computer on the Internet
Traceroute	List the hosts that used to transmit Internet data between the selected host and a computer on the Internet, as well as elapsed time
Ping	Test for a response from a computer on the Internet

To run network utilities from the Control Center

1. In the Control Center, click **Administration > Hosts > Utilities**.
2. Click the **Troubleshooting** tab.
3. From the **Host** drop-down list, select a host name.
4. Under **Select Utility** area, use the drop-down lists to specify a utility name and host name or IP address.

   If you select Nslookup in the **Utility** drop-down list, you must also specify a DNS query type in the **Record type** drop-down list.
5. Click **Run**.

   The results of the operation appear in the **Results** box.

Licensing your product

You obtain a license file from Symantec when you purchase Symantec Brightmail Gateway or renew an existing license. You must register the license for each

Scanner that you install and enable to use Symantec Brightmail Gateway features. You can use the same license file to register multiple Scanners.

In addition to this license, you can also obtain a separate license to enable Symantec Content Encryption. Content encryption lets you encrypt outbound messages for greater security and to track statistics for those messages through the Control Center.

When you obtain your license file from Symantec, save it to a location that you can access from the Control Center.

See "Viewing license statuses" on page 679.

See "Where to get more information" on page 41.

To license your product

1 In the Control Center, click **Administration > Hosts > Licenses**.

2 Do one of the following tasks:

   - In the **Provide a license file** field, type the full path and license file name.
   - Click **Browse** and locate the license file.

3 Click **Register License**.

Viewing license statuses

You can view the status of your licenses to determine which features are licensed for each Scanner and when a license expires.

The Licenses page in the Control Center contains all of the Symantec Brightmail Gateway features that require licenses. The page also list whether the feature is licensed and when the license expires.

Note: An alert is sent when a license approaches expiration. Another alert is sent when it expires. Contact your Symantec sales representative for assistance renewing licenses.

See "Licensing your product" on page 678.

To view license statuses

1 In the Control Center, click **Administration > Hosts > Licenses**.

2 In the **Host** drop-down list, select a host name.

   The status of the licenses and their expiration dates appear.

Enforcing strong passwords

You can enable or disable strong passwords. Strong passwords make access to the Control Center more secure. When you enable the strong password feature, the current passwords for all administrators expire. However, the password for the admin administrator does not expire. Administrators must set new strong passwords the next time that they access the Control Center.

If you disable the strong password feature, the password history is erased. If you later turn on strong passwords again, administrators may reuse their old passwords that would not have been allowed if strong passwords were enabled.

See “Strong password criteria” on page 681.

See “Resetting an administrator password” on page 688.

The ability to use strong passwords with the **Remember me** feature depends on the order in which the features are enabled, as follows:

Strong passwords are enabled first, then the administrator subsequently enables the **Remember me** feature.	When the **Remember me** feature is enabled, if the bypass duration is greater than 60 days (the maximum number of days permitted for strong passwords), then the login bypass duration is set to the same expiration date as the strong password. This change is not reflected in the Control Center settings.
The **Remember me** feature is enabled first, then the administrator subsequently enables strong passwords.	The **Remember me** feature is disabled in the Control Center.

See “Enabling users to bypass Control Center login credentials” on page 674.

To enforce strong passwords

1. In the Control Center, click **Administration > Users > Administrators**.
2. Do one of the following:
   - To enable strong passwords, check **Require strong passwords**.
   - To disable strong passwords, uncheck **Require strong passwords**.
     The new strong password policy takes effect immediately.

Password best practices

To create secure passwords, consider the following suggestions:

- Do not create a password that uses any of the following formats:
  - A word that is found in a dictionary (in any language or jargon)
  - A name (such as the name of a spouse, parent, child, pet, fantasy character, famous person, or location)
  - Any variation of your personal name or account name
  - Accessible information about you (such as your phone number, license plate, or social security number) or your environment
  - A birthday or a simple pattern (such as backwards, followed by a digit, or preceded by a digit)
- Create a password that is based on the following recommendations:
  - Use a mixture of upper and lower case letters, as well as digits or punctuation
  - Make sure the password is unrelated to any previous password
  - Use long passwords (eight characters or longer)
  - Consider using a pair of words with punctuation inserted
  - Consider using a pass phrase (an understandable sequence of words)
  - Consider using the first letter of each word in a pass phrase

See “Strong password criteria” on page 681.

After you reset an administrator's password, use a secure method (such as a phone call) to notify the administrator of the new password. Email and instant messaging are not typically secure methods.

Strong password criteria

Strong passwords must contain all of the following requirements:

- US-ASCII character encoding
- At least eight characters
- At least one uppercase character
- At least one lowercase character
- At least one number

See “Enforcing strong passwords” on page 680.

Strong passwords cannot be changed more frequently than once a day, but they must be changed every 60 days. They cannot be the same password as any of the user's last five passwords.

Adding administrators

When you add an administrator, you can specify the administrator's rights and which alerts and notifications the administrator receives.

See "Editing an administrator" on page 683.

See "Deleting an administrator" on page 684.

To add an administrator

1. In the Control Center, click **Administration > Users > Administrators**.
2. Click **Add**.
3. In the **User name** box, type the user name (in US ASCII characters).
4. In the **Password** box, type a password.
5. In the **Confirm password** box, type the password again to confirm it.
6. In the **Email address** box, type the email address of the administrator.
7. If this administrator is to receive system alerts, check **Receive Alert Notifications**.
8. Choose the administrative rights that you want to assign to the administrator as follows:

Full Administration Rights	Click **Full Administration Rights** to let the administrator view and modify all available rights.
Limited Administration Rights	Click **Limited Administration Rights** and choose the specific rights for this administrator as follows: ■ None Administrators do not have any rights on selected task. ■ View Administrators can view appropriate pages but cannot manage them. ■ Modify Administrators have full rights to view and modify selected tasks. See "Administrator rights" on page 684.

9. If you select Limited Administration Rights, beside **Content Incident Folders**, choose the specific rights for this administrator as follows:

None	Administrators can see content incident folder names but cannot see incidents in the Incident Management Overview page for a given folder.
View	Administrators can view incidents in the indicated content incident folder but cannot manage them. Administrators with View permissions cannot perform any actions on incidents in the indicated content incident folder.
Modify	Administrators can view and modify all incidents in the indicated content incident folder.

10. Check **Receive Incident Notifications** to indicate the administrator is to receive notifications for the incidents that are created for the specific content incident folder.
11. Click **Save**.

Editing an administrator

You can edit an administrator to modify any of the following:

- Administrator's name
- Administrator's email address
- Administrator's password
- Whether the administrator receives notifications
- Administrator's rights
  See "Administrator rights" on page 684.
- Whether the administrator receives notifications about content incident folder incidents

See "Adding administrators" on page 682.

To edit an administrator

1. In the Control Center, click **Administration > Users > Administrators**.
2. Select an Administrator from the list and click **Edit**.
3. Change the Administrator definition as needed.
4. Click **Save**.

Deleting an administrator

You can delete an administrator at any time. However, when an administrator is deleted, the settings cannot be retrieved. If the administrator might still need access to the Control Center, you may want to consider modifying the administrator's rights rather than delete the administrator.

See "Editing an administrator" on page 683.

To delete an administrator

1. In the Control Center, click **Administration > Users > Administrators**.
2. Check the box beside the administrator that you want to remove.
3. Click **Delete**.
4. In the confirmation dialog box, click **OK** to confirm the deletion.

Administrator rights

When you add or edit an administrator, you can assign that administrator full administration rights or limited administration rights. This section explains limited administration rights.

You can assign limited administration rights to an administrator on the **Add Administrator** or **Edit Administrator** page, both available from the **Administration > Users > Administrators** page. You can choose **None**, **View**, or **Modify** rights for each of the following:

- Manage Status and Logs
- Manage Reports
- Manage Policies
- Manage Settings
- Manage Administration
- Manage Quarantine

In addition, under **Content Incident Folders**, you can choose **None**, **View**, or **Modify** for each content incident folder that you create, and you can check **Receive Incident Notifications** for any content incident folder.

Any administrator without either **Full Administration Rights** or **Manage Administration** rights will see the **Administration > Users > Administrators** page but will only be able to change his or her own password on that page.

Each type of limited administrator rights grants the administrator the ability to view a subset of the pages of the Control Center.

Note: Although many of the types of limited administrator rights let you view the pages for all content incident folders, they do not let you actually see the content of those folders. To view or modify the contents of a content incident folder, you must have **View** or **Modify** rights for that folder under **Content Incident Folders** on the **Add Administrator** or **Edit Administrator** page.

Table 20-1 shows the pages available to an administrator with **Manage Status and Logs** rights.

Table 20-1 Manage Status and Logs pages

Tab	Menu	Pages
Status	**System**	■ **Dashboard** ■ **Hosts** ■ **Logs**
Status	**SMTP**	■ **Message Audit Logs** ■ **Message Queues**
Status	**Instant Messaging**	■ **Active Users** ■ **Network Status**
Content	**Incident Management**	■ **Folder Overview**
Administration	**Users**	■ **Administrators**
Administration	**Settings**	■ **LDAP** ■ **Logs**
Administration	**Hosts**	■ **Configuration** ■ **Licenses** ■ **Utilities**

Table 20-2 shows the pages available to an administrator with **Manage Reports** rights.

Table 20-2 Manage Reports pages

Tab	Menu	Pages
Reports	**View**	■ **Create a Report** ■ **Favorite Reports**
Administration	**Users**	■ **Administrators**
Administration	**Settings**	■ **Report**

Table 20-3 shows the pages available to an administrator with **Manage Policies** rights.

Table 20-3 Manage Policies pages

Tab	Menu	Pages
Protocols	**Instant Messaging**	■ **Network Access Control**
Reputation	**Policies**	■ **Bad Senders** ■ **Connection Classification** ■ **Good Senders**
Reputation	**Reputation Tools**	■ **Find Sender** ■ **IP Reputation Lookup**
Spam	**Policies**	■ **Email** ■ **Instant Messaging**
Spam	**Settings**	■ **Sender Authentication**
Virus	**Policies**	■ **Email** ■ **Instant Messaging**
Content	**Policies**	■ **Email**
Content	**Resources**	■ **Annotations** ■ **Attachment Lists** ■ **Dictionaries** ■ **Notifications** ■ **Patterns** ■ **Records**

Table 20-3 Manage Policies pages *(continued)*

Tab	Menu	Pages
Administration	**Users**	■ **Administrators** ■ **Find User** ■ **Policy Groups**

Table 20-4 shows the pages available to an administrator with **Manage Settings** rights.

Table 20-4 Manage Settings pages

Tab	Menu	Pages
Protocols	**SMTP**	■ **Address Masquerading** ■ **Aliases** ■ **Domains** ■ **Invalid Recipients** ■ **Settings**
Protocols	**Instant Messaging**	■ **Registered Users**
Spam	**Settings**	■ **Quarantine Settings** ■ **Scan Settings**
Virus	**Settings**	■ **LiveUpdate** ■ **Scan Settings** ■ **Suspect Virus Settings**
Content	**Settings**	■ **Archive** ■ **Content Encryption** ■ **Content Incident Folders** ■ **Symantec DLP Connect**
Administration	**Settings**	■ **Alerts** ■ **Certificates** ■ **Control Center** ■ **LDAP** ■ **Logs** ■ **Reports** ■ **SNMP** ■ **UPS**
Administration	**Hosts**	■ **Configuration**

Table 20-5 shows the pages available to an administrator with **Manage Administration** rights.

Table 20-5 Manage Administration pages

Tab	Menu	Pages
Administration	**Users**	■ **Administrators**
Administration	**Hosts**	■ **Licenses** ■ **Shutdown** ■ **Version**

Table 20-6 shows the pages available to an administrator with **Manage Quarantine** rights.

Table 20-6 Manage Quarantine pages

Tab	Menu	Pages
Spam	**Settings**	■ **Quarantine Settings**
Spam	**Quarantine**	■ **Email Spam**
Virus	**Settings**	■ **Suspect Virus Settings**
Virus	**Quarantine**	■ **Email Suspect Virus**
Content	**Incident Management**	■ **Folder Overview** ■ **Informational Incidents** ■ **Quarantine Incidents** ■ Custom content incident folders
Administration	**Users**	■ **Administrators**
Administration	**Settings**	■ **LDAP**

Resetting an administrator password

You may need to reset a Control Center administrator password if an administrator has forgotten the password.

Note: Only the admin administrator can change another administrator's password if the "Require strong passwords" setting is enabled and the one-day minimum password age is not met. In other cases, any administrator with Manage Administration rights can change another administrator's password.

See “Administrator rights” on page 684.

After resetting an administrator's password, use a secure method (such as a phone call) to notify the administrator of the new password. Email and instant messaging are not typically secure methods.

See “Enforcing strong passwords” on page 680.

See “Editing an administrator” on page 683.

To reset an administrator password

1. In the Control Center, click **Administration > Users > Administrators**.
2. Check the box beside the administrator whose password you want to change, and click **Edit**.
3. In the **Password** box, type the new password.
4. In the **Confirm password** box, type the password again.
5. Click **Save**.

About maintaining adequate disk space

Symantec Brightmail Gateway performs better with more available disk space. Certain features such as extended reporting data and spam quarantine can use a large quantity of disk space and tax system resources. Periodically compare the disk usage to the disk capacity to ensure that the Control Center and Scanners have adequate disk space.

See “Viewing the status of your hardware” on page 607.

See “Viewing information about your hardware” on page 608.

Check Scanners and Control Center for old or unnecessary log files and manually delete them.

See “Clear disk space checklist” on page 634.

See “Manually deleting log files” on page 632.

To mitigate the burden to the system from over-production of log data, configure the log database size, log levels, and set the purge frequency rate of older log data.

See “Managing the log database size” on page 630.

See “About log disk space alerts” on page 633.

Specify a smaller disk space usage for spam virus quarantine and content incident folders.

See “Modifying the disk space allotted for Suspect Virus Quarantine” on page 238.

See “About managing the size of content incident folders” on page 439.

Symantec Brightmail Gateway provides a number of Expungers which keep disk space from filling. You can schedule these to run at regular intervals.

See “Scheduled tasks types” on page 622.

Turning off an appliance

When you turn off an appliance, the process begins immediately. If you have emails in your message queues, those emails remain in the queues. Before you turn off the appliance, first stop the mail flow. As a precaution, you might also want to flush your inbound, outbound, and delivery message queues.

See “MTA and message queue behavior” on page 613.

See “Managing services and MTA operations” on page 110.

To turn off an appliance

1. In the Control Center, click **Administration > Hosts > Configuration**.
2. Check the box beside the Scanner that you want to turn off and click **Edit**.

   The **Edit Host Configuration page** appears, showing the **Services** tab.
3. Under **MTA Operation**, click **Do not accept incoming messages**.
4. Click **Save**.
5. On the **Status > SMTP > Message Queues** page, view each of the message queues to see if they are empty.
6. Click **Flush All** to flush any queues that contain messages.

   You can also click **Delete** to delete a message or **Delete All** to delete all messages in a queue.
7. Repeat the previous step until all queues are empty.
8. On the **Administration > Hosts > Shutdown** page, click the **Host** drop-down list to select the host to turn off.
9. Click **Shutdown**.

   Before you turn off power to the appliance, be sure the message **Power Down** appears on a locally connected video console or through a serial connection.

Restarting an appliance

You can restart an appliance when needed. Restarting an appliance entails the appliance turning itself off and then restarting itself.

To restart an appliance

1. In the Control Center, click **Administration > Hosts > Shutdown**.
2. Under **System Shutdown**, click the **Host** drop-down list and select the appliance that you want to restart.
3. Click **Reboot**.

   All connections close and the system restarts.

Resetting an appliance to its factory default

You can return an appliance to the factory default condition of the most recent version that is installed on the appliance. When you enable this feature, Symantec Brightmail Gateway does all of the following actions for the Host that you select:

- Stops the Scanner host
- Clears the Host from the host table
- Clears the logs from the database
- Clears the reports from the database
- Clears the status information from the database
- Resets all settings and policies to their default values
- Deletes all backup files stored on the appliance

Note: After you perform a factory reset on a particular Scanner, you must delete the Scanner through the Control Center and then add it again. See "Adding Scanners" on page 79.

After you reset the appliance, you are automatically logged out.

To reset an appliance to its factory default

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Factory Reset** tab, click the **Host** drop-down list and select a host.
3. Click **Reset**.

4 Click **OK** to confirm reset or **Cancel** to stop the process.

5 In the confirmation dialog box, click **Reset**.

Scheduling backups

Symantec Brightmail Gateway contains a backup program and scheduler in the Control Center. Use these features to back up of various types of Symantec Brightmail Gateway data. Restore backup data on the **Restore/Download** tab of the **Host Version** page.

See "Restoring an appliance from backups" on page 695.

You can check the status of your scheduled task from the **Status > Scheduled Tasks** page.

See "About scheduled tasks" on page 621.

Note: The `db-backup` command can back up your appliance with SCP. The Control Center does not offer this option.

See db-backup on page 746.

To schedule backups

1 In the Control Center, click **Administration > Hosts > Version**.

2 On the **Backup** tab, click **Add**.

The **Add Scheduled Backup** page appears.

3 Under **Backup description**, type a description of the scheduled backup.

4 Under **Backup Data**, choose from among the following backup types:

Full Backup	Backs up the complete database, as well as Suspect Virus Quarantine messages and content filtering messages that are stored on disk. The file format for this backup type is: `db-backup.brightmail.mm-dd-yr-hr-mm.full.manual.tar.bz2`
Backup only Configuration and Incidents	Backs up all configuration data in the database, as well as content filtering data and content filtering messages that are stored on disk. The file format for this backup type is: `db-backup.brightmail.mm-dd-yr-hr-mm.config.incidents.tar.bz2`
Backup only Configuration, Incidents, Logs and Reports	Backs up all configuration, incident, report, and log data in the database; as well as content filtering messages that are stored on disk. The file format for this backup type is: `db-backup.brightmail.mm-dd-yr-hr-mm.config.incidents.reports.logs.tar.bz2`

For all file types, "month" is expressed in standard three-letter format. The following example shows the file name of a full backup:

```
db-backup.brightmail.Aug-31-09-10-04.full.manual.tar.bz2
```

5 Under **Backup Schedule**, define the time and frequency to run the backup.

6 Under **Backup To**, specify whether to store backups on the local server or on a remote host using FTP.

The following shows sample values for specifying a remote backup through FTP:

Domain	host.symantecs.org (or 192.168.2.42)
Port	21
Path	/home/username/backups/

7 If you back up on the local server, indicate how many backup versions to keep.

The default is 3.

You only need to specify the number of backup versions that are retained when you store files locally. When you store backup data at a remote location, you must supply the necessary information for FTP transfer. You must also provide user authentication information when required by the remote location.

8 Click **Save**.

Editing a scheduled backup

You can modify a scheduled backup as needed.

To edit a scheduled backup

1 In the Control Center, click **Administration > Hosts > Version**.

2 On the **Backup** tab, check the box beside the backup that you want to edit.

3 Click **Edit**.

The Edit Scheduled Backup page appears.

4 Edit options for the scheduled backup.

5 Click **Save**.

Deleting a scheduled backup

You can delete a scheduled backup when it is no longer needed.

See “Scheduling backups” on page 692.

To delete a scheduled backup

1 In the Control Center, click **Administration > Hosts > Version**.

2 On the **Backup** tab, check the box beside the backup that you want to delete.

3 Click **Delete**.

Performing an on-demand backup

You can perform an on-demand backup of the appliance at any time. Restore backup data on the **Restore/Download** tab of the **Host Version** page.

See “Restoring an appliance from backups” on page 695.

To perform an on-demand backup

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Backup** tab, click **Backup Now**.
3. Under **Backup Data**, click the backup data type that you want.
4. Under **Backup To**, specify whether the backup data is to be stored on the local server or at a remote location through file transfer protocol.

   When you store backup data at a remote location, you must supply the necessary information for FTP transfer. You must also provide user authentication information when required by the remote location.
5. Click **Backup Now**.

Restoring an appliance from backups

If you have previously backed up your databases, it is possible to restore them from any of the available backup stores.

If you restore a backup from one appliance to a different physical appliance, make sure that the date is set correctly on the new appliance. This verification ensures that messages in quarantine at the time of the original backup are displayed correctly appear after the restore.

Note: After you restore an appliance that functions as your Control Center from a different IP address than the original IP address, you must reboot the appliance. If that appliance also hosts a Scanner, you should stop the Scanner first. If you restore the appliance from a backup taken on a different appliance, the restored appliance inherits the network configuration of the other appliance.

See “Restarting an appliance” on page 691.

See “Stopping and starting Scanners” on page 108.

To restore the appliance from a local backup

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Restore/Download** tab, click **Restore/Download backup from server**.
3. Under **Available Backups**, check the box beside the backup that you want to restore.
4. Click **Restore**.

To download a backup file to the appliance running the Control Center

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Restore/Download** tab, click **Restore/Download backup from server**.
3. Under **Available Backups**, check the box beside the backup that you want to download.
4. Click **Download**.

To restore your appliance from a remote backup

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Restore/Download** tab, click **Restore backup from a remote location**.

   Supply the protocol, domain (host name) or IP address, port, and fully qualified (absolute) path to the file. Supply authentication information if required.

   The following are sample values for restoring the system from a remote backup:

Host/IP address	host.symantecs.org (or 192.168.2.42)
Port	21
Path	/home/username/backups/

3. Click **Restore**.

To restore your system from a local file

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Restore/Download** tab, click **Upload a backup file from your local computer**.

   This step assumes that you have a local copy of the backup file, such as from backing up using FTP.

3. Click **Restore**.

About software updates

Symantec periodically releases new versions of the Symantec Brightmail Gateway software. New versions contain new features and fixes for software defects. Symantec recommends that you keep your appliance up to date with the latest version of the software. However, some software updates may contain changes to the underlying architecture of Symantec Brightmail Gateway. Always read the software update notes and release notes for specific instructions before you run the update.

See "Updating your software" on page 700.

See "Software update best practices" on page 698.

You must run software update on the Control Center and all Scanners. The order in which you apply the software updates to your appliances does not matter. For example, you can update the Control Center before you update Scanners or update a Scanner before the Control Center. Update one appliance at a time.

Table 20-7 Where to get more information about software updates

Document	Location	Description
Software update notes	In the Control Center, click **Administration > Version > Updates**. Click the button next to a new version and click **View Description**.	The software update notes contains abbreviated update information and a list of issues that are addressed in the new version.
Release notes	On the Internet, go to the following URL: www.symantec.com/business/support/overview.jsp?pid=53991 Click **Release Notes** to view the release notes for the most recent software version. To view release notes for previous releases, click **Documentation**.	The release notes contains more detailed information about updating to the new version. The release notes also describes new features and lists the known issues that are not fixed in the new version.
Late-breaking news	The URL for the late-breaking news is listed in the software update notes and release notes.	The late-breaking news Web page may contain information about the release that is not in the software update notes or release notes.

Determining which version of software is installed

You can determine what version of software is installed on your appliance.

See “Updating your software” on page 700.

To determine which version of software is installed

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Updates** tab, click the **Host** drop-down list and select a host.

   The version and status of your software appears.

Software update best practices

Before you run a software update, Symantec recommends that you prepare your system.

See “About software updates” on page 696.

See “Updating your software” on page 700.

Table 20-8 describes the steps to prepare and run a software update.

Note: Do not reboot or shut down an appliance while the software update runs. The appliance can become corrupted and require reinstallation.

Table 20-8 Software update procedure

Step	Action	Description
Step 1	Read software update documentation	Read the software update notes, release notes, and late-breaking news for the software update. The documentation may describe special steps to prepare for the software update. See “About software updates” on page 696.
Step 2	Delete messages in Spam Quarantine	If your site policies let you, delete all the messages in Spam Quarantine. See “Deleting spam messages in quarantine” on page 267.
Step 3	Delete log messages	If your site policies let you, delete all Scanner and directory data service log messages. See “Manually deleting log files” on page 632.

Table 20-8 Software update procedure *(continued)*

Step	Action	Description
Step 4	Delete content filtering incidents	If your site policies let you, delete all incidents from the content incident folders. See “Deleting incidents” on page 454.
Step 5	Perform a backup	Perform a backup to save your configuration settings. See “Performing an on-demand backup” on page 694.
Step 6	Stop incoming messages and drain the queue	If you are running software update on a Scanner or a combination Control Center and Scanner, set the MTA operation to **Do not accept incoming messages**. Incoming messages are temporarily rejected. After you stop incoming messages, flush messages to drain them from the queue. See “Managing services and MTA operations” on page 110. Do not stop the Scanner Services as described in that section.
Step 7	Monitor the update progress	Optionally, you can monitor the update process by running the `tail -f update.log` on the command line. Run the command before the software update so you can see the data being written to the log. See “Monitoring software update using the command line interface” on page 700.
Step 8	Run the software update	Run the software update either using the Control Center or the `update` command. See “Updating your software” on page 700. See update on page 817.

Table 20-8 Software update procedure *(continued)*

Step	Action	Description
Step 9	Restart incoming messages	After the update completes, exit from your Web browser, restart the Web browser, and log into the Control Center. If you ran the software update on a Scanner or a combination Control Center and Scanner, set the MTA operation to **Accept and deliver messages normally**. See "Managing services and MTA operations" on page 110.

Monitoring software update using the command line interface

The software update process can take several minutes or several hours. If you update using the Control Center, it is difficult to determine if the software update has encountered a problem. You can monitor the progress of the software update by viewing the `update.log` file as the update process writes to it. If you see information being written to `update.log` periodically, it usually means that the update is proceeding normally.

See "About software updates" on page 696.

You can start this process before you run the software update or after you have started the software update. You can use this process if you update using the Control Center or command line. If you update using an SSH client on the command line, run the `tail -f update.log` in a separate SSH client window than the `update install` command. If you update using the console to access the command line, it is not possible to run the `tail -f update.log` simultaneously with the update on that console.

See update on page 817.

To monitor the software update using the command line interface

1 Using an SSH client or the console, log into the appliance that you want to update.

2 Type the following command:

```
tail -f update.log
```

Updating your software

If a new version is available, you can update your Control Center and Scanners.

See "About software updates" on page 696.

See "Software update best practices" on page 698.

You can also update using the command line interface.

See update on page 817.

To update your software

1. In the Control Center, click **Administration > Hosts > Version**.
2. On the **Updates** tab, select a host.
3. If available, select an updated software version and click **View Description**.

   You must view the software update notes to display the **Update** option.
4. Click **Update**.

   The update begins.
5. After you install a software update, close and restart your Web browser before you log on to the Control Center.

Configuring bad message handling

In rare instances, a malformed email message may cause the Brightmail Engine to fail. If this situation occurs, Symantec Brightmail Gateway isolates the small number of messages that were being processed at the time of the failure and rescans them one at a time until the message that caused the failure is identified (causing the Brightmail Engine to fail again). The Brightmail Engine recovers from these failures quickly. The flow of mail is neither interrupted nor significantly delayed. When a bad message is detected, an alert is sent.

With Symantec Brightmail Gateway's bad message handling feature, you can specify how many times the system scans a potentially malformed message before it classifies it as such and places it in the bad message queue.

Note: For each retry (every time the system scans the malformed message) the Brightmail Engine fails. Use caution when setting this value.

You can then access and manage messages in the bad message queue with the `mta-control` command-line option.

Once the malformed message is identified, you have several options when you use the `mta-control` command as follows:

- List email messages in the bad message queue.

- View or export the message to a specified URL or to the screen.
- Delete a message from the bad message queue.
- Bypass the Brightmail Engine and deliver the message to its original recipient.
- Deliver the message to system administrator(s) by email.
- Resend the message through the Brightmail Engine.

See "Administering Symantec Brightmail Gateway through the command line" on page 711.

To configure bad message handling

1 Click **Protocols > SMTP > Settings**.

2 Under **Scan Settings**, check **Enable bad message handling**.

3 Type a value in the **Number of retries before classifying a message as bad** field.

4 Click **Save**.

Setting up your SMTP greetings and postmaster address

When Symantec Brightmail Gateway connects with other mail servers to initiate inbound or outbound messaging traffic, the SMTP protocol conversation begins with a greeting.

By default, Symantec Brightmail Gateway uses the following greeting:

Symantec Brightmail Gateway

You can change both the inbound SMTP greeting and the outbound SMTP greeting that Symantec Brightmail Gateway uses.

You can also change the postmaster address. Symantec Brightmail Gateway includes the postmaster address in bounce messages.

To set up your SMTP greetings and postmaster address

1 In the Control Center, click **Protocols > SMTP > Settings**.

2 Highlight the text in the **Inbound SMTP greeting** field and type the text you want to use.

3 Highlight the text in the **Outbound SMTP greeting** field and type the text you want to use.

4 Highlight the text in the **Postmaster address** field and type the text you want to use.

5 Click **Save**.

Customizing the date format and time format

Symantec Brightmail Gateway installs with default date and time formats. However, you can modify these formats. You must have Full Administration rights or Manage Settings rights to modify date and time formats.

Table 20-9 Configurable date and time formats

Format	Syntax and default value	Example
Time	hh:mm:ss a	02:29:34 PM
Short date	MM/dd/yy	2/25/10
Long date	MMMM dd, yyyy	February 25, 2010
Full date and time	EEEE, MMM dd, yyyy hh:mm:ss a zz	Thursday, Feb 25, 2010 01:23:45 PM PST
AM/PM symbols	AM/PM	AM

Symantec Brightmail Gateway applies the date format and time formats that you configure as follows:

Detail view	Full date and time
List view	Long date
Mailbox	Combination of long date format and time Uses the time format for same day's email; uses the long format for everything else
Reports	Short date and time

The maximum syntax length that you can specify is 32 characters. The 32-character limit does not include spaces.

See "Date format and time format pattern syntax" on page 704.

The option to customize the date format and time format is disabled by default. You must enable this feature to modify date and time formats.

Configuring the date format and time format

1. In the Control Center, click **Administration > Settings > Control Center.**
2. Click the **Locale** tab.
3. Check **Apply custom date and timestamp format**.
4. Specify the custom date and timestamp formats that you want to use.
5. Click **Save**.

Date format and time format pattern syntax

Table 20-10 provides the supported date format and time format syntax.

Table 20-10 Date and time format syntax

Letter	Component	Presentation	Example in US locale
G	Era	Text	AD
yyyy; yy	Year	Year	2010; 10
MMMM;MMM;MM	Month in year	Month	July; Jul; 07
w	Week in year	Number	27
W	Week in month	Number	2
D	Day in year	Number	189
d	Day in month	Number	10
F	Day of week in month	Number	2
EEEE; EEE	Day in week	Text	Tuesday; Tue
a	AM/PM mark	Text	PM
H	Hour in day 24H (0-23)	Number	0
k	Hour in day 24H (1-24)	Number	24
K	Hour in day 12H (0-11)	Number	0
h	Hour in day 12H (1-12)	Number	12

Table 20-10 Date and time format syntax *(continued)*

Letter	Component	Presentation	Example in US locale
m	Minute in hour	Number	20
s	Second in minute	Number	55
S	Millisecond	Number	978
zzzz; z	Time zone	Time zone	Pacific Daylight Time; PDT
Z	Time zone	RFC822	-0800

See "Customizing the date format and time format" on page 703.

Specifying what to include in diagnostic packages

A diagnostic package consist of the following, default contents:

- Configuration data
- Five latest core directories for each job under /data/scanner/jobs (excluding core files in those directories)
- Log data (up to 100,000 lines per file)

For more detailed information about what content is included in a default diagnostic package, refer to the command-line description for creating diagnostic packages.

See diagnostics on page 759.

You can also specify what additional data you want to include in the package.

If you select **All log data**, the report contains all of the lag data that is currently available. If you select **Message tracking**, the report contains all of the message tracking that is on disk. If you select any core file option, the report contains the most recent core files that are available.

After you specify what to include in the report package, generate the package and transfer or download it to the location that you choose. You can only generate one package at a time. To generate a diagnostic package, you must have Full Administration rights or rights to manage status.

See "Generating diagnostic packages and transferring them to a remote location" on page 709.

See "Generating and downloading diagnostic packages" on page 707.

See “About troubleshooting issues with Symantec Brightmail Gateway with diagnostics” on page 706.

Specifying what to include in diagnostic packages

1. In the Control Center, click **Administration > Hosts > Utilities**.
2. Click the **Diagnostics** tab.
3. Click the **Host** drop-down menu and select the host for which you want to run the diagnostic package.
4. Under **Select Components**, check the items that you want to include in the diagnostic package.

   If you select no components, a default diagnostic package is generated.

   If you select a host that is a Control Center only, then the only component options that are available are as follows:

   - **All log data**
   - **All core files - Agent core files**
   - **All core files - Other core files**
5. In the **Maximum number of cores included in diagnostics package** field, specify the number of cores that you want to include in the diagnostics report.

   The default value is 5.

   See “Generating diagnostic packages and transferring them to a remote location” on page 709.

About troubleshooting issues with Symantec Brightmail Gateway with diagnostics

If you experience trouble with Symantec Brightmail Gateway, Symantec Support may request that you create a diagnostic package to send to them. The diagnostic package helps Symantec Support troubleshoot issues with your product.

By default, all diagnostic packages contain certain components. However, you can also specify any additional components that you want to include in the package. If you do not specify any components, a default diagnostic package is generated.

See “Specifying what to include in diagnostic packages ” on page 705.

When you generate a diagnostic package, Symantec Brightmail Gateway creates the package on the host for which you are running the package. You can then transfer a copy of the package to the location that you specify: to a remote location through a transfer protocol or download it locally on the computer on which you run your browser. This feature only works with those hosts that use the same version of Symantec Brightmail Gateway as the Control Center. You can only generate one package at a time.

See "Generating diagnostic packages and transferring them to a remote location" on page 709.

See "Generating and downloading diagnostic packages" on page 707.

You can delete a package from the Control Center host when it is no longer needed, which frees up disk space on the Control Center.

See "Deleting diagnostic packages" on page 710.

You can also perform this task from the command line.

See diagnostics on page 759.

Generating and downloading diagnostic packages

When you generate a diagnostic package to download, Symantec Brightmail Gateway creates the package on the host for which you are running the package. After Symantec Brightmail Gateway creates the package, it appears in the **Available Diagnostic package** list in the Control Center. You can download the package to the Control Center computer immediately after you generate it or later on at your convenience.

You can only generate one diagnostic package at a time, so only one package appears in the **Available Diagnostic package** list at a time.

Note: Symantec Brightmail Gateway does not support downloading any diagnostic packages that exceed 2 GB. If the diagnostic package exceeds 2 GB, an error message appears. Regenerate the package and transfer it through one of the supported protocols to a remote location.

See "Generating diagnostic packages and transferring them to a remote location" on page 709.

When a package is successfully generated, it appears in the **Available Diagnostic package** list. If the generation does not complete successfully, an error appears on the Control Center. The error message only appears if you do not navigate from the page. The error is also logged to the Symantec Brightmail Gateway BrightmailLog.log log.

If you navigate away from the Diagnostics page and return, you may not see the diagnostic package in the **Available Diagnostic package** list.

One of the following events occurred:

The package generation is successful.	Refresh your browser. The package appears in the **Available Diagnostic package** list.
The package is still being generated.	No status appears on the Control Center. Do not click **Generate** again. Instead, refresh your browser to see if the package generation is finished. If it is finished, the package appears in the **Available Diagnostic package** list.
The package generation is unsuccessful.	No error message appears, but an error is logged to the Symantec Brightmail Gateway BrightmailLog.log log.

See "About troubleshooting issues with Symantec Brightmail Gateway with diagnostics" on page 706.

When you no longer need a diagnostic package, delete it from the Control Center to free up space.

See "Deleting diagnostic packages" on page 710.

To generate diagnostic packages, you must have Full Administration rights or rights to manage status.

Generating and downloading diagnostic packages

1. Specify what you want to include in a diagnostic package and the host.

   See "Specifying what to include in diagnostic packages " on page 705.

2. In the **Protocol type** drop-down list, select **Download to desktop**.
3. Click **Generate**.
4. Do one of the following tasks:

To download the package immediately after you generate it	In the **Save As** dialog box, type the location where you want the diagnostic package saved, and then click **Save**.
To download the package at a later time	Do all of the following tasks: ■ Click **Download**. If there are no diagnostic packages in the **Available Diagnostic package** list, the **Download** option is disabled. ■ In the **Save As** dialog box, type the location where you want the diagnostic package saved, and then click **Save**.

Generating diagnostic packages and transferring them to a remote location

When you generate a diagnostic package, Symantec Brightmail Gateway creates the package on the host for which you generate the package. After the package is created, you can transfer it to a remote location. You can only generate one diagnostic package at a time. To create diagnostic packages, you must have Full Administration rights or rights to manage status.

If you remain on the **Diagnostics** page after you generate a package, you receive feedback about whether the generation is successful. It is recommended that you remain on the **Diagnostics** page until you receive this confirmation.

If you navigate from the **Diagnostics**page after you generate a package, you receive no feedback about whether the generation is finished or if it was successful. If the package generation is still in progress and you try to regenerate a package, an error message appears that indicates a generation is already in progress.

If the generation is unsuccessful for any reason, an error is logged to the BrightmailLog.log.

See "About troubleshooting issues with Symantec Brightmail Gateway with diagnostics" on page 706.

Generating diagnostic packages and transferring them to a remote location

1. Specify what you want to include in a diagnostic package and the host.

   See "Specifying what to include in diagnostic packages " on page 705.

2. Click the **Protocol type** drop-down list and select one of the following options:

   - Download to desktop
   - **FTP**
   - **SCP**

   The default setting is **Download to desktop**.

   If you select FTP or SCP, you must also specify protocol parameters and credentials.

   If you selected **Download to desktop**, skip to step 8.

3. In the **Host** box, type the host name or IP address for the computer where you want to send the package .

4. In the **Port** box, type the port number for the computer where you want to send the package.

   The default port for FTP is 21 and for SCP is 22.

5. In the **Path** field, type the path.

   If file path ends with a forward slash, then the directory is assumed, and the default file name is appended to it. If it does not end with a forward slash, it assumes a full file path name.

6. In the **Username** box, type the user name.
7. In the **Password** box, type the password.
8. Click **Generate**.

Deleting diagnostic packages

When you create a diagnostic package and select the **Download to desktop** option, the package appears in **Available Diagnostic package** list. (Only one package can appear in the **Available Diagnostic package** list at a time.) When you no longer need this package, you can delete it. When you delete the package from the **Available Diagnostic package** list, you delete the file that is stored on Control Center.

To delete a diagnostic package, you must have Full Administration rights or rights to manage status.

See "Generating diagnostic packages and transferring them to a remote location" on page 709.

See "Generating and downloading diagnostic packages" on page 707.

See "About troubleshooting issues with Symantec Brightmail Gateway with diagnostics" on page 706.

To deleting diagnostic packages

1. In the Control Center, click **Administration > Hosts > Utilities**.
2. Click the **Diagnostics** tab.
3. Under **Available Diagnostic package**, click **Delete**.

Converting 8-bit MIME messages to 7-bit MIME

If you encounter issues with MTAs that cannot handle 8-bit MIME, you can force the sending MTA to convert the 8-bit MIME messages to 7-bit MIME.

To convert 8-bit MIME messages to 7-bit MIME

1. In the Control Center, click **Protocols > SMTP > Settings**.
2. Under **Mime Handling**, check **Do not advertise 8BITMIME**.
3. Click **Save**.

Administering Symantec Brightmail Gateway through the command line

Each appliance (real or virtual) has a set of commands that you can use to configure, optimize, and administer your product. You can execute these commands from an SSH session or from the system console. The help for these commands is presented in Linux man page format.

See “Command line interface access methods” on page 713.

These help pages use the following Linux man page conventions:

- Square brackets ([]) indicate that a statement is optional
- The pipe character (|) indicates that one of two statements can be specified.
- Text in *italics* indicates that the text should be replaced with the text that you specify.

The Symantec Brightmail Gateway man pages contain the following sections:

- Synopsis
  A description of the options and arguments available for the command.
- Description
  General information about the command.
- Options
  Options that you can use to control the behavior of a command. Options always begin with one or two dashes, such as `-s` or `--status`. Use two dashes for the full term; one dash for the abbreviated term.
  Some options have arguments. For example, `--log` *level*. Square brackets mean that element of the command is optional.
  Not all commands have options.
- Arguments
  Some commands require arguments. Arguments are names of files, host names, IP addresses, and so on that you specify to control the behavior of the command. Not all commands have arguments. Unlike options, you do not precede arguments with dashes.

- Examples
  This section provides sample command usage. Not all commands have examples.
- See Also
  This section lists related commands. Not all commands have see also references.

Use the following commands to navigate through the man pages:

- f or SPACE
  Forward one window
- b
  Backward one window
- /pattern
  Search for a word or pattern
- <
  Go to the beginning of the document
- >
  Go to the end of the document
- q
  Quit
- h
  Display more help with man pages

Type `help` *`command_name`* to get information about a specific command. Type `help` to get general information about command-line man pages.

The following commands are available:

- agent-config
- cat
- cc-config
- clear
- db-backup
- db-restore
- delete
- diagnostics
- dns-control
- grep

- help
- ifconfig
- iostat
- ldapsearch
- list
- mallog
- malquery
- monitor
- more
- mta-control
- netstat
- nslookup
- password
- ping
- reboot
- route
- rpmdb
- service
- show
- shutdown
- sshd-config
- tail
- telnet
- traceroute
- update

Command line interface access methods

You can log into the command line interface on each Symantec Brightmail Gateway appliance. Some of the commands duplicate functions in the Control Center. Some of the commands provide functions that are not available in the Control Center.

See "Administering Symantec Brightmail Gateway through the command line" on page 711.

Table 20-11 and Table 20-12 describe the methods that you can use to access the command line interface. After connecting to the command line interface, type **admin** at the `login as:` prompt and type the `admin` password at the `password:` prompt.

Table 20-11 Command line interface access methods for physical appliances

Access method	How to connect
System console using directly attached keyboard and VGA monitor	You must have physical access to the appliance to access the command line interface with a keyboard and VGA monitor. Connect a keyboard to the keyboard port on the appliance. Connect a VGA-compatible monitor to the D-sub 15 VGA port on the appliance. You can also connect the keyboard and VGA ports on the appliance to a KVM switch.
System console using a serial cable	You must have physical access to the appliance to access the command line interface with serial cable. Connect a null modem cable from the DB9 serial port on the appliance to the serial port on another computer. Use a terminal emulation software on the computer to access the appliance through the serial port. On a Windows computer, ensure that the terminal emulation software is set to use the correct COM port. Configure the terminal emulation software on the computer to the following settings: ■ 9600 bps ■ 8 data bits ■ No parity bit ■ 1 stop bit
Remote access using an SSH client	Using an SSH client lets you access the command line interface from any computer on your network, unless firewall rules prohibit access. For a Windows computer, use an SSH client such as PuTTY. On a UNIX computer you can use the `ssh` command that is typically included in the operating system. The host name or IP address to connect to using the SSH client is the name you specified when you initially configured the appliance. For a Control Center appliance, the host name is also the name in the URL that you use to access the Control Center.

Table 20-12 Command line interface access methods for virtual appliances

Access method	How to connect
VMware Virtual Machine Console	You can use the VMware Virtual Machine Console to log into the virtual appliance. Refer to the VMware Virtual Machine Console documentation for more information.
Remote access using an SSH client	If you configured the virtual appliance with a host name or IP address that resolves on your network, you can use an SSH client to access the command line interface. You can access the virtual appliance from any computer on your network, unless firewall rules prohibit access. For a Windows computer, use an SSH client such as PuTTY. On a UNIX computer you can use the `ssh` command that is typically included in the operating system.

Appendix A

Message filtering options

This appendix includes the following topics:

About filtering

Although Symantec Brightmail Gateway provides default settings for dealing with spam, spim, and viruses, you can configure the actions taken on filtered messages to meet your organizational requirements. Content filtering policies offer further methods of managing mail flow into and out of your organization. You can also use content filtering policies to monitor and enforce compliance with regulatory and organizational requirements.

Symantec Brightmail Gateway provides a variety of actions for filtering email and instant messages (IM), and allows you to set identical options for all users, or specify different actions for distinct policy groups.

You can assign members to a policy group based on email addresses, domain names, or LDAP groups. For each policy group, you can specify an action or group of actions to perform when a message results in a particular verdict. You specify actions when you create or edit a spam, spim, virus, or content filtering policy. Each of these policies is a filtering policy.

See "Adding members to a policy group" on page 317.

Verdicts are the conclusions reached on a message by the filtering process. When you create or edit a filtering policy, you specify the conditions that you are looking for in messages. Each condition or set of conditions corresponds to a type of verdict that Symantec Brightmail Gateway can assign to a message.

Symantec Brightmail Gateway performs actions on an email or IM based on the verdict applied to the message and the policy groups that include the message recipient or sender as a member, as follows:

- For inbound email filtering, the policy groups that impact message filtering are those policy groups that include the message recipient.
- For outbound email filtering, the policy groups that impact message filtering are those policy groups that include the message sender.

See "Verdicts by verdict category" on page 731.

Verdicts and actions for email messages

The following describes the filtering actions available for each verdict category.

Note: The Sender Groups column includes both good and bad sender groups, but does not include Fastpass. You cannot specify any actions for the Fastpass verdict.

In general, messages from senders in the good sender groups bypass spam filtering, but do not bypass virus filtering or content filtering.

See "Verdicts by verdict category" on page 731.

See "Verdict combinations" on page 733.

Table A-1 Verdicts and actions for email messages

Action	Description	Attack verdict	Virus verdict	Spam verdict	Content Filtering verdict	Sender Group verdict
Add a header	Add an email header.	x	x	x	x	x
Add annotation	Insert predefined text (a disclaimer, for example).	x	x	x	x	x
Add BCC recipients	Blind carbon copy to the designated SMTP address(es).	x	x	x	x	x
Archive the message	Forward a copy to the designated SMTP address and, optionally, host.	x	x	x	x	x
Bypass content filtering policy	Do not filter spam messages for content filtering policies. You can choose all content filtering policies or specify the policies to bypass.			x		x
Bypass spam scanning	Do not scan messages that meet this policy for spam. Cannot be added to the list of approved or rejected actions.				x	
Clean the message	Repair repairable virus infections and delete unrepairable virus infections. Only available for the virus verdict.		x			
Create an incident	Create a record of a content filtering policy incident. Optionally, hold for review and defer certain actions.				x	
Defer SMTP connection	Using a 4xx SMTP response code, tell the sending MTA to try again later. Cannot be used with the Local Bad Sender Domains or Local Good Sender Domains groups.	x				x
Delete message	Delete the message.	x	x	x	x	x
Deliver message normally	Deliver the message. Viruses and mass-mailing worms are neither cleaned nor deleted.	x	x	x	x	x

Table A-1 Verdicts and actions for email messages *(continued)*

Action	Description	Attack verdict	Virus verdict	Spam verdict	Content Filtering verdict	Sender Group verdict
Deliver the message to the recipient's Spam folder	Deliver to end-user Spam folder(s). Requires use of the Symantec Spam Folder Agent for Exchange or the Symantec Spam Folder Agent for Domino. **Note:** Symantec no longer provides technical support for the Symantec Spam Folder Agent for Exchange and the Symantec Spam Folder Agent for Domino.	x	x	x	x	x
Deliver message with content encryption	Deliver via the designated encryption host over a mandatory TLS channel.				x	
Deliver message with TLS encryption	Send the message over an encrypted channel.				x	
Forward a copy of the message	Copy the message to designated SMTP address(es), and also deliver the original message to the recipient.	x	x	x	x	x
Hold message in Spam Quarantine	Send to the Spam Quarantine.	x	x	x	x	x
Hold message in Suspect Virus Quarantine	Hold in the Suspect Virus Quarantine for a configured number of hours (default is six), then refilter for viruses only, using the latest virus definitions. Only available for the suspicious attachment verdict.		x			
Modify the Subject line	Add a tag to the message's `Subject:` line.	x	x	x	x	x
Reject messages failing bounce attack validation	If a message fails bounce attack validation, reject the message. Only available for the Failed bounce attack validation verdict.			x		

Table A-1 Verdicts and actions for email messages *(continued)*

Action	Description	Attack verdict	Virus verdict	Spam verdict	Content Filtering verdict	Sender Group verdict
Reject SMTP connection	Using a 5xx SMTP response code, notify the sending MTA that the message is not accepted. Cannot be used with the Local Bad Sender Domains or Local Good Sender Domains groups.	x				x
Remove unresolved recipients (for Directory Harvest Attacks only)	If a directory harvest attack is taking place, remove each unresolved recipient rather than sending a bounce message to the sender.	x				
Route the message	Deliver via the designated SMTP host.	x	x	x	x	x
Send a bounce message	Return the message to its `From:` address with a custom response and deliver it to the recipient, with or without attaching the original message.	x	x	x	x	x
Send notification	Deliver the original message and send a predefined notification to designated SMTP address(es) with or without attaching the original message.	x	x	x	x	x
Strip and Delay in Suspect Virus Quarantine	Remove all non-text content and deliver the stripped message immediately. Hold the complete message in Suspect Virus Quarantine for a configured number of hours (default is six hours), then release and rescan. Only available for the Suspicious Attachment verdict.		x			
Strip attachments	Remove all attachments according to a specific attachment list. Cannot be used with sender authentication.		x	x	x	
Treat as a bad sender	Process using the action(s) specified in the Local Bad Sender Domains group. Applies even if the Local Bad Sender Domains group is disabled.				x	
Treat as a mass-mailing worm	Process using the action(s) specified in the associated worm policy.				x	

Table A-1 Verdicts and actions for email messages *(continued)*

Action	Description	Attack verdict	Virus verdict	Spam verdict	Content Filtering verdict	Sender Group verdict
Treat as a good sender	Process using the action(s) specified in the Local Good Sender Domains group. Applies even if the Local Good Sender Domains group is disabled. When used in a content filtering policy, messages that match the policy will not be scanned for spam.				x	
Treat as a virus	Process using the action(s) specified in the associated virus policy.				x	
Treat as spam	Process using the action(s) specified in the associated spam policy.				x	
Treat as suspected spam	Process using the action(s) specified in the associated suspected spam policy.				x	

When using Table A-1 consider the following limitations:

- By default, inbound and outbound messages containing a virus are cleaned of the virus. Inbound and outbound messages containing a mass-mailing worm, and unscannable messages, including malformed MIME messages, are deleted. If you are concerned about losing important messages, you may want to create a different filter policy for unscannable messages and apply that new filter policy to some or all of your policy groups.
  See "Virus categories and default actions" on page 321.
- The Send a bounce message action returns the message to its `From:` address with a custom response, and also delivers the original message to the recipient. Symantec does not recommend using the Send a bounce message action for virus or spam policies. Virus and spam messages often use falsified headers. An attempt to deliver a bounce message to an uninvolved party named in the `From:` address is basically sending a spam message. This could result in your domains being blacklisted by other domains or by third-party black lists.
- If a message matches a content filtering policy that invokes the Bypass spam scanning action and also matches a spam policy that invokes the Bypass content filtering policy action, neither spam nor content filtering policy actions are applied to that message.
- When you select certain actions, you can encode an optional archive tag for Western European (ISO-8859-1, default), Unicode (UTF-8), Japanese

(ISO-2022-JP, EUC-JP, or Shift_JIS), Simplified Chinese(GB2312 or GB18030), Traditional Chinese (Big5), or Korean (KS_C_5601-1987).

Verdicts and actions for instant messages

The following filtering actions are available for each IM verdict category.

Table A-2 IM filtering actions by verdict

Action	Description	Virus verdict	Spim verdict
Add annotation	Add an annotation to an IM message informing the recipient that the message contains spim or suspected spim.		x
Allow file transfer	Allow an infected, encrypted, or unscannable file to be transferred to its recipient.	x	
Block file transfer	Block an infected, encrypted, or unscannable file from being transferred to its recipient.	x	
Delete the message	Delete an IM message that contains spim or suspected spim.		x
Deliver message normally	Deliver an IM message that contains spim or suspected spim.		x
Send notification	Send a predefined notification to the sender of an IM message that contains spim or an infected file informing the sender of the action.	x	x

Multiple actions per verdict

Within a filtering policy, you can create compound actions, performing multiple actions for a particular verdict.

An example follows:

1. Defining a virus policy, the administrator selects the Virus verdict and then assigns the actions, Clean, Add annotation, and Send notification to the policy.
2. Defining a policy group, the administrator assigns members then selects the new virus policy.

3 An email message is received whose recipients include someone in the new policy group.

4 Symantec Brightmail Gateway cleans the message, annotates it, delivers it, then sends a notification to its intended recipients.

When more than one filtering policy applies to a message, Symantec Brightmail Gateway uses special logic to combine actions from different filtering policies.

See "Verdict and action combinations" on page 724.

Verdict and action combinations

Symantec Brightmail Gateway offers the ability to combine multiple actions for different verdicts on the same message. This capability provides advantages over a model in which only one verdict for a message can result in actions. For example, suppose a spam message also contains a virus and your policies specify quarantining of spam messages and cleaning of viruses. Instead of cleaning the virus and delivering the spam to user inboxes, Symantec Brightmail Gateway cleans the virus and holds the cleaned spam message in Spam Quarantine. Or, if your policies specify modification of the subject line of spam messages and cleaning of viruses, Symantec Brightmail Gateway cleans the virus from the message and modifies the subject line.

Other types of messages can be affected by more than one filtering policy. A message can meet the criteria for two different content filtering policies. Or, the same spam message could contain a virus and meet the criteria for several content filtering policies. Symantec Brightmail Gateway combines the various filtering policies to determine which actions should be taken on the message.

In order to implement multiple actions, Symantec Brightmail Gateway includes sophisticated processing logic that automatically resolves potential conflicts between actions. In general, there is no need to worry about how actions will combine between your filtering policies. However, because a particular message can match multiple filtering policies, the resulting actions may not match your expectations. This section explains the basics of how actions from different policies can combine.

What happens to a message depends on the particular combination of actions applied to that message by the one or more policies that affect the message. In other words, actions combine with each other (or not, in some cases) based solely on action types. The kind of policy that called for the action has no impact on processing. The order in which actions are listed in the Control Center has no impact on processing.

For example, you create a content filtering policy to take action on messages that contain two or more words from your Profanity custom dictionary in the subject, body, or attachments of the message. You only use this policy for your Sales group. The action that you specify for these messages is Delete message. Your default virus policy specifies the action Clean the message, and your default spam policy specifies the action Modify the subject line, placing [SPAM] before the subject line text. Your Sales group uses the default virus and spam policies. A spam message addressed to a member of your Sales group arrives containing three words from your Profanity dictionary and also containing a virus. What happens to that message?

Because one of the actions specified is Delete message, Symantec Brightmail Gateway deletes the message and does not apply the other actions. In most cases, the Delete message action prevents other actions from occurring. However, what if the content filtering policy did not apply because the message contained only one word from your Profanity dictionary? In that case, the message is cleaned and delivered to the user's inbox with [SPAM] prepended to the subject line.

Many types of actions from different policies can be combined for the same message.

See "Limits on combining actions" on page 725.

About multiple content filtering policies

When more than one content filtering policy applies to a message, some of the actions specified may not take place as follows.

- The order of policies on the **Email Content Filtering Policies** page determines content filtering policy priority. Higher priority content filtering policies appear higher up in the list.
- Actions specified for the highest priority content filtering policy that applies to a message are triggered according to the rules for combining actions.
  See "Limits on combining actions" on page 725.
- For the other content filtering policies that apply to the message, the only actions that can happen are the Send notification and Create an incident (without holding for review) actions.

Limits on combining actions

Symantec Brightmail Gateway includes two kinds of limitations on action combinations:

Limitations imposed by the Control Center user interface	These limitations apply when you are creating a policy. When more than one policy applies to a message, additional combinations can result that exceed these limitations. See “User interface action combinations” on page 728.
Limitations imposed by message processing	These limitations apply at all times. Even if actions are invoked by multiple policies, the actions on a message cannot exceed these limitations. See “Action processing combinations” on page 726.

Action processing combinations

In general, the actions invoked by each policy determine how a message is processed, not the policies themselves. However, there are exceptions to this rule. Actions invoked by end-user Good and Bad Senders groups override all actions except those actions invoked by content filtering policies.

See “Verdict combinations” on page 733.

Note: If two actions are both the result of content filtering policies, additional restrictions on their combination apply.

See “About multiple content filtering policies” on page 725.

Table A-3 shows the processing category for each action. Use the processing category to find the action inTable A-4

Table A-3 Action processing categories

Processing Category	Actions
Firewall	Defer SMTP Connection, Reject SMTP Connection, Reject messages failing bounce attack validation
Event	Archive the message, Create an incident (without holding for review), Forward a copy of the message, Send a bounce message, Send notification

Table A-3 Action processing categories *(continued)*

Processing Category	**Actions**
Delay	Create an incident and hold message for review, Hold message in Suspect Virus Quarantine, Strip and Delay in Suspect Virus Quarantine
Delete	Delete message
Modify	Add a header, Add annotation, Add BCC recipients, Clean the message, Deliver the message to the recipient's Spam folder, Deliver message with TLS encryption, Deliver message with content encryption (this action exhibits both route and modify category behaviors), Modify the subject line, Remove unresolved recipients, Strip attachments
Route	Hold message in Spam Quarantine, Route the message, Deliver message with content encryption (this action exhibits both route and modify category behaviors)
No action	Deliver message normally

Table A-4 shows how actions in processing categories combine. Actions are listed according to their processing category. Match an action-processing category in the left-hand column with an action-processing category in the top row to see how actions in those processing categories combine. Actions whose names begin with Treat as are processed according to the action specified by the policy to which they refer. For example, the **Treat as a virus** action is processed according to the action specified by your virus policy. Actions whose names begin with Bypass are subject only to the user interface limitations.

See "User interface action combinations" on page 728.

Special rules apply when combining routing and modify actions, as follows:

- If one action is Hold message in Spam Quarantine and one is Route the message, the message is moved to Spam Quarantine.
- If both actions are Hold message in Spam Quarantine, the message is moved to Spam Quarantine.
- If both actions are Route the message, the message is routed to the routing address that appears first in an alphanumeric sort.
- When combining Route and Modify actions, if one action is Hold message in Spam Quarantine and one action is Deliver message with TLS encryption, the TLS encryption does not take place.

Table A-4 Processing action combination matrix

	Firewall	Event	Delay	Delete	Modify	Route	No action
Firewall	Firewall	Firewall	Firewall	Firewall	Firewall	Firewall	Firewall
Event	Firewall	Event + event	Event + delay	Delete + event	Modify + event	Route + event	Event
Delay	Firewall	Event + delay	Delay	Delay, defer deletion	Delay, defer modification	Delay, defer routing	Delay
Delete	Firewall	Delete + event	Delay, defer deletion	Delete	Delete	Delete	Delete
Modify	Firewall	Modify + event	Delay, defer modification	Delete	Modify + modify	Modify + route	Modify
Route	Firewall	Route + event	Delay, defer routing	Delete	Modify + route	One route wins	Route
No action	Firewall	Event	Delay	Delete	Modify	Route	No action

User interface action combinations

Table A-5 describes the limitations on combining actions within a filtering policy. These limitations are imposed by the Control Center user interface. Note that the Treat as actions are those listed in the last 6 rows of the table.

When you create a policy, the Control Center user interface imposes the limitations described in Table A-5. More than one policy can impact a message. Therefore, these limitations can be exceeded for a particular message. However, message processing imposes additional limitations that cannot be exceeded.

See "Action processing combinations" on page 726.

Table A-5 Compatibility of filtering actions by verdict

Action	Compatibility with other actions	Can add multiple times?
Add a header	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes
Add annotation	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Add BCC recipients	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes

Table A-5 Compatibility of filtering actions by verdict *(continued)*

Action	Compatibility with other actions	Can add multiple times?
Archive the message	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Bypass content filtering policy	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes. One per content filtering policy.
Bypass spam scanning	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Clean the message	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Create an incident (without Hold message for review option)	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes
Create an incident (with Hold message for review option checked)	Can only be combined with Send notification.	No
Defer SMTP Connection	Cannot be used with other actions	No
Delete message	Send a bounce message, Send notification, Archive the message, Create an incident, Forward a copy of the message, Bypass content filtering policy, Bypass spam scanning	No
Deliver message normally	Any except Hold message in Suspect Virus Quarantine, Delete message, Hold message in Spam Quarantine, Strip and Delay in Suspect Virus Quarantine, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Deliver the message to the recipient's Spam folder	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Deliver message with TLS encryption	Any except Delete message, Deliver message normally, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Forward a copy of the message	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes

Table A-5 Compatibility of filtering actions by verdict *(continued)*

Action	Compatibility with other actions	Can add multiple times?
Hold message in Spam Quarantine	Any except Hold message in Suspect Virus Quarantine, Deliver the message normally, Delete message, Route the message, Strip and Delay in Suspect Virus Quarantine, Defer SMTP Connection, Reject SMTP Connection, Treat as actions If used with Deliver the message to the recipient's spam folder, affected messages are quarantined. If released from Spam Quarantine, messages are delivered to the recipient's Spam folder.	No
Hold message in Suspect Virus Quarantine	Any except Delete message, Deliver message normally, Hold message in Spam Quarantine, Strip and Delay in Suspect Virus Quarantine, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Modify the subject line	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	One for prepend and one for append
Reject messages failing bounce attack validation	Cannot be used with other actions	No
Reject SMTP Connection	Cannot be used with other actions	No
Remove unresolved recipients	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Route the message	Any except Delete message, Defer SMTP Connection, Hold message in Spam Quarantine, Reject SMTP Connection, Treat as actions	No
Send a bounce message	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Send notification	Any except Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Strip and Delay in Suspect Virus Quarantine	Any except Delete message, Deliver the message normally, Hold message in Spam Quarantine, Hold message in Suspect Virus Quarantine, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	No
Strip attachments	Any except Delete message, Defer SMTP Connection, Reject SMTP Connection, Treat as actions	Yes
Treat as a bad sender	Cannot be used with other actions.	No

Table A-5 Compatibility of filtering actions by verdict *(continued)*

Action	Compatibility with other actions	Can add multiple times?
Treat as a mass-mailing worm	Cannot be used with other actions.	No
Treat as a good sender	Cannot be used with other actions. When used in a content filtering policy, messages that match the policy will not be scanned for spam.	No
Treat as a virus	Cannot be used with other actions.	No
Treat as spam	Cannot be used with other actions.	No
Treat as suspected spam	Cannot be used with other actions.	No

Verdicts by verdict category

Table A-6 Verdicts by verdict category

Verdict Category	Verdict	Description
Bad Sender Policies	Directory harvest attack	An attempt is underway to capture valid email addresses. A directory harvest attack is accomplished by emailing to your domain with a specified number of non-existent recipient addresses sent from the same IP address.
	Email virus attack	A specified quantity of infected email messages has been received from a particular IP address.
	Bad Sender Groups	An email message, domain, or IP address is a member of one of the following groups: ■ Local Bad Sender Domains ■ Local Bad Sender IPs ■ Third Party Bad Senders ■ Symantec Global Bad Senders See “About blocking and allowing messages using sender groups” on page 174.

Table A-6 Verdicts by verdict category *(continued)*

Verdict Category	Verdict	Description
Good Sender Policies	Good Sender Groups	An email message, domain, or IP address is a member of one of the following groups: ■ Local Good Sender Domains ■ Local Good Sender IPs ■ Third Party Good Senders ■ Symantec Global Good Senders See "About blocking and allowing messages using sender groups" on page 174.
	Fastpass	Allows most email messages from verified good senders to bypass spam filtering. You cannot specify any actions for the Fastpass verdict.
Sender authentication	Sender authentication	An email message has failed either SPF or Sender ID authentication. See "Enabling SPF and Sender ID authentication" on page 135.
Virus	Virus	An email or IM message contains a virus, based on current Symantec virus filters.
	Mass-mailing worm	An email or IM message contains a mass-mailing worm, based on current Symantec virus filters.
	Unscannable for viruses	An email or IM message exceeds the container limits configured on the Scanning Settings page or is unscannable for other reasons. For example, a message or an attachment that contains malformed MIME cannot be scanned for viruses.
	Encrypted attachment	An email or IM message contains an attachment that is encrypted or password-protected and therefore cannot be scanned.
	Spyware or adware	An email or IM message contains any of the following types of security risks: spyware, adware, hack tools, dialers, joke programs, or remote access programs. See "Spyware or adware verdict details" on page 207.
	Suspicious attachment	An email or IM message either shows virus-like signs or because suspicious new patterns of message flow involving this attachment have been detected.

Table A-6 Verdicts by verdict category *(continued)*

Verdict Category	Verdict	Description
Spam	Spam	An email message is spam, based on current spam filters from Symantec.
	Suspected spam	An email message is suspected spam, based on a configurable Suspected Spam Threshold.
	Failed bounce attack validation	An email message is part of a bounce attack, based on bounce attack validation filtering. See "About defending against bounce attacks" on page 182.
Spim	Spim	An IM message contains spim, based on current spim filters from Symantec.
Content Filtering	Text in the Subject, Body, or Attachments	An email message contains keywords in your configurable dictionary, matches/does not match a regular expression or pattern, or matches data in a record.
	Text in this specific part of the message	Text in any of 13 message parts contains or matches in one of several ways a specific string, or matches/does not match a regular expression or pattern.
	Text in this specific part of the message header	Text in the envelope recipient or envelope sender contains/does not contain an email address, domain, or country code from a specific dictionary.
	Message size	The message size is equal to/greater than/less than a specific number of bytes, KB, or MB.
	File metadata	An attachment is in an attachment list, has a specific filename or MIME type, or contains/does not contain a filename or file extension from specific dictionary.
	For all messages	All email is flagged. You can create a content filtering rule that applies to all messages, for example to universally attach an annotation to all inbound or outbound email messages.

Verdict combinations

In general, messages from senders in Good Sender groups bypass spam filtering but do not bypass virus filtering or content filtering. This section provides a more detailed explanation of how this process works.

Messages arriving at the gateway first undergo connection-time processing.

Messages from certain domains or IP addresses can be deferred, rejected, or accepted, based on the following:

- Locally (within your system) collected reputation information
- Globally collected reputation information
- Local Good and Bad Sender Groups
- Global Good and Bad Sender Groups

These actions occur before the Brightmail Engine processes the messages.

See "How Symantec Brightmail Gateway works" on page 35.

During connection processing, certain steps can be skipped, as shown in Table A-7.

Table A-7 Connection-time good sender processing

If any of these match:	Symantec Brightmail Gateway does not check for any of these:
Local Good Sender IPs Third Party Good Senders Symantec Global Good Senders	Third Party Bad Senders
Local Good Sender IPs Symantec Global Good Senders	Symantec Global Bad Senders
Local Good Sender IPs	Local Bad Sender IPs

See "About blocking and allowing messages using sender groups" on page 174.

See "Verdicts by verdict category" on page 731.

During Brightmail Engine message processing, certain steps may be skipped, as shown in Table A-8 and Table A-9.

Table A-8 Brightmail Engine inbound message processing

If any of these match:	Symantec Brightmail Gateway does not check for any of these:
Third Party Good Senders	Third Party Bad Senders Spam Suspected spam Sender authentication failure

Table A-8 Brightmail Engine inbound message processing *(continued)*

If any of these match:	Symantec Brightmail Gateway does not check for any of these:
Local Good Sender Domains Local Good Sender IPs Content filtering policy with a Bypass spam scanning action Content filtering policy with a Treat as good sender action Symantec Global Good Senders End user Good Senders List	Third Party Bad Senders Spam Suspected spam Sender authentication failure Symantec Global Bad Senders
Local Good Sender IPs Content filtering policy with a Treat as a good sender action End user Good Senders List	Local Bad Sender IPs Local Bad Sender Domains
Local Good Sender Domains	Local Bad Sender Domains
Spam	Suspected spam
When the associated policy specifies a Bypass content filtering policy action: ■ Local Good Sender Domains ■ Local Good Sender IPs ■ Third Party Good Senders ■ Local Bad Sender Domains ■ Local Bad Sender IPs ■ Third Party Bad Senders ■ Symantec Global Good Senders ■ Symantec Global Bad Senders ■ Spam ■ Suspected spam ■ Sender authentication failure	All content filtering policies or those content filtering policies specified in the policy

Table A-9 Brightmail Engine outbound message processing

If any of these match:	Symantec Brightmail Gateway does not check for any of these:
Content filtering policy with a Bypass spam scanning action	Spam Suspected spam
When the associated policy specifies a Bypass content filtering policy action: ■ Spam ■ Suspected spam	All content filtering policies or those content filtering policies specified in the policy
Spam	Suspected spam

Appendix B

Command reference

This appendix includes the following topics:

- more
- mta-control
- netstat
- nslookup
- password
- ping
- reboot
- route
- rpmdb
- service
- show
- shutdown
- sshd-config
- tail
- telnet
- traceroute
- update

agent-config

`agent-config` – configures the agent that connects hosts to the Control Center

SYNOPSIS

```
agent-config [--norestart] [--force] --add | --delete ip
agent-config --help | --status
agent-config [--norestart] --log level
```

DESCRIPTION

The `agent-config` command lets you edit the allowed IP configuration for the Scanner. Use this command when you change the IP address of the Control Center. You must run this command on every host to re-allow the new Control Center IP to connect to the hosts. The Agent restarts when you add or delete an IP address to or from the allowed IP list, unless you include `--norestart` in the command.

OPTIONS

`--add, -a`

Add an IP address to the agent-allowed IP address list.

`--delete, -d` *`ip`*

Delete an IP address from the agent-allowed IP address list. Specify an IP address in dotted quad format. For example, `192.168.2.1`.

`--log, -l` *`level`*

Set the log level. The log levels are listed below from least verbose to most verbose and each level includes the previous level. For example, if you specify the `errors` level, only the most urgent log messages are stored. If you specify the `notices` level, `errors`, `warnings`, and `notices` level log messages are stored.

Specify one of the following log levels:

- `errors`
- `warnings`
- `notices`
- `information`
- `debug`

`--force, -f`

Used with `--delete` option to bypass the deletion warning.

`--help, -h`

Display this message.

`--norestart, -n`

Do not restart the agent after modifying the IP address list or log level.

`--status, -s`

Display the allowed IP address list and current log level.

cat

`cat` – standard Linux command to view a file

DESCRIPTION

The `cat` command displays the contents of plain text files. The `more` command can be more useful than `cat` for listing long files or multiple files.

Type `help cat` on the command line for more information about the options available for `cat`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

The `cat` command is a standard Linux command that has been modified to only display the files that the `list` command shows.

SEE ALSO

See list on page 776.

See more on page 791.

cc-config

`cc-config` – configures the logging and network access to the Control Center

SYNOPSIS

```
cc-config ( --help | --status )
cc-config cclog --level level
cc-config compliancelog --days days
cc-config database ( --status | --check [tableName] | --repair
[tableName] | --optimize [tableName] )
cc-config http ( --on | --off )
cc-config port-443 ( --on | --off )
```

DESCRIPTION

The `cc-config` command lets you modify the selected settings that the Control Center uses. These settings include Content Filtering Audit logs, port 443 access, and more.

ARGUMENTS

`cclog`

Change the log level of the main Control Center log, BrightmailLog.log.

When you apply this to the Control Center log, `cc-config` writes the command-line parameters to the log4j properties file. It then restarts the Control Center.

`compliancelog`

Change the rollover frequency of the Content Filtering log.

`database`

List, optimize, validate, or repair the database tables that the Control Center uses.

`http`

Turn on or off access to the Control Center through port 80 (the standard, unsecure port for Web servers).

If http access is off, you cannot access the Control Center with a URL that starts with http://. If http access is on, you can access the Control Center with a URL that starts with http://. To access the Control Center using http, append

:41080 to the URL. Regardless of the http setting, you can always access the Control Center with a URL that starts with https://.

`port-443`

Turn on or off access to the Control Center through port 443 (the standard, ssl-secured port for Web servers).

When port 443 access is off, you must append :41443 to the URL when you use an https:// URL to access the Control Center. When port 443 access is enabled, you do not need to append the port number for an https:// URL to access the Control Center.

OPTIONS

`--check, -c`

Check the given database table. If no table name is specified, then check all tables.

`--days, -d`

Set the number of days to keep logs before they roll over.

`--help, -h`

Display this message.

`--level, -l`

Set the log level. The log levels are listed below from least verbose to most verbose and each level includes the previous level. For example, if you specify the `errors` level, only the most urgent log messages are stored. If you specify the `debug` level, `errors`, `warnings`, `information` and `debug` level log messages are stored.

Specify one of the following log levels:

- errors
- warnings
- information
- debug

`--off`

Disable a feature.

`--on`

Enable a feature.

`--optimize, -o`

Optimize the table so it takes less space on disk. If no table name is specified, then optimize all tables.

`--repair, -r`

Repair the given database table. If no table name is specified, then attempt a repair operation on all damaged tables.

`--status, -s`

Display the current log settings and port statuses.

clear

`clear` – standard Linux command to clear the screen

SYNOPSIS

```
clear
```

DESCRIPTION

The `clear` command erases all of the text on the screen and displays the command prompt at the top of the screen.

This command is a standard Linux command that has not been modified.

db-backup

`db-backup` – back up the Control Center database

SYNOPSIS

```
db-backup [options]
```

DESCRIPTION

The `db-backup` command backs up the Brightmail databases, such as configuration settings, report data, log data, and incidents. You can store backups on the appliance or on a remote server. Only run this command on the appliance that contains the Control Center. This command does not function on a Scanner-only appliance. Only one instance of `db-backup` can run at a time.

By default, backup files are compressed before they are written to disk to minimize the size of backup files. The `db-backup` command calculates the amount of disk space the backup file requires. It does not run if at least twice this amount is available on the `/data` partition.

Use `db-restore` or the Control Center restore feature to restore a backup on the appliance or a backup on a remote remote computer. If you specify `--file` *`path`* for a backup to the appliance, you can only restore the backup using the `db-restore` command, not the Control Center restore feature.

You can also create backups using the Control Center. In the Control Center, click **Administration > Hosts > Version > Backup**.

OPTIONS

`--backup, -b` *`number`*

The number of backups to store on the appliance. If you have more backups stored than *number*, then older backups are deleted. Each unique combination of type and schedule is retained separately. If you do not specify `--backup` *`number`*, the default is 5 for each type and schedule combination. See examples 4 and 6.

`--file, -f` *`path`*

The name and, optionally, location to save the backup. Use the `--file` option to specify an alternate file name for the backup file or to save the backup file to a remote computer. If you do not specify `--file` *`path`*, the backup is saved to the appliance as

`db-backup.brightmail.`*`Mon-Day-Year-Hour-Min`*`.full.manual.tar.bz2.`

You can save the backup to a remote computer using either FTP (file transfer protocol) or SCP (secure copy protocol). If the path ends with `/` the backup is saved in that directory using the default file name. If the path ends with a file name the backup is saved with that name in the specified path. When you save the backup to a remote computer, `db-backup` temporarily stores the backup file on the appliance, checks the file for data integrity, copies the file to the remote computer, and checks to ensure that the file was successfully copied.

Use one of the following two path formats to save the backup to a remote server:

`FTP`

Use the following format: `ftp://'`*`user`*`':'`*`password`*`'@`*`host`*`[:`*`port`*`]/`*`path`*. If special characters are included in the password, you must enclose the password in single quotes (‘). If the special characters in a password include a single quote, you can use the double quote instead (“). Passwords containing single and double quotes are not valid. If no user name and password are specified, an anonymous login is used.

`SCP`

Use the following format: `scp://'`*`user`*`'@`*`host`*`/`*`path`*. You must specify a user name. The `db-backup` command prompts you for the password.

`--gzip, -g`

Use the gzip compression algorithm instead of the default bzip2 compression algorithm. The gzip algorithm performs less efficient compression than bzip2.

`--list, -l`

List existing backups on the appliance.

`--help, -h`

Display this message.

`--nocompress, -n`

Do not compress the backup file. Use this option if you want to visually scan the file contents.

`--purge, -p`

Purge backups. Use the `--purge` option to delete old backup files that match the parameters that you specify. To delete all but the *number* most recent backups of a type and schedule combination, specify `--purge --backup` *`number`* along with the type and schedule. Specify `--purge --backup 0` to delete all backups of a type and schedule combination. To delete a specific file, specify `--file` *`file`* along with `--purge`. See examples 5 and 6.

`--schedule, -s` *`schedule`*

The schedule name to include in the backup file name. If you specify a schedule name, `db-backup` does not create automatic backups at that interval. The schedule that you specify only names the backup file with that name. The schedule names differentiate backups. See `--backup` and `--purge` for more information. Use the backup feature in the Control Center to create automatic scheduled backups. The following schedules are available:

`manual`

Label this backup a manual backup. This option is the default.

`daily`

Label the backup a daily, manual backup.

`weekly`

Label the backup a weekly, manual backup.

`monthly`

Label the backup a monthly, manual backup.

`--type, -t` *`type`*

The type of backup to create. Each backup type has two aliases that are alternate short versions of the backup type. See example 4. The following types are available:

`full`

Perform a full backup (aliases: `f`, `1`). This option is the default.

`config-incidents`

Back up configuration and content filtering incident data (aliases: `ci`, `2`).

`config-incidents-reports-logs`

Back up configuration, content filtering incident, report and log data (aliases: `cirl`, `3`).

EXAMPLES

Example 1

Save a full backup on the appliance with the default schedule of `manual` and the default type of `full`. The newest five backups with a schedule of `manual` and type of `full` are kept (including the backup just created) and the rest of the backups matching that combination are deleted.

```
db-backup
```

Example 2

Save a full backup on a remote server with SCP. The database backup file in the format `db-backup.brightmail.`*`date-time`*`.full.manual.tar.bz2` is copied to 192.168.2.42 in the /tmp directory through SCP. Log on to the SCP server with the `support` user account. The `db-backup` command prompts for the password for the `support` user account.

```
db-backup --file scp://support@192.168.2.42/tmp/
```

Example 3

Save a full backup on a remote server with FTP. The database backup file `db-backup.brightmail.`*`date-time`*`.full.manual.tar.bz2` is copied to `host.symantecexample.org` in the `/user/jmuir` directory. Log on to the FTP server with the `jmuir` user account and `secret` password.

```
db-backup -f ftp://jmuir:secret@host.symantecexample.org/user/jmuir/
```

Example 4

Backup configuration and content filtering incident data to the appliance and include the word `weekly` in the backup file name. In addition to the newly created backup, keep one additional existing backup with `config-incidents` and `weekly` in the file name.

```
db-backup --backup 2 --schedule weekly --type ci
```

Example 5

Delete a single backup file.

```
db-backup --purge --file
db-backup.brightmail.Jan-21-10-19-26.config-incidents.weekly.tar.bz2
```

Example 6

Delete all but the one most recent backup file of type `config-incidents` and schedule `manual`.

```
db-backup --purge --backup 1 --type config-incidents --schedule manual
```

SEE ALSO

See db-restore on page 750.

db-restore

`db-restore` – restores the Brightmail databases to an appliance from previously created backups on the appliance or from remote locations with FTP, SCP, and HTTP

SYNOPSIS

```
db-restore [--force --list --help] file
```

DESCRIPTION

The `db-restore` command restores Brightmail databases to an appliance from a single, previously created backup. These are the backups that you have previously generated and saved on the appliance or from remote locations with FTP, SCP, and HTTP. If you attempt to run more than one instance of `db-restore` at a time, an error results. If any part of the operation fails, `db-restore` fails, and an explanatory message appears on the command line. You must be on the Control Center host to use the `db-restore` command.

The backup that you create from the Control Center may not be the same as the backup that you create from the command line. If you want to restore system policies and system databases, restore the databases from the Control Center, not this command-line option. The `db-restore` command expects the compressed backup files that `db-backup` creates.

When you restore a database backup on a different appliance than it was created, keep in mind the following considerations:

- If you restore the appliance from a backup taken on a different appliance, the restored appliance inherits the network configuration of the other appliance. For example, assume that you create a backup on Computer A. You restore the backup on Computer B. Computer A's configuration settings overwrite the configuration settings of Computer B.
- If you attempt to restore a backup to an appliance other than the one on which it was created, you must reboot the appliance.

Stop the Control Center while this operation runs. Restart it when the restore has completed.

OPTIONS

`--force, -f`

Force a restore even when the version of appliance software in the backup file differs from the software that is currently on the appliance.

`--list, -l`

List the backup files that are stored on the appliance.

`--help, -h`

Display this message.

ARGUMENTS

Specify *file* with one of the following formats. If the file is stored on a remote computer, specify the directory path to the file.

file

Type the file name without the FTP, HTTP, or SCP prefix to specify a backup that is stored locally.

ftp://*user*:*password*@[:*port*] /*path*

Copy files from their remote location with FTP.

Logon is attempted with the user name and password credentials that you provide on the command line. If special characters are included in the password, enclose the password in single quotes (‘). If the special characters in a password include a single quote, you can use the double quote instead (“). If no credentials are specified, anonymous logon is used. Error checking ensures that the copies are complete.

http://*host*[:*port*]/*path*

Allow for Web-based transfer of a restore file from the Control Center. With this mode, backups that are stored on your local appliance can be retrieved at either of the following addresses:

```
http://host.domain.com:41080/brightmail/backups/file.bz2
```

```
https://host.domain.com/brightmail/backups/file.bz2
```

If special characters are included in the password, enclose the password in single quotes (‘). If the special characters in a password include a single quote, you can use the double quote instead (“). Error checking ensures that the copies are complete.

Note: If you use Internet Explorer, ensure that **Do not save encrypted pages to disk** is unchecked. This option is found in the **Internet Explorer Tools** menu > **Internet Options** menu, expanded **Security** view.

`scp://`*`username@host/path`*

Copy the backup file from its remote location with SCP. A complete path, file name, and user name are required when you specify a backup file through SCP. You are prompted for a password for the user name that you specify. Return codes are checked to ensure that the entire backup file is copied from the remote host. The script exits with non-zero status on failure. If the script fails, an error message appears. Error checking ensures that the copies are complete.

SEE ALSO

See db-backup on page 746.

See diagnostics on page 759.

See “Restarting an appliance” on page 691.

See “Stopping and starting Scanners” on page 108.

delete

`delete` – clear logs, configuration information, and data

SYNOPSIS

```
delete [--purge num] component component ...
delete file file
```

DESCRIPTION

Use the `delete` command to delete logs, configuration information, and other data. You may want to delete data if disk space is low or to clear configuration data to correct or diagnose a problem. The `delete` command restarts the Brightmail Engine if necessary after you run the `delete` command.

OPTIONS

`--purge, -p` *`num`*

Delete all database backup files except for the *num* most recent files. This option is only valid with the `database` component.

ARGUMENTS

You can delete individual files or you can specify one or more components to delete logical groups of files.

`file` *`file`*

Delete the file that you specify. You can only delete the files that you can view with the `list` command. Specify the entire path to the file as shown by the `list` command.

Symantec recommends that you delete items by specifying a component instead of deleting individual files. If you delete individual files, you may change the effectiveness or performance of Symantec Brightmail Gateway. If you delete log files or temporary files with the `delete file` *`file`* command, some log data may be lost. To delete log files, specify one of the components in the log components group.

If you do delete individual log files with the `delete file` *`file`* command, restart the service that applies to the log file that you deleted. For example, if you delete the Control Center log file `Brightmaillog.log`, restart the Control

Center service. Use the `service` command or the Control Center to restart a service.

The following components are available and are listed in groups of similar behavior.

Log components:

`alllogs`

Delete all logs in the log component group.

`bcclogs`

Delete all Control Center logs.

`ddslogs`

Delete all directory data service logs.

If you delete `ddslogs`, the `bmclient_log` and `bmserver_log` log files may contain many `Could not connect: Connection refused` errors. These errors are normal.

`imlogs`

Delete all IM logs.

`mallogs`

Delete all Message Audit Logs.

`oslogs`

Delete all operating system logs.

`scannerlogs`

Delete all Scanner logs.

Configuration components:

`allconfig`

Delete all configuration data in the configuration component group.

`bccconfig`

Delete all Control Center configuration files.

`clearsockets`

Delete all socket files in the `/var/tmp` directory.

`ddsconfig`

Delete all directory data service configuration files.

`imconfig`

Delete IM configuration files.

`osconfig`

Delete operating system configuration files.

`scannerconfig`

Delete all of the Scanner configuration files for a given Scanner (including support sieve scripts). It does not affect the Scanner configuration information that is stored in the Control Center.

When you run `delete scannerconfig`, it restarts the appliance on which the command is run. After you run `delete scannerconfig`, you must recommit Scanner configuration information from the Control Center to disk and relicense your Scanner.

You can recommit the Scanner information to disk unchanged or edit the information to correct potential problems before you save this information to disk. To do either of these tasks, access **Administration > Hosts > Configuration** in the Control Center, select the Scanner, and click **Edit**. To recommit the information unchanged, click **Save**. Alternatively, edit any settings for this Scanner as necessary to correct a problem in the configuration and click **Save**.

You can delete the Scanner configuration if you change the Scanner configuration of an independent Scanner appliance. Then you can re-add it with the Add Scanner Wizard. This option is not available for an appliance that hosts both a Control Center and Scanner.

Symantec recommands that you do not use `delete scannerconfig`.

Data components:

`alldata`

Delete all data in the data component group.

`bccdata`

Delete all Control Center data including any license files. Afterwards, your configuration is the same as an out-of-the-box the Control Center configuration.

`ddsdata`

Delete all directory data service data.

`imdata`

Delete all IM data.

`keystore`

Delete Control Center HTTPS certificates from the keystore.

`scannerdata`

Delete mail from MTA queues and the following file:

```
/data/scanner/rules/matchEngine/tmp/data_match_engine_jce_keystore
```

`sudata`

Delete all of the files that are related to software updates.

Quarantine components:

`allquarantine`

Delete all messages from all quarantines.

`contentquarantine`

Delete all messages from the content quarantine.

`spamquarantine`

Delete all messages from Spam Quarantine.

`virusquarantine`

Delete all messages from Suspect Virus Quarantine.

Rule components:

`allrules`

Delete all rules and replace them with the factory default rules.

`avrules`

Delete all antivirus rules and replace them with the factory default rules.

`dayzerorules`

Delete all day zero rules and replace them with the factory defaults rules.

`fastpassrules`

Delete all Fastpass rules.

`gatekeeperrules`

Delete gatekeeper antispam rules and replace with factory default rules.

`intsigrules`

Delete all intsig rules and replace them with the factory default rules.

`ipfreqrules`

Delete IP frequency rules.

`regexrules`

Delete regex filter rules.

`spamhunterrules`

Delete all spam hunter rules and replace them with the factory default rules.

`spamsigrules`

Delete spamsig rules and replace them with the factory default rules.

Note: The `delete` command may take half a minute to delete rules. Wait for the command prompt to return before you run additional commands. Do not press **Ctrl+C** to stop the `delete` command while it is running.

Miscellaneous components:

`all`

Delete all logs, configuration data, passwords, support sieve scripts, Scanner data, cores, diagnostic packages, rules, queue data, and backup files to restore your appliance to the original factory configuration.

`bcchostacl`

Delete the Scanner access controls made on the **Administration > Settings > Control Center** page to permit access from all Scanners.

`cores`

Delete all core directories.

`database`

Delete all backups of the Control Center database that were created with `db-backup`.

`diagnostics`

Delete all diagnostic packages.

`help`

Display a summary of components that you can delete.

`monitor`

Delete the files made by the `monitor` command.

`oldqueuedata`

Delete all old Postfix email queue data.

EXAMPLES

Example 1

Delete the `BrightmailLog.log` file.

```
delete file /data/logs/bcc/BrightmailLog.log
```

Example 2

Delete all messages in the Spam Quarantine.

```
delete spamquarantine
```

Example 3

Delete all Control Center database backup files that are stored on the appliance except for the three most recent backup files.

```
delete --purge 3 database
```

SEE ALSO

See cat on page 741.

See list on page 776.

See more on page 791.

See “Clear disk space checklist” on page 634.

diagnostics

`diagnostics` – generate diagnostics package

SYNOPSIS

```
diagnostics [options] url
```

DESCRIPTION

The `diagnostics` command generates a diagnostic package that Symantec Support can use to analyze problems with the product.

You should specify a valid URL unless you use the `--find-other-cores` option. If you specify a valid URL but do not specify the data collection options, `diagnostics` uses the following parameters by default:

```
--config --crash-info 5 --logs 100000
```

When the user name or password are part of the URL, write them in quotes if they have any special shell characters in them. The password can be specified in the URL or at the password prompt. An example of the URL syntax is as follows:

```
scp://'user':'password'\@host[:port]/path
```

If you specify a path that ends with a forward slash, the diagnostics file is written to the path that you specify with the default file name. If you specify a path that does not end with a forward slash, the backup file is written with the file name specified in the path.

The default diagnostics file name is in the following format:

```
diagnostics.yy-mmm-dd-hh-mm.hostname.tar.gz
```

For example:

```
diagnostics.09-Sep-10-15-42.host9902.symantecexample.com.tar.gz
```

An option cannot be specified more than once whether it is in its long form or short form. For the `--cores` option, a component cannot be specified more than once either with the component name or convenient string `all`. If you attempt to, an error message appears along with the appropriate usage text.

OPTIONS

`--config, -c`

Collect only the configuration data.

`--cores, -o` *`component n`*

Collect the latest *n* core directories, including core files for a component. The valid range for *n* is 1 through 9,999.

The list of components include the following:

- `--cores mta` *`n`* collects MTA core packages
- `--cores bmagent` *`n`* collects Brightmail Agent core packages
- `--cores imrelay` *`n`* collects IM Relay core packages
- `--cores bmserver` *`n`* collects Brightmail Server core packages
- `--cores conduit` *`n`* collects Conduit core packages
- `--cores jlu-controller` *`n`* collects Java LiveUpdate core packages
- `--cores other` *`n`* collects the other core files that are not collected with other options.
- `--cores all` *`n`*

  `all` is a convenient identifier that means all components.

`--crash-info` *`n`*`, -a`

Collect the latest *n* core directories (excluding the core files in those directories) for the following processes:

- mta
- bmagent
- imrelay
- bmserver
- conduit
- jlu-controller

The valid range for *n* is 1 through 9,999.

`--edm, -e`

Collect the exact data match (EDM) record sets.

`--find-other-cores, -d`

Discover any core file outside of `/data/scanner/jobs` and move them to `/data/scanner/jobs/other`.

If Symantec Brightmail Gateway discovers and moves any core files, an email notification is sent to the administrators that are specified to receive alerts. If not, no email notification is sent.

You can use this option with the `delete cores` command to clean up core files on your product. Run this command first to move the core files that are not in the jobs directory to the jobs directory. Then use `delete cores` to delete the core files.

If `--find-other-cores` is the only data collection option specified, a URL is not required. No diagnostics package is generated.

`--force, -f`

Force diagnostics to run even if a diagnostics collection that is started from the user interface is still in progress. If a package creation is in progress, the existing diagnostics collection fails.

`--help, -h`

Display this message.

`--include-old-queues, -i`

Collect queue data from old postfix queues.

This command is only useful on configurations in which Symantec Brightmail Gateway is migrated from version 7.7 or earlier. This command is not applicable for Symantec Brightmail Gateway version 8 or higher.

`--ldap, -p`

Collect legacy ldapsync data.

This command is only useful on configurations in which Symantec Brightmail Gateway is migrated from version 8 or earlier. This command is not applicable for Symantec Brightmail Gateway version 9 or higher.

`--logs all, -l`

Collect the entire log file.

`--logs` *`n`*`, -l`

Collect log data that is limited to *n* lines per log file.

The valid range for *n* is 1 through 2,147,483,647.

`--monitor, -m`

Collect a snapshot output of the following `monitor` command: `monitor -c 6 --proc bmserver --proc mta system database disk mta p_all` and existing monitor logs under `/data/monitor`.

`--rules, -r`

Collect all rules that are present on the Scanner, except exact data match data.

`--tracking, -t`

Collect Message Audit Log files.

`--verbose, -v`

Show the command process in verbose mode.

ARGUMENTS

The syntax for the URL paths referenced by this command is as follows:

- scp://'user':'password'\@host[:port]/path
  Copies the diagnostics package remotely through SCP.
- ftp://'user':'password'\@host[:port]/path
  Copies the diagnostics package remotely through FTP.

Logon is attempted with the user name and password credentials that are provided on the command line. If special characters are included in the password, you must enclose the password in single quotes (‘). If the special characters in a password include a single quote, you can use the double quote instead (“). If no credentials are specified, anonymous logon is used.

EXAMPLES

Create a diagnostics file and transfer it with the SCP protocol. The diagnostics file (in the format: `diagnostics.yy-mmm-dd-hh-mm.hostname.tar.gz`) is transferred to the SCP destination.

```
diagnostics scp://'support'@10.160.248.128/tmp/
```

Note: The month is expressed in the three-letter format, not two-digit format.

dns-control

`dns-control` – control the local DNS cache

SYNOPSIS

```
dns-control command
```

DESCRIPTION

The `dns-control` command manages local caching for the name server.

All `dns-control` command outputs end with either a completion message or a failure message. Examples are: "Command cmdname completed successfully" and "Command cmdname failed."

Some commands require the DNS cache to be running before they can be executed. In these cases, the only output is: "The DNS Cache is currently stopped." Start the cache with the `dns-control start` command before you run those commands.

ARGUMENTS

The command components are as follows:

`start`

Start the local caching name server.

`stop`

Stop the local caching name server.

`restart`

Restart the local caching name server.

`status`

Display the status of the local caching name server.

`flush`

Flush the cache.

`list`

List the locally configured name servers for the resolver.

`trace`

Increment the tracing (debug) level by +1.

`notrace`

Disable tracing (debug).

`reconfig`

Reconfigure name server.

`help`

Display this page.

grep

`grep` – a standard Linux command to search in files for text or a regular expression

DESCRIPTION

The `grep` command searches in the files that you specify for text or regular expressions.

Type `help grep` on the command line for more information about the options available for `grep`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

help

`help` – Display help for individual commands or display all available commands.

SYNOPSIS

```
help [ --list | command ]
```

DESCRIPTION

The `help` command displays a list of available commands on the product. If you specify a command name, the `help` command displays help for that command.

The help for commands is presented in Linux man page format. These help pages use the following Linux man page conventions. Do not type the brackets, parenthesis, or pipe symbol when you run a command.

`Brackets [ ]`

The options and the arguments that are listed within square brackets are optional. The options and the arguments that are not listed within square brackets are required.

`Parenthesis ( )`

The options and the arguments that are listed within parenthesis are required but are mutually exclusive. A pipe symbol separates the mutually exclusive options or arguments.

`Pipe |`

The pipe symbol indicates the options or arguments that are mutually exclusive. For example `[ -e` *`pattern`* `| -f` *`file`* `]` means that you can specify `-e` *`pattern`* or `-f` *`file`*, but not both.

`Colored, italic, or underlined text`

Text that is italic, colored, or underlined indicates that you should substitute that text with specific text. When you type `help` *`command`*, the terminal or terminal software that you use to access the command line determines how this text appears. When you view help pages in a PDF or in the online help, this type of text is italic.

`--option, -o`

Some command options are available in long and short versions. The long version and short version produce the same behavior. Use whichever version is most convenient for you. In the OPTIONS section, these options are displayed with the long version first, followed by a comma, and then the short version. The long version is preceded with two dashes and the short version is preceded with one dash. Some options have required parameters that you specify after the option, like a log level or IP address.

The help pages contain the following sections:

`SYNOPSIS`

A description of the options and arguments available for the command.

`DESCRIPTION`

General information about the command.

`OPTIONS`

Options that you can use to control the behavior of a command. Options always begin with one or two dashes, such as `-s` or `--status`. If an option is listed in square brackets in the synopsis, the options are optional. If not, the option is required.

Some options have arguments. For example, `--log` *level*. Square brackets indicate optional arguments.

Not all commands have options.

`ARGUMENTS`

Some commands require arguments. Arguments are names of files, host names, IP addresses, and so on that you specify to control the behavior of the command. Not all commands have arguments.

`EXAMPLES`

The EXAMPLES section provides sample command usage. Not all commands have examples.

`SEE ALSO`

The SEE ALSO section lists related commands. Not all commands have see also references.

Use the following commands to navigate through the help pages:

`f or SPACE`

Forward one screen

`b`

Backward one screen

`/pattern`

Search for a word or pattern

`<`

Go to the beginning of the document

`>`

Go to the end of the document

`q`

Exit from the document and display the command prompt

`h`

Display additional information about navigating the help pages

OPTIONS

`--list, -l`

Display a list of all the available commands.

ARGUMENTS

`command`

Display help for the specified command.

If you do not specify a command, help for the `help` command is displayed (this page). Specify one of the following commands:

`agent-config`

Configures the agent that connects hosts to the Control Center

`cat`

Standard Linux command to view a file

`cc-config`

Configures the logging and network access to the Control Center

`clear`

A standard Linux command to clear the screen

`db-backup`

Back up the Control Center database

`db-restore`

Restores the Brightmail databases to an appliance from previously created backups on the appliance or from remote locations with FTP, SCP, and HTTP

`delete`

Clear logs, configuration information, and data

`diagnostics`

Generate diagnostics package

`dns-control`

Control the local DNS cache

`grep`

a standard Linux command to search in files for text or a regular expression

`help`

Display help for individual commands or display all available commands

`ifconfig`

A standard Linux command to configure network interfaces

`iostat`

A standard Linux command to display CPU and device load

`ldapsearch`

A standard Linux command to query an LDAP directory

`list`

Display the file names of all files that certain commands can act on

`mallog`

List, backup, or restore Message Audit Logs

`malquery`

Query Message Audit Logs

`monitor`

View and record information about Brightmail-specific processes

`more`

A standard Linux command to page through a text file

`mta-control`

Control the MTA processes and backup and restore mail queues

`netstat`

A standard Linux command to view network connections

`nslookup`

A standard Linux command to query DNS servers

`password`

Change your administrative password

`ping`

A standard Linux command to test for a response from a remote computer

`reboot`

Reboot the appliance

`route`

A standard Linux command to show and manipulate the IP routing table

`rpmdb`

Manage and repair the RPM database

`service`

A standard Linux command to start or stop services

`show`

Display system information

`shutdown`

Shut down the appliance without rebooting

`sshd-config`

Configure which addresses can SSH to the appliance

`tail`

A standard Linux command to view the end of a file

`telnet`

A standard Linux command to connect to a remote computer

`traceroute`

A standard Linux command to view the path that network packets take

`update`

Update the appliance software

HISTORY

In Symantec Brightmail Gateway version 9.0, some commands that existed in version 8.0 and previous versions were renamed, incorporated into other commands, or removed. The following commands were changed in version 9.0:

`agentconfig`

Replaced with `agent-config`.

`clear`

Replaced with `delete`. In version 9.0, the `clear` command clears the screen.

`crawler`

Part of `diagnostics`.

`date`

Replaced with `show --date`.

`deleter`

Replaced with `delete cores`.

`dn-normalize`

The functionality of the `dn-normalize` command is not available in version 9.0.

`eula`

Replaced with `show --eula`.

`http`

Replaced with `cc-config http`.

`install`

Replaced with `update install`.

`ls`

Replaced with `list`.

`mta-stats`

Replaced with `monitor mta`.

`passwd`

Replaced with `password`.

`pause-mode`

Replaced with `mta-control pause-mode`.

`rebuildrpmdb`

Replaced with `rpmdb --repair`.

`rm`

Replaced with `delete files`.

`set-control-center-port-443`

Replaced with `cc-config port-443`.

`sshdctl`

Replaced with `sshd-config`.

`sshdver`

Replaced with `sshd-config --version`.

`sys-info`

Replaced with `show --info`.

`system-stats`

Replaced with `monitor system`.

`tls-ca-cert-control`

The functionality of the `tls-ca-cert-control` command is not available in version 9.0.

ifconfig

`ifconfig` – a standard Linux command to configure network interfaces

DESCRIPTION

The `ifconfig` command displays the status and configuration of network interfaces and can make temporary changes to interface configurations.

Type `help ifconfig` on the command line for more information about the options available for `ifconfig`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

iostat

`iostat` – a standard Linux command to display CPU and device load

DESCRIPTION

The `iostat` command monitors system input/output device loading by observing the time devices are active in relation to their average transfer rates.

Type `help iostat` on the command line for more information about the options available for `iostat`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

ldapsearch

`ldapsearch` – a standard Linux command to query an LDAP directory

DESCRIPTION

The `ldapsearch` command searches in the LDAP source that you specify and displays matching records.

Type `help ldapsearch` on the command line for more information about the options available for `ldapsearch`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

list

list – display the file names of all files that certain commands can act on

SYNOPSIS

```
list [--all] [--cores] [--diagnostics] [--logs] [--monitor] [--temp]
[--top]
list --help
```

DESCRIPTION

The list command displays the file names of all of the files that can be acted upon by certain commands. The following commands can act upon the files that are listed with list:

cat

Display the contents of one or more files.

delete

Delete one or more files.

more

Display the contents of one or more files and pause at the end of each screen.

tail

Show the last 50 lines of the named log file.

OPTIONS

If list does not list any files when you specify an option, there are no files in that category.

--all, -a

List all files.

--cores, -c

List all core files.

--diagnostics, -d

List all diagnostic packages.

--help, -h

Display this message.

`--logs, -l`

List all log files.

`--monitor, -m`

List all monitor files.

`--temp, -p`

List all temporary files.

`--top, -t`

List the largest files that the administrator can delete and their sizes.

EXAMPLES

Example 1

List all the files that can be viewed with `cat` (except core files and diagnostic files) or deleted with `delete`.

```
list --all
```

Example 2

List the largest files that you can delete. You can use the `delete` command to delete large files if you do not need them.

```
list --top
```

SEE ALSO

See cat on page 741.

See delete on page 753.

See more on page 791.

mallog

`mallog` – list, backup, or restore Message Audit Logs

SYNOPSIS

```
mallog [ --list | --help ]
mallog [ --backup | --restore ] url
```

DESCRIPTION

The `mallog` command backs up and restores Message Audit Log data that resides on the Scanner. The `mallog` command also lists the Message Audit Log files on the Scanner. To view message activity in the Message Audit Logs, use the Control Center or the `malquery` command.

Available log files include the following:

- `/data/logs/scanner/audit_bmengine_log*`
- `/data/logs/scanner/audit_mte_log*`
- `/data/logs/scanner/audit_mta_log*`

Note: When you run `mallog --backup` or `mallog --restore`, email processing stops while these commands run. No inbound email or outbound email is delivered during this time. If your organization's email availability policies are strict, it may be appropriate to only run these commands during off hours.

OPTIONS

`--backup` *url*

Create a backup of all of the message tracking logs that are in tar.gz format, and upload the resulting file to the specified URL.

Note: This option suspends mail processing while the command is executed.

`--help`

Display this message.

`--list`

List individual message tracking logs on the file system and their timestamps and sizes.

`--restore` *`url`*

Restore message tracking logs from the specified URL. Existing logs are overwritten.

Note: This option suspends mail processing while the command is executed.

URLs may have a scheme of either FTP, SCP, or, HTTP (for restore only).

If you specify a path that ends with a forward slash, the diagnostics file is written to the path that you specify with the default file name. If you specify a path that does not end with a forward slash, the backup file is written with the file name specified in the path. The `--restore` option requires a full path name which includes a file name. The entire URL should be taken in double quotes. If any part of the URL contains special characters, such as full or double quotes, escape the special characters with a backslash. When the password is part of the URL, it should be written in quotes if it has any special shell characters in it.

`url`

Transmit the package to the *url* location by SCP or FTP.

The entire URL should be taken in double quotes. If any part of the URL contains special characters, such as full or double quotes, escape the special characters with a backslash. When the password is part of the URL, it should be written in quotes if it has any special shell characters in it.

SEE ALSO

See malquery on page 780.

malquery

`malquery` – query Message Audit Logs

SYNOPSIS

```
malquery (-l start,end | -g start,end)
(-u uid [-u uid ...] | -e event[,arg_num]<=|*>string [-e ...] | -q
event[,arg_num]<=|*>quoted-printable-string [-q ...])
[-m max_results] [-I index_max] [-o output_file] [-d] [-v]
```

DESCRIPTION

You can track messages in the Control Center by querying the Message Audit Logs. Alternatively, you can use the `malquery` command-line command to track messages. Use `malquery` instead of the Control Center for complex queries or queries where you expect voluminous data. The `malquery` command only returns data for the Scanner that you are logged into.

Enabling Message Audit Logging results in approximately 800 bytes of audit logs per message. Message Audit Logging can cause performance and storage problems if your site receives more than 1,000,000 messages per day.

The output from `malquery` is in .xsd format, for example:

```
<malResults count="message result count">
        <message UID="uid">
            <events>
                <event time="utc" name="event id">parameters</event>
                <event time="utc" name="event id">parameters</event>
                <event time="utc" name="event id">parameters</event>
                <event time="utc" name="event id">parameters</event>
            </events>
        </message>
</malResults>
```

OPTIONS

`-e ...`

Find email messages that contain the events that match the specified criterion.

Examples:

-e RCPTS=dale@company.com

RCPTS is recipient. In this example, the recipient is dale@company.com.

-e SUBJECT*"my flowers"

SUBJECT is the subject of the email message. In this example, the subject contains the words 'my flowers'.

`-g` *`start,end`*

Find messages by the GMT date range to search in UNIX time (the number of time units that have elapsed since the epoch time 1/1/1970). For example,

July 4, 2008, 11:59 P.M. = 1215212340.

Separate the start date and end date by a comma with no space.

`-I` *`index_max_n`*

Use the index (.idx file) if the number of matching results is less than or equal to *index_max_n*. Otherwise, the index is ignored. This option searches a flat file, which saves time when you want to look up large numbers of events.

The default for *index_max_n* is 1000.

`-l` *`start,end`*

Find messages based on the specified date range. The date format is YYYYMMDDhhmm. For the hours and minutes, use a 24-hour clock. For example:

July 4, 2008, 11:59 P.M. = 200807042359.

Separate the start date and end date by a comma with no space.

`-m` *`max_results`*

Return the *max_results* number of messages. The default is 1000.

`-o` *`file`*

Output data that matches results to *file*.

`-q ...`

Find email messages that contain the events that match the specified criterion in quoted-printable encoding. For example:

-q SUBJECT*"red =3D rose" -- subject contains 'red = rose'

`-u` *`audit_id`*

Find the email message with the specified audit ID.

`-v`

Show the command process in verbose mode.

EXAMPLES

Example 1

Search for an email based on the following criteria:

- Start date is between July 4, 2008, 2:00 P.M. and date of July 4, 2008, 11:59 P.M. in GMT time
- Recipient is “dale@company.com”
- Subject contains the words "check this out"
- Maximum output is 500 results
- Write results to file `/tmp/results.xml`

```
malquery -g 1215140340,1215212340 -e RCPTS=dale@company.com -e
SUBJECT*"check this out" -m 500 -o /tmp/results.xml
```

Example 2

Search for an email based on the following criteria:

- Start date is between July 4, 2009, 11:00 P.M. and date of July 4, 2009, 11:59 P.M.
- Recipient is “dale@company.com”
- The audit ID of *uid number*
- Maximum output is 500 results
- Write results to file `/tmp/results.xml`

```
malquery -l 200907042300,200907042359 -e RCPTS=dale@company.com -uid
uid number -m 500 -o /tmp/results.xml
```

Example 3

Search for an email based on the following criteria:

- Start date is between July 4, 2009, 11:00 P.M. and date of July 4, 2009, 11:59 P.M.
- Recipient is “dale@company.com”
- Subject contains the quoted-printable encoding words that return the email messages that match the word “red” in the subject. For example: “red roses”.
- Maximum output is 500 results
- Write results to file `/tmp/results.xml`

```
malquery -l 200907040000,200907090000 -e RCPTS=dale@company.com -q
SUBJECT*"red " -m 500 -o /tmp/results.xml
```

Example 4

Search for an email based on the following criteria:

- Start date between July 4, 2009, 11:00 P.M. and date of July 4, 2009, 11:59 P.M.
- Recipient is “dale@company.com”
- Subject contains the words "check this out"
- Number of matching results for this command is the default index_max of 1000
- Write results to file `/tmp/results.xml`

```
malquery -l 200907040000,200907090000 -e RCPTS=dale@company.com -e
SUBJECT*"check this out" -I -o /tmp/results.xml
```

SEE ALSO

See mallog on page 778.

See “About message audit logging” on page 646.

See “Audit log events” on page 648.

monitor

`monitor` – view and record information about Brightmail-specific processes

SYNOPSIS

```
monitor options [--proc name] [identifier ...]
monitor list
monitor stop ( pid | all )
```

DESCRIPTION

The `monitor` command lets you view and record detailed information about the Brightmail system and its processes.

OPTIONS

`--count, -c` *`num`*

Produce *num* samples.

The default is 1. The upper limit is 2^31-1 (roughly, 2.1 billion).

`--help, -h`

Display this message.

`--interval, -i` *`num`*

Take a sample at the *num* interval (measured in seconds).

The default is 10 seconds. For any long-running monitor jobs that are written to disk, you should increase this interval (to 60 or more). If the disk space fills up, the `monitor` process stops. Increase the interval time to avoid this issue.

`--output, -o` *`file`*

Save the output to a file instead of printing it to the console. The file is saved as `/data/monitor/`*`file`*.

When you use this option, `monitor` runs in the background and returns the process ID (PID) of the monitor process. Use `cat`, `more`, or `tail` to view the file. The file name can contain ASCII characters.

`--proc, -p` *`name`*

Collect data for one of the following Brightmail processes and its children. The valid process names and the programs that they represent are as follows:

`afasnmpd`

The `afasnmpd` process provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`bmagent`

The Brightmail Agent facilitates communicating configuration information between the Control Center and each Scanner.

`bmserver`

The `bmserver` process filters email messages.

`conduit`

The Conduit retrieves updated email filters and manages statistics.

`controlcenter`

The Control Center provides centralized Web administration, collects statistics, and hosts quarantines.

`im`

The IM process filters instant messaging.

`liveupdate`

LiveUpdate downloads virus definitions from Symantec Security Response to the Scanner.

`lsisnmpd`

The `lsisnmpd` process provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`monitor`

The `monitor` process displays or saves information about Symantec Brightmail Gateway processes.

`mta`

The mail transfer agent routes inbound and outbound messages to the Brightmail Engine for processing and delivers filtered messages.

`mysql`

The MySQL database on the Control Center stores settings and message information.

`percsnmpd`

The `percsnmpd` process provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`snmpd`

The `snmpd` process waits for requests from SNMP management software.

`stunnel`

The `stunnel` process provides secure encrypted connections.

`--quiet, -q`

Suppress any warnings from the monitor program.

`--tab, -t`

Produce data in a tabular format. Use the `--tab` option with the `--output` option to create output to import into a spreadsheet. The `--tab` does not format text correctly for the screen. For example, on the screen the column headings are not aligned with the column data.

When you format data for tabular output `--tab`, the column headings for each identifier are prefaced with the process name. For example, `controlcenter_p_%user`.

ARGUMENTS

`list`

Produce a list of all monitor processes, their PIDs, and the options that were used at runtime. The `monitor list` command always shows the `monitor list` command as one of the monitor processess that is running. This behavior is normal.

`stop ( `*`pid`*` | all )`

Stop the specified monitor processes. Type a PID to stop a single process. Type the word `all` to stop all monitor processes.

`identifiers`

The information that is displayed or saved depends on the identifiers that you specify. If you do not specify one or more identifiers, then the default of `system` is used. Some identifiers represent multiple identifiers and are provided for convenience. Five groups of identifiers are available: system, database, disk, MTA, and process.

System identifiers are as follows:

- `%user` - Percent of the available CPU time that is spent in user mode
- `%nice` - Percent of the available CPU time that is spent running as nice
- `%sys` - Percent of the available CPU time that is spent in system mode
- `%wait` - Percent of the available CPU time that is spent in IO wait
- `%idle` - Percent of the available CPU time that is spent idling
- `memt` - Total memory (k)

- memu - Memory in use (k)
- pageout - The number of memory pages that are swapped out to disk
- system - A convenience identifier that includes the following system identifiers: %user %sys %wait memt memu memf

Database Identifiers - These identifiers denote the size of the Control Center database, the size of its various quarantines, and how many messages they contain. The identifiers are as follows:

- db_size - The total size of the Control Center database in kilobytes

  The db_size includes not only the messages in the various quarantines, but other data that the Control Center maintains. All of the sizes do not add up, but inasmuch db_size should be equal to or larger than the sum of all the quarantine sizes.
- db_qsize - The size of the Spam Quarantine directory, in kilobytes
- db_qqty - The number of messages in the Spam Quarantine
- db_vsize - The size of the Suspect Virus Quarantine directory, in kilobytes
- dv_vqty - The number of messages in the Suspect Virus Quarantine
- db_csize - The size of the content incident directories
- db_cqty - The number of messages in the content incident quarantine
- database - A convenience identifier that includes all the database identifiers.

Disk identifiers - The disk identifiers provide information on disk utilization on the partitions that the administrator controls. The identifiers are as follows:

- data_used - The percentage of /data that is being used
- data_free - The amount of free space in /data, in kilobytes
- opt_used - The percentage of /opt that is being used
- opt_free - The amount of free space in /opt, in kilobytes
- other_used - The percentage of /opt that is being used (for example, 20%)
- other_free - The amount of space available on /opt in kilobytes
- disk - A convenience identifier that includes all the above disk data.

MTA identifiers - These identifiers report MTA statistics. The identifiers are as follows:

- i_conn - Number of inbound connections

- i_qmsgs - Number of queued inbound messages
- i_dmsgs - Number of deferred inbound messages
- i_qsize - Size of the inbound queue (k)
- i_drate - Inbound listener data rate
- i_mrate - Inbound listener message rate
- mta_in - All of the inbound statistics (the identifiers that begin with i_)
- o_conn - Number of outbound connections
- o_qmsgs - Number of queued outbound messages
- o_dmsgs - Number of deferred outbound messages
- o_qsize - Size of the outbound queue (k)
- o_drate - Outbound listener data rate
- o_mrate - Outbound listener message rate
- mta_out - All of the outbound statistics (the identifiers that begin with o_)
- d_conn - Number of delivery connections
- d_qmsgs - Number of queued delivery messages
- d_dmsgs - Number of deferred delivery messages
- d_qsize - Size of the delivery queue (k)
- d_drate - Delivery listener data rate
- d_mrate - Delivery listener message rate
- mta_del - All of the delivery statistics (the identifiers that begin with d_)
- mta - A convenience identifier that includes all of the MTA identifiers.

  The information that is collected depends on the identifiers that are provided. If none are provided, then the default of "system" is used. Some identifiers represent multiple identifiers and are provided for convenience.

  This command does not give any indication about the average load or amount of work that is done between one sample and the next. Each sample is a snapshot of the MTA status at that point in time.

Process identifiers - The --proc option lets you monitor statistics for groups of Brightmail processes. If the --proc flag is used without any p_* identifiers, the following default value is used: p_%user p_%sys p_memv p_memr p_mems. Identifiers for use with --proc include:

- p_%user - Percent of the available CPU time that is spent in user mode
- p_%sys - Percent of the available CPU time that is spent in system mode
- p_memv - Virtual memory that the processes use (k)
- p_memr - Resident memory in use by the processes (k)
- p_mems - Highest amount of the shared memory that any of the processes use (k)
- p_all - All of the proc identifiers

EXAMPLES

The following examples describe some ways that you can use the `monitor` command. These examples include a mix of the long and short forms of some of the option names, such as `-o` and `--output`.

Example 1

Check one time the percent of available CPU time and memory that the conduit service consumes. Save the result to file `/data/monitor/conduit_mon`.

```
monitor --proc conduit --output conduit_mon
```

Example 2

Collect the average load of the MTA service on the system every 3 seconds 1000 times. Display the average load on the system from the MTA service in a tabbed format and written out to file `/data/monitor/mta_mon`.

```
monitor --proc mta --interval 3 --count 1000 --tab --output mta_mon
```

Example 3

Collect the average load of the afasnmpd service on the system every 3 seconds 1000 times. Display the average load on the system from the MTA process in a tabbed format and saved to file `/data/monitor/snmp_mon`. Normally, no output appears on the screen because of the `-q` option.

```
monitor --proc afasnmpd --interval 3 --count 1000 --tab -q -o snmp_mon
```

Example 4

Check one time the percent of available CPU time and the memory that the LiveUpdate service uses. Save the result to file `/data/monitor/liveupdate_mon`.

```
monitor --proc liveupdate --output liveupdate_mon
```

Example 5

Check one time the percent of available CPU time and the memory that the monitor service consumes. Save the result to file `/data/monitor/monitor_mon` in tabbed format.

```
monitor --proc monitor --output monitor_mon --tab
```

Example 6

Check the percent of available CPU time and the memory the stunnel service consumes. Save the result to file `/data/monitor/stunnel _mon` in tabbed format. Normally, no output appears on the screen because of the `--quiet` option.

```
monitor --proc stunnel --output stunnel_mon --tab --quiet
```

SEE ALSO

See cat on page 741.

See delete on page 753.

See list on page 776.

See more on page 791.

See tail on page 812.

more

`more` – a standard Linux command to page through a text file

DESCRIPTION

The `more` command displays the contents of plain text files one screen at a time. Press **Space** to view the next screen. Use the `list` command to list the files that `more` can display.

You can run the output of another command to `more` to view the output one screen at a time. After the command that you are running, type the pipe symbol and then `more`. See the example below.

Type `help more` on the command line for more information about the options available for `more`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

The `more` command is a standard Linux command that has been modified to only display the files that the `list` command shows.

EXAMPLES

Example 1

Display `BrightmailLog.log` one screen at a time.

```
more /data/logs/bcc/BrightmailLog.log
```

Example 2

Examine the output of `list --top` one screen at a time.

```
list --top | more
```

SEE ALSO

See list on page 776.

mta-control

`mta-control` – control the MTA processes and backup and restore mail queues

SYNOPSIS

```
mta-control queue command
mta-control pause-mode mode
```

DESCRIPTION

The `mta-control` command lets you query MTA queues, and control specific elements within MTA message processing. For example, you can flush message queues.

Note: Do not use the ~ (tilde) character when you specify output file names, paths, passwords, email addresses, and user names (for exporting). Specify the full path name.

ARGUMENTS

Specify one of the following MTA queues:

- inbound
- outbound
- delivery
- all

The following components are available:

- start – Start the queue.
- stop – Stop the queue.
- status – Display the current status. The status can be: running, not running, enabled or disabled.
- restart – Restart the queue.
- flush – Reattempt delivery for all queued messages.
- delete-msgs-by-sender *regexp* – Delete from the queue all messages with Envelope Sender that matches the given Perl regular expression (case insensitive).

- delete-msgs-by-rcpt *regexp* – Delete from the queue all messages with an Envelope Recipient that matches the given Perl regular expression (case insensitive).

  Note: This deletes the entire message, not just the recipient.

- delete-msg-by-id *queue-ID* – Delete the message with the given queue-ID from the queue.
- delete-all-msgs – Delete all messages from the queue.
- active-routes – Print all active routes and the number of messages for each route.
- num-messages-in-route *route* – Print the number of messages for the given route.
- list-msgs *route* – Print the messages for the given route.
- list-msg-details *msgid* – Given a message ID, print details about that message.
- route-info *route* – Display DNS lookup information, destination, and number of messages for a route.
- reroute *src-routedst-route* – Reroute messages from *src-route* to *dst-route*.
- delete-msgs-by-sender *perl regexp* – Delete from the queue all messages with an envelope sender that matches the given Perl regular expression (case insensitive).
- delete-msgs-by-rcpt *perl regexp* – Delete from the queue all messages with an envelope recipient that matches the given Perl regular expression. Note that this deletes the entire message, not just the recipient (case insensitive).
- delete-msg-by-id *queue-ID* – Delete the message with the given queue-ID from the queue. Note that the ID is only unique per queue.
- delete-all-msgs – Delete all messages from the queue.
- import-queues *url* – Import an entire mail queue from backup. Specify `all` for the queue. Ensure that the MTA is running before importing a mail queue. To start the MTA, run `mta-control all start`. Specify the URL as described for the export-msg-by-id component.
- export-queues *url* – Back up the mail queue to a URL. Specify `all` for the queue. Ensure that the MTA is stopped before exporting the mail queue. To stop the MTA, run `mta-control all stop`. Specify the URL as described for the export-msg-by-id component.

- export-msg-by-id *queue-ID* [*url*] – Export the message with the given queue-ID from the queue and save it to the specified URL. If you do not specify a URL, the message data is displayed on the screen. If you do not specify the FTP password, `mta-control` prompts you for the password. If you specify a path that ends with '/', Symantec Brightmail Gateway stores the file in that location using a default file name. Otherwise, Symantec Brightmail Gateway stores the file with the file name that you specified in the path. The URL syntax is as follows:
  scp://'user'\@host/path (user is prompted for password)
  ftp://'user':'password'\@host[:port]/path
  ftp://'user'\@host[:port]/path
  Put a double-quote character before and after the URL. If any part of the URL contains special characters, such as full or double quotes, put a backslash before each special character.
- query-queue – Query the message queue.

  The following optional parameters are accepted:

  - sender_match=*perl regexp*
  - rcpt_match=*perl regexp*
  - deferred - selects the messages that are deferred
  - include_subject
  - start=N
  - limit=N
  - format=*neat|xml*

  The parameters sender_match, rcpt_match and deferred are logically ANDed together if present. The intermediate result set after applying these matches is sorted by date, and then the start and limit are applied: \$start messages are skipped and then \$limit messages are returned. The default is to show all messages in 'neat' format, which is meant to be human readable.
- query-queue – Query the message queue based on one or more provided parameters.

  The following optional parameters are accepted:

  - sender_match=*perl regexp*
  - rcpt_match=*perl regexp*
  - deferred - selects the messages that are deferred
  - include_subject

  - start=N
  - limit=N
  - format=*neat|xml*

  sender_match, rcpt_match and deferred are logically ANDed together if present. The intermediate result set after applying these matches is sorted by date, and then the start and limit are applied: $start messages are skipped and then $limit messages are returned.
  The default is to show all messages in 'neat' format, which is meant to be human readable.

  Note: For the following bad message queue commands, use `all` instead of a *queue-ID* to apply the command to all bad messages in the queue.

- bad-msg-list – List the times and IDs of messages in the bad message queue. The queue is either inbound or outbound.
- bad-msg-export *queue-ID* [*url*] – Export or display the message. See export-msg-by-id for URL format.
  To display the message on the screen, type `mta-control` *`queue`* `bad-msg-export` *`queue-ID`*.
  Specify the URL as described for the export-msg-by-id component.
- bad-msg-delete *queue-ID* – Delete the message.
- bad-msg-bypass *queue-ID* – Submit the message for delivery to the original recipients and bypass scanning.
- bad-msg-forward *queue-IDaddress* – Submit a copy of the message for delivery to the given address and bypass scanning. The original bad message remains in the bad message queue.
- bad-msg-retry *queue-ID* – Retry scanning the message as if it were new.

The six pause modes affect email scanning (`scan`), acceptance (`accept`), and delivery (`delivery`). Each pause mode sets scanning, acceptance, and delivery to a particular state as described below, regardless of the previous state of `scan`, `accept`, and `delivery`. Pause modes are as follows:

- `status` – Display the current pause mode status. If you type `mta-control pause-mode`, `mta-control` displays the pause mode status.
- `pause-accept` – Set `scan` to running and set `accept` to paused. The `delivery` state is not affected by `pause-accept`.

- pause-deliver – Set delivery to paused. The accept and scan states are not affected by pause-deliver. This is equivalent to mta-control delivery stop.
- pause-scan – Set scan to paused and set accept to running. The delivery state is not affected by pause-scan.
- resume-accept – Set scan to running and set accept to running. The delivery state is not affected by resume-accept
- resume-deliver – Set delivery to running. The accept and scan states are not affected by resume-deliver. This is equivalent to mta-control delivery start.
- resume-scan – Set scan to running and set accept to running. The delivery state is not affected by resume-scan.

EXAMPLES

Example 1

Show the status of the MTA (inbound, outbound, and delivery queues and whether they are running or not).

```
mta-control pause-mode status
```

Example 2

Do not accept any new mail on the appliance but scan mail in the queue. This command does not affect the delivery of email.

```
mta-control pause-mode pause-accept
```

Example 3

Accept email on the appliance, but do not scan it. This command does not affect the delivery of email.

```
mta-control pause-mode pause-scan
```

Example 4

Do not deliver email on the appliance.

```
mta-control pause-mode pause-deliver
```

Example 5

Accept and scan email on the appliance. This command does not affect the delivery of email.

```
mta-control pause-mode resume-accept
```

Example 6

Accept and scan email on the appliance. This command does not affect the delivery of email.

```
mta-control pause-mode resume-scan
```

Example 7

Deliver email on the appliance.

```
mta-control pause-mode resume-deliver
```

Example 8

Display the queue-id of messages in delivery queue.

```
mta-control delivery query-queue
```

Example 9

View a raw message in the delivery queue with a message queue-id.

```
mta-control delivery export-msg-by-id 00/00-25597-EFD46794
```

Example 10

Export a specific message from the delivery queue with a message queue-id. The message queue-id is 00/00-25597-EFD46794. Export it to the 192.168.159.99 SCP server in the /tmp directory with the support account. `mta-control` queries for the password.

```
mta-control delivery export-msg-by-id 00/00-25597-EFD46794
"scp://support\@192.168.159.99/tmp/"
```

Example 11

Export all message queues. Export the message queue file to the 192.168.159.99 FTP server in the /tmp directory with the sysadmin account. Since a password is not specified, `mta-control` queries for the password.

```
mta-control all export-queues "ftp://sysadmin\@192.168.159.99/tmp/"
```

Example 12

```
mta-control all query-queue
```

Show all messages currently in the inbound queue, the outbound queue, and the delivery queue.

netstat

`netstat` – a standard Linux command to view network connections

DESCRIPTION

The `netstat` command prints network connections, routing tables, interface statistics, masquerade connections, and multicast memberships.

Type `help netstat` on the command line for more information about the options available for `netstat`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

EXAMPLES

Example 1

Display network connections.

```
netstat -an
```

Example 2

Display routing table.

```
netstat -r
```

nslookup

`nslookup` – a standard Linux command to query DNS servers

DESCRIPTION

The `nslookup` command performs a DNS lookup of the given hostname or IP address.

Type `help nslookup` on the command line for more information about the options available for `nslookup`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is part of the standard Linux command set. It has been modified for use by Symantec Brightmail Gateway, but this modification does not affect its functionality.

EXAMPLES

Look up MX records for a domain (yahoo.com, for example):

```
nslookup -querytype=mx yahoo.com
```

password

`password` – change your administrative password

SYNOPSIS

```
password [--help] [--reset]
```

DESCRIPTION

The `password` command changes the password that you use to logon to the command line. You are prompted to type your old password, and to type your new password twice.

OPTIONS

`--help, -h`

Display this message.

`--reset, -r`

Set the administrative password to the factory default.

ping

`ping` – a standard Linux command to test for a response from a remote computer

DESCRIPTION

The `ping` command tests, through data packet, the transfer of that data between the appliance and the hostname or IP address that you specify.

Type `help ping` on the command line for more information about the options available for `ping`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

reboot

`reboot` – reboot the appliance

SYNOPSIS

```
reboot [--force]
```

DESCRIPTION

The `reboot` command stops all services and then restarts the appliance.

Note: If you reboot the appliance while you run software update on Symantec Brightmail Gateway, you can corrupt the appliance software.

OPTIONS

`--force, -f`

Reboot the appliance, even if software update is running (not recommended). The appliance can become corrupted and require reinstallation. Contact Symantec Technical Support for information about reinstalling the appliance software.

`--help, -h`

Display this message.

SEE ALSO

See shutdown on page 809.

route

`route` – a standard Linux command to show and manipulate the IP routing table

DESCRIPTION

The `route` command lets you view routing tables or add entries to a routing table temporarily. Its primary use is for viewing the routing tables.

Type `help route` on the command line for more information about the options available for `route`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

rpmdb

`rpmdb` – manage and repair the RPM database

SYNOPSIS

```
rpmdb [--verify] [--repair]
```

DESCRIPTION

The `rpmdb` command lets you verify the current RPM database and rebuild it. This command can be useful in the event the database is corrupted and you want to repair it. Software updates for Symantec Brightmail Gateway are stored as RPM packages.

OPTIONS

`--repair, -r`

Rebuild the RPM database.

`--verify, -v`

Verify the current RPM database.

service

`service` – a standard Linux command to start or stop services

SYNOPSIS

```
service name command
service name help
```

DESCRIPTION

Start, stop, and check the status of Symantec Brightmail Gateway services with the `service` command. Services are programs that run continuously to perform specific tasks. During normal operation, you do not have to stop or start services. You may need to stop or start services to diagnose or resolve a problem with Symantec Brightmail Gateway.

The `service` command is a standard Linux command that has been modified to work with services available on Symantec Brightmail Gateway.

ARGUMENTS

Specify a service *name* and *command* when you run `service`.

name

Specify one of the following service names:

`afasnmpd`

The `afasnmpd` service provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`agent`

The Brightmail Agent facilitates communicating configuration information between the Control Center and each Scanner.

`connector`

The Conduit and LiveUpdate services download spam and virus definitions.

`controlcenter`

The Control Center provides centralized Web administration, collects statistics, and hosts quarantines.

`dds`

Directory data service interfaces with LDAP to provide authentication, email address validation, message routing, and policy groups.

If you restart the `dds` service, the `bmclient_log` and `bmserver_log` log files may contain many `Could not connect: Connection refused` errors. These errors are normal.

`imrelayd`

The IM service filters instant messaging.

`lsisnmpd`

The `lsisnmpd` service provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`mta`

The mail transfer agent processes, routes, and delivers email messages in cooperation with the Brightmail Engine.

`mysql`

The MySQL database on the Control Center stores settings and message information.

`osconfig`

The `osconfig` service manages network interfaces and related services.

`percsnmpd`

The `percsnmpd` service provides SNMP information for some Dell PowerEdge Expandable RAID Controllers.

`smsswapfile`

The `smsswapfile` service manages secondary swap file space.

`snmpd`

The `snmpd` service waits for requests from SNMP management software.

`stunnel`

The `stunnel` service provides secure encrypted connections.

`command`

The following commands are available. Some commands do not apply to certain commands. Type `service` *`name`* `help` to display the commands that apply to a service.

`condrestart`

Restart the service only if it is currently running. This command is available only for the `controlcenter`, `snmpd`, `mta`, and `stunnel` services.

`delete`

Delete the swap file on the appliance. This command is available only for the `smsswapfile` service.

`help`

Display the commands available for the service that you specify.

`reload`

This command is available only for the `mysql` and `stunnel` services.

`restart`

Stop the service and then start the service.

`status`

Display the status of a service.

`start`

Start the service.

`stop`

Stop the service.

EXAMPLES

Example 1

Display the commands that are available for the `mta` service.

```
service mta help
```

Example 2

Display the status of the `mta` service.

```
service mta status
```

Example 3

Stop the `mta` service.

```
service mta stop
```

Example 4

Stop the Conduit, LiveUpdate, and jlu-controller.

```
service connector stop
```

show

`show` – display system information

SYNOPSIS

```
show [--date] [--eula] [--info] [--version]
show --help
```

DESCRIPTION

The `show` command displays the following information:

- Current date and time
- End User License Agreement
- System information
- Product version number

OPTIONS

`--date, -d`

Show the current date and time.

`--eula, -e`

Show the End User License Agreement.

`--help, -h`

Display this message.

`--info, -i`

Show the system hardware information.

`--version, -v`

Show the product version number and installation date.

shutdown

`shutdown` – shut down the appliance without rebooting

SYNOPSIS

```
shutdown [--help | --force]
```

DESCRIPTION

The `shutdown` command turns off the appliance immediately. The appliance is not restarted. Shutdown occurs immediately and email messages remain in the queues. To start an appliance after you run the `shutdown` command, you must press the appliance power button, unless you have configured remote access to the appliance hardware.

Note: If you shut down the appliance during the software update process, you can corrupt the appliance software.

OPTIONS

`--help, -h`

Display this message.

`--force, -f`

Shut down the appliance, even if software update is running (not recommended). The appliance can become corrupted and require reinstallation. Contact Symantec Technical Support for information about reinstalling the appliance software.

SEE ALSO

See reboot on page 802.

See “Turning off an appliance ” on page 690.

sshd-config

`sshd-config` – configure which addresses can SSH to the appliance

SYNOPSIS

```
sshd-config (--list | --help)
sshd-config --add (allow|deny) address
sshd-config --delete (allow|deny) rule#
sshd-config --version [1|2]
```

DESCRIPTION

The `sshd-config` command lets you specify which addresses can access the appliance through SSH.

OPTIONS

`--add, -a`

Add a new rule.

`--delete, -d`

Delete an active rule.

`--help, -h`

Display this message.

`--list, -l`

Display the active rules and the current protocol number.

`--version, -v`

Set the version number of the protocol to use (1 or 2).

ARGUMENTS

`allow/deny`

When an SSH client connects, the client address is compared to the allow list and deny list in the following order:

- If the client address matches any allow rules, then the connection is allowed.
- If the client address matches any deny rules, then the connection is rejected. The client is allowed.

`rule`

Each rule is a list of one or more addresses and wildcards that are separated by commas, as follows:

- some.hostname.com
  Matches a specific host
- .hostname.com
  Matches some.hostname.com and other.hostname.com
- 1.2.3.4
  Matches a specific IP address
- 1.2.
  Matches any IP address starting with 1.2
- 1.2.3.0/255.255.255.0
  Matches any IP address within the 1.2.3.* subnet
  The EXCEPT keyword can be used to exclude a subset of addresses. For example, hostname.com EXCEPT forbidden.hostname.com.

You can specify one of the following keywords instead of a host name or IP address for the address parameter. Use the KNOWN and UNKNOWN keywords with care since they depend on DNS service.

- ALL
  Matches any address
- LOCAL
  Matches any host whose name does not contain a dot character
- KNOWN
  Matches any host whose name and address are known
- UNKOWN
  Matches any host whose name or address are unknown

tail

`tail` – a standard Linux command to view the end of a file

SYNOPSIS

```
tail [-f | --help ] log_name
```

DESCRIPTION

The `tail` command is part of the standard Linux command set which shows the last 50 lines of the named log file.

However, this command is modified in the following ways:

- Only the `-f` and `--help` options that are described here are available.
- If a character in a log file is not printable or is not ASCII, the sequence \xAB is displayed instead of that character. AB is the hexadecimal value of the character. For example, a character with a decimal value of 128 is displayed as \x80.
- This command is restricted to the file names that are obtainable from the `list` command. The `list` command displays the file names of all of the files that can be acted upon by certain commands. In addition to the `tail` command, the following commands can act upon the files that are listed with `list`:

  `cat`

  Display the contents of one or more files.

  `delete`

  Delete one or more files.

  `more`

  Display the contents of one or more files and pause at the end of each screen.

OPTIONS

`-f`

Follow the file as new text is added to it. The `tail -f` command prints the last 10 lines of the file but does not exit. As new text lines are added to the file, `tail` displays the new text lines. The `-f` option is useful for monitoring a log file as additional information is added to the log file. If you type `tail`

`-f` *`log_name`* and nothing seems to happen, the file is empty, the file is not being written to, or both.

To stop monitoring a file, press **Ctrl+C**.

`--help, -h`

Display this message.

ARGUMENTS

`log_name`

log_name can be any of the following:

- `agent_log`
- `bmclient_log`
- `bmserver_log`
- `boot.log`
- `BrightmailLog.log`
- `conduit_log`
- `cron`
- `db-migration.log`
- `dds.log`
- `dmesg`
- `imlinkage_log`
- `imrelay_log`
- `jlu-controller_log`
- `liveupdt.log`
- `maillog`
- `messages`
- `named.run`
- `secure`
- `update.log`

EXAMPLES

Example 1

Display the last 50 lines of the `BrightmailLog.log` log file.

```
tail BrightmailLog.log
```

Example 2

During an update, monitor the `update.log` log file. If you see information being written to `update.log` periodically, it usually means that the update is proceeding normally.

```
tail -f update.log
```

SEE ALSO

See list on page 776.

telnet

`telnet` – a standard Linux command to connect to a remote computer

DESCRIPTION

The `telnet` command lets you log into the command line of another computer on your network from the appliance.

Type `help telnet` on the command line for more information about the options available for `telnet`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

traceroute

`traceroute` – a standard Linux command to view the path taken by network packets

DESCRIPTION

The `traceroute` command displays the network route to the given hostname or IP address.

Type `help traceroute` on the command line for more information about the options available for `traceroute`. The information that is displayed may contain references to commands that are not available on Symantec Brightmail Gateway.

This command is a standard Linux command that has not been modified.

update

`update` – update the appliance software

SYNOPSIS

```
update check | list
update ( download | install | notes ) [--version number]
update --help
```

DESCRIPTION

You can perform the following tasks with the `update` command:

- Check for new software updates
- Download software updates
- Install software updates
- List the available software updates for download or installation

Before you update the software, ensure that your appliance does not perform any tasks that if disrupted could cause problems after you reset the system.

OPTIONS

`--help, -h`

Display this message.

`--version, -v`

Specify a software update version number for the `download`, `install`, or `notes` arguments. Use `update list` to determine what versions are available for the `--version` option.

ARGUMENTS

`check`

Perform a test update. The test update demonstrates what happens if you choose to perform a software update. Running `update check` does not update your appliance software.

`download`

Download but do not install a software update. After you download a software update, you can install it by typing `update install`. If you do not specify a version, the latest software update is downloaded.

If your Internet connection to the Symantec software update servers is not reliable, try downloading as a separate step from installing.

`install`

Download and install a software update. If you do not specify a version, the latest software update is installed on your appliance.

`list`

Display the available software updates.

`notes`

Display the software update notes. If you do not specify a version, the latest software update notes are displayed.

EXAMPLES

Download but do not install a software update. After you download a software update, you can install it by typing `update install`.

```
update download
```

SEE ALSO

See “Software update best practices” on page 698.

Content Filtering templates

This appendix includes the following topics:

- U.S. regulatory policy templates
- Confidential data-protection policy templates
- Acceptable use policy templates
- Customer and employee data-protection templates
- Network security policy templates
- UK and international regulatory policy templates

U.S. regulatory policy templates

This section describes the US regulatory policy templates that Symantec Brightmail Gateway provides.

See "About content filtering policy templates" on page 336.

Table C-1 U.S. regulatory policy templates

Name	Description	Described Content	Structured Data
Export Administration Regulations (EAR) See "Export Administration Regulations (EAR)" on page 821.	The Export Administration Regulations (EAR) are enforced by the US Department of Commerce. These regulations primarily cover technologies and technical information with both commercial and military applications. They are also known as dual use technologies (for example, chemicals, satellites, software, computers). This policy detects violations based on countries and controlled technologies designated by the EAR.	X	X

Table C-1 U.S. regulatory policy templates *(continued)*

Name	Description	Described Content	Structured Data
Gramm-Leach Bliley See "Gramm-Leach-Bliley" on page 822.	The Gramm-Leach-Bliley (GLB) Act gives consumers the right to limit some sharing of their information by financial institutions. This policy detects transmittal of customer data.	X	X
HIPAA (including PHI) See "HIPAA" on page 823.	This policy strictly enforces the US Health Insurance Portability and Accountability Act (HIPAA) by searching for data concerning prescription drugs, diseases, and treatments in conjunction with Protected Health Information (PHI). This policy can be used by organizations that are not subject to HIPAA but want to control PHI data.	X	X
International Traffic in Arms Regulations (ITAR) See "International Traffic in Arms Regulations (ITAR)" on page 825.	The International Traffic in Arms Regulations (ITAR) are enforced by the US Department of State. Exporters of defense services or related technical data are required to register with the federal government and may need export licenses. This policy detects potential violations based on countries and controlled assets designated by the ITAR.	X	X
NASD Rule 2711 and NYSE Rules 351 and 472 See "NASD Rule 2711 and NYSE Rules 351 and 472" on page 825.	NASD Rule 2711 and NYSE Rules 351 and 472 protect the name(s) of any companies involved in an upcoming stock offering, internal project names for the offering, and the stock ticker symbols for the offering companies.	X	
NASD Rule 3010 and NYSE Rule 342 See "NASD Rule 3010 and NYSE Rule 342" on page 826.	NASD Rule and NYSE Rule 342 require brokers-dealers to supervise certain brokerage employee's communications. This policy monitors the communications of registered principals who are subject to these regulations.	X	
NERC Security Guidelines for Electric Utilities See "NERC Security Guidelines for Electric Utilities" on page 826.	Detects information outlined in the North American Electric Reliability Council (NERC) security guidelines for protecting and securing potentially sensitive information about critical electricity infrastructure.	X	
Office of Foreign Assets Control (OFAC) See "Office of Foreign Assets Control (OFAC)" on page 826.	The Office of Foreign Assets Control of the US Department of the Treasury administers and enforces economic and trade sanctions based on US foreign policy and national security goals against certain countries, individuals, and organizations. This policy detects communications involving these targeted groups.	X	

Table C-1 U.S. regulatory policy templates *(continued)*

Name	Description	Described Content	Structured Data
Payment Card Industry Data Security Standard See "Payment Card Industry Data Security Standard" on page 827.	The Payment Card Industry (PCI) data security standards are jointly determined by Visa and MasterCard to protect cardholders by safeguarding personally identifiable information. Visa's Cardholder Information Security Program (CISP) and MasterCard's Site Data Protection (SDP) program both work toward enforcing these standards. This policy detects credit card number data.	X	X
Sarbanes-Oxley See "Sarbanes-Oxley" on page 828.	The US Sarbanes-Oxley Act (SOX) imposes requirements on financial accounting, including the preservation of data integrity and the ability to create an audit trail. This policy detects sensitive financial data.	X	
SEC Fair Disclosure Regulation See "SEC Fair Disclosure Regulation" on page 829.	The US SEC Selective Disclosure and Insider Trading Rules prohibit public companies from selectively divulging material information to analysts and institutional investors prior to its general release to the public. This policy detects data indicating disclosure of material financial information.	X	
State Data Privacy See "State Data Privacy" on page 830.	Many states in the US have adopted statutes mandating data protection and public disclosure of information security in which confidential data of individuals is compromised. This policy detects these breaches of confidentiality.	X	X

Export Administration Regulations (EAR)

Described Content condition

EAR Commerce Control List and Recipients - A compound rule that looks for both a country code in the recipient from the EAR Country Codes dictionary and a keyword from the EAR CCL Keywords dictionary.

Structured Data conditions

Indexed EAR Commerce Control List Items and Recipients – A compound condition that looks for both a country code in the recipient from the "EAR Country Codes" dictionary and for a specific "SKU" from a Record resource View.

Included dictionaries

The included dictionaries are as follows:

- EAR Country Codes
- EAR CCL Keywords

Gramm-Leach-Bliley

Described Content condition

The Described Content conditions are as follows:

- US Social Security Numbers – Looks for social security numbers. For this rule to match, there must be both a number that fits the Valid Social Security Number premium pattern, and a keyword or phrase that indicates the presence of a US SSN with a keyword from the US SSN Keywords dictionary. The keyword condition is included to reduce false positives with numbers that may match the SSN format.
- ABA Routing Numbers – Looks for a match to the ABA Routing number regex rule and a keyword from the dictionary "ABA Routing Number Keywords.
- Credit Card Numbers, All - Looks for credit card numbers. Similar to the first rule, this rule requires that there be both a number that fits a credit card system pattern ccn and a keyword or phrase that indicates the presence of a credit card number from the Credit Card Number Keywords dictionary. The keyword condition is included to reduce false positives with numbers that may match the credit card format.

Structured Data conditions

The Structured Data conditions are as follows:

- Username/Password Combinations – Looks for user names and passwords in combination from a Record resource View.
- 3 or more critical customer fields – Looks for any three fields that can identify a customer uniquely from a Record resource View, except for combinations of phone, email, and first or last name.
- Exact SSN or CCN – Looks for SSN or Bank Card Number from a Record resource View.
- Customer Directory – Looks for Phone or Email from a Record resource View.

Included dictionaries

The included dictionaries are as follows:

- US SSN Keywords
- Credit Card Number Keywords
- ABA Routing Number Keywords

HIPAA

Described Content condition

The Described Content conditions are as follows:

- SSN and Drug Keywords – Looks for the Social Security number (SSN) with the Valid Social Security Number premium pattern, in combination with a keyword from US SSN Keywords dictionary and a keyword from the Prescription Drug Names dictionary.
- SSN and Treatment Keywords – Looks for the Social Security number (SSN) with the Valid Social Security Number premium pattern, in combination with a keyword from US SSN Keywords dictionary and a keyword from the Medical Treatment Keywords dictionary.
- SSN and Disease Keywords – Looks for the Social Security number (SSN) with the Valid Social Security Number premium pattern, in combination with a keyword from US SSN Keywords dictionary and a keyword from the Disease Names dictionary.
- SSN and Drug Codes – Looks for the Social Security number (SSN) with the Valid Social Security Number premium pattern, in combination with a keyword from US SSN Keywords dictionary and a drug code using the Drug Code regular expression.

Structured Data conditions

The Structured Data conditions are as follows:

- Patient Data and Drug Codes – Any part of the message matches the NDC Drug Code regular expression and any part of the message matches text in a Record resource View.
- Patient Data and Drug Keywords – Any part of the message matches a Prescription Drug Names dictionary entry and any part of the message matches an entry in a Record resource View
- Patient Data and Treatment Keywords – Any part of the message matches an entry in the Medical Treatment Keywords dictionary and any part of the message matches text in a Record resource View.
- Patient Data and Disease Keywords – Any part of the message matches an entry in the Disease Names dictionary and any part of the message matches text in a Record resource View.

Exception conditions:

- TPO Exception – Looks for a recipient email address matching one from the TPO Email Addresses dictionary. If a match is found, the policy is not triggered even if the other conditions are met.

  Note: TPOs (Treatment, Payment, or health care Operations)—companies that partner with the health care organization—have a specific carve-out for the HIPAA information restrictions. This exception in the rules does not trigger the policy if the protected information is sent to one of these allowed companies. The template requires that the customer enter the allowed email addresses of these companies.

Included dictionaries

The included dictionaries are as follows:

- US SSN Keywords
- Prescription Drug Names
- Medical Treatment Keywords
- Disease Names
- TPO Email Addresses

International Traffic in Arms Regulations (ITAR)

Described Content condition

ITAR Munitions List and Recipients - A compound rule that looks for both a recipient country code from the ITAR Country Codes dictionary and a keyword from the ITAR Munition Names dictionary.

Structured Data conditions

Indexed ITAR Munition Items and Recipients – A compound rule that looks for a recipient country code from the ITAR Country Codes dictionary and for a specific Stock Keeping Unit (SKU) number from a Record resource View.

Included dictionaries

The included dictionaries are as follows:

- ITAR Country Codes
- ITAR Munition Names

NASD Rule 2711 and NYSE Rules 351 and 472

Described Content condition

NASD Rule 2711 and NYSE Rules 351 and 472 - Compound rule that contains a sender condition and a keyword condition. The sender condition requires editing by the user and is a list of email addresses of research analysts at the user’s company (Analysts' Email Addresses dictionary). The keyword condition works for any upcoming stock offering, internal project names for the offering, and the stock ticker symbols for the offering companies (NASD 2711 Keywords dictionary). Like the sender condition, it requires editing by the user.

Included dictionaries

The included dictionaries are as follows:

- Analysts' Email Addresses
- NASD 2711 Keywords

NASD Rule 3010 and NYSE Rule 342

Described Content condition

The Described Content conditions are as follows:

- Stock Recommendation - Compound rule that looks for both a keyword in the NASD 3010 Stock Keywords dictionary and the NASD 3010 Buy/Sell Keywords keyword dictionary. This rule requires that there is evidence of both a stock recommendation of some sort in combination with a recommendation for a specific buy or sell action.
- NASD Rule 3010 and NYSE Rule 342 Keywords - Looks for keywords in the NASD 3010 General Keywords dictionary. These keywords look for any general stock broker activity.

Included dictionaries

The included dictionaries are as follows:

- NASD 3010 Stock Keywords
- NASD 3010 Buy/Sell Keywords
- NASD 3010 General Keywords

NERC Security Guidelines for Electric Utilities

Described Content condition

A compound rule that looks for any keyword matches from the "Sensitive Keywords" dictionary and the "Vulnerability Keywords" dictionary.

Included dictionaries

The included dictionaries are as follows:

- Sensitive Keywords
- Vulnerability Keywords

Office of Foreign Assets Control (OFAC)

There are two primary conditions in the OFAC policy template. The first deals with the Specially Designated Nationals (SDN) list, and the second deals with general OFAC policy restrictions.

Described Content condition

The Described Content conditions are as follows:

- OFAC Special Designated Nationals List and Recipients - Looks for a recipient with a country code matching entries from the OFAC SDN Country Codes dictionary in combination with a match of a keyword from the SDN List dictionary.

  The SDN list refers to specific people or organizations that are subject to trade restrictions. The U.S. Treasury Department provides text files with specific names, last known addresses, and known aliases for these individuals and entities. However, the Treasury Department stipulates that the addresses may not be correct or current, and different locations do not change the restrictions against these people and organizations.
- Communications to OFAC countries - Looks for a recipient with a country code matching entries from the OFAC Country Codes dictionary.

  This condition provides guidance around the restrictions the U.S. Treasury Department has placed on general trade with specific countries. This is distinct from the SDN list, since individuals and organizations are not specified. The template looks for recipients that have the listed countries as the designated country code.

Included dictionaries

The included dictionaries are as follows:

- SDN List
- OFAC SDN Country Codes
- OFAC Country Codes

Payment Card Industry Data Security Standard

Described Content condition

Credit Card Numbers, All - Looks for a match to the Valid Credit Card pattern and a keyword from the Credit Card Number Keywords dictionary.

Structured Data conditions

Searches record resource view for a match.

Included dictionaries

Credit Card Number Keywords

Sarbanes-Oxley

Described Content condition

The Described Content conditions are as follows:

- SEC Fair Disclosure Regulation – Mirrors the rule in the SEC Fair Disclosure policy; looks for three different conditions, and all must be satisfied: any keywords in the SEC Fair Disclosure Keywords dictionary, any keywords in the Company Name Keywords dictionary and any commonly used documents in the spreadsheet or document writing file types.

  The SEC Fair Disclosure keywords indicate possible disclosure of advance financial information. The company name keywords require editing by the user. This can include any name, alternate name, or abbreviation that might indicate a reference to the company. Specifically, the file type groups detected are: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, wordperfect, and pdf.
- Financial Information – Three different conditions that must be satisfied, including: a word from the Financial Keywords dictionary, a word from the Confidential/Proprietary Words dictionary, and a spreadsheet file type. The spreadsheet file types required are excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, and 123.

Included dictionaries

The included dictionaries are as follows:

- SEC Fair Disclosure Keywords
- Company Name Keywords
- Financial Keywords
- Confidential/Proprietary Words

SEC Fair Disclosure Regulation

Described Content condition	The Described Content conditions are as follows: ■ SEC Fair Disclosure Regulation - Mirrors the rule in the SEC Fair Disclosure policy; looks for three different conditions, and all must be satisfied: any keywords in the SEC Fair Disclosure Keywords dictionary, any keywords in the Company Name Keywords dictionary and any commonly used documents in the spreadsheet or document writing file types. The SEC Fair Disclosure keywords indicate possible disclosure of advance financial information. The company name keywords require editing by the user. This can include any name, alternate name, or abbreviation that might indicate a reference to the company. Specifically, the file type groups detected are: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, wordperfect, and pdf.
Included dictionaries	The included dictionaries are as follows: ■ SEC Fair Disclosure Keywords ■ Company Name Keywords

State Data Privacy

Described Content condition

The Described Content conditions are as follows:

- US Social Security Numbers - Looks for a word from the US SSN Keywords dictionary and a hit from the Valid Social Security Number premium pattern.
- ABA Routing Numbers - Looks for a word from the ABA Routing Number Keywords dictionary and a hit from the ABA Routing Number regular expression.
- Credit Card Numbers, All - Looks for a word from Credit Card Number Keywords and the credit card number system pattern.
- CA Drivers License Numbers - Looks for a match for the CA driver's license number pattern, a match for a regular expression for terms relating to driver's license, and a keyword from the California Keywords dictionary.
- NY Drivers License Numbers - Looks for a match for the NY driver's license number pattern, a match for a regular expression for terms relating to driver's license, and a keyword from the New York Keywords dictionary.
- Letter + 12 Digits Drivers License Numbers - Looks for a match for the stated driver's license number pattern, a match for a regular expression for terms relating to driver's license, and a keyword from the Letter/12 Num. DLN State Words dictionary (namely, Florida, Minnesota, and Michigan).
- IL Drivers License Numbers - Looks for a match for the IL driver's license number pattern, a match for a regular expression for terms relating to driver's license, and a keyword from the Illinois Keywords dictionary.
- NJ Drivers License Numbers - Looks for a match for the NJ driver's license number pattern, a match for a regular expression for terms relating to driver's license, and a keyword from the New Jersey Keywords dictionary.

Exception condition:

- Email to Affiliates - An exception for email messages to affiliates who are legitimately allowed to receive information covered under the State Data Privacy regulations. The Affiliate Domains dictionary requires editing by the user.

Structured Data conditions	State Data Privacy, Consumer Data – Searches Record resource View for any three matches with the exception of First name, last name, pin and First name, last name, password. Exception condition: ■ Email to Affiliates - An exception for email messages to affiliates who are legitimately allowed to receive information covered under the State Data Privacy regulations. The Affiliate Domains dictionary requires editing by the user.
Included dictionaries	The included dictionaries are as follows: ■ US SSN Keywords ■ ABA Routing Number Keywords ■ Credit Card Number Keywords ■ California Keywords ■ New York Keywords ■ Letter/12 Num. DLN State Words ■ Illinois Keywords ■ New Jersey Keywords ■ Affiliate Domains

Confidential data-protection policy templates

This section describes the confidential data protection policy templates that Symantec Brightmail Gateway provides.

See "About content filtering policy templates" on page 336.

Table C-2 Confidential data-protection policy templates

Name	Description	Described Content	Structured Data
Confidential Documents See "Confidential Documents" on page 833.	This policy detects company-confidential documents at risk of exposure.	X	

Table C-2 Confidential data-protection policy templates *(continued)*

Name	Description	Described Content	Structured Data
Defense Message System (DMS) GENSER Classification See "Defense Message System (DMS) GENSER Classification" on page 834.	DMS General Service categories for messaging classify national security information according to access controls, which limit message distribution to authorized recipients. This policy template differs from US Intelligence Control Markings (CAPCO & DCID 1/7) in that it includes terms used to identify unclassified but sensitive information.	X	
Design Documents See "Design Documents" on page 834.	This policy detects various types of design documents, such as CAD/CAM, at risk of exposure.	X	
Encrypted Data See "Encrypted Data" on page 835.	This policy detects the use of encryption by a variety of methods including S/MIME, PGP, GPG, and file password protection.	X	
Financial Information See "Financial Information" on page 835.	This policy detects financial data and information.	X	
Mergers and Acquisitions Data See "Mergers and Acquisition Data" on page 836.	This policy detects information and communications about upcoming merger and acquisition activity. It may be modified with company-specific code words to detect specific deals.	X	
Price Information See "Price Information" on page 836.	This policy detects specific SKU and/or pricing information at risk of exposure.		X
Project Data See "Project Data" on page 836.	This policy detects discussions of sensitive projects.	X	
Publishing Documents See "Publishing Documents" on page 836.	This policy detects various types of publishing documents, such as FrameMaker files, at risk of exposure.	X	
Resumes See "Resumes" on page 837.	This policy detects active job searches.	X	X
Source Code See "Source Code" on page 837.	This policy detects various types of source code at risk of exposure.	X	

Table C-2 Confidential data-protection policy templates *(continued)*

Name	Description	Described Content	Structured Data
US Intelligence Control Markings (CAPCO & DCID 1/7) See "US Intelligence Control Markings (CAPCO) and DCID 1/7" on page 838.	This policy detects authorized terms to identify classified information in the US Federal Intelligence community as defined in the Control Markings Register, which is maintained by the Controlled Access Program Coordination Office (CAPCO) of the Community Management Staff (CMS). The register was created in response to the Director of Central Intelligence Directive (DCID) 1/7.	X	

Confidential Documents

Described Content condition

The Described Content conditions are as follows:

- Confidential Documents - A compound condition that looks for a combination of keywords from the Confidential Keywords dictionary and the following file types: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, pdf, and ppt.
- Proprietary Documents - A compound condition that looks for a combination of keywords from the Proprietary Keywords dictionary and the following file types: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, pdf, and ppt.
- Internal Use Only Documents - A compound condition that looks for a combination of keywords from the Internal Use Only Keywords dictionary and the following file types: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, pdf, and ppt.
- Documents Not For Distribution - A compound condition that looks for a combination of keywords from the Not For Distribution Words dictionary and the following file types: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, 123, doc, pdf, and ppt.

Included dictionaries

The included dictionaries are as follows:

- Confidential Keywords
- Proprietary Keywords
- Internal Use Only Keywords
- Not For Distribution Words

Defense Message System (DMS) GENSER Classification

Described Content condition

Looks for any keywords in the Secret, Top Secret, Classified or Restricted, or Other Sensitive Information dictionaries. Keywords and phrases other than those indicated in the titles of the Secret, Top Secret, and Classified or Restricted dictionaries are user-defined. The Other Sensitive Information dictionary includes phrases used to categorize sensitive but unclassified information.

Included dictionaries

The included dictionaries are as follows:

- Secret
- Top Secret
- Classified or Restricted
- Other Sensitive Information

Design Documents

Described Content condition

The Described Content conditions are as follows:

- Design Document Extensions - Looks for specified file name extensions are found in the Design Document Extensions dictionary.
- Design Documents - Looks for specified file types: cad_draw and dwg.

Included dictionaries

Design Document Extensions

Encrypted Data

Described Content condition

The Described Content conditions are as follows:

- Password Protected Files - Looks for the following file type extensions in the Password Protected Files attachment list resource: encrypted_zip, encrypted_doc, encrypted_xls, or encrypted_ppt.
- PGP Files - Looks for the following file type: pgp
- GPG Files - Looks for a keyword from the GPG Encryption Keywords dictionary.
- S/MIME - Looks for a match with the S/MIME regular expression.
- PGP8 Header Keywords – Looks for characteristic keywords in PGP8 files headers.
- PGP8 Keywords – Looks for characteristic strings in PGP8 encrypted files.
- PGP Encrypted Documents – Looks for .pgp and .aex.message or file-attachment extensions in the PGP file extension dictionary.

Included dictionaries

The included dictionaries are as follows:

- GPG Encryption Keywords
- PGP file extensions
- PGP8 Keywords

Financial Information

Described Content condition

Financial Information - Looks for the combination of specified file types, keywords from the Financial Keywords dictionary, and keywords from the Confidential/Proprietary Words dictionary. The specified file types are: excel_macro, xls, works_spread, sylk, quattro_pro, mod, csv, applix_spread, and 123.

Included dictionaries

The included dictionaries are as follows:

- Financial Keywords
- Confidential/Proprietary Words

Mergers and Acquisition Data

Described Content condition	M & A Activity - Looks for any keywords from the M & A Project Code Names dictionary, which is user-defined. Merger and acquisition activity is extremely customer and project specific; there are no general terms across all customers and users that suffice.
Included dictionaries	M & A Project Code Names

Price Information

Described Content condition	Looks for the combination of user-specified Stock Keeping Unit (SKU) numbers and the price for that SKU number in a Record resource View.

Project Data

Described Content condition	Project Activity - Looks for any keywords in the Sensitive Project Code Names dictionary, which is user-defined.
Included dictionaries	Sensitive Project Code Names

Publishing Documents

Described Content condition	The Described Content conditions are as follows: ■ Publishing Documents - Looks for the specified file types: qxpress, frame, aldus_pagemaker, and publ (Microsoft Publisher). ■ Publishing Documents, extensions - Looks for specified file name extensions in the Publishing Document Extensions dictionary.
Included dictionaries	Publishing Document Extensions

Resumes

Described Content condition

Resumes, All - Looks for files of a specified type (.doc) that are less than 50 kB and match at least one keyword from each of the following dictionaries: Job Search Keywords, Education, Job Search Keywords, Work, and Job Search Keywords, General.

Structured Data conditions

Resumes, Employee – Matches files that are less than 50 KB of a specified type (e.g., .doc) in the Resumes, Employee attachment list resource and that fulfill the following conditions: (1) matches at least one keyword from each of the following dictionaries: Job Search Keywords, Education; Job Search Keywords, Work; and Job Search Keywords, General; and (2) matches first and last names of employees in Record resource View.

Included dictionaries

The included dictionaries are as follows:

- Job Search Keywords, Education
- Job Search Keywords, Work
- Job Search Keywords, General

Source Code

Described Content condition

The Described Content conditions are as follows:

- Source Code Extensions - Looks for file name extensions from the Source Code Extensions dictionary.
- Java Source Code - Looks for the Java Import Statements or Java Class Files regular expression.
- C Source Code - Looks for the C Source Code regular expression.
- VB Source Code - Looks for the VB Source Code regular expression.
- PERL Source Code - Looks for the three different PERL-related system patterns and regular expressions.

Included dictionaries

Source Code Extensions

US Intelligence Control Markings (CAPCO) and DCID 1/7

Described Content condition

Looks for any keywords in the Secret, Top Secret, or Classified or Restricted dictionaries. Keywords and phrases other than those indicated in the dictionary titles are user-defined.

Included dictionaries

The included dictionaries are as follows:

- Secret
- Top Secret
- Classified or Restricted

Acceptable use policy templates

This section describes the acceptable use policy templates that Symantec Brightmail Gateway provides.

See "About content filtering policy templates" on page 336.

Table C-3 Acceptable use policy templates

Name	Description	Described Content	Structured Data
Competitor Communications See "Competitor Communications" on page 839.	This policy detects communications with competitors.	X	
Gambling See "Gambling" on page 839.	This policy detects any reference to gambling.	X	
Illegal Drugs See "Illegal Drugs" on page 840.	This policy detects conversations about illegal drugs and controlled substances.	X	
Media Files See "Media Files" on page 840.	This policy detects various types of video and audio files (including mp3).	X	
Offensive Language See "Offensive Language" on page 840.	This policy detects the use of offensive language.	X	
Racist Language See "Racist Language" on page 841.	This policy detects the use of racist language.	X	

Table C-3 Acceptable use policy templates *(continued)*

Name	Description	Described Content	Structured Data
Restricted Files See "Restricted Files" on page 841.	This policy detects various file types that are generally inappropriate to send out of the company such as MS Access and executable files.	X	
Restricted Recipients See "Restricted Recipients" on page 841.	This policy detects communications with any specified recipients such as former employees.	X	
Sexually Explicit Language See "Sexually Explicit Language" on page 841.	This policy detects vulgar and sexually explicit pornographic language content.	X	
Violence & Weapons See "Violence and Weapons" on page 842.	This policy detects violent language and discussions about weapons.	X	

Competitor Communications

Described Content condition	Competitor List - Looks for keywords (domains) from the Competitor Domains dictionary, which is user-defined.
Included dictionary	Competitor Domains

Gambling

Described Content condition

The Described Content conditions are as follows:

- Suspicious Gambling Keywords - Looks for five instances of keywords from the Gambling Keywords, Confirmed dictionary.
- Less Suspicious Gambling Keywords - Looks for ten instances of keywords from the Gambling Keywords, Suspect dictionary.

Included dictionary

The included dictionaries are as follows:

- Gambling Keywords, Confirmed
- Gambling Keywords, Suspect

Illegal Drugs

Described Content condition

The Described Content conditions are as follows:

- Street Drugs - Looks for five instances of keywords from the Street Drug Names dictionary.
- Mass Produced Controlled Substances - Looks for five instances of keywords from the Manufd. Controlled Substances dictionary.

Included dictionary

The included dictionaries are as follows:

- Street Drug Names
- Manufd. Controlled Substances

Media Files

Described Content condition

The Described Content conditions are as follows:

- Media Files - Looks for the following files types: qt, riff, macromedia_dir, midi, mp3, mpeg_movie, quickdraw, realaudio, wav, video_win, and vrml.
- Media Files Extensions - Looks for file name extensions from the Media Files Extensions dictionary.

Included dictionary

Media Files Extensions

Offensive Language

Described Content condition

The Described Content conditions are as follows:

- Offensive Language, Explicit - Looks for any single keyword in the Offensive Language, Explicit dictionary.
- Offensive Language, General - Looks for any three instances of keywords in the Offensive Language, General dictionary.

Included dictionary

The included dictionaries are as follows:

- Offensive Language, Explicit
- Offensive Language, General

Racist Language

Described Content condition	Racist Language - Looks for any single keyword in the Racist Language dictionary.
Included dictionary	Racist Language

Restricted Files

Described Content condition	MSAccess Files and Executables - Looks for files of the specified types: access, exe, and exe_unix.
Included dictionary	None

Restricted Recipients

Described Content condition	Restricted Recipients - Looks for messages to recipients with email addresses in the Restricted Recipients dictionary.
Included dictionary	Restricted Recipients

Sexually Explicit Language

Described Content condition	The Described Content conditions are as follows: ■ Sexually Explicit Keywords, Confirmed - Looks for any single keyword in the Sex. Explicit Keywords, Confirmed dictionary. ■ Sexually Explicit Keywords, Suspected - Looks for any three instances of keywords in the Sex. Explicit Words, Suspect dictionary. ■ Sexually Explicit Keywords, Possible - Looks for any three instances of keywords in the Sex. Explicit Words, Possible dictionary.
Included dictionary	The included dictionaries are as follows: ■ Sex. Explicit Words, Confirmed ■ Sex. Explicit Words, Suspect ■ Sex. Explicit Words, Possible

Violence and Weapons

Described Content condition: Violence & Weapons - A compound condition that looks for a keyword from the Violence Keywords dictionary and a keyword from the Weapons Keywords dictionary.

Included dictionary: The included dictionaries are as follows:

- Violence Keywords
- Weapons Keywords

Customer and employee data-protection templates

This section describes the customer and employee data-protection templates that Symantec Brightmail Gateway provides.

See "About content filtering policy templates" on page 336.

Table C-4 Customer and employee data-protection templates

Name	Description	Described Content	Structured Data
Canadian Social Insurance Numbers See "Canadian Social Insurance Number" on page 844.	This policy detects patterns indication Canadian social insurance numbers (SINs) at risk of exposure.	X	
Credit Card Numbers See "Credit Card Numbers" on page 844.	This policy detects patterns indicating credit card numbers at risk of exposure.	X	
Customer Data Protection See "Customer Data Protection" on page 844.	This policy detects customer data at risk of exposure.	X	X
Employee Data Protection See "Employee Data Protection" on page 845.	This policy detects employee data at risk of exposure.	X	X
Individual Taxpayer Identification Numbers (ITIN) See "Individual Taxpayer Identification Numbers (ITIN)" on page 846.	An Individual Taxpayer Identification Number (ITIN) is a tax processing number issued by the US Internal Revenue Service (IRS). The IRS issues ITINs to track individuals that are not eligible to obtain a Social Security Number (SSNs).	X	

Table C-4 Customer and employee data-protection templates *(continued)*

Name	Description	Described Content	Structured Data
SWIFT Codes See "SWIFT Codes" on page 846.	The Society for Worldwide Interbank Financial Telecommunications (SWIFT) is a cooperative organization under Belgian law and is owned by its member financial institutions. The SWIFT code (also known as a Bank identifier Code, or ISO 9362) has a standard format to identify a bank, location, and the branch involved. The codes are used when transferring money between banks, particularly across international borders.	X	
UK Drivers License Numbers See "UK Drivers License Numbers" on page 847.	This policy detects UK Drivers License Numbers using the official specification of the UK Government Standards of the UK cabinet Office.	X	
UK Electoral Roll Numbers See "UK Electoral Roll Numbers" on page 847.	This policy detects UK Electoral Roll Numbers using the official specification of the UK Government Standards of the UK Cabinet Office.	X	
UK National Insurance Number See "UK National Insurance Number" on page 847.	The National Insurance Number is issued to individuals by the UK Department for Work and Pensions and Inland Revenue (DWP/IR) for administering the national insurance system.	X	
UK Passport Numbers See "UK Passport Numbers" on page 847.	This policy detects valid UK passports using the official specification of the UK Government Standards of the UK Cabinet Office.	X	
UK Tax ID Numbers See "UK Tax ID Numbers" on page 848.	This policy detects UK Tax ID Numbers using the official specification of the UK Government Standards of the UK Cabinet Office.	X	
US Social Security Numbers See "US Social Security Numbers" on page 848.	This policy detects patterns indicating social security numbers at risk of exposure.	X	

Canadian Social Insurance Number

Described Content conditions	Canadian Social Insurance Numbers - Looks for a match to the Canadian Social Insurance Number regular expression and a keyword from the Canadian Social Ins. No. Words dictionary.
Included dictionaries	Canadian Social Ins. No. Words

Credit Card Numbers

Described Content conditions	Credit Card Numbers, All - Looks for a match to the credit card number system pattern and a keyword from the Credit Card Number Keywords dictionary.
Included dictionaries	Credit Card Number Keywords

Customer Data Protection

Described Content conditions	The Described Content conditions are as follows: ■ US Social Security Number Patterns - Looks for a match to the Social Security number regular expression and a keyword from the US SSN Keywords dictionary. ■ Credit Card Numbers, All - Looks for a match to the credit card number system pattern and a keyword from the Credit Card Number Keywords dictionary. ■ ABA Routing Numbers - Looks for a match to the ABA Routing number regular expression and a keyword from the ABA Routing Number Keywords dictionary.

Structured Data conditions

The Structured Data conditions are as follows:

- Username/Password Combinations – Looks for user names and passwords in combination in Record resource View.
- 3 or more customer data fields – Looks for combination of data from any 3 of the following fields in Record resource View: SSN, Phone, Email, First Name, Last Name, Bank Card number, Account Number, ABA Routing Number, Canadian Social Insurance Number, and UK National Insurance Number, Date of Birth. Exceptions: combination of phone, email, and first or last names; email or phone and first and last names.
- Exact SSN or CCN – Looks for SSN or Bank Card Number in Record resource View.
- Customer Directory – Looks for Phone or Email in Record resource view

Included dictionaries

The included dictionaries are as follows:

- US SSN Keywords
- Credit Card Number Keywords
- ABA Routing Number Keywords

Employee Data Protection

Described Content conditions

The Described Content conditions are as follows:

- US Social Security Number Patterns - Looks for a match to the Social Security number regular expression and a keyword from the US SSN Keywords dictionary.
- Credit Card Numbers, All - Looks for a match to the credit card number system pattern and a keyword from the Credit Card Number Keywords dictionary.
- ABA Routing Numbers - Looks for a match to the ABA Routing number regular expression and a keyword from the ABA Routing Number Keywords dictionary.

Structured Data conditions	The Structured Data conditions are as follows: ■ Username/Password Combinations – Looks for user names and passwords in combination in Record resource View. ■ 3 or more employee data fields – Looks for combination of data from any 3 of the following fields in Record resource View: SSN, Phone, Email, First Name, Last Name, Bank Card Number, Account Number, ABA Routing Number, Canadian Social Insurance Number, and UK National Insurance Number, employee number, medical insurance number, salary, direct deposit account, and Date of Birth. ■ Employee Directory: Looks for Phone or Email in Record resource view
Included dictionaries	The included dictionaries are as follows: ■ US SSN Keywords ■ Credit Card Number Keywords ■ ABA Routing Number Keywords

Individual Taxpayer Identification Numbers (ITIN)

Described Content conditions	US ITIN - Looks for a match to the US ITIN regular expression and a keyword from the US ITIN Keywords dictionary.
Included dictionaries	US ITIN Keywords

SWIFT Codes

Described Content conditions	SWIFT Code Regex - Looks for a match to the SWIFT code regular expression and a keyword from the SWIFT Code Keywords dictionary.
Included dictionaries	SWIFT Code Keywords

UK Drivers License Numbers

Described Content conditions

Contains a single compound condition with three parts: a single keyword from the UK Keywords dictionary, the pattern matching that of the UK driver's license regular expression, and different combinations of the phrase driver's license using a regular expression.

Included dictionaries

UK Keywords

UK Electoral Roll Numbers

Described Content conditions

Contains a single compound condition with three parts: a single keyword from the UK Keywords dictionary, a pattern matching the UK Electoral Roll Number regular expression, and a single keyword from the UK Electoral Roll Number Words dictionary.

Included dictionaries

The included dictionaries are as follows:

- UK Keywords
- UK Electoral Roll Number Words

UK National Insurance Number

Described Content conditions

UK National Insurance Numbers - Looks for a match to the UK National Insurance number regular expression and a keyword from the UK NIN Keywords dictionary.

Included dictionaries

UK NIN Keywords

UK Passport Numbers

Described Content conditions

The Described Content conditions are as follows:

- UK Passport Numbers (Old Type) - Looks for a keyword from the UK Passport Keywords dictionary and a pattern matching the regular expression for UK Passport Numbers (Old Type).
- UK Passport Numbers (New Type) - Looks for a keyword from the UK Passport Keywords dictionary and a pattern matching the regular expression for UK Passport Numbers (New Type).

Included dictionaries	UK Passport Keywords

UK Tax ID Numbers

Described Content conditions	UK Tax ID Numbers - Looks for a match to the UK Tax ID number regular expression and a keyword from the UK Tax ID Number Keywords dictionary.
Included dictionaries	UK Tax ID Number Keywords

US Social Security Numbers

Described Content conditions	US Social Security Number Patterns - Looks for a match to the social security number regular expression and a keyword from the US SSN Keywords dictionary.
Included dictionaries	US SSN Keywords

Network security policy templates

The following table provides an overview of the premium network security enforcement pre-built policy templates that Symantec Brightmail Gateway provides.

See "About content filtering policy templates" on page 336.

Table C-5 Network security policy templates

Name	Description	Described Content	Structured Data
Network Diagrams See "Network Diagrams" on page 849.	This policy detects computer network diagrams at risk of exposure.	X	
Network Security See "Network Security" on page 849.	This policy detects evidence of hacking tools and attack planning.	X	
Password Files See "Password Files" on page 849.	This policy detects password file formats, such as SAM, /etc/password, and /etc/shadow.	X	

Network Diagrams

The Described Content conditions are as follows:

- Network Diagrams with IP Addresses - Looks for a Visio file type in combination with an IP address regular expression.
- Network Diagrams with IP Address Keyword - Looks for a Visio file type in combination with phrase variations of IP address with a regular expression.

Network Security

Described Content conditions	The Described Content conditions are as follows: ■ GoToMyPC Activity - Looks for a GoToMyPC command format with a regular expression. ■ Hacker Keywords - Looks for a keyword from the Hacker Keywords dictionary. ■ KeyLoggers - Looks for a keyword from the Keylogger Keywords dictionary.
Included dictionary	The included dictionaries are as follows: ■ Hacker Keywords ■ Keylogger Keywords

Password Files

Described Content conditions	The Described Content conditions are as follows: ■ Password Filenames - Looks for the file names passwd or shadow. ■ /etc/passwd Format - Looks for a pattern with /etc/passwd format regular expression. ■ /etc/shadow Format - Looks for a pattern with /etc/shadow format regular expression. ■ SAM Passwords - Looks for a pattern with SAM format regular expression.
Included dictionary	Password Filenames

UK and international regulatory policy templates

This section describes the UK and international regulatory templates that Symantec Brightmail Gateway provides.

See “About content filtering policy templates” on page 336.

Table C-6 UK and international regulatory policy templates

Name	Description	Described Content	Structured Data
Caldicott Report See “Caldicott Report” on page 851.	The Caldicott Report (December 1997) was a review commissioned by the UK Chief Medical Officer to make recommendations to improve the way the National Health service handles and protects patient information. The Caldicott Committee was set up to review the confidentiality and flows of data throughout the NHS for purposes other than direct care, medical research or where there is a statutory requirement for information. Its recommendations are now being put into practice throughout the NHS and in the Health Protection Agency.	X	X
Data Protection Act 1998 See “Data Protection Act 1998” on page 853.	The Data Protection Act 1998 (replacement of Data Protection Act 1984) set standards which must be satisfied when obtaining, holding, using or disposing of personal data in the UK. The Data Protection Act 1998 covers anything with personal identifiable information (for example, data about personal health, employment, occupational health, finance, suppliers, and contractors).	X	X
EU Data Protection Directives See “EU Data Protection Directives” on page 854.	Directives 95/46/EC of the European Parliament deal with the protection of individuals with regard to the processing and free movement of personal data. This policy detects personal data specific to the EU directives.		X
Human Rights Act 1998 See “Human Rights Act 1998” on page 854.	The Human Rights Act 1998 allows UK citizens to assert their rights under the European Convention on Human Rights in UK courts and tribunals. The Act states that "so far as possible to do so, legislation must be read and given effect in a way which is compatible with convention rights." This policy enforces Article 8 by ensuring the private lives of British citizens stay private.	X	X

Table C-6 UK and international regulatory policy templates *(continued)*

Name	Description	Described Content	Structured Data
PIPEDA See "PIPEDA" on page 855.	Canada's Personal Information Protection and Electronic Documents Act (PIPEDA) protects personal information in the hands of private sector organizations and provides guidelines for the collection, use and disclosure of that information.	X	X

Caldicott Report

Described Content conditions

The Described Content conditions are as follows:

- UK National Insurance Number and Drug Keywords - Looks for a keyword from UK NIN Keywords dictionary in combination with a pattern matching the UK NIN regular expression and a keyword from the Prescription Drug Names dictionary.
- UK National Insurance Number and Disease Keywords - Looks for a keyword from UK NIN Keywords dictionary in combination with a pattern matching the UK NIN regular expression and a keyword from the Disease Names dictionary.
- UK National Insurance Number and Treatment Keywords - Looks for a keyword from UK NIN Keywords dictionary in combination with a pattern matching the UK NIN regular expression and a keyword from the Medical Treatment Keywords dictionary.

Structured Data conditions

The Structured Data conditions are as follows:

- Patient Data and Drug Keywords – Looks for any match to the following data in a Record Resource view: NIN (National Insurance Number), account number, last name, ID card number, email, phone, and UK NHS (National Health Service) number. This data must appear in combination with a keyword from the Prescription Drug Names dictionary.
- Patient Data and Disease Keywords – Looks for any match to the following data in a Record resource View: NIN (National Insurance Number), account number, last name, ID card number, email, phone, and UK NHS (National Health Service) number. This data must appear in combination with a keyword from the Disease Names dictionary.
- Patient Data and Treatment Keywords – Looks for any match to the following data in a Record resource View: NIN (National Insurance Number), account number, last name, ID card number, email, phone, and UK NHS (National Health Service) number. This data must appear in combination with a keyword from the Medical Treatment Keywords dictionary.

Included dictionary

The included dictionaries are as follows:

- Prescription Drug Names
- Disease Names
- Medical Treatment Keywords
- UK NIN Keywords

Data Protection Act 1998

Described Content conditions

The Described Content conditions are as follows:

- UK Electoral Roll Numbers - Looks for a single compound condition with three parts: a single keyword from the UK Keywords dictionary, the pattern matching that of the UK Electoral Roll Number regular expression, and single keyword from the UK Electoral Roll Number Words dictionary.
- UK National Insurance Numbers - Looks for a match to the UK National Insurance number regular expression and a keyword from the UK NIN Keywords dictionary.
- UK Tax ID Numbers - Looks for a match to the UK Tax ID number regular expression and a keyword from the UK Tax ID Number Keywords dictionary.
- UK Drivers License Number - Looks for a single compound condition with three parts: a single keyword from the UK Keywords dictionary, a pattern matching that of the UK driver's license regular expression, and different combinations of the phrase driver's license using a regular expression.
- UK Passport Numbers (Old Type) - Looks for a keyword from the UK Passport Keywords dictionary and a pattern matching the regular expression for UK Passport Numbers (Old Type).
- UK Passport Numbers (New Type) - Looks for a keyword from the UK Passport Keywords dictionary and a pattern matching the regular expression for UK Passport Numbers (New Type).

Structured Data conditions

UK Data Protection Act, Personal Data – Looks for three of the following columns of data in a Record resource View: NIN (National Insurance Number), account number, pin, bank card number, first name, last name, drivers license, password, tax payer ID, UK NHS number, date of birth, mother's maiden name, email address, and phone number. Combinations of first and last names with pin, password, email, phone, or mother's maiden name are excepted.

Included dictionary

The included dictionaries are as follows:

- UK Keywords
- UK Electoral Roll Number Words
- UK NIN Keywords
- UK Tax ID Number Keywords
- UK Passport Keywords

EU Data Protection Directives

The Structured Content conditions are as follows:

- EU Data Protection Directives

  Searches EU Country Codes dictionary for country codes that do not match recipient part of the message header and looks for any two of the following data columns in a record resource view: last name, bank card number, driver's license, account number, pin, medical account number, and ID card number, username, password, ABA routing number, email, phone, and mother's maiden name. Combinations of last name with email, phone, account number, and username data are excepted.
- EU Data Protection, Contact Info

  Searches EU Country Codes dictionary for country codes that do not match recipient part of the message header and looks for any two of the following data columns: last name, phone, account number, username, and email.

Human Rights Act 1998

Described Content conditions	UK Electoral Roll Numbers - Looks for a single compound condition with four parts: a single keyword from the UK Keywords dictionary, the pattern matching that of the UK Electoral Roll Number regular expression, a single keyword from the UK Electoral Roll Number Words dictionary, and a single keyword from the UK Personal Data Keywords dictionary.
Structured Data conditions	UK Data Protection Act, Personal Data – A compound rule that looks for last name and electoral roll number in a Record resource View in combination with a keyword from the UK Personal Data Keywords dictionary.
Included dictionary	The included dictionaries are as follows: ■ UK Personal Data Keywords ■ UK Keywords ■ UK Electoral Roll Number Words

PIPEDA

Described Content conditions

The Described Content conditions are as follows:

- Canadian Social Insurance Numbers - Looks for a match to the Canadian Social Insurance Number regular expression and a keyword from the dictionary Canadian Social Ins. No. Words.
- ABA Routing Numbers - Looks for a word from the ABA Routing Number Keywords dictionary and a match from the ABA Routing Number regular expression.
- Credit Card Numbers, All - Looks for a word from Credit Card Number Keywords dictionary and a match from the credit card number system pattern.

UK Electoral Roll Numbers - Looks for a single compound condition with four parts: a single keyword from the UK Keywords dictionary, the pattern matching that of the UK Electoral Roll Number regular expression, a single keyword from the UK Electoral Roll Number Words dictionary, and a single keyword from the UK Personal Data Keywords dictionary.

Structured Data conditions

The Structured Data conditions are as follows:

- PIPEDA – Looks in a Record resource View for any two of the following data columns: last name, bank card, medical account number, medical record, agency number, account number, PIN, username, password, SIN, ABA routing number, email, phone, mother's maiden name. Combinations of last names with email, phone, account number, or username are excepted.
- PIPEDA, Contact Info – Looks for in Record resource View any two of the following data columns: last name, phone, account number, username, email

Included dictionary

The included dictionaries are as follows:

- Canadian Social Ins. No. Words
- ABA Routing Number Keywords
- Credit Card Number Keywords

Appendix D

Symantec Brightmail Gateway support for VMware Tools

This appendix includes the following topics:

- Symantec Brightmail Gateway support for VMware tools

Symantec Brightmail Gateway support for VMware tools

Symantec Brightmail Gateway virtual appliances provide support for a limited set of VMware tools.

Only the following tools are supported:

- The second generation vmxnet Virtual NIC driver: This tool loads automatically at virtual appliance boot time. No action is required to activate this support.
- The vmtoolsd daemon: This tool starts automatically during virtual appliance boot time. No action is required to activate this support. The vmtoolsd daemon supports automatic power down of the virtual appliance from the vSphere4 Client dashboard. The vmtoolsd daemon also supports the Guest Information Service.

No other VMware tools functionality is supported.

Glossary

Adware

Adware is a type of program that displays an advertisement of some sort. Adware is usually related to a specific Web site that is cached in the Web browser. In some cases, it changes the home page of your Web browser to point to a specific Web site. Because adware is not malicious in nature, it is not considered a virus.

Agent

A component that facilitates communicating configuration information between the Control Center and each Scanner.

annotation

A phrase or paragraph placed at the beginning or end of the body of an email message. Up to 1000 distinct annotations are allowed for use in specific categories of messages for specific policy groups of recipients. You can use this feature to automate email disclaimers.

archive

An action that can be performed on email messages which consists of forwarding the messages to a specific SMTP address.

attachment list

A list of attachment types for use in filtering. You can create attachment lists based on file naming (for example, based on the file extension), or on the true type of each file, or you can use a pre-filled list.

Audit ID

A unique identifier included as a message header in all processed messages.

authentication

The process of determining the identity of a user attempting to access a network. Authentication occurs through challenge/response, time-based code sequences, or other techniques. Authentication typically involves the use of a password, certificate, PIN, or other information that can be used to validate identity over a computer network.

bad sender

A sender from whom you do not want to accept email messages. A bad sender is a member of at least one of the following groups: Local Bad Sender Domains, Local Bad Sender IPs, Third Party Bad Senders, or Symantec Global Bad Senders.

bounce

An action that can be performed on an email message by an email server, which consists of returning the message to its `From:` address with a custom response.

Bounce Attack Prevention

A feature of Symantec Brightmail Gateway that eliminates bounced messages that are a result of redirection, while still allowing legitimate bounce message notification.

Brightmail Engine

The component of Symantec Brightmail Gateway that scans email and attachments, instant messages, and file transfers for viruses, spam, and content filtering according to polices that you configure.

broadcast address A common address that is used to direct (broadcast) a message to all systems on a network. The broadcast address is based upon the network address and the subnet mask.

CA (Certificate Authority) A trusted third-party organization or company that issues digital certificates used to create digital signatures and public-private key pairs. The role of the CA in this process is to guarantee that the entity granting the unique certificate is, in fact, who it claims to be. This means that the CA usually has an arrangement with the requesting entity to confirm a claimed identity. CAs are a critical component in data security and electronic commerce because they guarantee that the two parties exchanging information are really who they claim to be.

certificate A file that is used by cryptographic systems as proof of identity. It contains a user's name and public key.

Certificate Authority-signed SSL A type of Secure Sockets Layer (SSL) that provides authentication and data encryption through a certificate that is digitally signed by a Certificate Authority.

CIDR (classless interdomain routing) A way of specifying a range of addresses using an arbitrary number of bits. For instance, a CIDR specification of 206.13.1.48/25 would include any address in which the first 25 bits of the address matched the first 25 bits of 206.13.1.48.

clean An action that consists of deleting unrepairable virus infections and repairing repairable virus infections.

Conduit A component that retrieves new and updated filters from Symantec Security Response through secure HTTPS file transfer. Once retrieved, the Conduit authenticates filters, and then alerts the Brightmail Engine that new filters are to be received and implemented. Finally, the Conduit manages statistics for use by Symantec Security Response and for generating reports.

content filtering A set of features that enable administrators to enforce corporate email policies, reduce legal liability, and ensure compliance with regulatory requirements.

Control Center A Web-based configuration and administration center. Each site has one Control Center. The Control Center also houses Spam Quarantine and supporting software. You can configure and monitor all of your Scanners from the Control Center.

DDS (directory data service) A Symantec Brightmail Gateway service that lets you use the information that is stored in your Lightweight Directory Access Protocol (LDAP) directories for features throughout Symantec Brightmail Gateway.

defer An action that an MTA receiving an email message can take, which consists of using a 4xx SMTP response code to tell the sending MTA to try again later.

DMZ (de-militarized zone) A network added between a protected network and an external network to provide an additional layer of security. Sometimes called a perimeter network.

dictionary A list of words and phrases against which email messages can be checked for non-compliant content. Symantec Brightmail Gateway allows you to create content filtering policies that screen email against a specific dictionary. You can use the provided dictionaries, add terms to the provided dictionaries, or add additional dictionaries.

directory data source A set of configuration data that includes host connection parameters and the set of functions enabled for that source.

directory harvest attack A tactic that spammers use to determine valid email addresses. A directory harvest attack occurs when a spammer sends a large quantity of possible email addresses to a site. An unprotected mail server rejects messages that are sent to invalid addresses. This behavior lets spammers know which email addresses are valid by checking the rejected messages against the original list.

DNS (Domain Name Server) proxy An intermediary between a workstation user and the Internet that allows the enterprise to ensure security and administrative control.

DNS (Domain Name System) A hierarchical system of host naming that groups TCP/IP hosts into categories. For example, in the Internet naming scheme, names with .com extensions identify hosts in commercial businesses.

DNS server A repository of addressing information for specific Internet hosts. Name servers use the Domain Name System (DNS) to map IP addresses to Internet hosts.

downstream At a later point in the flow of email. A downstream email server is an email server that receives messages at a later point in time than other servers. In a multiple-server system, inbound mail travels a path from upstream mail servers to downstream mail servers. Downstream can also refer to other types of networking paths or technologies.

email virus attack A series of virus-infected emails from a specific domain. Symantec Brightmail Gateway allows you to choose an action to perform on these messages.

encrypted attachment A message attachment that has been converted into a form that is not easily understood by unauthorized persons. Symantec Brightmail Gateway does not scan encrypted attachments, but allows you to choose an action to take when an encrypted attachment is detected.

Expunger A component of Spam Quarantine and content filtering quarantine incidents, which resides on the Control Center computer in Symantec Brightmail Gateway. The Expunger can be configured to periodically remove older or unwanted messages or incidents from the Spam Quarantine or content filtering Quarantine incident folders.

false positive A piece of legitimate email that is mistaken for spam and classified as spam by Symantec Brightmail Gateway.

Fastpass A feature of Symantec Brightmail Gateway that lets most email messages that are from verified good senders bypass spam filtering. Fastpass conserves resources by providing a temporary exemption from spam scanning for senders with a demonstrated history of sending no spam messages. Fastpass reduces the processing time required for messages from legitimate sources.

filter A method for analyzing email messages, used to determine what action to take on each message. Symantec Brightmail Gateway uses a variety of types of filters to process messages. A filter can be provided by Symantec, created by a local administrator, created by an end user, or provided by a third party.

filter policy In Symantec Brightmail Gateway, a set of actions that apply to a category of messages. The actions specified in a filter policy are only applied to users who are members of a policy group that includes the filter policy. There are three types of filter policies: spam, virus, and content filtering policies. Filter policies can also make use of policy resources. See also policy group, policy resources.

firewall A program that protects the resources of one network from users from other networks. Typically, an enterprise with an intranet that allows its workers access to the wider Internet will want a firewall to prevent outsiders from accessing its own private data resources.

gateway A network point that acts as an entrance to another network. A gateway can also be any computer or service that passes packets from one network to another network during their trip across the Internet.

good sender A sender from whom you want to accept email messages. A good sender is a member of at least one of the following groups: Local Good Sender Domains, Local Good Sender IPs, Third Party Good Senders, or Symantec Global Good Senders.

heuristic Filters that pro-actively target patterns common in spam and viruses.

host 1. In a network environment, a computer that provides data and services to other computers. Services might include peripheral devices, such as printers, data storage, email, or Web access. 2. In a remote control environment, a computer to which remote users connect to access or exchange data.

IP (Internet Protocol) The method or protocol by which data is sent from one computer to another on the Internet. Each computer (known as a host) on the Internet has at least one address that uniquely identifies it to all other computers on the Internet.

IP address A unique number that identifies a workstation on a TCP/IP network and specifies routing information. Each workstation on a network must be assigned a unique IP address, which consists of the network ID, plus a unique host ID assigned by the network administrator. This address is usually represented in dot-decimal notation, with the decimal values separated by a period (for example, 123.45.6.24).

LDAP (Lightweight Directory Access Protocol) A software protocol that enables anyone to locate organizations, individuals, and other resources such as files and devices in a network, whether on the Internet or on a corporate intranet. LDAP is a lightweight (smaller amount of code) version of Directory Access Protocol (DAP), which is part of X.500, a standard for directory services in a network.

LDAP group
A group stored in your LDAP directory server that references users or other groups through LDAP distinguished name (DN) on a multivalued attribute.

LDIF (LDAP Data Interchange Format)
An Internet Engineering Task Force (IETF) standard format for representing directory information in a flat file, specified in RFC 2849.

Local Bad Sender Domains
Domains of senders from whom you do not want to accept messages. Specify Local Bad Sender Domains in the Control Center.

Local Bad Sender IPs
IP addresses of senders from whom you do not want to accept messages. Specify Local Bad Sender IPs in the Control Center.

Local Good Sender Domains
Domain addresses of senders that are permitted by default and bypass spam filtering. Specify Local Good Sender Domains in the Control Center.

Local Good Sender IPs
IP addresses of senders that are permitted by default and bypass spam filtering. Specify Local Good Sender IPs in the Control Center.

malware
Programs and files that are created to do harm. Malware includes computer viruses, worms, and Trojan horses.

MIME (Multipurpose Internet Mail Extensions)
A protocol used for transmitting documents with different formats via the Internet.

MTA (Mail Transfer Agent)
A generic term for programs such as Sendmail, postfix, or qmail that send and receive mail between servers using SMTP. The MTA in each Symantec Brightmail Gateway Scanner routes the inbound messages and outbound messages to the Brightmail Engine for processing. Then the MTA delivers filtered messages to their internal destinations or to remote destinations.

Each Scanner MTA has the following queues for the temporary storage of email:

Delivery queue
The queue that temporarily holds inbound and outbound messages that have already been filtered before the MTA delivers the messages to their required destinations.

Inbound queue
The queue that temporarily holds inbound messages before the MTA forwards them to the Brightmail Engine for processing.

Outbound queue
The queue that temporarily holds outbound messages before the MTA forwards them to the Brightmail Engine for processing.

Administrator actions that affect the outbound queue also affect SMTP authentication.

name server

See DNS (Domain Name Server) proxy.

notification

1. In Symantec Brightmail Gateway, a separate email that can be automatically sent to the sender, recipients, or other email addresses when a specified condition is met. For example, if you have a policy that strips .exe attachments from incoming messages, you may want to also notify the sender that the attachment has been stripped. 2. In Symantec Brightmail Gateway, a periodic email summary sent by Spam Quarantine to users, listing the newly quarantined spam messages, and including links for users to immediately release messages to their inbox or to log in to their personal quarantines. See also Notifier.

Notifier

A component of Spam Quarantine, which resides on the Control Center in Symantec Brightmail Gateway. Notifier sends periodic email messages to users, providing a digest of their spam. The Notifier message (notification) is customizable; it can contain a list of the subject lines and senders of all spam messages.

packet

A unit of data that is formed when a protocol breaks down messages that are sent along the Internet or other networks. Messages are broken down into standard-sized packets to avoid overloading lines of transmission with large chunks of data. Each of these packets is separately numbered and includes the Internet address of the destination. Upon arrival at the recipient computer, the protocol recombines the packets into the original message.

phishing

An attempt to illegally gather personal and financial information by sending a message that appears to be from a well known and trusted company. A phishing message typically includes at least one link to a fake Web site, designed to mimic the site of a legitimate business and entice the recipient to provide information that can be used for identity theft or online financial theft.

policy

A set of message filtering instructions that Symantec Brightmail Gateway implements on a message or set of messages. See also filter policy, policy group.

policy group

In Symantec Brightmail Gateway, a group of users to which you can apply a unique set of filter policies. Users can be specified by email address or domain. See also filter policy.

policy resources

In Symantec Brightmail Gateway, sets of data that enable customization of email filtering and the actions taken on filtered email. You can employ policy resources when you create filter policies. Policy resources include annotations, archive, attachment lists, dictionaries, and notifications. See also filter policy, annotation, archive, attachment list, dictionary, and notification (definition 1).

port

1. A hardware location used for passing data into and out of a computing device. Personal computers have various types of ports, including internal ports for connecting disk drives, monitors, and keyboards, and external ports, for connecting modems, printers, mouse devices, and other peripheral devices. 2. In TCP/IP and UDP networks, the name given to an endpoint of a logical connection. Port numbers

identify types of ports. For example, both TCP and UDP use port 80 for transporting HTTP data.

probe account An invalid email address used to attract spam. Probe accounts are created and added to the Symantec Probe Network by using tools in the Control Center.

Probe Network A network of email accounts provided by Symantec's Probe Network Partners. Used by Symantec Security Response for the detection of spam, the Probe Network has a statistical reach of over 300 million email addresses and includes over 2 million probe accounts.

Probe Network Partners ISPs or corporations that participate in the Probe Network.

protocol A set of rules for encoding and decoding data so that messages can be exchanged between computers and so that each computer can fully understand the meaning of the messages. On the Internet, the exchange of information between different computers is made possible by the suite of protocols known as TCP/IP. Protocols can be stacked, meaning that one transmission can use two or more protocols. For example, an FTP session uses the FTP protocol to transfer files, the TCP protocol to manage connections, and the IP protocol to deliver data.

proxy server A server that acts on behalf of one or more other servers, usually for screening, firewall, or caching purposes, or a combination of these purposes. Also called a gateway. Typically, a proxy server is used within a company or enterprise to gather all Internet requests, forward them out to Internet servers, and then receive the responses and in turn forward them to the original requester within the company.

reject An action that an MTA receiving an email message can take, which consists of using a 5xx SMTP response code to tell the sending MTA that the message is not accepted.

release In Symantec Brightmail Gateway, an action that end users or administrators can take on messages in Spam Quarantine. Releasing removes the message from Spam Quarantine and returns the message to the end user's inbox. See also Spam Quarantine.

Scanner The component in Symantec Brightmail Gateway that filters mail. Each site can have one or many Scanners. The configuration of each Scanner is managed through the Control Center.

sender group A category of email senders that Symantec Brightmail Gateway manages using the Brightmail Adaptive Reputation Management (Brightmail ARM) feature. Sender groups can be based upon IP addresses, domains, third party lists, or Symantec lists. You can configure the Brightmail ARM to take a variety of actions on messages from each sender group.

Sender ID A set of standard practices for authenticating email. If the sender's domain owner participates in Sender ID, the recipient MTA can check for forged return addresses.

Symantec Brightmail Gateway allows you to specify an action for messages that fail Sender ID authentication.

signature 1. A state or pattern of activity that indicates a violation of policy, a vulnerable state, or an activity that may relate to an intrusion. 2. Logic in a product that detects a violation of policy, a vulnerable state, or an activity that may relate to an intrusion. This can also be referred to as a signature definition, an expression, a rule, a trigger, or signature logic. 3. Information about a signature including attributes and descriptive text. This is more precisely referred to as signature data.

SMTP (Simple Mail Transfer Protocol) The protocol that allows email messages to be exchanged between mail servers. Then, clients retrieve email, typically through the POP or IMAP protocol.

spam 1. Unsolicited commercial bulk email. 2. An email message identified as spam by Symantec Brightmail Gateway, using its filters.

Spam Quarantine A file directory that stores email messages separately from the normal message flow, and allows access to those messages. In Symantec Brightmail Gateway, Spam Quarantine is located on the Control Center computer and provides users with Web access to their spam messages. Users can browse, search, and delete their spam messages and can also redeliver misidentified messages to their inbox. An administrator account provides access to all quarantined messages. Spam Quarantine can also be configured for administrator-only access.

spam scoring The process of grading messages when filtering email for spam. Symantec Brightmail Gateway assigns a spam score to each message that expresses the likelihood that the message is actually spam. See also suspected spam.

SPF (Sender Policy Framework) A set of standard practices for authenticating email. If the sender's domain owner participates in SPF, the recipient MTA can check for forged return addresses. Symantec Brightmail Gateway allows you to specify an action for messages that fail SPF authentication.

spyware Stand-alone programs that can secretly monitor system activity and detect passwords and other confidential information and relay the information back to another computer.

SSH (Secure Shell) A program that allows a user to log on to another computer securely over a network by using encryption. SSH prevents third parties from intercepting or otherwise gaining access to information sent over the network.

SSL (Secure Sockets Layer) A protocol that allows mutual authentication between a client and server and the establishment of an authenticated and encrypted connection, thus ensuring the secure transmission of information over the Internet. See also TLS.

subnet mask Used to subdivide an assigned network address into additional subnetworks by using some of the unassigned bits to designate local network addresses. Subnet masking facilitates routing by identifying the network of the local host. The subnet mask is a required configuration parameter for an IP host.

A local bit mask (set of flags) that specifies which bits of the IP address specify a particular IP network or a host within a subnetwork. Used to “mask” a portion of an IP address so that TCP/IP can determine whether any given IP address is on a local or remote network. Each computer configured with TCP/IP must have a subnet mask defined.

suspected spam

A message that Symantec Brightmail Gateway deems could potentially be spam based on scores derived from pattern matching and heuristic analysis. Through policies, you can specify different actions for the messages that are identified as suspected spam. See also spam.

Suspect Virus Quarantine

A file directory that temporarily holds messages suspected of containing viruses. Messages with suspicious attachments can be held in Suspect Virus Quarantine for a number of hours, then filtered again, with updated filters, if available. This processing delay capability enables Symantec Brightmail Gateway to more effectively deal with new virus threats as they emerge.

suspicious attachment

A message attachment that Symantec Brightmail Gateway has determined may contain a virus. You can choose what action to take when a suspicious attachment is detected.

Symantec Global Bad Senders

A list of IP addresses collected by Symantec, based on global spam data from mail servers protected by Symantec. One of the sender groups in Symantec Brightmail Gateway.

Symantec Global Good Senders

A list of IP addresses collected by Symantec, based on global legitimate sender data from mail servers protected by Symantec. One of the sender groups in Symantec Brightmail Gateway.

Symantec Security Response

Symantec Security Response is a team of dedicated intrusion experts, security engineers, virus hunters, threat analysts, and global technical support teams that work in tandem to provide extensive coverage for enterprise businesses and consumers. Symantec Security Response also leverages sophisticated threat and early warning systems to provide customers with comprehensive, global, 24x7 Internet security expertise to proactively guard against today’s blended Internet threats and complex security risks.

· Security Response covers the full range of security issues to provide complete protection for customers including the following areas:

· Viruses, worms, Trojan horses, bots and other malicious code

· Hackers

· Vulnerabilities

· Spyware, adware, and dialer programs

· Spam

· Phishing and other forms of Internet fraud

Security Response keeps Symantec and its customers ahead of attackers by forecasting the next generation of threats using its worldwide intelligence network and unmatched insight. The team delivers the bi-annual Internet Security Threat Report that identifies critical trends & statistics for the entire security community, placing Symantec at the forefront of the rapidly shifting landscape.

With the steadily increasing sophistication of today's threats, a holistic approach to defending your digital assets is the key to repelling attackers. With a unified team covering the full range of security issues, Symantec Security Response helps provide its customers with fully integrated protection as it combines the collective expertise of hundreds of security specialists to bring updates and security intelligence to the full range of Symantec's products and services. Symantec has research and response centers located around the world.

Third Party Bad Senders A sender group in Symantec Brightmail Gateway that allows administrators to add multiple lists of bad senders compiled by third-party services.

Third Party Good Senders A sender group in Symantec Brightmail Gateway that allows administrators to add multiple lists of good senders compiled by third-party services.

threat A circumstance, event, or person with the potential to cause harm to a system in the form of destruction, disclosure, modification of data, or denial of service.

TLS (Transport Layer Security) A protocol that provides communications privacy over the Internet by using symmetric cryptography with connection-specific keys and message integrity checks. TLS provides some improvements over SSL in security, reliability, interoperability, and extensibility. See also SSL.

Transformation Engine A component of a Symantec Brightmail Gateway Scanner that performs actions on messages.

true file type recognition A technology that identifies the actual type of a file, whether or not the file extension matches that type. In Symantec Brightmail Gateway, you can specify filtering actions based on the true file type or true file class of a file, or you can filter based on the file name or extension.

unscannable In Symantec Brightmail Gateway, a message can be unscannable for viruses for a variety of reasons. For example, if it exceeds the maximum file size or maximum scan depth configured on the Scanning Settings page, or if it contains malformed MIME attachments, it may be unscannable. Compound messages such as zip files that contain many levels may exceed the maximum scan depth. You can configure how unscannable messages are processed.

virus A piece of programming code inserted into other programming to cause some unexpected and, for the victim, usually undesirable event. Viruses can be transmitted by downloading programming from other sites or present on a diskette. The source of the file you are downloading or of a diskette you have received is often unaware of the virus. The virus lies dormant until circumstances cause the computer to execute its code. Some viruses are playful in intent and effect, but some can be harmful, erasing data or causing your hard disk to require reformatting.

Symantec Network Prevent

A component of Symantec Data Loss Prevention which discovers, monitors, and protects confidential data wherever it is stored or used. Symantec Brightmail Gateway integrates with Symantec Network Prevent to deliver, route, hold, or block email traffic.

worm

A special type of virus. A worm does not attach itself to other programs like a traditional virus, but creates copies of itself, which create even more copies.

Index