SGOS Administration (Admin Console Edition)

ProxySG Admin Console version 1.1.3.1
SGOS 6.7.4 and later

Symantec
A Division of Broadcom

Legal Notice

Broadcom, the pulse logo, Connecting everything, and Symantec are among the trademarks of Broadcom. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Copyright © 2020 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Wednesday, November 11, 2020

Table of Contents

- "About the ProxySG Admin Console " on the next page
  - "Navigate the Admin Console" on page 6
  - "Documentation Conventions" on page 16
  - "SGOS Documentation" on page 21
- "Configure the ProxySG Appliance" on page 23
  - "Manage ProxySG Policy" on page 24
  - "Configure Networking" on page 58
  - "Manage ProxySG Services" on page 149
  - "Configure SSL and Device Authentication" on page 203
  - "Authenticate Administrators and Users" on page 230
- "Monitor the Appliance" on page 312
  - "Advanced Statistics" on page 313
- "View System Reports " on page 314
  - "View Active Sessions (Per-Session Connection Statistics)" on page 315
- "Administer the Appliance" on page 316
  - "Configure General ProxySG Settings" on page 317
  - "Configure Data Services" on page 323
  - "Manage System Images" on page 370
  - "Manage Licensing" on page 373
  - "Configure ProxySG Access Logs" on page 383
  - "Configure Malicious Content Scanning Services" on page 416
  - "Verify Service Health and Status" on page 446
  - "Service Information" on page 467

About the ProxySG Admin Console

The Blue Coat ProxySG Admin Console is the next-generation web interface for the ProxySG appliance, designed to help you manage and monitor the appliance more efficiently. The Admin Console is the successor to the Java-based Management Console and can be accessed through the latest browsers.

To access the Admin Console, you require Symantec Management Center. See "Compatibility Information" below for supported software versions to access the Admin Console.

> **Note:** This early release of the Admin Console is a preview of the new design and functionality, and as such does not include all of the areas of the Management Console. Some features can't be configured in the Admin Console; you must use the Management Console or the CLI to complete some tasks.
>
> To access the Management Console, use the following URL:
> https://*Proxy_SG_IP_address*:8082
>
> Refer to the *SGOS Administration Guide*, *Web Visual Policy Manager Reference*, and the *Command Line Interface Reference* as needed. See "SGOS Documentation" on page 21.

Compatibility Information

The Admin Console supports the following product and browser versions:

Supported Versions

- SGOS 6.7.4 and later
- Advanced Secure Gateway 6.7.4 and later

  > **Tip:** Use the Admin Console to manage ProxySG configuration only; for Content Analysis configuration, use the Advanced Secure Gateway management console.

- Management Center version 2.4.x and later

  > **Tip:** Users require appropriate read/write permissions to view and work in the Admin Console via Management Center. For information on role-based administration in Management Center, and to download the Admin Console installation package, refer to Management Center documentation.

Supported Browsers

You can use the Admin Console in the latest versions of the following web browsers:

- Apple Safari
- Google Chrome
- Microsoft Edge
- Mozilla Firefox

Getting Started

If this is your first time using the ProxySG appliance or the Admin Console, Symantec recommends following these steps:

1. Refer to the *SGOS Release Notes* for your version to learn about features, known issues, and bug fixes.
2. (If upgrading) Refer to the *SGOS Upgrade/Downgrade WebGuide* to learn about important upgrade information, behavior changes, and removals.
3. Create a myBroadcom account. You need an account to register your appliance and retrieve the license key.

   Refer to Knowledge Base article 145885: https://knowledge.broadcom.com/external/article?articleId=145885
4. (Recommended) Review the following topics in this help system:
   - "Navigate the Admin Console" on the next page
   - "Documentation Conventions" on page 16
     - " Abbreviations" on page 18
   - "SGOS Documentation" on page 21

Navigate the Admin Console

Refer to the following topics to learn about navigating the Admin Console, including this help system.

Functional Area

The Admin Console divides functions into four areas, represented by icons in the navigation pane:

- **Dashboard**: Allows you to quickly view important device data.
- **Reports**: Look at statistics about the appliance
- **Configuration**: Configure all aspects of the appliance, such as network, services, policy, and more.
- **Administration**: Manage the appliance license, upgrades, data services, and diagnostics.

To collapse or expand the navigation menu, select the **< >** at the bottom of the pane.

See the "Example" on page 9 below.

Default View and Compact View

Areas in the Admin Console are displayed by default as follows:

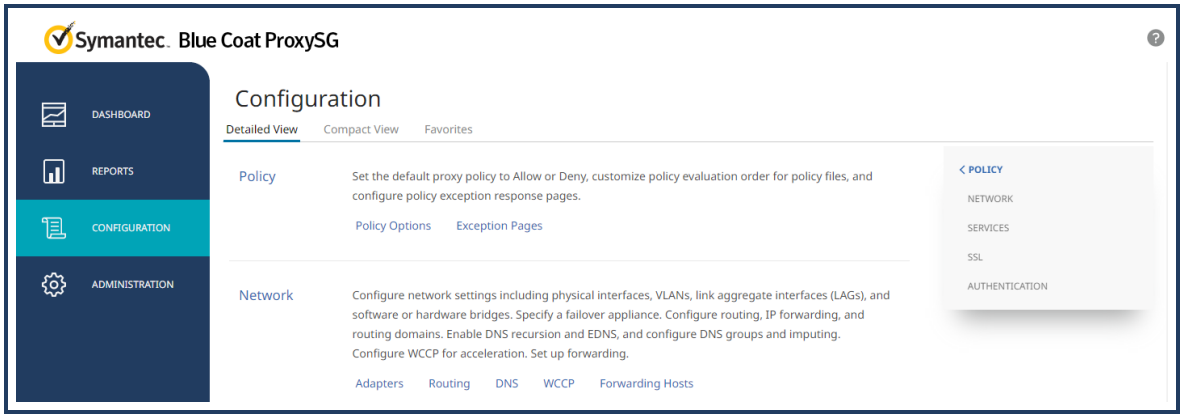

To display links in compact groups without descriptions, select **Compact View**:

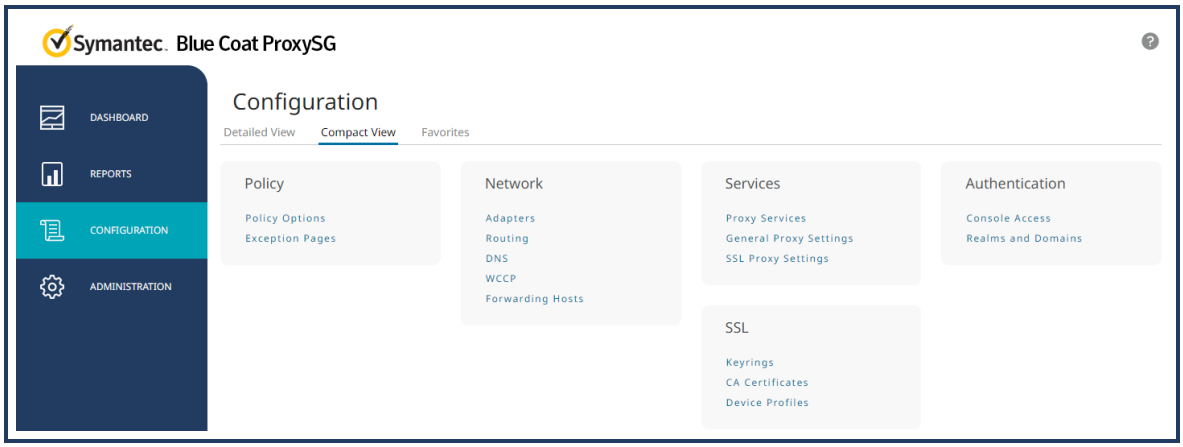

Saving Favorites

In both Detailed View and Compact View, you can save work areas that you often use for quick access later. Hover over an option and select the star that appears, as shown in Compact View below:

You can also save frequently-used work areas at the top of each page:

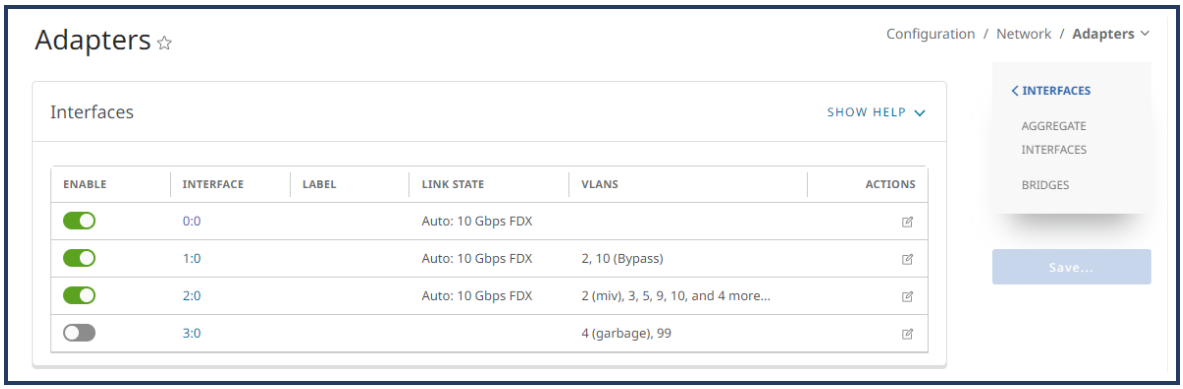

After you save a favorite, it appears under Favorites:

To remove a favorite, select the star beside its name.

Breadcrumbs

Pages in the Admin Console are organized hierarchically. Use the breadcrumb links at the top right of each page to identify your location in a process. See the "Example" on the facing page below.

Navigation

You can navigate pages within the Admin Console using several methods:

- Move down to more specific levels: Select the link on a page, depending on context. In the following example, you can select the **Policy Options** link.

  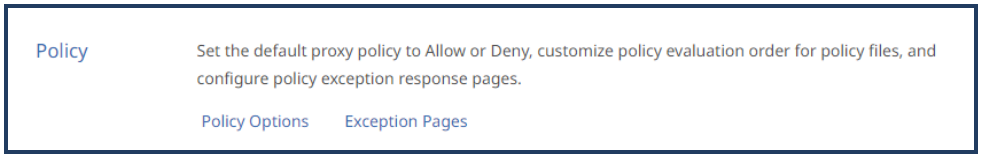

  The following shows the links in Compact View:

  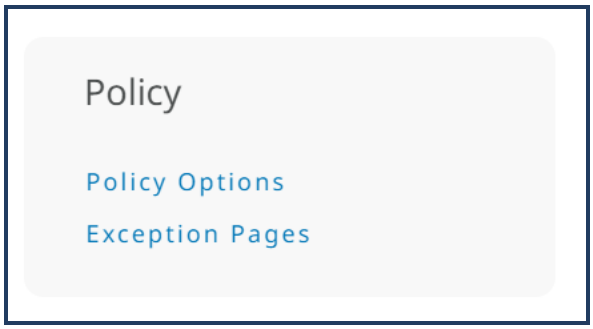

- Move up to more general levels: Select a link in the breadcrumbs on the top right. In the following example, you can

select the **Policy** link in breadcrumbs to go up one level. Select **Configuration** to go up two levels.

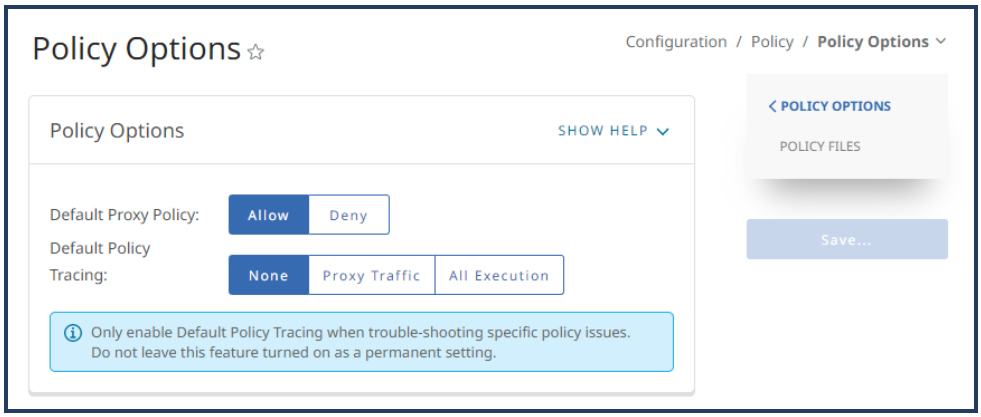

- (If applicable) Move to another page at the same level as the current page: If the breadcrumbs shows the current page with an arrow, select it to see other pages. In the preceding example, select **Policy Options** to go to **Exception Pages**, which is also under **Configuration > Policy**.

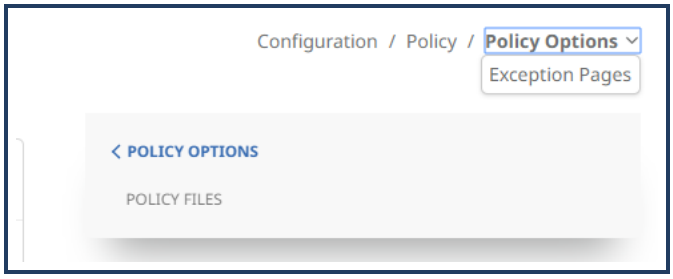

Example

In the following example:

- The navigation menu is expanded.
- In the menu, **Configuration** is selected.
- Breadcrumbs show the navigation path **Configuration > General > Time**.
- You can select **Time** in the breadcrumbs to go to another page.

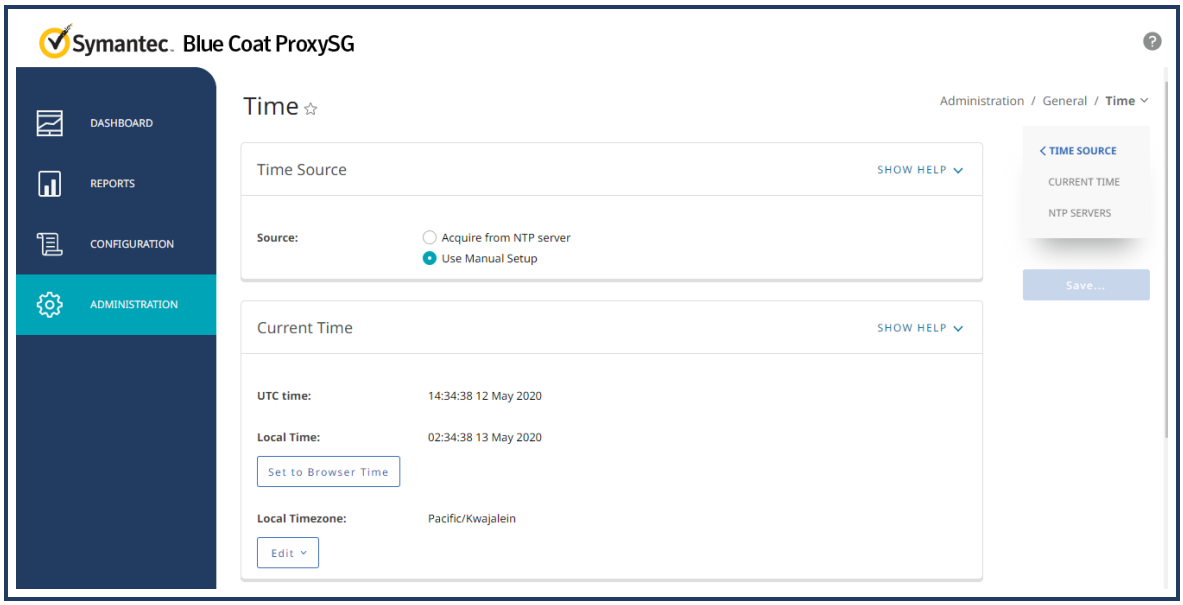

Sorting and Filtering

Some columns in the Admin Console can be sorted; if you hover over a column name, it appears underlined to indicate it can be sorted.

Select a column name to sort the rows in ascending (*null*, 0-9, A-Z) or descending order (Z-A, 9-0, *null*). An arrow appears in the column header to indicate the sort order.

The following example shows the **Interface** column sorted in ascending order.

Column headers with a filter icon indicate that the column fields can be filtered. Select the filter icon and enter the filter text in the dialog.

The filter icon turns blue to indicate that a filter is active and not all of the data is visible. Clear the filter by selecting the filter icon and deleting the filter text.

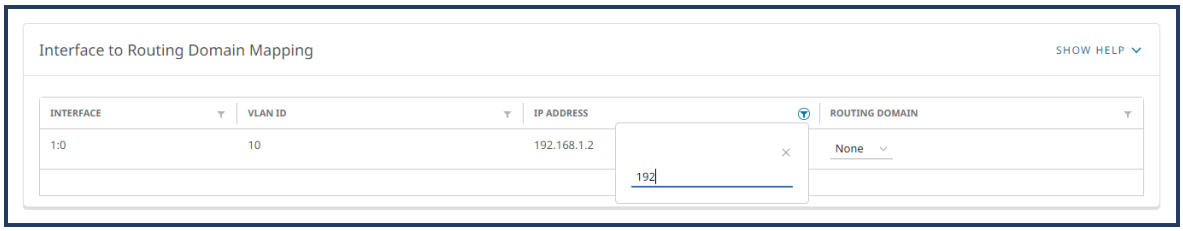

Changing Settings

The **Actions** column in some sections of the Admin Console displays options for editing and deleting settings or entries in a section. Select the appropriate icon to edit or delete an entry. Sometimes, system-defined or default entries can't be deleted; for these, only edit icons are available.

The following section shows that the default routing domain can be edited, but not deleted. The user-created domain1 can be edited or deleted.

Some defaults can't be edited or deleted, such as system-defined service groups. For these items, no icons are available in the **Actions** column.

Saving Changes

The Admin Console indicates if there are pending changes while you work:

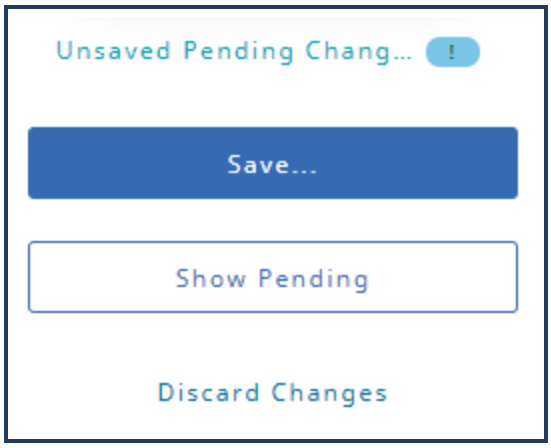

Select **Save** to save changes. Alternatively, you can select **Show Pending** to display pending CLI changes; see "Save Changes Confirmation" below for an example.

Otherwise, select **Discard Changes** to clear your changes. The settings revert to the previous configuration.

If there are no changes to save, the console indicates that you have unsaved pending changes. In addition, the **Save** option is unavailable (grayed out) and there is no **Discard** option.

If you navigate away from a page or dialog while changes are in progress, the console warns you to save changes or discard them:

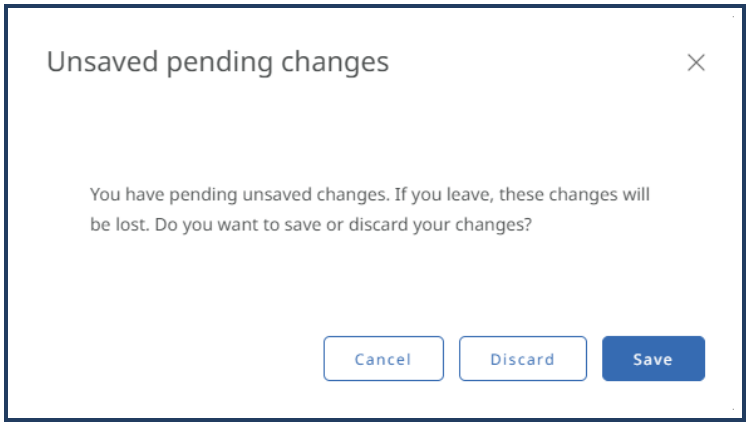

Save Changes Confirmation

When you save changes, the Admin Console prompts you to confirm the action. The Save Changes dialog displays the configuration changes you are about to make, including the underlying CLI commands that are executed.

In the following example, the CLI shows that saving changes will enable policy tracing for all proxy transactions. Selecting **Copy** copies the CLI commands to the clipboard.

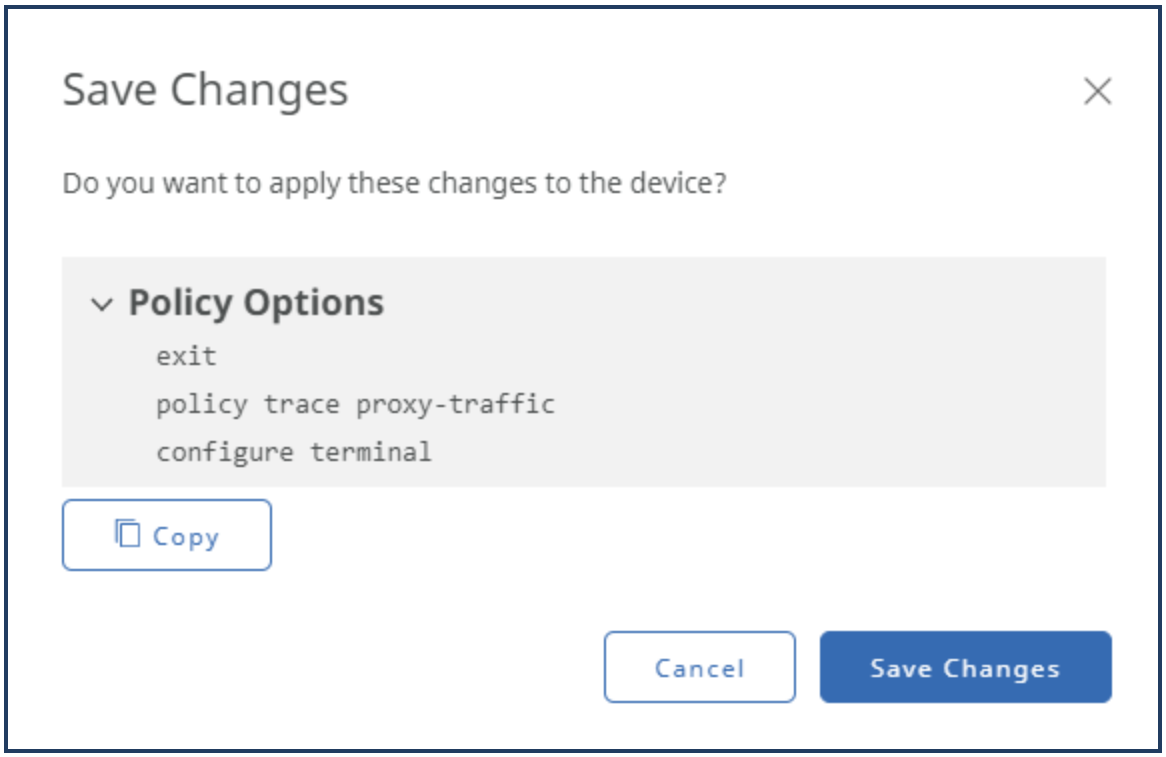

Note: This dialog only confirms that changes are saved, and doesn't check the data you enter (for example, whether a specified server URL is reachable or not).

Automatically Saved Changes

The **Instant Save** feature is available in some areas of the Admin Console when you make configuration changes:

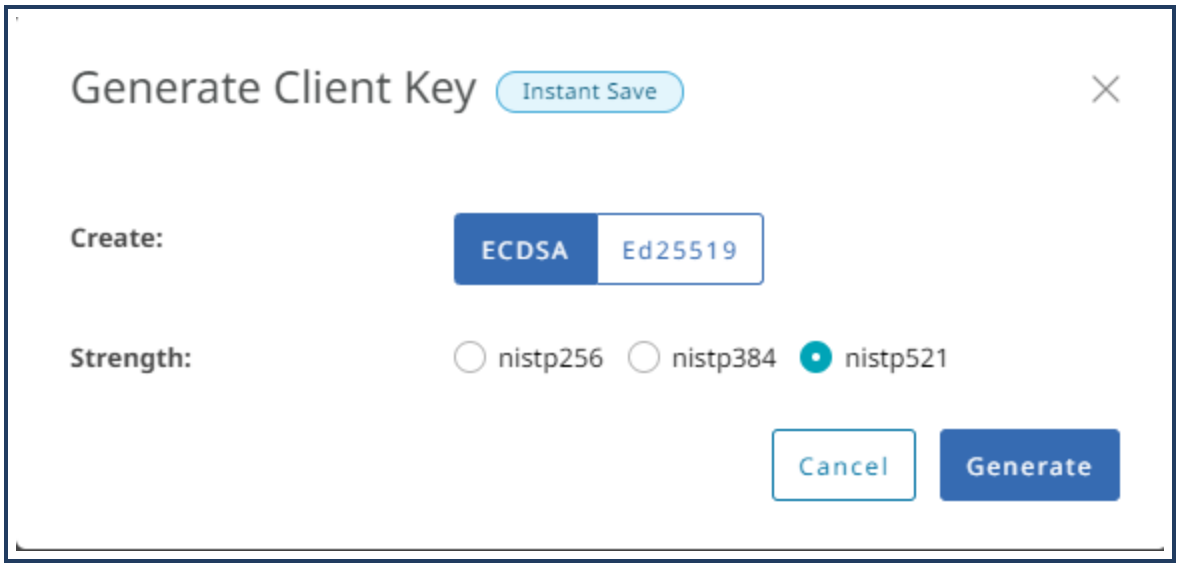

These changes are applied immediately without confirmation from you.

Help Topics

Two kinds of help are available in the Admin Console:

- Select the **?** icon at the top right of each Admin Console page to open a general help topic.
- To read a brief description about a section, select **SHOW HELP**. For full information on a topic, select the **Learn more** link. The following example shows an expanded description.

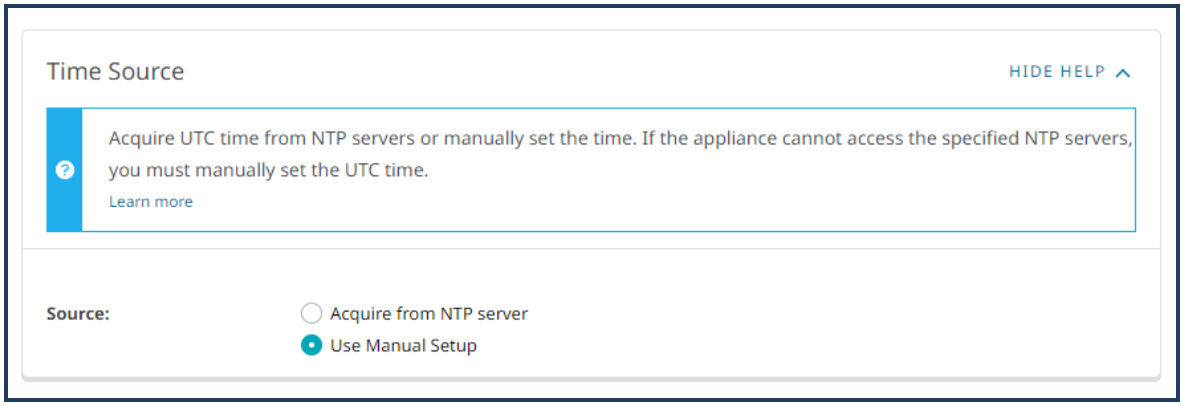

Whenever you open a help topic, it is displayed within the hierarchy of the full Admin Console help system. For other ways to navigate the help system:

- Browse the table of contents in the navigation pane.
- Type text into the **Search** field at the top of help pages.

- Select links in the breadcrumbs at the top of help pages.
- Use cross-reference links throughout the documentation.

Documentation Conventions

This help system provides comprehensive information to help you configure and use the ProxySG appliance. Refer to the following to understand the help system's conventions.

Help System Version

The Admin Console has a single-version help system. This means that the same documentation is available for version 6.7.x and 7.1. Any feature or behavior differences between versions are described in the help.

Typography

The following table lists the typographical conventions as well as CPL and CLI syntax conventions used in this guide.

Convention	Usage	Example
bold text	Labels, such as page and dialog names, in the Admin Console	Select **Allow** or **Deny**.
bold text separated by >	Navigation path in the Admin Console	Select **Configuration > Policy > Policy Options**.

`fixed-width text`	CLI input/output, CPL, log output, other output	`Download log:`
		`Local database download at:`
		`2008/08/11 17:40:42-0400`
		`...`

`fixed-width text, bold`	In CLI or CPL, a literal string to be entered as shown.	`# (config proxy-services)` **`dynamic-bypass`**
`fixed-width text, italics`	CLI or CPL variable to be substituted with a literal name or value	`# (config dynamic-bypass)` **`max-entries`** *`number`*
regular text, italics	Variable to be substituted with a literal name or value	
`{ }`	One of the CLI or CPL parameters enclosed within the braces must be supplied.	`# (config dynamic-bypass) {`**`enable`** `|` **`disable`**`}`
|	The CLI or CPL parameter before or after the pipe character can or must be entered.	
`[ ]`	Optional CLI or CPL parameter.	`# (config dynamic-bypass) [`**`no`**`]` **`trigger all`**

Procedures

When you follow a procedure in this document, you might not want to save the changes immediately. For this reason, all procedures omit the final step of saving or applying changes.

If you attempt to navigate away from a page with unsaved settings, the Admin Console prompts you with an **Unsaved pending changes** dialog.

See "Navigate the Admin Console" on page 6 for more information.

Abbreviations

This help system includes abbreviated terms. See " Abbreviations" on the next page for the expanded forms and definitions of the terms.

Abbreviations

This help system includes abbreviated terms. Refer to the following for the expanded forms and definitions of the terms.

Abbreviation	Expanded Term	Definition
ACL	access control list	Text file that includes lists of IP addresses that can access the console.
AD	Active Directory	Microsoft services that manage user authentication and authorization, as well as user directory information.
ADN	Application Delivery Network	Defines the framework that enables application acceleration between corporate offices in a WAN.
BCAAA	Blue Coat Authentication and Authorization Agent	Software installed on a domain server, acting as an intermediary between the ProxySG appliance and the domain.
CA	Certificate Authority	Organization that issues digital certificates.
CCL	CA Certificate List	List of trusted certificates.
CIFS	Common Internet File System	Protocol based on the SMB protocol used for file sharing, printers, serial ports, and other communications. The ProxySG appliance supports CIFS proxy.
CLI	Command Line Interface	One of two ways to access the ProxySG appliance; a command line tool where you can configure SGOS and execute administrative commands. Set up the CLI admin account and password during the initial creation of your appliance.
CPL	Content Policy Language	The language in which ProxySG appliance policy is written. Policies can be customized to an organization's specific set of users and unique enforcement needs.
CRL	certificate revocation list	List of certificates that have been revoked and should not be trusted.
DC	domain controller	Server that responds to authentication requests and verifies users on computer networks
DNS	Domain Name System	System that translates domain names to IPv4 or IPv6 addresses.
EDNS	extension mechanisms for DNS	Allows DNS requesters to receive DNS UDP messages longer than the default 512 bytes.
ELFF	Extended Log File Format	Log file format defined by the W3C.
FTP	File Transfer Protocol	Commonly used protocol used to transfer files between client and server.
GRE	Generic Routing Encapsulation	Protocol that can encapsulate multiple network layer protocols inside virtual point-to-point links or point-to-multipoint links over an IP network.

Abbreviation	Expanded Term	Definition
ICAP	Internet Content Adaptation Protocol	With the ProxySG appliance, used for content scanning and ICAP servers.
ICMP	Internet Control Message Protocol	Protocol used by routers to generate error messages to the source IP address when network problems prevent delivery of IP packets.
IWF	Internet Watch Foundation	A third-party content filter provider supported on the appliance
MAPI	Messaging Application Programing Interface	Protocol used by Microsoft Outlook (client) to communicate with Microsoft Exchange (server).
NTP	Network Time Protocol	Protocol used for synchronizing clocks in a network.
OCS	origin content server	The source of the internet content that the client is requesting. A web server is an example of an OCS.
OCSP	Online Certificate Status Protocol	Protocol that provides a secure means of checking certificate revocation status in real time.
PKI	public key infrastructure	System of processes, technologies, and policies that allows you to encrypt and sign data.
RIP	Routing Information Protocol	Protocol that uses hop count to select the fastest route to a destination.
RTS	Return-to-Sender	Allows the appliance to send response packets to interface that received the request packet.
RTSP	Real Time Streaming Protocol	Protocol used for communication with streaming media servers.
SCP	Secure CoPy	SCP protocol allows you to transfer data securely over an SSH connection.
SMB	Server Message Block	Client-server communication protocol used for sharing access to resources on a network. The ProxySG appliance supports SMBv1 and SMBv2.
SPN	Service Principal Name	Unique identifier of a service instance, used by Kerberos to associate a service instance with a service login.
SSL	Secure Sockets Layer	Cryptographic protocol used to secure internet communications.
STP	Spanning Tree Protocol	Protocol that prevents bridge loops in a network with redundant paths.
SWG VA	Secure Web Gateway virtual appliance	Virtual edition of ProxySG appliance. Supported deployments include VMWare, AWS, and Hyper-V environments.
TTL	time-to-live	Counter within an ICMP packet. As a packet goes through each router, the router decrements this TTL value by 1.

Abbreviation	Expanded Term	Definition
TLS	Transport Layer Security	Newer version of SSL.
UTC	Coordinated Universal Time	International time standard.
VA	virtual appliance	The virtual machine image.
VAP	virtual appliance package	The zip file that contains the OVF file and the virtual disk files (.vmdk) required for creating the ProxySG VA.
VPM	Visual Policy Manager	A graphical policy editor. In the VPM, you can define policies without having to manually edit policy files.
WAF	Web Application Firewall	Symantec's WAF and reverse proxy solution that provides OWASP Top 10 protection and more.
WCCP	Web Cache Communication Protocol	Cisco-developed protocol that allows certain Cisco routers and switches to transparently redirect traffic to a cache engine such as a ProxySG appliance.

SGOS Documentation

The following lists the latest SGOS 7.2.x documentation.

In addition, Symantec provides other deployment guides targeted for specific solutions. For these guides and earlier versions of documentation, refer to: https://techdocs.broadcom.com/content/broadcom/techdocs.html

Document	Overview
SGOS Upgrade/Downgrade Guide https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/upgrade-downgrade-guide.html	Steps for upgrading or downgrading SGOS,. Also covers behavior changes and policy deprecations.
SGOS Administration Guide https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/admin-guide.html	Detailed information for configuring and managing the ProxySG appliance using the Java-based Management Console.
Command Line Interface Reference https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/command-line-interface-reference.html	Commands available in the ProxySG appliance CLI and how to use them to perform configuration and management tasks.
ProxySG Web Visual Policy Manager WebGuide https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/web-vpm-policy.html	How to create and implement policy in the ProxySG appliance's web-based Visual Policy Manager, including layer interactions, object descriptions, and advanced tasks.
Legacy Visual Policy Manager Reference https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/legacy-vpm-policy.html	How to create and implement policy in the ProxySG appliance's legacy Visual Policy Manager.
Content Policy Language Reference https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/content-policy-language.html	CPL gestures available for writing the policy by which the ProxySG appliance evaluates web requests.
Reverse Proxy WebGuide https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/reverse-proxy-deployment-guide.html	How to deploy a ProxySG appliance as a front-end for Internet-based users to access secure application, content, and web servers.

Document	Overview
Web Application Firewall Solutions Guide https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/web-application-firewall-solutions-guide.html	How to configure Symantec WAF solution to protect your web servers, accelerate web content, and simplify operation.
ProxySG Log Fields and CPL Substitutions Reference https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/proxysg-log-fields-and-substitutions.html	Describes fields available for creating access log formats on the ProxySG appliance.
Secure Web Gateway - Content Analysis Policy Best Practices Improvement https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/content-analysis-best-practice-policies.html	Provides a secure and customizable policy model for bypassing content scanning to improve the user experience where needed or save resources by excluding low-risk/high-volume traffic. This document identifies weak policy conditions and is intended to reduce risk by using different sets of policy conditions.
SGOS 7.2.x Documentation https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2.html	Full list of documentation published for SGOS 7.2.x.
SGOS Release Notes SGOS Release Notes are available on the Broadcom Support Documents and Downloads page.	Changes, issues, fixes, and limitations pertaining to SGOS releases. Also includes any related security advisory (SA) fixes.

Configure the ProxySG Appliance

Refer to the following topics to learn to configure the appliance:

- "Manage ProxySG Policy" on the next page
- "Configure Networking" on page 58
- "Manage ProxySG Services" on page 149
- "Configure SSL and Device Authentication" on page 203
- "Authenticate Administrators and Users" on page 230

Manage ProxySG Policy

Refer to the following topics to learn about managing policy on the ProxySG appliance:

- "Configure Policy Options" on the facing page
- "Manage Policy Files" on page 27
- "About Policy Exception Pages" on page 27
  - "Exception Page Types" on page 29
  - " Exception Definition Overview" on page 35
  - "Use Substitution Variables in Exceptions" on page 37
  - "Manage Policy Exceptions" on page 40
    - "Examples: Create and Update Exception Pages" on page 42
  - "Include Exceptions in Policy" on page 50

See Also:

- *Web Visual Policy Manager Reference*
- *Content Policy Language Reference*

Configure Policy Options

Policy options allow you to change the default transaction setting, configure policy tracing, and re-order policy evaluation.

Transaction Settings: Deny and Allow

The default proxy transaction policy is to either deny proxy transactions or to allow proxy transactions.

A default proxy transaction policy of **Deny** prohibits proxy-type access to the appliance; you must then create policies to explicitly grant access on a case-by-case basis.

A default proxy transaction policy of **Allow** permits most proxy transactions; however, if protocol detection is enabled, the appliance allows HTTP CONNECT for both port 443 and other ports—provided the appliance detects a known protocol. If protocol detection is disabled, HTTP CONNECT is only allowed on port 443. If your policy is set to Allow, you must create policies to explicitly deny access on a case-by-case basis.

> **Note:** The default proxy policy does not apply to admin transactions. By default, admin transactions are denied unless you log in using console account credentials or if explicit policy is written to grant read-only or read-write privileges.

What is the default proxy policy?

- If you installed SGOS using the front panel or a serial console port, the default setting is Deny.
- If you upgraded SGOS from a previous version, the default remains whatever it was for the previous policy.

> **Note:** Changing the default proxy transaction policy affects the basic environment in which the overall policy is evaluated. It is likely that you must revise policies to retain expected behavior after such a change. Changing the default proxy transaction policy does not affect the evaluation of cache and admin transactions.

To set the default proxy policy, select **Allow** or **Deny**.

Configure Policy Tracing

Tracing enabled with the Admin Console or CLI is global; that is, it records every policy-related event in every layer. It should be used only while troubleshooting. For information on troubleshooting policy, refer to the *Content Policy Language Reference*. Turning on policy tracing of any kind is expensive in terms of system resource usage and slows down the appliance's ability to handle traffic.

Caution: Tracing is best used temporarily, such as for troubleshooting. Do not enable policy tracing during normal operation of the appliance, as doing so can cause performance issues.

Configure policy tracing:

1. In the Admin Console, select **Configuration > Policy > Policy Options**.
2. In **Policy Options**, for Default Policy Tracing, select one of the options:
    - **None**: No policy tracing.
    - **Proxy Traffic**: Trace only proxy traffic.
    - **All Execution**: Trace all traffic.

Manage Policy Files

The order in which the appliance evaluates policy rules is significant. Changes to the evaluation order can result in different effective policy, as the order of policy evaluation defines general rules and exceptions. While this order is configurable, the default and recommended order is:

- VPM
- Local
- Central
- Forward

This order prevents policies in the Central file that block virus signatures from being inadvertently overridden by allow (access-granting) policy rules in the and Local files.

On an upgraded appliance, the policy evaluation order is the order that existed on the appliance before the upgrade.

To change policy order, use the arrows in the **Actions** column to move files up and down on the list.

Tip: When changing the policy file evaluation order, remember that decisions from files later in the order can override decisions from earlier files.

Manage Policy Files

The order in which the appliance evaluates policy rules is significant. Changes to the evaluation order can result in different effective policy, as the order of policy evaluation defines general rules and exceptions. While this order is configurable, the default and recommended order is:

- VPM
- Local
- Central: The Central policy file is updated when needed by Symantec. The file can be updated automatically or you can request e-mail notification. You can also configure the path to point to your own custom Central policy file.

  > **Note:** If you import policy that was created externally (for example, by a third-party tool) into a policy file, the ProxySG appliance validates the policy file's contents for syntax errors but cannot check for formatting or typographical mistakes. Such errors may result in unintended behavior after the policy is loaded. Symantec strongly recommends that you review the CPL in imported policy for correctness before installing the file.

- Forward

This order prevents policies in the Central file that block virus signatures from being inadvertently overridden by allow (access-granting) policy rules in the and Local files.

On an upgraded appliance, the policy evaluation order is the order that existed on the appliance before the upgrade.

Change the Policy Evaluation Order

To change policy order, use the arrows in the **Actions** column to move files up and down on the list.

> **Tip:** When changing the policy file evaluation order, remember that decisions from files later in the order can override decisions from earlier files.

For more information on policy files, refer to the *Content Policy Language Reference*.

About Policy Exception Pages

Exception pages are web pages displayed to users under specified conditions, and are intended to provide feedback to users about errors and policy actions. When users request a URL that triggers such an error or action, exception pages are returned to the web browser instead of users' requested content.

The appliance ships with a number of predefined exceptions, which you can use with or without customization. In addition, you can create custom exception pages. See "Exception Page Types" on the facing page.

Exception definitions consist of specific fields, such as `details` and `summary`, which provide information about the exception. The `format` field in an exception definition contains valid HTML to format the exception page (applicable to HTTP responses). You can also include substitution variables in any part of the exceptions page. See " Exception Definition Overview" on page 35.

Implement Exception Pages

Implementing exception pages consists of two steps:

1. Creating or editing exception pages:
   - Creating and installing user-defined exception pages.
   - Updating existing built-in or user-defined exception pages.

   See "Examples: Create and Update Exception Pages" on page 42.
2. Referencing the exception pages in policy.

   While configuring policy, you can also override the exception's current definition for the specific policy rule. See "Include Exceptions in Policy" on page 50.

> **Note:** Exception pages are returned to users whose requests are intercepted and match the policy conditions. See "Change the State of a Proxy Service" on page 154 for details on setting services to Intercept.

> **Tip:** Enable access logging to record when exceptions are returned and users' ensuing actions. See "Configure ProxySG Access Logs" on page 383.

See Also:

- "Manage Policy Exceptions" on page 40
- "Exception Page Types" on the facing page
- " Exception Definition Overview" on page 35
- "Include Exceptions in Policy" on page 50

Exception Page Types

Exceptions are sent in response to certain ProxySG client requests, such as denial by policy, failure to handle the request, and authentication failure. Exceptions are returned to users based on policy rules defined by the ProxySG administrator. For example, if a client sends a request for content that is not allowed, an exception HTML page (for HTTP connections) or an exceptions string (for non-HTTP connections) is returned, informing the user that access is denied.

Two types of exceptions can be used on the appliance: "Built-in Exceptions" below and "User-Defined Exceptions" on page 34.

Built-in Exceptions

Built-in exceptions are a set of pre-defined exceptions. They send information back to the user under operational contexts that are known to occur, such as `policy_denied` or `invalid_request`.

Built-in exceptions are always available and can also have their contents customized; however, they can't be deleted. To edit a built-in exception, see "Configure Exception Pages" on page 41.

The table below lists the built-in exceptions and the context under which they are issued.

Exception Type	HTTP Response Code	Issued When...
authentication_ failed	401	The transaction cannot be authenticated, usually because the credentials were incorrect. `authentication_failed` is a synonym for `deny.unauthorized`.
bad_credentials	400	The username or password were sent using an invalid/ unrecognized format. This can have two causes: ■ The username or password contains non-ASCII characters, and the appliance is not configured to use the same authentication character encoding as is being used by the web browser. ■ The username or password is too long. (The limits for the username and password are 64 bytes each,after being translated to UTF-8.)
client_failure_ limit_ exceeded	503	Too many requests from your IP address (`$(client.address)`) have failed.
configuration_ error	403	A configuration error on the appliance was detected, and the requested operation could not be handled because of the configuration error. This exception is a likely indicator that the administrator of the ProxySG must intervene to resolve the problem.
connect_ method_denied	403	A user attempted an CONNECT method to a nonstandard port when explicitly proxied.Symantec does not allow CONNECT methods to non-standard ports by default because it is considered a security risk to do so.
content_ encoding_error	502	A Web site presented a content encoding header of one type but encoded the data differently

Exception Type	HTTP Response Code	Issued When...
content_filter_ denied	403	A particular request is not permitted because of its content categorization.
content_filter_ unavailable	403	An external content-filtering service could not be contacted, and the appliance is failing closed in such a situation.
dns_server_ failure	503	The request could not be processed because the appliance was unable to communicate with the DNS server in order to resolve the destination address of the request.
dns_ unresolved_ hostname	404	The request could not be processed because the appliance was unable to resolve the hostname in the request with DNS.
dynamic_ bypass_reload	200	The dynamic_bypass policy action is matched.
gateway_error	504	There was a network error while attempting to communicate with the upstream gateway.
icap_ communication_ error	504	A network error occurred while the appliance was attempting to communicate with an external ICAP server.
icap_error	504	A network problem occurred, the ICAP service might be misconfigured, or the ICAP server might have reported an error.
internal_error	500	The appliance encountered an unexpected error that resulted in the inability to handle the current transaction.
invalid_auth_ form	403	The submitted authentication form is invalid. The form data must contain the username, password, and valid original request information.
invalid_request	400	The request received by the appliance was unable to handle the request because it detected that there was something fundamentally wrong with the syntax of the request.
invalid_ response	502	The server's response could not be processed because of a malformed response or a misconfiguration.
license_ exceeded	403	Access is denied because a license has been exceeded on the proxy, and the request is not permitted.
license_expired	403	The requested operation cannot proceed because it would require the usage of an unlicensed feature.

Exception Type	HTTP Response Code	Issued When…
method_denied	403	The requested operation utilizes a method that has been explicitly denied because of the service properties associated with the request.
not_ implemented	501	The protocol cannot handle the requested operation because it utilizes a feature that is not currently implemented.
notify	200	Used internally by the VPM. You do not need to customize the text of this exception, since in this case the entire HTML response is generated by VPM and is not taken from the exception definition.
notify_missing_ cookie	403	This exception is returned when a Notify User action is being used to notify the user, and the user has disabled cookies in the Web browser.
policy_denied	403	policy_denied is a synonym for deny.
policy_redirect	302	A redirect action is matched in policy.
radius_splash_ page	200	The user is authorized. Click the refresh button on the browser to proceed to the requested site. The user/ session ID is $(x-radius-splash-username)/$(x-radiussplash- session-id)
redirected_ stored_ requests_ not_ supported	403	This applies to forms authentication with POST requests only): The origin server returned a redirect for the request. The appliance is configured to not allow stored requests to be redirected.
refresh	200	A refresh (using the HTTP Refresh: header) is required. The refresh exception (by default) refreshes the originally requested URL (or in some cases, its post-imputed form).
server_request_ limit_ exceeded	503	Too many simultaneous requests are in progress to $(url.host).
silent_denied	403	An exception(silent_denied) is matched in policy. This exception is pre-defined to have no body text, and is silent in that it results in only the status code being sent to the client.
server_ authentication_ error	500	Internal error. The appliance encountered an internal error while preparing to send the username/password upstream. This error can only occur when the appliance “server authentication” feature is enabled.
ssl_client_cert_ expired: Expired SSL Client Certificate	503	A web site presents an incorrect or invalid certificate or a configuration error has occurred.
ssl_client_cert_ ocsp_check_ failed OCSP Error On Client Certificate	503	An error occurred while checking the revocation status of the certificate.

Exception Type	HTTP Response Code	Issued When...
ssl_domain_ invalid: SSL Certificate Host Mismatch	409	There was a failure contacting a web site through HTTPS because the certificate has a common name that does not match the web site's domain name.
ssl_failed: SSL Certificate Verification Error	503	A secure connection could not be established to an web site. This typically occurs when a web site that is not configured to accept SSL connections.
ssl_server_cert_ expired: Expired SSL Server Certificate	503	A Web site presents an incorrect or invalid certificate or a configuration error has occurred.
ssl_server_cert_ ocsp_check_ failed OCSP Error On Server Certificate	503	An error occurred while checking the revocation status of the certificate.
ssl_client_cert_ revoked: Revoked SSL Client Certificate	503	The client presents a revoked certificate or a configuration error has occurred.
ssl_client_cert_ ocsp_status_ unknown: Unknown OCSP Status of Client Certificate	503	An OCSP check returned unknown status for a client certificate.
ssl_client_cert_ untrusted_ issuer Untrusted SSL Client Certificate	503	A Web site presents an incorrect or invalid certificate or a configuration error has occurred.
ssl_server_cert_ ocsp_status_ unknown Unknown OCSP Status of Server Certificate	503	The server certificate revocation status is unknown. This is caused by a certificate revocation check for which the server does not have a status.

Exception Type	HTTP Response Code	Issued When...
ssl_server_cert_revoked: Revoked SSL Server Certificate	503	A Web site presents a revoked certificate or a configuration error has occurred.
ssl_server_cert_untrusted_issuer: Untrusted SSL Server Certificate	503	A Web site presents an incorrect or invalid certificate or a configuration error has occurred.
tcp_error	503	A network error occurred attempting to communicate with an upstream host.
transformation_error	403	The server sends an unknown encoding and the appliance is configured to do content transformation.
unsupported_encoding	406	The client makes a request with an Accept- Encoding: Identity;q=0, … header. Only uncompressed content is available in cache, the appliance is not configured to compress the content, or the compression license is expired, or the client request results in to Accept-Encoding: Identity;q=0 because of the combination of request and configured policy.
unsupported_protocol	406	The protocol used in the request is not understood.
upstream_407_rejected	407	An authentication challenge (HTTP status code 407 "Proxy authentication required") from an upstream OCS was blocked.
virus_detected	200	Virus was detected in the content.
data_leak_detected	200	A violation of DLP policy was detected in the content. The ICAP response must contain the following: ■ `X-Violations-Found` in the HTTP header ■ `dlp` string in the Server header

Some built-in exceptions can be initiated directly through the `exception()` action (or in the VPM, the **Return Exception** object). Some require additional states, which makes initiating them either problematic or out of context. The following are exceptions that can't be initiated though policy:

- `authentication_failed`
- `authentication_failed_password_expired`
- `authentication_redirect_from_virtual_host`
- `authentication_redirect_to_virtual_host`

- authentication_success
- dynamic_bypass_reload
- license_expired
- ssl_domain_invalid
- ssl_failed

To view the contents of a built-in exception, enter the `show exceptions` *`exception_id`* at the `#(config)` prompt, as follows:

```
#(config) exceptions
#(config exceptions) show exceptions configuration_error
configuration_error exception:
all protocols:
summary text:
     SG configuration error
details text:
     Your request could not be processed because of a configuration error: $(exception.last_error)
help text:
     The problem is most likely because of a configuration error, $(exception.contact) and provide
them with any pertinent information from this message.
http protocol:
     code: 403
```

User-Defined Exceptions

User-defined exceptions are created and deleted by the administrator. If a user-defined exception is referenced by policy, it can't be deleted. The default HTTP response code for user-defined exceptions is 403. To create and edit user-defined exceptions, see "Configure Exception Pages" on page 41.

See Also:

- "About Policy Exception Pages" on page 27
- " Exception Definition Overview" on the facing page
- "Examples: Create and Update Exception Pages" on page 42

Exception Definition Overview

In ProxySG appliance policy, an exception definition is a file that specifies the contents and appearance of exception pages that are returned to users. Exception definitions are either built-in or user-defined.

You can deploy exceptions with their default values (some fields are undefined by default), which means that the browser presents plain exception pages with minimal information to the user. To provide specific error details, organizational information, and custom styling, define the appropriate exception fields.

An exception definition file shows the fields in a structured data list format and allows you to define each field:

```
(exception.exception_id
  (contact "") ; defines the $(exception.contact) variable <---- start of fields that apply to all protocols
  (details "") ; the exception reason using the $(exception.details) variable
  (format "") ; defines format of the HTML
  (help "") ; defines the $(exception.help) variable
  (summary "") ; the exception message using the $(exception.summary) variable
  (http ; specific settings for HTTP/S exceptions <---- start of fields that apply to just HTTP
    (code "") ; valid HTTP response code to return to the client
    (contact "") ; defines the $(exception.contact) text
    (details "") ; HTML exception detail text
    (format "") ; HTML exception format with variables
    (help "") ; defines the $(exception.help) text
    (summary "") ; short name for $(exception.summary)
  ) <---- closing tag for HTTP fields
) <---- closing tag for exception definition
```

To view exception definition files, in the Admin Console, select **Configuration > Policy > Exception Pages**.

Exception Fields

Refer to the following descriptions of exception fields:

Field name	Description
`exception.`*`exception_id`*	Specifies the exception identifier: ■ A built-in exception is called `exception.`*`exception_id`*, where *`exception_id`* is the exception type. See "Exception Page Types" on page 29 for the built-in exception types. ■ A custom exception is called `exception.user-defined.`*`exception_id`*, where *`exception_id`*is a name you specify.

Field name	Description
`contact`	Sets site-specific contact information that can be used in all exceptions. It is possible to customize contact information per exception, but setting the top-level contact information for all exceptions is sufficient in most environments.
`format`	Controls the appearance of the exception page. For an HTTP exception response, the format is an HTML file. For other protocol responses that can't render HTML, the exception page returns a single line.
`summary`	Gives a reason for the exception. For example, the default `policy_denied` exception summary is “Access Denied”.
`details`	Describes in more detail the reason for the exception. For example, the default `policy_denied` exception (for the HTTP protocol) detail is `Your request has been denied by system policy.`
`help`	Provides possible causes and potential steps for users to take. For example, to review URL categories, append the `$(exception.category_review_url)` and `$(exception.category_review_message)` substitutions to the `$(exception.help)` definition. See "Use Substitution Variables in Exceptions" on the facing page and *SGOS Administration Guide* for details.
`http`	Specifies the HTTP response code to use when the exception occurs. For example, the `policy_denied` exception by default returns the 403 Forbidden HTTP response code.

Note: Fields other than `format` must contain fewer than 8000 characters. Characters exceeding this limit are not shown to the user.

Recommendations for Modifying Exception Files

- To prevent issues from occurring with exception pages in production, download the existing exception file, modify it locally, and then install the new file.
- For ensure correct syntax, begin each file with a definition for `exception.all`. End each file with the `exception.all` closing parenthesis:

```
(exception.all
(exception.policy_denied)
...
)
```

- Before deleting a user-defined exception, check if it is referenced in installed policy. Remove the reference from policy, reinstall policy, and then delete the exception.

How Fields Inherit Definitions

Exceptions are stored on the appliance in a hierarchical model. Two parent exceptions, `exception.all` and `exception.user-defined.all`, provide definitions that all child exceptions inherit. Child exception files don't need each field to be defined as long as they inherit definitions from one of the parent exceptions.

The `user-defined.all` exception is the parent of all user-defined exceptions, but it is also a child of `exception.all`. Configuring `exception.user-defined.all` is only necessary if you want certain fields to be common for all user-defined exceptions, but not common for built-in exceptions.

> **Note:** By default, the built-in exceptions don't define the `format` field; however, because each built-in and user-defined exception derives its default values from `exception.all`, built-in exception `format` fields inherit the `exception.all`'s `format` field definition. To change the format text for all built-in and user-defined exceptions, customize the `format` field in `exception.all`.

See Also:

- "Use Substitution Variables in Exceptions" below
- "Examples: Create and Update Exception Pages" on page 42

Use Substitution Variables in Exceptions

When defining exception fields, you can use substitution variables to generate details specific to the request. Refer to the *ProxySG Content Policy Language Reference* for information on substitution variables.

Some exception fields can be configured as substitutions:

- `$(exception.id)`
- `$(exception.summary)`
- `$(exception.details)`
- `$(exception.help)`
- `$(exception.contact)`

Additionally, the `format`, `summary`, `details`, `help`, and `contact` fields can be configured specifically for HTTP in the `http` section of the format, or configured commonly for all protocols. See "Use Substitution Variables in Exceptions" above for an example.

The `format` field contains the HTML body of the exception. It is not available as a substitution itself, but it can include substitutions. The following is a simple HTML format that uses substitution variables:

```
$(exception.id): $(exception.summary)
Request: $(method) $(url)
Details: $(exception.details)
Help: $(exception.help)
Contact: $(exception.contact)
```

> **Tip:** Other useful substitutions related to exceptions:
>
> - `$(exception.last_error)`: Includes additional exception details when available.
> - `$(exception.reason)`: Shows the reason when the appliance terminates a transaction. For example, a transaction that matches a DENY rule in policy has its `$(exception.reason)` set to "Either 'deny' or 'exception' was matched in policy".

Specify Substitution Key-Value Pairs

You can specify substitution variables in key-value pairs in the exception file:

```
(exception.all
   (format "This is an exception: $(exception.details)")
   (details "")
   (exception.policy_denied
   (format "")
   (details "your request has been denied by system policy")
   )
)
```

In this example, `policy_denied` defines the `$(exception.details)` substitution as "Your request has been denied by system policy". Because the `policy_denied``format` field is undefined, it inherits the `format` field from its parent exception (`exception.all`). When the `policy_denied` exception is issued, the resulting text is:

```
This is an exception: your request has been denied by system policy.
```

Configure a Substitution Variable Using the Command Line

Use the `# (config exceptions)` **`inline`** command to configure substitution variables. The following sets the `$(exception.contact)` substitution for every HTTP exception:

```
# (config exceptions) inline http contact EOF
For assistance, contact <a href="mailto:sysadmin@example.com">sysadmin</a>EOF
```

To configure a different `$(exception.contact)` substitution for every HTTP exception:

```
# (config exceptions) user-defined inline http contact EOF
For assistance, contact <a href="mailto:policyadmin@example.com">policyadmin</a>EOF
```

Customize a Substitution Variable in the Exception File

To customize the `$(exception.contact)` substitution variable for every HTTP exception, edit the `exception.all` element:

```
(exception.all
   (contact "For assistance, contact your network support team.")
   (details "")
   (format "$(exception.id): $(exception.details)")
   (help "")
   (summary "")
   (http
     (code "200")
     (contact "")
     (details "")
     (format <<EOF
<format_contents>
EOF
    )
    (help "")
    (summary "")
    )
<built_in_exceptions>
    )
)
```

To add the `$(exception.contact)` information, modify the contact substitution variable under the `http` element:

```
(exception.all
   (contact "For assistance, contact your network support team.")
   (details "")
   (format "$(exception.id): $(exception.details)")
   (help "")
   (summary "")
   (http
     (code "200")
     (contact "For assistance, contact <a href="mailto:sysadmin@example.com">sysadmin</a>")EOF
     (details "")
     (format <<EOF
<format_contents>
EOF
      )
      (help "")
      (summary "")
<built_in_exceptions>
      )
)
```

See "Examples: Create and Update Exception Pages" on page 42 for detailed instructions on modifying exception files and examples of custom exception pages.

Manage Policy Exceptions

You can create new custom exceptions and edit existing built-in and custom exceptions. The Admin Console allows you to edit and install exceptions, and also provides details on the contents of the exceptions files.

To manage policy exceptions:

1. In the Admin Console, select **Configuration > Policy > Exception Pages**.
2. Select the following as needed:

Task	Field or Value
Display current exceptions	In the **Source** section, select **Current - Source**. The browser opens the current exceptions in a new browser tab.
Display default exceptions definition	In the **Source** section, select **Default - Source**. The browser opens the default source in a new browser tab.
Install new and edit existing policy exceptions For instructions, see "Configure Exception Pages" on the facing page. For help with editing the structured data language format of the exceptions file, see " Exception Definition Overview" on page 35.	In the **Install** section, select an installation method: ■ Text editor, which allows you to enter settings (or copy and paste the contents of an existing file) directly onto the appliance. ■ Local file, to which the appliance can browse and install. ■ Remote URL, where you put an already-created file on an FTP or HTTP server to be downloaded to the appliance. Use the `# (config exceptions) path` *URL* command to set the path and the `# (config exceptions) load exceptions` command to load the exceptions. When the exceptions file is edited, it updates the existing exceptions already on the appliance. The configuration remains in effect until it is overwritten by another update.

Task	Field or Value
View most-recently installed policy definitions, including any substitution variables, and preview policy exception pages in HTML format	In the **Install** section, select **Most recent install**. A new browser tab shows the following information: ■ **Common Configuration**: Fields in `exception.all` and `exception.user-defined.all` that have default values and are not overridden in child exceptions. ■ **User-Defined Exceptions**: Installed user-defined exceptions. ■ **Built-in Exceptions**: Pre-defined exceptions. In both user-defined and built-in exceptions lists, you can select the following options: ■ **View Sample HTML**: Exception page as it appears to users. ■ **View Settings**: Overview of exception fields. ■ **View Expanded Settings**: Detailed view of exception fields and Format contents.
Verify results of most recent exception page installation	In the **HTML Pages and Configuration** section, select **Most recent install**. The console opens the results of the last installation, including any errors or warnings.

Configure Exception Pages

You can edit existing exceptions and create user-defined exception pages.

> **Tip:** Before configuring exception pages, review " Exception Definition Overview" on page 35 for details on the structure of an exception file and the use of each element in the file.

> **Note:** You can't create user-defined exceptions for patience pages.

To configure exception pages:

1. In the Admin Console, select **Configuration > Policy > Exception Pages**.
2. Prepare the exception page file and select an installation method in **Install**:
   - If using a local file, create the file (for example, select **Install from text editor**, copy and paste the contents to a new file, and make your changes).

- If using a remote URL, use the same method as above, host the exception file on a server that the ProxySG appliance can access, and specify the URL.
- If using text editor, open the text editor and make changes within the file.

For example, to change the default contact details on all built-in exception pages, make the following edit under `exception.all`:

```
(contact "For assistance, email support@yourcompany.org.")
```

To create a new exception page, see "Create a User-Defined Exception Page " below.

3. Save the exception page and check the results of the installation (in **HTML Pages and Configuration**, select **Most Recent Install**).

   You can also preview the page as it would appear to users. In **Install**, select **Most Recent Install**. In the list of user-defined exceptions, select **View Sample HTML** for the exception you created.

After you create or edit exceptions (following the "Examples: Create and Update Exception Pages" below), they can be selected in the **Return Exception** action in the VPM.

Next, include the exception page in a policy rule. See "Include Exceptions in Policy" on page 50 for instructions; these instructions also include steps on using the VPM to customize single exception pages for specific use cases.

See Also:

- " Exception Definition Overview" on page 35
- "Examples: Create and Update Exception Pages" below
- "Include Exceptions in Policy" on page 50

Examples: Create and Update Exception Pages

The following examples illustrate how to create and update exception pages:

- "Create a User-Defined Exception Page " below
- "Edit an Existing Exception Page Definition" on page 47

Create a User-Defined Exception Page

This example uses the text editor to create a user-defined exception page called `user-defined.legal_reasons` for internal compliance.

To create a user-defined exception page:

1. In the Admin Console, select **Configuration > Policy > Exception Pages**.
2. Under Install, select **Install from text editor**.

3. In the dialog that opens, look for `exception.user-defined.all`. Go to the bottom of the exception definition and enter a new line before the last closing parenthesis:

```
(exception.user-defined.all   <---- opening exception.user-defined.all tag
   (contact "")
   (details "")
   (format "")
   (help "")
   (summary "")
   (http
      (code "403")
      (contact "")
      (details "")
      (format "")
      (help "")
      (summary "")
   )
                        <---- put the new exception definition here
)   <---- closing exception.user-defined.all tag
```

4. Copy an exception definition and paste it in the new line, omitting the fields preceding `http`. Then, edit the `http` fields as follows:

> **Note:** All user-defined exceptions are named `user-defined.`*`exception_id`*.

```
(exception.user-defined.all <---- opening exception.user-defined.all tag
   (contact "")
   (details "")
   (format "")
   (help "")
   (summary "")
   (http
      (code "403")
      (contact "")
      (details "")
      (format "")
      (help "")
      (summary "")
   )
(exception.user-defined.legal_reasons <---- opening user-defined exception tag, start of new user-defined exception definition
   (http <---- opening http tag
      (code "404")
      (contact "")
      (details "")
```

```
      (format <<--legal_reasons--   <---- opening format tag; the <<--exception_id-- is required
      <!DOCTYPE html>
      <html>
      <head><!--Specify the text that appears in the browser toolbar and any metadata attributes-
->
      <title>Denied due to legal policy </title>
      <meta name="author" content="organization_name">
      <meta name="description" content="description">
      <meta name="category" content="$(exception.category)">
      </head>
      <body>
      <center>
      <img src="path_to_logo_image"/><br> <!--Make sure that the appliance can access the URL-->
      <p><font face="Arial, Helvetica, sans-serif" color="blue"><b>You are about to access the
internet<br>
      INTERNET USAGE IS MONITORED AND LOGGED.</font></p> <!--Enter any text as required-->
      <p><font face="Arial, Helvetica, sans-serif" size="4">Your IP address: $(client.address)
<br>Your username: $(user.name)<br>
      YOU HAVE BEEN DENIED ACCESS TO THIS SITE. PLEASE READ THE COMPLIANCE POLICY AT: <a
href="url">http://intranet.example.com/up.html</a><br>
      This has been reported by: $(proxy.name)<br>
      <!--Use any substitution variables that are useful-->
      If you have questions about this decision, email: <a
href="mailto:YourCompany@example.com?subject=Monitored web page $(url), IP address:
$(client.address), User ID: $(user)">Compliance Team</a></font></p>
      </center>
      </body>
      </html>
      --legal_reasons--   <----the --exception_id-- is required
      ) <---- closing format tag
     ) <---- closing http tag
  ) <---- closing user-defined exception tag, end of new user-defined exception definition
) <---- closing exception.user-defined.all tag
(exception.coach
  ...
```

5. Select **Install** to save the exception page. If there are no syntax errors, the exception installs successfully:

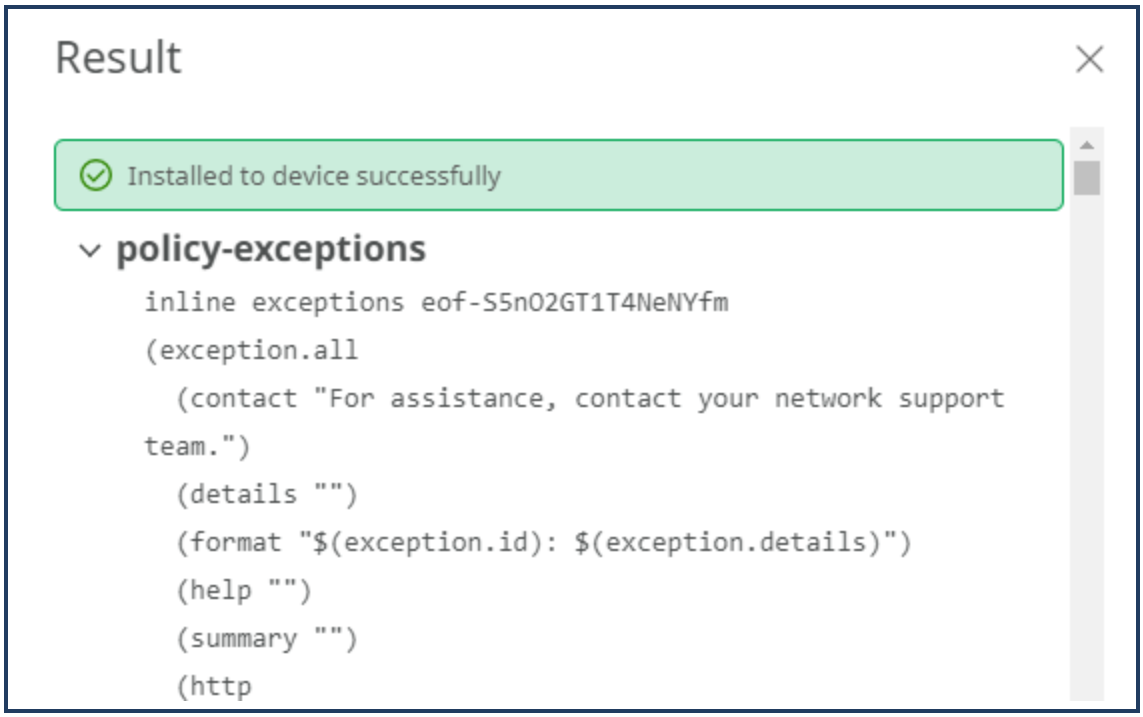

If the appliance reports syntax errors, review and correct the file before installing again.

6. Check the results of the installation (in **HTML Pages and Configuration**, select **Most Recent Install**).

   You can also preview the page as it would appear to users. In **Install**, select **Most Recent Install**. In the list of user-defined exceptions, select **View Sample HTML** for the exception you created.

7. Include the exception in a policy rule. See "Include Exceptions in Policy" on page 50 for instructions.

   The exception page is available in the **Return Exception** VPM object as follows:

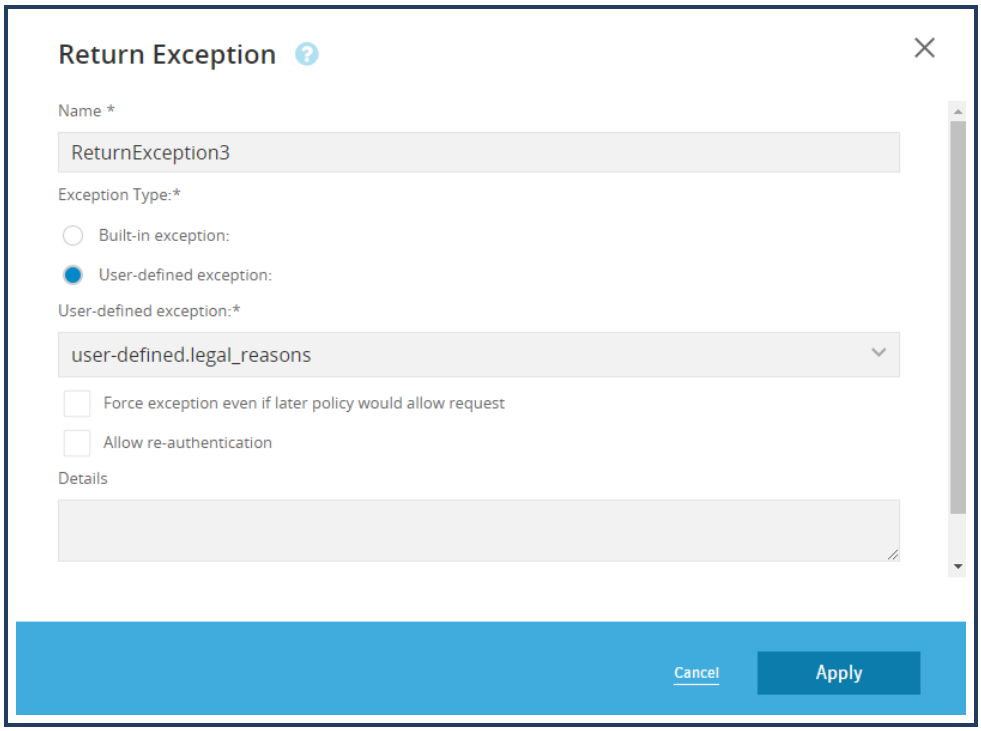

When a user request matches the policy rule referencing this exception, the appliance returns the `user-defined.legal_reasons` exception page as follows:

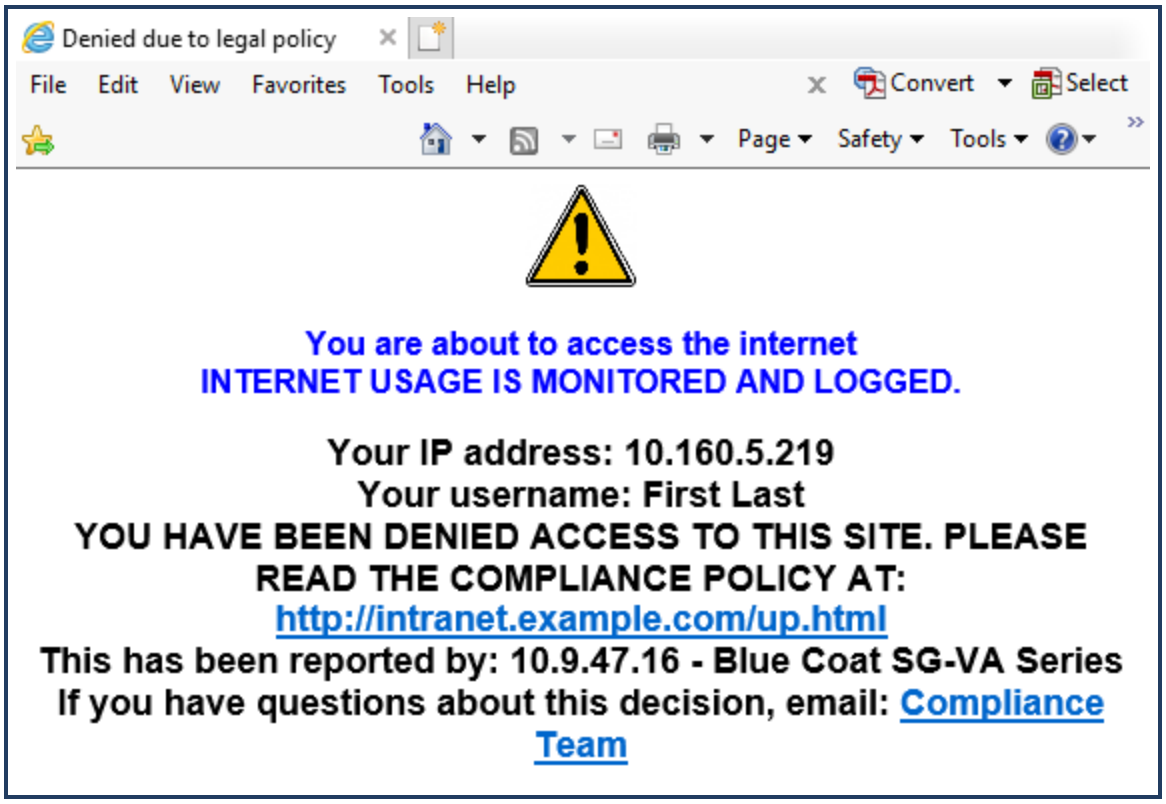

Edit an Existing Exception Page Definition

You can edit any existing built-in or user-defined exception page. This example uses the text editor to edit the contact text and logo on the `exception.all` page. It also adds a email link for user feedback.

To edit the `exception.all` exception page:

1. In the Admin Console, select **Configuration > Policy > Exception Pages**.
2. In **Install**, select **Install from text editor**.
3. In the dialog that opens, `exception.all` is at the top. Edit or add the definition as follows:

```
(exception.all
  (contact "For assistance, contact your local IT help desk.") <---- modify the contact message
  (details "")
  (format "$(exception.id): $(exception.details)")
  (help "")
  (summary "")
  (http
    (code "200")
    (contact "")
    (details "")
    (format <<--3ed3acdf.0--
```

```
<HTML><HEAD>
<TITLE>$(exception.summary)</TITLE>
</HEAD>
<BODY>
<center> <---- center the contents
<img src="path_to_logo_image"/><br> <!--Make sure that the appliance can access the URL--> <----
 add company logo
<FONT face="Helvetica">
<big><strong>$(exception.company_name)</strong></big><BR>
</FONT>
<blockquote>
<TABLE border=0 cellPadding=1 width="80%">
<TR><TD>
<FONT face="Helvetica">
<big>$(exception.summary) ($(exception.id))</big>
<BR>
<BR>
</FONT>
</TD></TR>
<TR><TD>
<FONT face="Helvetica">
$(exception.details)
</FONT>
</TD></TR>
<TR><TD>
<FONT face="Helvetica">
$(exception.help)
</FONT>
</TD></TR>
<TR><TD>
<BR>
<FONT face="Helvetica">Transaction ID: $(x-bluecoat-transaction-uuid)</FONT>
</TD></TR>
<TR><TD>
<FONT face="Helvetica" SIZE=2>
<BR>
$(exception.contact)
</FONT>
</TD></TR>
</TABLE>
</blockquote>
<!--Use any substitution variables that are useful--> <---- add user feedback option
If you have questions about this decision, email: <a
href="mailto:YourCompany@example.com?subject=Monitored web page $(url), IP address:
$(client.address), User ID: $(user)">Compliance Team</a></font></p>
</FONT>
```

```
</center>
</BODY></HTML>
--3ed3acdf.0--
    )
    (help "")
    (summary "")
)
```

4. Select **Install** to save the exception page. If there are no syntax errors, the exception installs successfully.

   If the appliance reports syntax errors, review and correct the file before installing again.

5. Check the results of the installation (in **HTML Pages and Configuration**, select **Most Recent Install**).

   You can also preview the page as it would appear to users. In **Install**, select **Most Recent Install**. In the list of user-defined exceptions, select **View Sample HTML** for the exception you created.

6. Include exceptions in a policy rule. See "Include Exceptions in Policy" on the next page for instructions.

   When a user request matches the policy rule referencing an exception page that inherits the defaults specified in `exception.all` (that is, the settings are not overridden), the appliance returns an exception page as follows:

See Also:

- " Exception Definition Overview" on page 35
- "Include Exceptions in Policy" below

Include Exceptions in Policy

Create exception policies in the VPM by configuring a Web Access Layer that includes a trigger (**Destination** object) and a selected exception (**Action** object). Other settings, such as **Source** and **Time**, are optional.

To include an exception in policy:

1. In the VPM, add or edit a **Web Access** policy rule.
2. Specify conditions that will trigger the exception, for example, using a **Destination** object. If needed, specify properties such as a **Time** object to make the rule effective only at certain times or on certain days.
3. Add the **Return Exception** action and select a built-in or custom exception type:
    a. From Exception Type, select **Built-in exception** or **User-defined exception**. See "Exception Page Types" on page 29 for descriptions of exceptions.
    b. Select an exception type from the drop-down menu. Note that custom exceptions must be defined on the system before they are available from the menu.
    c. (Optional) Select **Force exception even if later policy would allow request** to supersede other policy that applies to this request.
    d. (Optional) Add a custom message to the exception page:
        a. In **Details**, enter the message text.
        b. Select **Edit** to add ELFF and CPL strings to the exception message.

    By default, exception pages show the exception type, a general message about the policy denial, and the transaction ID:

If you specify a custom message, it replaces the default message (see the image at end of "Example: Create Exception Policy for Request Categories" below).

Example: Create Exception Policy for Request Categories

This example creates policy to deny access to shopping websites.

> **Note:** This example uses the web VPM. Steps for the legacy VPM should be similar. Use the Management Console to launch the VPM.

To create exception policy for request categories:

1. In the Management Console, select **Configuration > Policy > Visual Policy Manager**.
2. Select **Launch VPM**. The web VPM launches in a new browser tab.
3. Select **Add Layer > Web Access** to add the policy layer.

4. In the Web Access layer, add a rule and specify the request categories that will trigger the exception:

   a. In **Destination**, select **Set** and add a **Request URL Category** object.

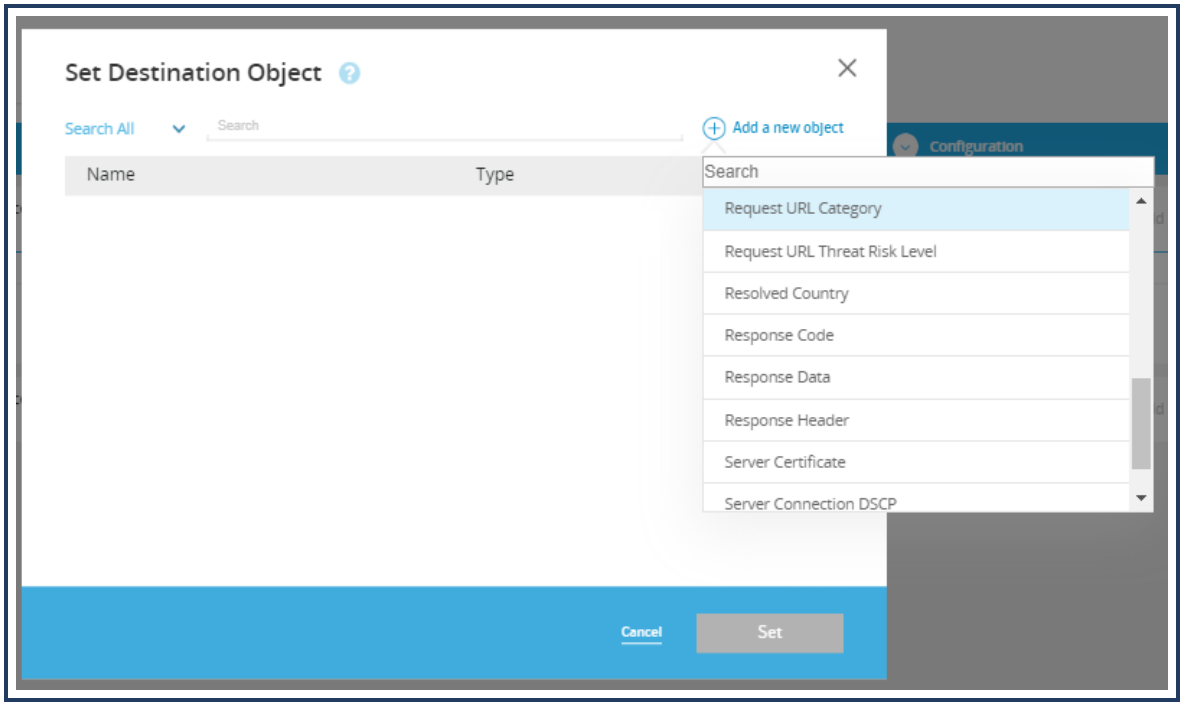

   b. In the **Request URL Category** dialog, select **Add Category**. The Add Category dialog opens.

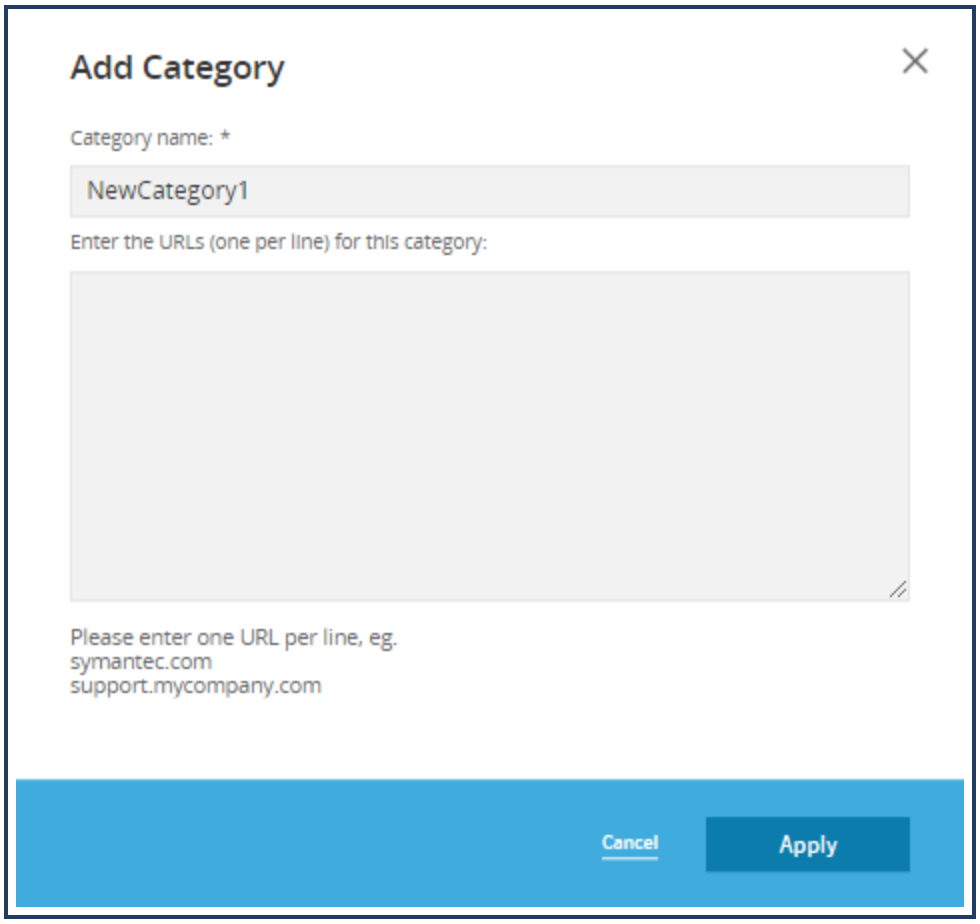

c. Name the object "Shopping sites" and add a list of URLs, one per line.

After you apply changes, the category is saved as a Policy category. Select it as shown below and then save the object.

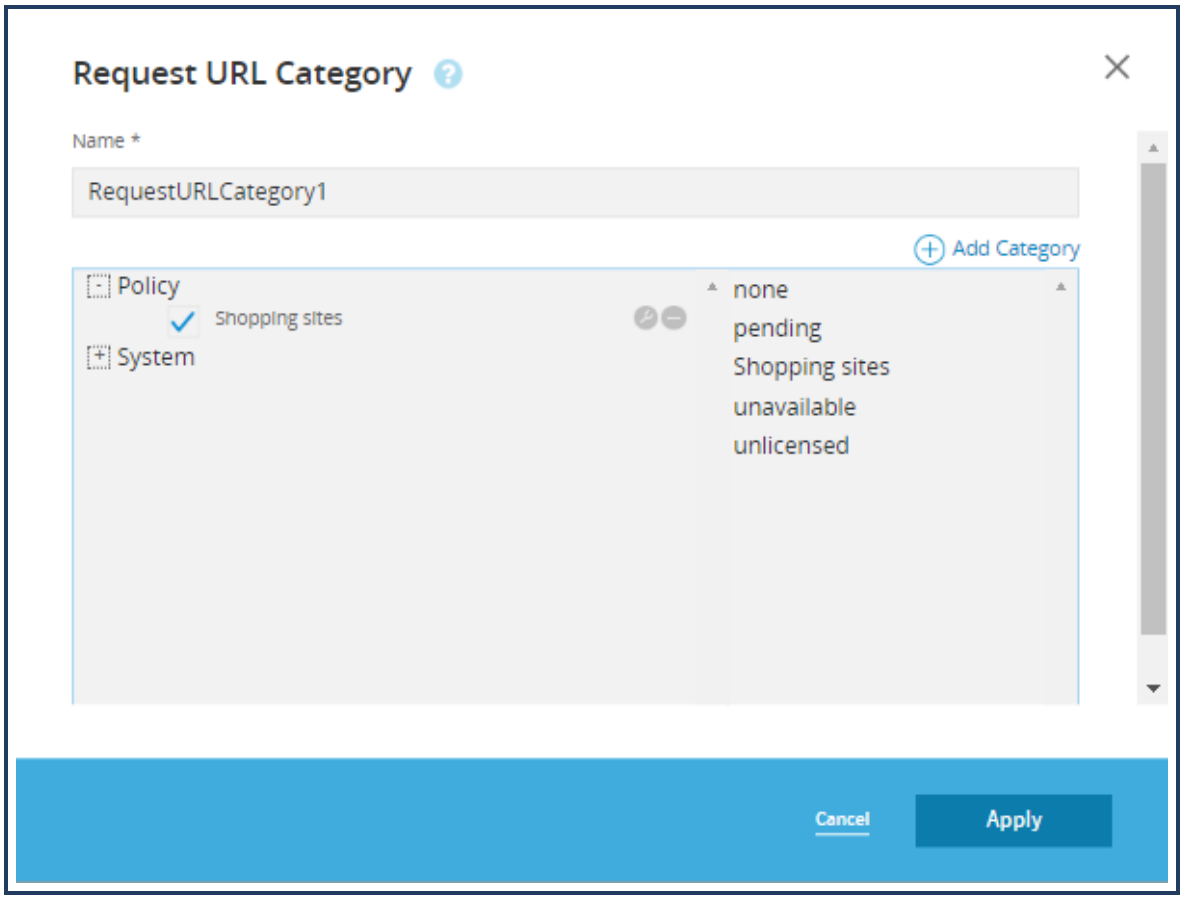

5. In the same policy rule, configure the exception to return when a requested URL matches the categories specified in the previous step:

a. In **Action**, select **Set** and add a **Return Exception** object.

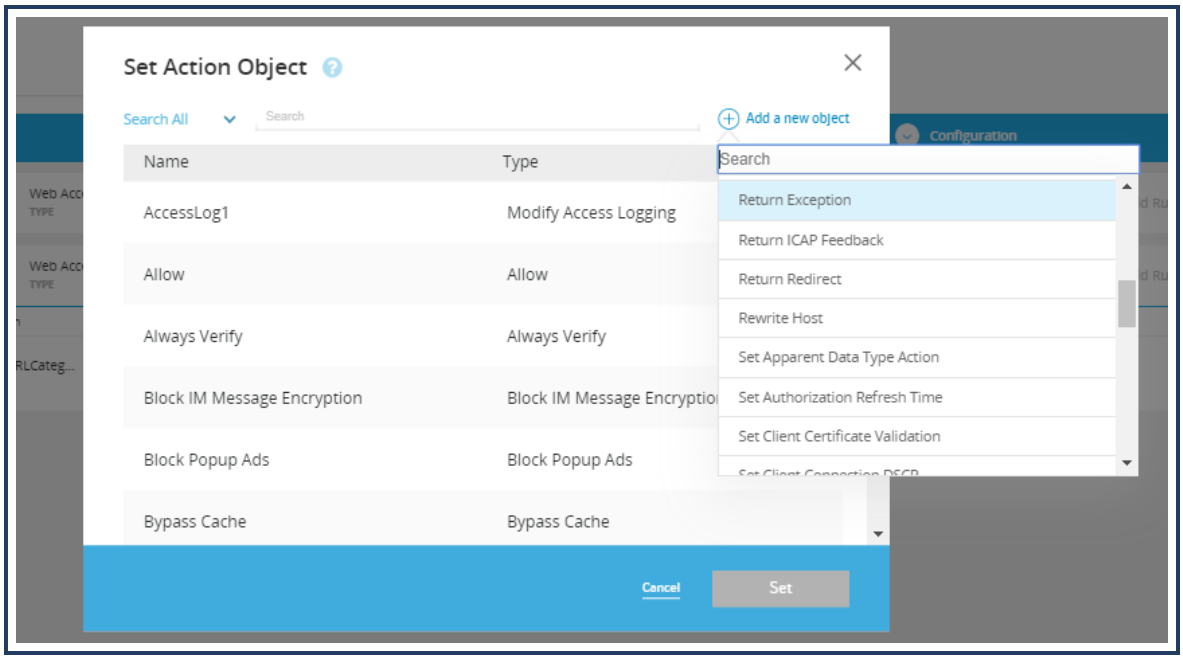

b. From Exception Type, select **Built-in exception**.

c. Select **policy_denied** from the drop-down menu.

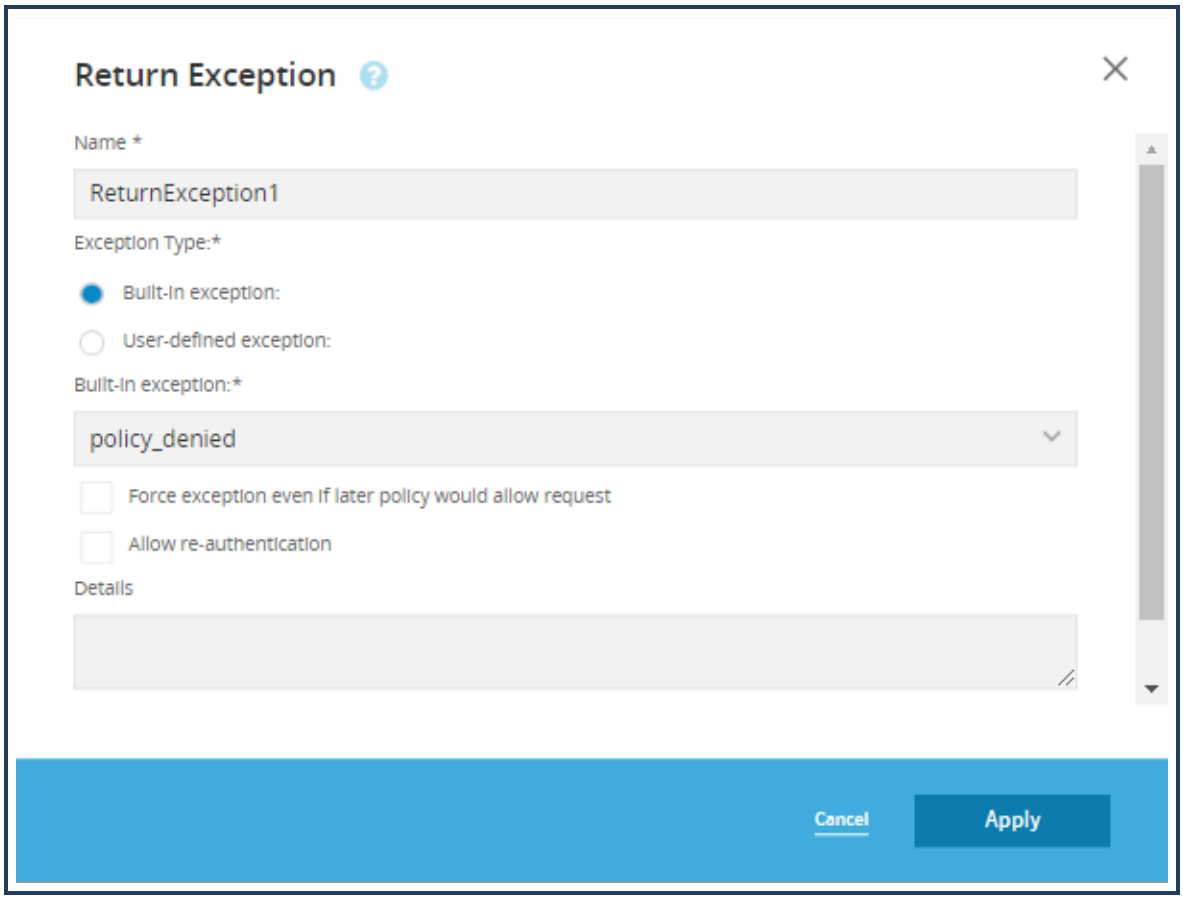

6. (Optional) Override the exception page's default settings for this policy rule only:
    a. In the **Details** field, enter the message text: "You are not permitted to access".
    b. To add an ELFF substitution to the message, select **Edit**. Select the `$(c-uri-hostname)` ELFF variable and

then apply changes.

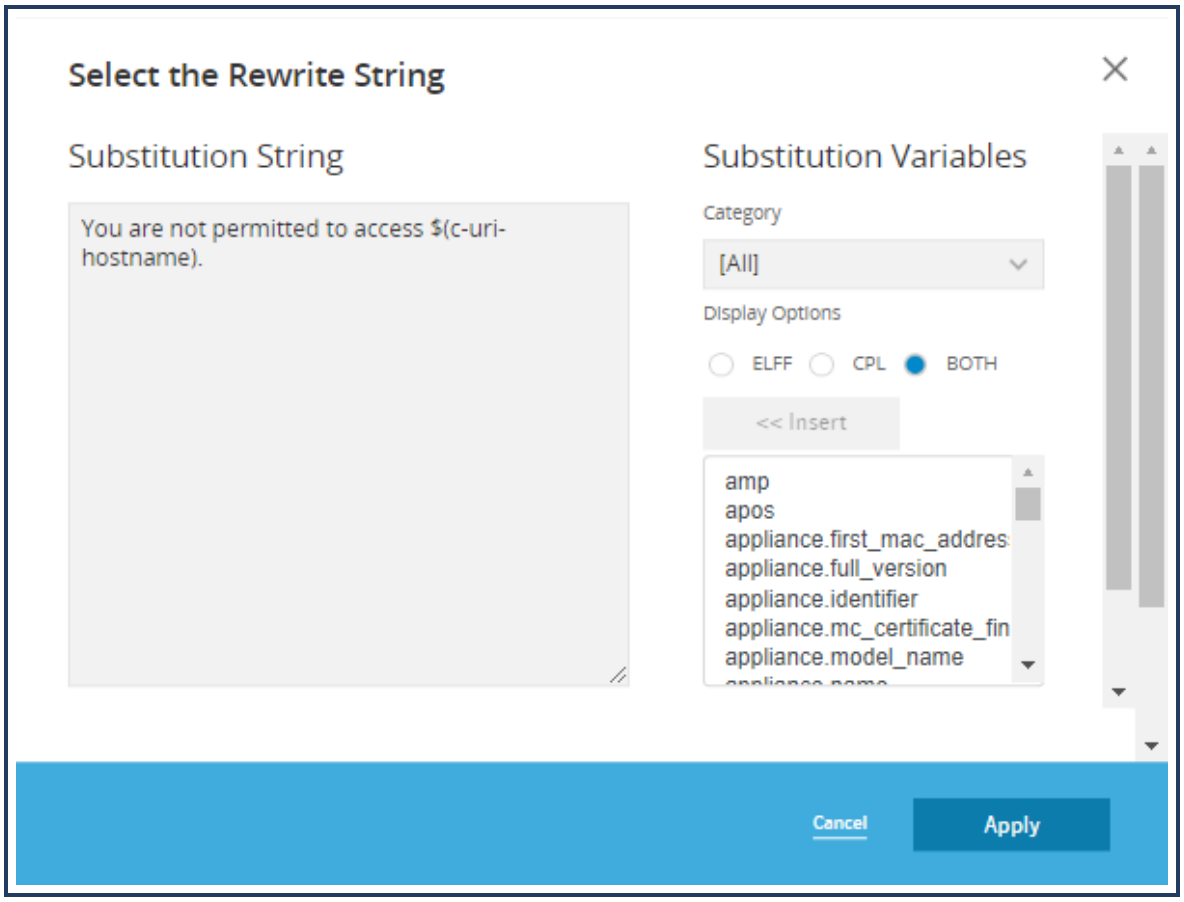

7. Install policy. Select **Apply Policy > OK**.

When users request sites that match the policy rule, they receive the following exception page:

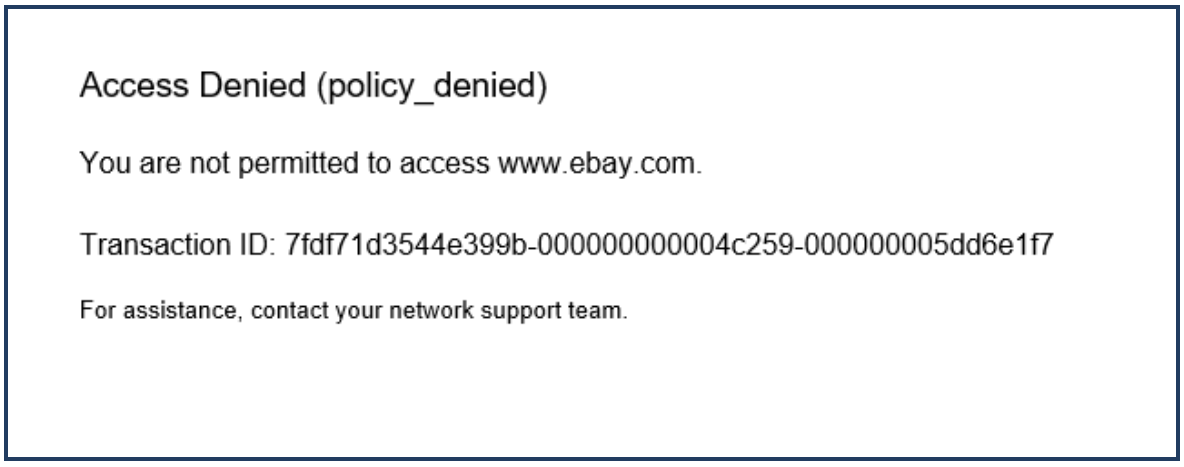

Configure Networking

Refer to the following topics to learn about network configuration on the ProxySG appliance:

- "Configure Adapters and Bridges" on page 60
  - Configure a Network Adapter
  - "Configure a VLAN" on page 66
    - "About VLAN Configuration" on page 68
  - "Interface Settings" on page 72
  - "Configure an Aggregate Interface" on page 74
  - "About Bridging" on page 76
    - "About Bridging" on page 76
    - "Bridge Settings" on page 82
    - "About Failover Mode" on page 84
- "Configure Routing" on page 87
  - "Distribute Traffic Through Multiple Default Gateways" on page 91
  - "Routing in Transparent Deployments" on page 93
  - "About Static Routes" on page 100
  - "Implement a Routing Domain" on page 103
  - "Configure Routing Information Protocol " on page 105
    - "RIP Configuration File Commands" on page 108
- "Configure DNS" on page 112
  - "Recursive and Non-Recursive DNS" on page 114
  - "Configure DNS Forwarding Groups" on page 117
  - "Configure DNS Name Imputing Suffixes" on page 121
- "About WCCP" on page 122
  - "Requirements for Configuring WCCP" on page 126
  - "Configure WCCP" on page 127

- "Configure Service Groups" on page 130

- "Configure a Forwarding System" on page 140
  - "Configure Forwarding Hosts " on page 142
  - "Configure Forwarding Lists" on page 143
  - "Set Global Defaults for Load Balancing and Host Affinity" on page 145
  - "Set the Forwarding Default Sequence" on page 147

Configure Adapters and Bridges

This section describes:

- ProxySG appliance network adapters, the adapter interfaces, and how to configure the appliance to function within a VLAN environment. Although you most likely have performed initial configuration tasks to get the appliance live on the network, you might require additional conceptual information to ensure the configuration matches the deployment requirement.
- The appliance's hardware and software bridging capabilities. Network bridging through the ProxySG appliance provides transparent proxy pass-through and failover support.

How Appliance Adapters Interact on the Network

Each appliance ships with multiple network adapters installed on the system, each with one or more interfaces (the number of available interfaces varies by appliance model).

In Symantec documentation, the convention for the interface is *adapter:interface*, for example, 0:0.

About WAN and LAN Interfaces

Recent appliance models have labels next to the physical interfaces (on the appliance backplate) that identify the WAN and LAN links. These interface labels are hard-coded and displayed in the respective interface graphics in the Admin Console. Based on your deployment type (the appliance directly in-path between users and a router or the appliance connected to a router that resides in-path, virtually in-path, and explicit), verify the following connections:

- The appliance is deployed in-path with bridging.

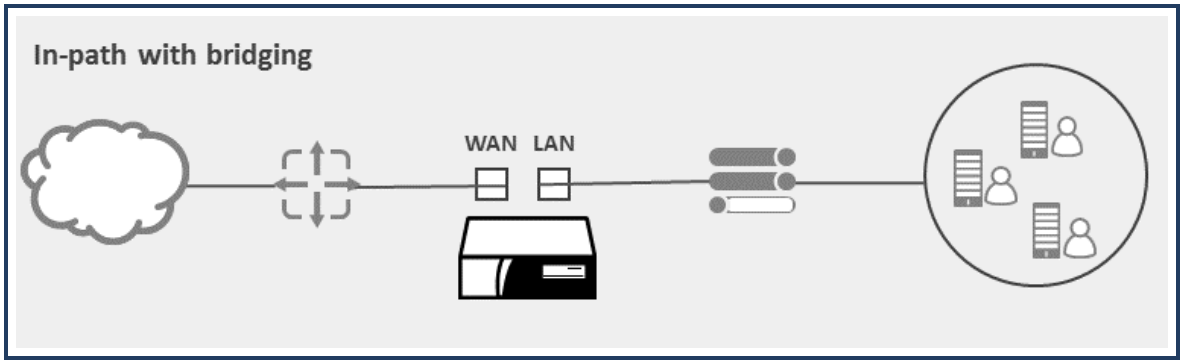

- Clients and WAN links connect to the appliance transparently through a router with WCCP.

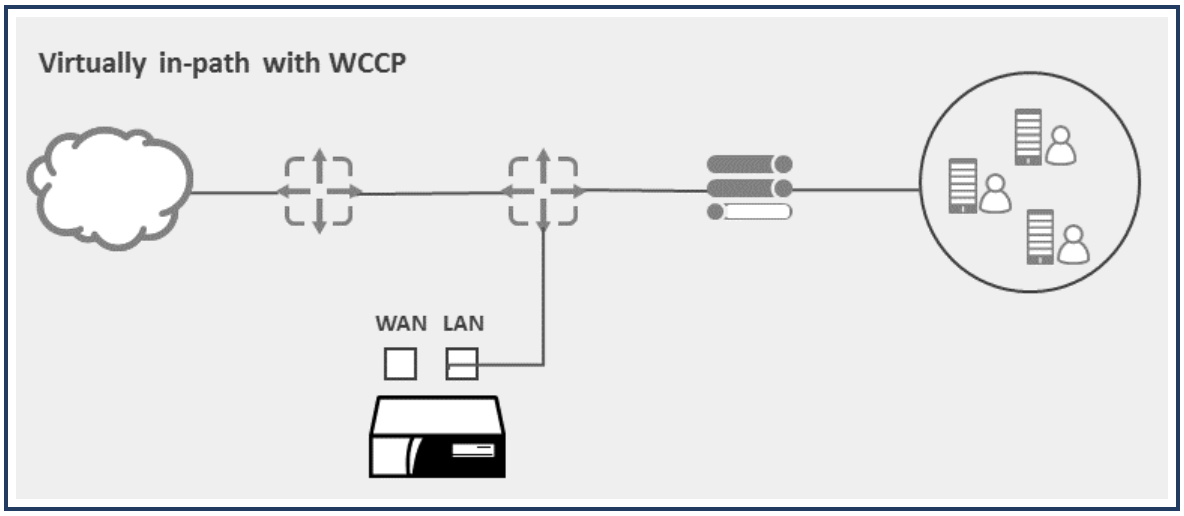

See Also:

- "Configure a Network Adapter" on page 63
- "Interface Settings" on page 72
- "About VLAN Configuration" on page 68
- "Configure a VLAN" on page 66

Improve Resiliency or Create a Bigger Pipe with an Aggregate Interface

Multiple physical interfaces may be bundled into one logical multi-gigabit aggregate interface using standard 1 GB or 10 GB physical interfaces. This provides increased throughput and network resiliency. If an interface which is part of an aggregated link goes down, its traffic will move to the other interfaces within the aggregate interface. When the interface comes back up, the traffic is redistributed across all of the links.

An aggregate interface is created on the fly when the first physical interface is added to it. Settings from the first member are applied to the aggregate link (or group). Consequent members take their common settings from the parent aggregate link. Common settings include:

- MTU size
- Reject inbound
- Allow intercept
- VLAN trunking
- Native VLAN

- Spanning tree
- IPv6 auto-linklocal

Editing VLAN settings will update the settings for that specific VLAN on all member interfaces.

Link aggregation is accomplished using the industry-standard IEEE 802.1AX Link Aggregation standard. Switch support and switch configuration are required. The switch and appliance must be cabled port-to-port.

See Also:

- "Configure an Aggregate Interface" on page 74

About Bridging

A bridge is a network device that interconnects multiple computer networks. Unlike a hub, a bridge uses the Ethernet frame's destination MAC address to make delivery decisions. Because these decisions are based on MAC addressing, bridges are known as Layer 2 devices. This Layer 2 functionality is similar to that used by switches. Bridging is especially useful in smaller deployments in which explicit proxies or L4 switches are not feasible options.

Bridging functionality allows each appliance to be easily deployed as a transparent redirection device, without requiring the additional expense and maintenance of L4 switches or WCCP-capable routers. Transparent bridges are deployed in-path between clients and routers—all packets must pass through hem, though clients are unaware of their presence.

A branch office that would take advantage of a bridging configuration is likely to be small; for example, it might have only one router and one firewall in the network, as shown below.

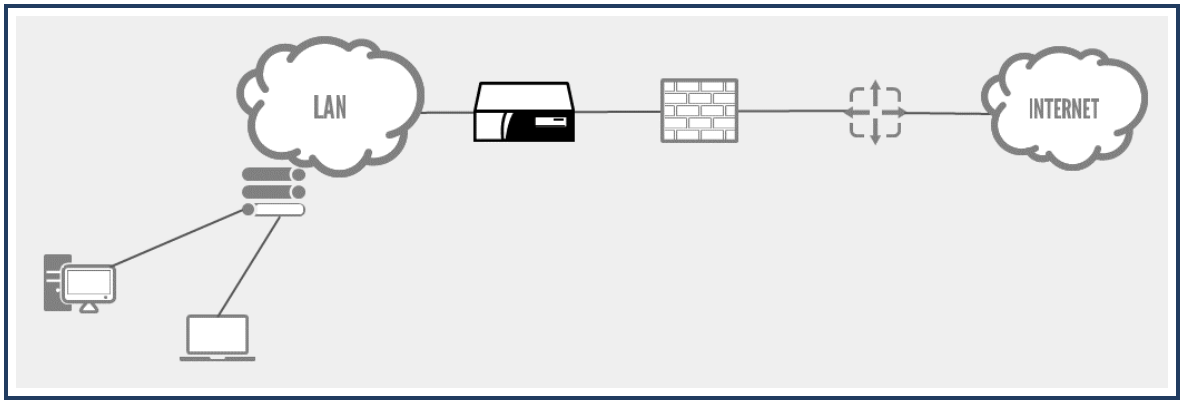

To ensure redundancy, the appliance supports both serial and parallel failover modes. See "About Failover Mode" on page 84 for more information about serial and parallel failover configurations.

See Also:

- "Configure a Software Bridge" on page 80
- "Bridge Settings" on page 82

Configure a Network Adapter

The following procedure describes how to disable, enable, or change the default adapter and interface settings for site-specific network requirements. These include inbound connection restrictions, link settings, browser/PAC file settings, and VLAN settings. Repeat the process if the system has additional adapters. By default, the appliance:

- Allows the transparent interception of inbound connections.
- Auto-negotiates link settings with the connected switch or router. Symantec recommends using auto-negotiation except under special circumstances.

> **Note:** Rejecting inbound connections improperly or manually configuring link settings improperly might cause the appliance to malfunction. Ensure that you know the correct settings before attempting either of these. If the appliance fails to operate properly after changing these settings, contact Symantec Support.

Identify Interface Details

The Interfaces section (**Configuration > Network > Adapters**) lists all configured interfaces on the appliance and displays information about each one:

Setting or field	Description
Enable	Enable or disable the interface. If you disable an interface, the console displays a warning message prompting you to confirm the action. If the interface you are disabling is in use by the Admin Console, you will see a warning message. If you lose connection to the Admin Console, reconnect with an active IP address.
Interface	Physical interface.
Label	Connection purpose or other identifying information.

Setting or field	Description
Link State	Displays information: ■ Auto/Manual: *speed* FDX/HDX ◦ Auto/Manual: Whether or not the appliance auto-negotiates with the router. ◦ *speed*: The maximum transfer speed available through the interface, depending on the type of Ethernet technology. The values are: 10 megabit/sec, 100 megabit/sec, 1 gigabit/sec, and 10 gigabit/sec. **Note:** An N/A status might indicate a network connectivity issue. ■ FDX/HDX: ◦ FDX: Full Duplex–the interface can simultaneously send and receive at the defined maximum speed (previous bullet). For example, a 100 Mbps full duplex link can send up to 100 Megabits per second (Mbps) of data and simultaneously receive up to 100 Mbps of data. ◦ HDX: Half Duplex–the interface can only send data in one direction at a time. For example, a 100 Mbps half duplex link can only send and receive a combined maximum of 100 Mbps of data.
VLANs	VLANs configured for the interface.

To configure a network adapter:

1. In the Admin Console, select **Configuration > Network > Adapters**.
2. In the Interfaces section, select the adapter and interface to configure.

   The Configure Interface page displays interface settings.
3. Enter or edit the following settings as needed:

Setting or field	Description
Label	Associate the appliance interface with the connection purpose. For example, label an interface "wansfodc" to indicate a WAN-OP connection to a datacenter concentrator in San Francisco.

Setting or field	Description
When receiving packets on this interface	Select behavior when receiving packets on inbound connection: ■ Allow transparent interception (default): The appliance intercepts the appropriate traffic based on settings configured in **Configuration > Services**; all other traffic is bridged or forwarded. ■ Bypass transparent interception: The appliance bridges or forwards all inbound traffic on this interface, regardless of the services configuration. ■ Firewall incoming traffic: The appliance drops all inbound connections on this interface, regardless of the services configuration. The default is **Allow transparent interception**. The appliance performs normal proxy interception, as configured in **Configuration > Services**, for the traffic arriving on the interface. If you require this appliance to perform interception on traffic from a specific interface or set of interfaces, set the other interfaces to either bypass the traffic (pass it through but not intercept it) or firewall it (block it completely). For more information, see " Interception Options Per Interface" on page 72.
Link settings	Select the link settings: ■ Automatically negotiate link settings (default, recommended): The appliance auto-negotiates the link settings for this interface. ■ Manually configure link settings: Select the options that meet your network requirements. This method requires a consistent configuration on the router or switch connected to this appliance. Half is not available for an aggregate interface. For more information, see " Link Settings" on page 72.

4. (Optional) Assign, change, or bind multiple IP addresses to the interface. See "Configure a VLAN" on the next page for instructions.

Configure a VLAN

You can assign multiple IP addresses to a single interface using VLANs. For more information, see "Bind Multiple IP Addresses to an Interface" on page 73.

To configure a VLAN:

1. In the Admin Console, select **Configuration > Network > Adapters**.
2. In **Interfaces**, select the adapter and interface to configure.

   The Configure Interface page displays interface settings.
3. In the **VLAN on *interface*** section, edit the Physical Interface.
4. On the **Edit Native VLAN for Interface** dialog, specify the IP address (IPv4 or IPv6) and subnet mask (for IPv4) or prefix length (for IPv6). An IPv6 address can be link-local or global.
5. Select **Add IP Address** to add the entry.
6. (Optional) Configure VLAN options:
7. a. In the **VLAN on *interface*** section, select **Add VLAN**.

   b. Enter or edit the following settings as needed:

Setting or field	Description
Label	Unique name, such as connection purpose.
VLAN ID	VLAN identifier. By default, the native VLAN ID for any appliance interface is 1, as most switches by default are configured to have their native VLAN IDs as 1. Only change the Native VLAN for Interface value if the native VLAN ID of the switch or router connected to this interface is a value other than 1; match that value here.
When receiving packets on this VLAN	The receiving packet and browser behavior is the same as for physical interfaces with the exception of **Use physical interface**, which applies the same configuration to the VLAN as was set on the physical interface. The default setting is to inherit the behavior from the physical interface.
IP Addresses	VLANs on the interface. To add more VLANs to the interface, enter the IP address and subnet mask.

> **Note:** If you attempt to edit a VLAN with an IP address being used by the Admin Console you will see a message warning of loss of connectivity. If you lose connection to the Admin Console, reconnect with an active IP address.

See Also:

- Configure a Network Adapter

About VLAN Configuration

Virtual LANs (VLANs) are logical network segments that allow hosts to communicate, regardless of physical network location. The benefit to this is that clients can be separated logically—based on organizational unit, for example— rather than based on physical connectivity to interfaces. The appliance treats VLAN interfaces identically to traditional physical LAN interfaces.

VLAN segments are defined on the switch. The network administrator specifies which ports belong to which VLANs. The following diagram illustrates a port-based VLAN configuration. Clients on network segments attached to switch ports 1 and 2 belong to VLAN 1, which has the network address 10.0.1.x; network segments attached to switch ports 14 and 15 belong to VLAN 2, which has the network address 10.0.2.x.

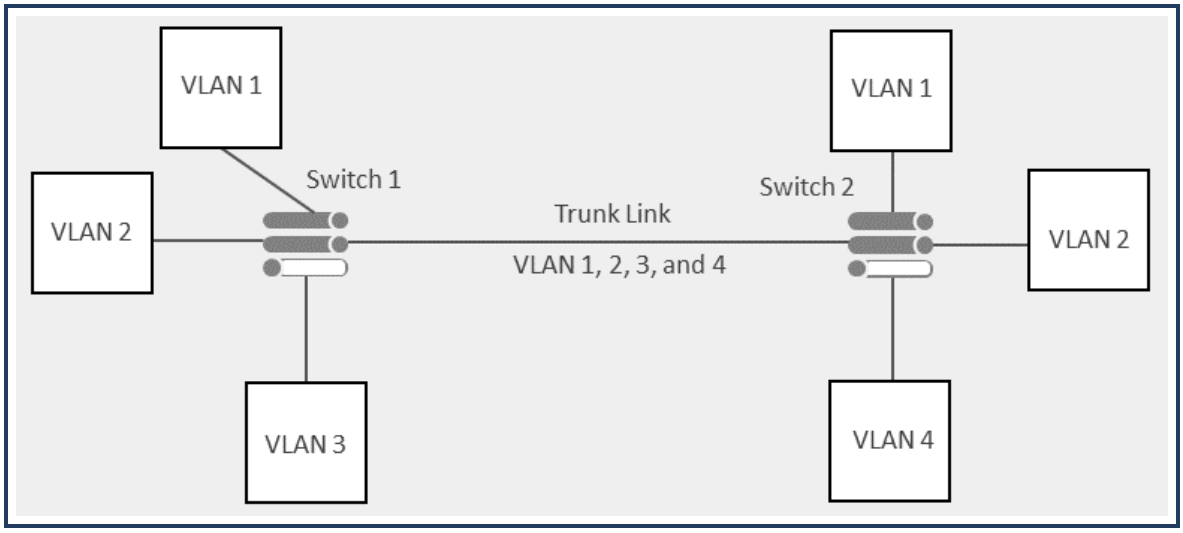

As also illustrated in the diagram, clients of different OS types can reside within a VLAN. Not all clients are able to detect (send or receive) VLAN-tagged packets.

VLAN Trunking

Trunk ports are ports that carry traffic for more than one VLAN. They tag each packet with the VLAN ID in the packet header. Trunk ports are commonly used between switches and routers that must switch or route traffic from or to multiple VLANs. By default, VLAN trunking is enabled on the appliance.

In the following diagram, multiple VLANs are connected by a trunk link between two switches.

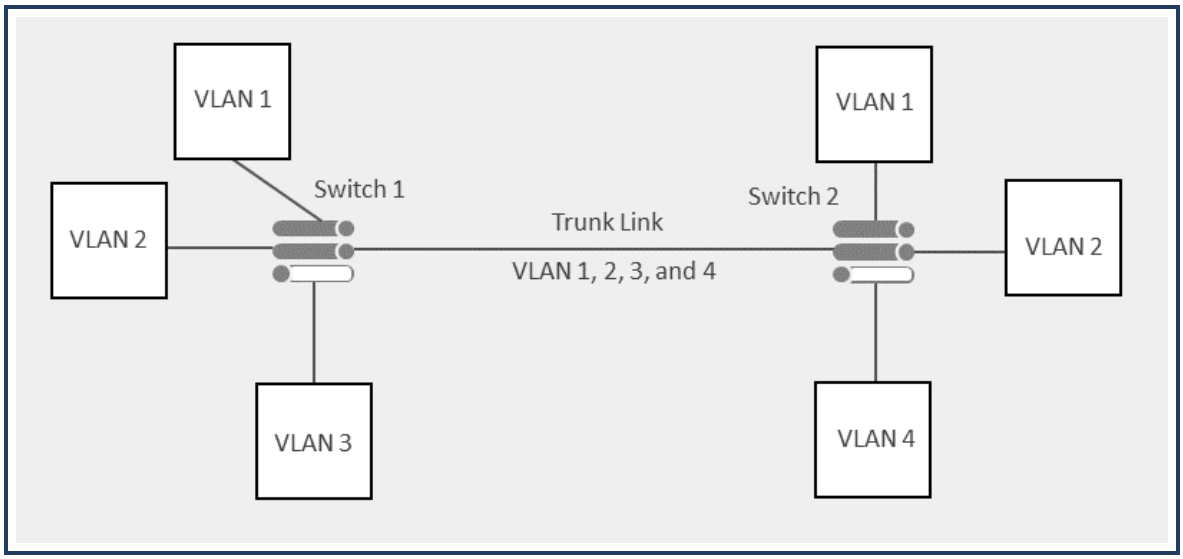

Native VLANs

Each switch port has a designated native VLAN. Traffic on the port associated with the native VLAN is not tagged. Traffic destined for VLANs other than the native VLAN is tagged.

The trunk link carries both the native VLAN and all other VLAN (tagged) packets, as illustrated in the following diagram.

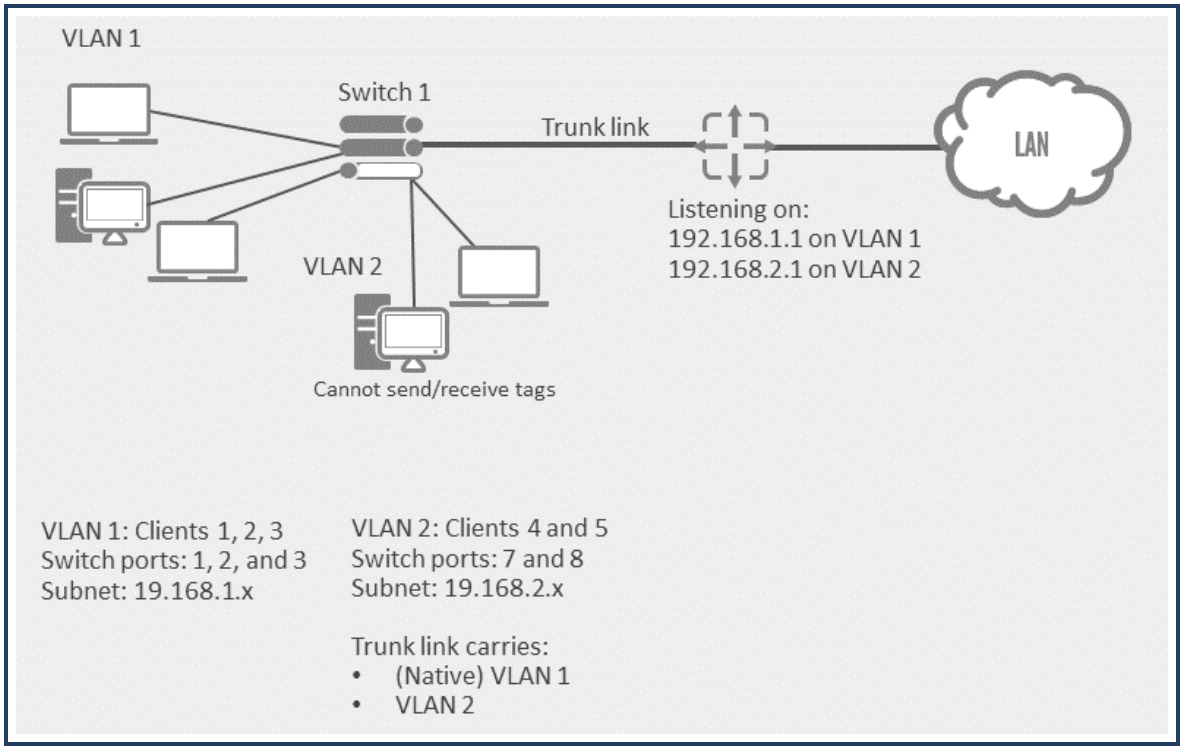

In this example, the client attached to port 7 belongs to VLAN 2. Even though port 7 is part of VLAN 2, it does not set tags or receive VLAN-tagged packets. The switch associates the traffic with VLAN 2 and tags it accordingly when appropriate. Conversely, it strips the VLAN 2 tag on the response. The trunk link carries VLAN 1 (the native) and 2 traffic to a router that forwards traffic for those VLANs.

Deployment complications arise when a device (other than a router) is required between switches. Any network device without VLAN-tagging support might drop or misinterpret the traffic.

As a best practice, do not deploy a device that is not configured to recognize VLAN-tagged traffic in-path of a trunk link.

> **Note:** In Symantec documentation, the convention for VLAN is adapter:interface.VLAN_ID. Example: 1:0.10 refers the VLAN ID 10 on adapter 1, interface 0.

ProxySG VLAN Support

The appliance supports VLAN tagging and it is enabled by default; therefore, an appliance can be deployed in-path with switches that are exchanging VLAN tagged traffic. This allows for uninterrupted VLAN service, plus enables benefits gained with the proxy features.

The Admin Console enables you to configure VLAN interfaces the same way you configure physical interfaces. After a VLAN is added, it appears in the list of network interfaces. Settings such as allow-intercept and reject-inbound are applicable to VLAN interfaces.

The most common deployment is an appliance residing between two switches or a switch and a router; in these cases, preserving tagged packets is essential to correct network operation.

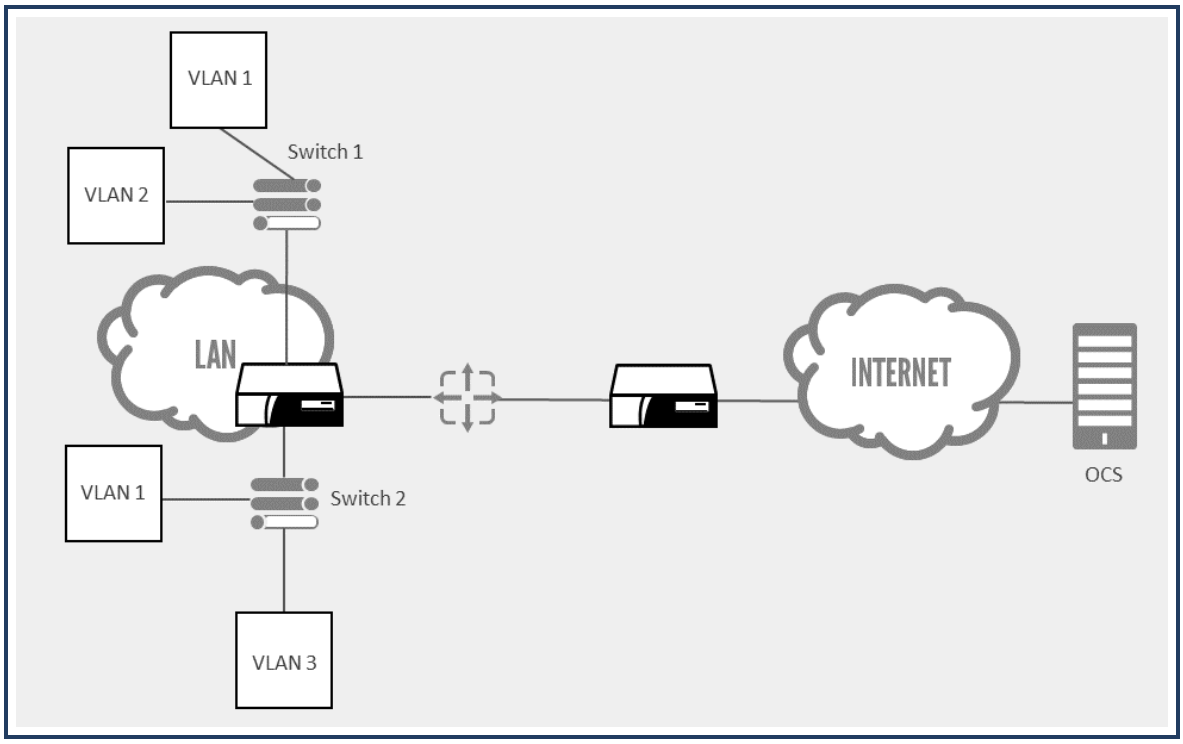

Based on this deployment:

- If configuration and policy allow, the appliance accepts all packets regardless of their VLAN tag and passes them from one interface to the other with the original VLAN tag preserved.
- If a packet arrives on one interface tagged for VLAN 2, it remains on VLAN 2 when it is forwarded out on another interface. If a packet arrives untagged and the destination interface has a different native VLAN configured, the appliance adds a tag to ensure the VLAN ID is preserved. Similarly, if a tagged packet arrives and the VLAN ID matches the native VLAN of the destination interface, the appliance removes the tag before transmitting the packet.
- The appliance strips the native VLAN tag on all outgoing traffic.

About Bridging and VLANs

On the appliance, bridges can be created between two physical interfaces only. If you have configured virtual interfaces (VLANs), all the VLANs on the selected physical interfaces will be bridged.

Although VLANs are supported on bridges, the appliance does not support creating a bridge group between VLANs when bridging or bypassing traffic. For example, when bridging you cannot send packets from VLAN 0:0.2 to 0:1.3.

For information on bridging, see "Configure a Software Bridge" on page 80.

See Also:

- "Configure a VLAN" on page 66

Interface Settings

Refer to the following information when configuring an interface:

Interception Options Per Interface

You can execute one of three actions upon intercepting traffic on a per-interface basis:

- Allow: Bridge/forward traffic and intercept appropriate traffic as defined by Proxy Services.
- Bypass: Bridge/forward all traffic without interception.
- Firewall: Drop (silently block) any traffic not related to established appliance connections.

The following table describes what effect each interception option setting has on different traffic types.

Option	reject-inbound	allow-intercept	ProxySG management and console connections	Explicit proxy service traffic	Transparent proxy service traffic	Other traffic
Allow	Disabled	Enabled	Intercepted	Intercepted	Intercepted	Forwarded
Bypass	Disabled	Disabled	Intercepted	Intercepted	Forwarded	Forwarded
Firewall	Enabled	Enabled/Disabled	Silently dropped	Silently dropped	Silently dropped	Silently dropped

By default, the intercept behavior is to bypass transparent interception.

Link Settings

By default, the appliance auto-negotiates the interface speed and duplex settings with the switch or router to which it is connected.

- The appliance supports multiple Ethernet modes. The speed setting is the maximum transfer speed, in Megabits or Gigabits per second (Mbps/Gbps), the interface supports.
- The duplex setting designates two-way traffic capabilities. In Full duplex mode, both devices may transmit to and from each other simultaneously, allowing each direction to use the maximum transfer speed without affecting the other direction. In Half duplex mode, only one device may transmit at any one time, effectively sharing the maximum transfer speed of the interface.

The appliance's health monitoring capability provides alerts if interface use reaches warning and critical capacity levels. In Full duplex mode, the appliance reports the larger percentage value of the sending and receiving values. For example, if the appliance is receiving 20 Mbps and sending 40 Mbps on a 100 Mbps-capable interface, the reported value is 40%. If the same interface was set to half duplex, the reported value is 60%, or the aggregated values.

Symantec strongly recommends using the (default) auto-negotiation feature. The key issue is the appliance settings must match the settings on the switch; therefore, if you manually change the settings on the appliance, you must also match the settings on the router or switch.

> **Note:** When the 100 Mbps Ethernet interfaces on the appliance 210 are connected to Gigabit Ethernet capable devices, they might incorrectly auto-negotiate when fail-open pass-through is used. If both the interfaces on these appliances are connected to Gigabit capable switches or hubs, Symantec recommends that you configure the link settings manually to 100 Mbps. To configure the link settings, see "Configure a Network Adapter" on page 63.

The following table lists the results of various appliance and router link settings for 100 Mbps speeds. Result values are listed in format speed/duplex.

Router/switch auto-negotiation result	Router/switch interface settings	ProxySG interface setting	ProxySG auto-negotiation result
100/Full Duplex	Auto	Auto	100/Full Duplex
N/A	100/Full Duplex	Auto	100/Half Duplex
N/A	100/Full Duplex	100/Full Duplex	N/A
100/Half Duplex	Auto	100/Full Duplex	N/A

The following table lists the results of various appliance and router link settings for 1 Gbps speeds. The values are listed in format speed/duplex.

Router/switch auto-negotiation result	Router/switch interface settings	ProxySG interface setting	ProxySG auto-negotiation result
No link	Auto	Gig/full duplex	No link
Gig/full duplex	Auto	Auto	Gig/full duplex
No link	Gig/full duplex	100/Full Duplex	No link

Bind Multiple IP Addresses to an Interface

The appliance allows you to bind multiple IP addresses to an interface, and typically, the assigned IP addresses are on the same subnet. Multiple IP addresses on an interface allows for managing one service under a specific IP and another service under a different IP. For example, you can assign one IP address for management services/console access and another IP address for managing proxy traffic. In addition, you could assign unique IP addresses to manage different services, that is have HTTP traffic on one and native FTP on another.

When using the appliance in a mixed IPv4/IPv6 environment, you should assign IPv4 and IPv6 addresses to each interface. The IPv6 address can be link-local or global.

Configure an Aggregate Interface

In the Admin Console (**Configuration > Network > Adapters**), the **Aggregate Interface** section lists all configured aggregate interfaces on the appliance and displays information about each one:

Setting or field	Description
Enable toggle	If disabled, the aggregate link interface will disable the physical interface of each member of the aggregate group. Individual physical interfaces within the group can be enabled or disabled.
ID	Identifier for the aggregate interface.
Label	Label for the aggregate interface.
Interface	Interfaces that are members of the aggregate interface.
LACP State	Possible states: ■ Up: The member is healthy and operates normally from LACP perspective. ■ Synchronizing: Peers are out of sync with the port, or unable to exchange LACP PDUs. ■ Negotiating: Exchanging key information with the peer. If it persists, this might indicate that the peer is not in the correct link aggregation, or that other configuration on the switch port differs from the rest of the aggregation. ■ Suspended: The port is not being used by LACP. Potential reason is port is in half-duplex. ■ Disabled: The physical interface is disabled. ■ Down: The physical interface has no link. Link aggregation should work on any switch configured to use LACP; consult the documentation from your switch vendor when configuring link aggregation. LACP standby link selection and dynamic key management are not supported.
VLANs	VLANs that have been added to the aggregate interface.

To add an aggregate interface or edit an existing aggregate interface:

1. In the Admin Console, select **Configuration > Network > Adapters**.
2. In Aggregate Interfaces:
    a. Select **Add Aggregate** to add an aggregate.
    b. Select an existing aggregate in the list or select Edit under Actions.

    The Add New Interface/Configure Interface page displays aggregate interface settings.

3. Enter or edit the following settings as needed:

Setting or field	Description
Label	Label for the aggregate interface.
Identifier	Identifier for the aggregate interface.
Select from available interfaces	Add or remove aggregate members. Note the following: ■ Only available interfaces are displayed. ■ Up to 32 interfaces can be added to an aggregate interface. ■ The MAC address of the first physical interface assigned to the aggregate interface becomes the MAC address of the aggregate link. ■ Interfaces from different adapters may belong to the same aggregate link. ■ Interfaces with different speeds are allowed in an aggregate interface. ■ The Link State column provides link information such as Enable requested, Disabled, Auto <negotiated speed, duplex>, and so on.
VLANs on interface	This information is for viewing only. To configure VLANs, see "Configure a Network Adapter" on page 63. When a physical link is added to an aggregate interface, the VLAN configurations are merged. A VLAN on the new physical interface is created on the aggregate, and the existing configuration is copied to the aggregate and all other group members; if a VLAN exists on the aggregate interface, it will be created on the new physical interface. If both a new member and the aggregate interface have a VLAN configuration, all VLAN settings are copied from the aggregate VLAN to the member VLAN, except the IP address.

Notes

- Packets for any given connection will always be transmitted over the same physical link, limiting the burst speed for the connection to the capacity of that specific link rather than the sum of all the links.
- Link aggregation and bridging can’t be configured on the same physical interface, though they can both be used in a single deployment.

About Bridging

Review the following bridging concepts:

Bridging Methods

The appliance provides bridging functionality by two methods:

- Software–A software, or dynamic, bridge is constructed using a set of installed interfaces. Within each logical bridge, interfaces can be assigned or removed. In the event of failure, software bridges fail closed–traffic is not passed. This behavior can be desirable if you want to pass traffic to a redundant appliance and/or link.
- Hardware–A hardware, or pass-through, bridge uses a dual interface Ethernet adapter. This type of bridge provides pass-through support–in the event of failure, traffic passes through the appliance.

See "About the Pass-Through Adapter" on the facing page for more information.

> **Note:** If you want to use an L4 switch or an explicit proxy instead of bridging, you must disable the pass-through card.

Traffic Handling

Bridges are used to segment Ethernet collision domains, thus reducing frame collisions. To make efficient delivery decisions, the bridge must discover the identity of systems on each collision domain. The bridge uses the source MAC address of frames to determine the interface that the device can be reached from and stores that information in the bridge forwarding table. When packets are received, the bridge consults the forwarding table to determine which interface to deliver the packet to. The only way to bypass the bridge forwarding table lookup is to define a static forwarding entry. For more information on static forwarding entries, see "About Static Forwarding Table Entries" on page 83.

Trust Destination MAC

When the appliance is in transparent bridging mode, the appliance always “trusts” the destination MAC address of inbound packets and does not consult its routing table. Trust Destination MAC is enabled by default (when the appliance is in transparent bridging mode) and cannot be disabled. For more information on Trust Destination MAC, see "About Trust Destination MAC" on page 93.

About Bridging and Policy

Because the bridge intercepts all traffic, you can take advantage of the powerful proxy services and policies built into the appliance to control how that traffic is handled. If the appliance recognizes the intercepted traffic, you can apply policy to it. Unrecognized traffic is forwarded out. The following diagram illustrates this traffic handling flow.

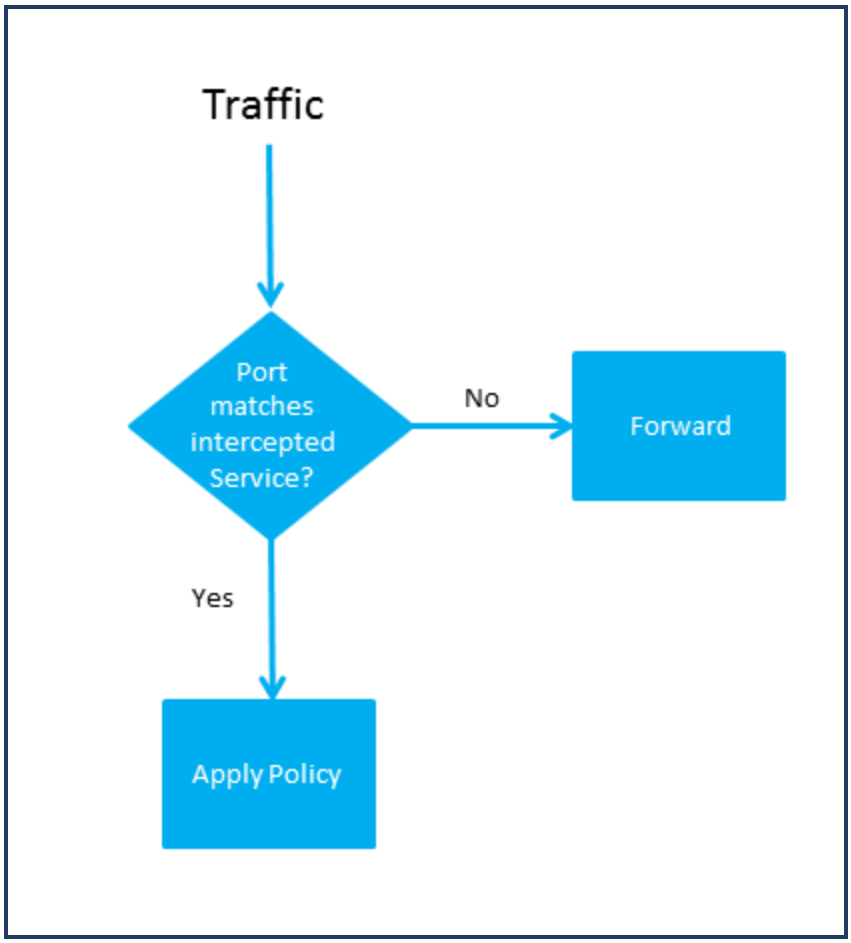

Because policy can be applied only to recognized protocols, it is important to specify port ranges that will capture all traffic, even that operating on lesser known ports.

About the Pass-Through Adapter

A pass-through adapter is a dual interface Ethernet adapter designed to provide an efficient fault-tolerant bridging solution. If this adapter is installed on an appliance, SGOS detects the adapter on system bootup and automatically creates a bridge–the two Ethernet interfaces serve as the bridge ports. If the appliance is powered down or loses power for any reason, the bridge fails open; that is, network traffic passes from one Ethernet interface to the other. Thus, network traffic is uninterrupted, but does not route through the appliance.

Caution: This scenario creates a security vulnerability.

After power is restored to the appliance, the bridge comes back online and network traffic is routed to the appliance and thus is subject to that appliance's configured features, policies, content scanning, and redirection instructions. Bridging supports only failover; it does not support load balancing.

> **Note:** To view adapter state, in the Admin Console, select **Configuration > Network > Adapters**.

Deployment Recommendations

Rack and cable the appliance while it is powered off. This enables you to confirm that the pass-through adapter is functioning and that traffic is passing through the appliance. If traffic is not being passed, confirm that you have used the correct cabling (crossover or straight).

Reflecting Link Errors

When the appliance is deployed transparently with bridging enabled, link errors that occur on one interface can be reflected to the other bridge interface. This allows a router connected to the appliance on the healthy link to detect this failure and recompute a path around this failed segment. When the interface with the original link error is brought back up, the other interface is automatically restarted as part of the health check process.

Reflecting link errors requires that two interfaces be available and connected in a bridging configuration; it also requires that **Propagate Failure** be enabled. By default, it is disabled.

> **Note:** This feature is only applicable to a two-interface hardware or software bridge. **Propagate Failure** sets itself to disabled in any other scenario.

If the link goes down while **Propagate Failure** is disabled, the previous link state is immediately reflected to the other interface if **Propagate Failure** is enabled. See "Configure a Software Bridge" on page 80 for details.

Bypass List Behavior

Static and dynamic bypass operate depending on how the appliance intercepts the traffic, as follows:

- When the appliance is installed in a bridging deployment, bridging is used for bypass.
- When the appliance is installed as a router or external layer 4 load balancers are used to redirect traffic to the appliance, routing is used for bypass, but only if IP Forwarding is enabled. Otherwise, traffic is dropped instead of being bypassed.
- When the appliance is installed in a WCCP deployment, either GRE or layer 2 redirection is used for bypass.

SGOS uses WCCP packet return to redirect bypassed traffic back to the router, supporting the following combination of packet forwarding and return options:

Packet forwarding	Packet return
GRE	GRE

Packet forwarding	Packet return
L2	GRE
L2	L2

To set these options in the Admin Console, see "Configure Service Groups" on page 130.

Configure a Software Bridge

This section describes how to link adapters and interfaces to create a network bridge.

Before configuring a software bridge, ensure that your adapters are of the same type and use the same settings. Although the software does not restrict you from configuring bridges with adapters of different speeds and MTU configurations (for example, ports speeds of 10/100 Mbit/s and 1 GigE combined with an MTU of 1400 and 1500, respectively), the resulting behavior is unpredictable.

> **Note:** The Bridges section (**Configuration > Network > Adapters**) displays the bandwidth class used for all bridges. After you have created and configured a bandwidth management class for bridging, you can configure bridging bandwidth allocation. See "Bandwidth Management for Bridging" on page 82 for details.

To configure a software bridge:

1. In the Admin Console, select **Configuration > Network > Adapters**.
2. In **Bridges**:
   - To add a bridge, select **Add Bridge**.
   - To edit an existing bridge, select it.

     The Add New Bridge/Configure Bridge page displays bridge settings.
3. Enter or edit the following settings as needed:

Setting or field	Description
Name	Name of the bridge.
Failover Group	(Optional) Assign the bridge to a failover group. Failover groups are configured in the CLI. See "About Failover Mode" on page 84.
Propagate Failure	If you have a two-interface bridge, select to enable link error propagation.
Clear Statistics	(Editing an existing bridge only) Select to reset the traffic history of the bridge, which includes packet and byte counts, to 0.

Setting or field	Description
Interfaces	Configure the interface settings to: ■ Allow transparent interception. It is bypassed by default. You must configure the WAN interface to allow transparent interception. ■ Firewall incoming traffic. Firewalls must be specifically configured. See "Configure a Network Adapter" on page 63 for details. To configure bridge interface settings: 1. Select **Add** to add an interface to the bridge, or edit an existing bridge interface. 2. Enter or edit the bridge settings: ■ **Spanning Tree**: Select to enable bridging loop avoidance. See "About Bridging Loop Detection" on page 83 for details. ■ **Static Forwarding Table**: If you are using firewall configurations that require the use of static forwarding table entries, add a static forwarding table entry that defines the next hop gateway that is on the correct side of the bridge. See "About Static Forwarding Table Entries" on page 83 for details. ■ **MAC Address**: MAC addresses for the next hop gateways.

Bridge Settings

Refer to the following information when configuring a bridge:

> **Note:** For failover settings, see "About Failover Mode" on page 84.

Bandwidth Management for Bridging

After you have created and configured a bandwidth management class for bridging, you can manage the bandwidth used by all bridges.

> **Note:** This setting only controls the bandwidth class used by bypassed traffic. To manage intercepted traffic, you must define bandwidth management policy (using VPM or CPL).

To set bandwidth management for bridging:

1. In the Admin Console, select **Configuration > Network > Adapters**.
2. In **Bridges**, select one of the following from the Bandwidth Class menu:
   - A bandwidth class for bridging.
   - **None** to disable bandwidth management for bridging.

Programmable Pass-Through/NIC Adapters

Some ProxySG appliances ship (when ordered) with a network adapter card that can be used as a pass-through adapter or as a NIC, depending on the configured mode. If the network adapter mode is set to disabled, the adapter interfaces can be used as NICs or as part of a software bridge.

If your appliance includes a programmable adapter card, the following programmable adapter modes are available:

- Disabled: Disables the bridge and allows the adapter interfaces to be reused as NICs or as part of another bridge.
- Fail Open: If the appliance fails, all traffic passes through the bridge so clients can still receive data.
- Fail Closed: If the appliance fails, all traffic is blocked and service is interrupted. This mode provides the same functionality as a user-configured software bridge.

> **Note:** If you create a software bridge, the programmable bridge card mode is implicitly Fail Closed (if the appliance fails, the software bridge is non-functional).

About Bridging Loop Detection

Bridging supports the Spanning Tree Protocol (STP). STP is a link management protocol that prevents bridge loops in a network that has redundant paths that can cause packets to be bridged infinitely without ever being removed from the network. STP ensures that a bridge, when faced with multiple paths, uses a path that is loop-free. If that path fails, the algorithm recalculates the network and finds another loop-free path. The administrator can enable or disable spanning tree participation for the interface.

About Static Forwarding Table Entries

Certain firewall configurations require the use of static forwarding table entries. These firewall failover configurations use VIP addresses and virtual MAC (VMAC) addresses. When a client sends an ARP request to the firewall VIP, the firewall replies with a VMAC (which can be an Ethernet multicast address); however, when the firewall sends a packet, it uses a physical MAC address, not the VMAC. The solution is to create a static forwarding table entry that defines the next hop gateway that is on the correct side of the bridge.

About Failover Mode

In failover mode, two appliances are deployed, a master and a slave. The master sends keepalive messages (advertisements) to the slave appliance. If the slave does not receive advertisements at the specified interval, the slave takes over for the master. When the master comes back online, the master takes over from the slave again.

The SGOS bridging feature allows two different types of failover modes, parallel and serial. Hardware and software bridges allow different failover modes:

- Software bridges allow serial or parallel failover. Note that if the appliance fails, serial failover also fails.
- Hardware bridges allow serial or parallel failover.

Parallel Failover

In parallel failover mode, two systems are deployed side by side on redundant paths. In parallel failover, the slave does not actively bridge any packets unless the master fails. If the master fails, the slave takes over the master IP address and begins bridging. A parallel failover configuration is shown in the following figure.

Because of the redundant paths, you must enable **Spanning Tree** to avoid bridge loops. See "About Bridging Loop Detection" on the previous page for more information about STP.

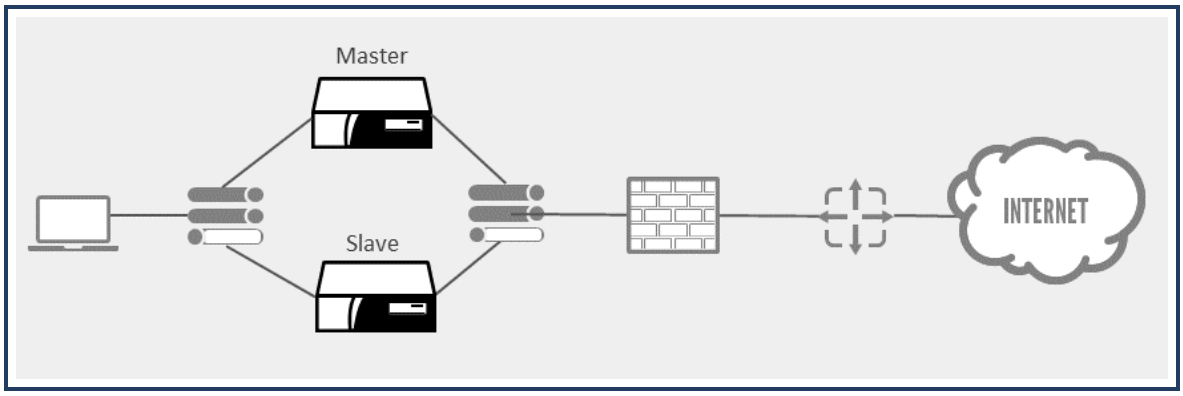

Serial Failover

In serial failover mode, the slave is in-path and continuously bridges packets, but does not perform any other operations to the bridged traffic unless the master fails. If the master fails, the slave takes over the master IP address and applies policy, etc. A serial configuration is shown in the following figure.

If you are relying on a hardware bridge for serial failover, you must configure the pass-through bridge to be in fail open mode. See "Programmable Pass-Through/NIC Adapters" on page 82 for more information about configuring bridge modes.

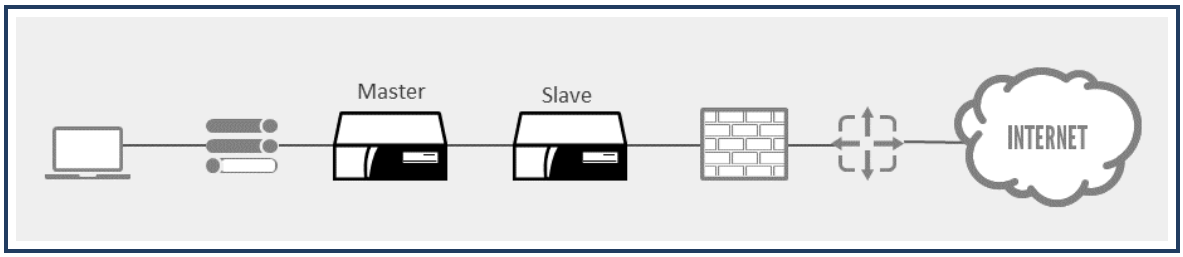

Configure Failover

Failover is accomplished by doing the following:

- Creating virtual IP addresses on each proxy.
- Creating a failover group.
- Attach the failover group to the bridge configuration.
- Selecting a failover mode (parallel or serial - this can only be selected using the CLI).

Both proxies can have the same priority (for example, the default priority). In that case, priority is determined by the local IP address—the appliance with the highest local IP will assume the role of master.

Example

The following example creates a bridging configuration with one bridge on standby.

- Appliance A—software bridge IP address: 10.0.0.2. Create a virtual IP address and a failover group, and designate this group the master.

  ```
  # (config) virtual-ip address 10.0.0.4
  # (config) failover
  # (config failover) create 10.0.0.4
  # (config failover) edit 10.0.0.4
  # (config failover 10.0.0.4) master
  # (config failover 10.0.0.4) enable
  ```

  The preceding commands create a failover group called 10.0.0.4. The priority is automatically set to 254 and the failover interval is set to 40.

- Appliance B—software bridge IP address: 10.0.0.3. Create a virtual IP address and a failover group.

  > **Note:** This deployment requires a hub on both sides of the bridge or a switch capable of interface mirroring.

```
# (config) virtual-ip address 10.0.0.4
# (config) virtual-ip address 10.0.0.4
# (config) failover
# config failover) create 10.0.0.4
# (config failover) edit 10.0.0.4
# (config failover 10.0.0.4) enable
```

On both appliances, attach the bridge configuration to the failover group:

```
# (config bridge bridge_name) failover group 10.0.0.4
```

- On both appliances, specify the failover mode:

```
# (config bridge bridge_name) failover mode serial
```

Configure Routing

Because it does not participate in a network routing protocol, the ProxySG appliance must be configured to reach clients and servers. To reach devices outside the network, you must configure a primary packet delivery path. This path is known as the default route (or the default gateway) and is configured during initial setup when you specify a default gateway. The appliance sends all traffic to the default gateway unless another route is specified. These alternate routes are called static routes and they list the IP addresses of other gateways that can be used to reach clients and servers in other parts of the network. Static routes are discussed in "About Static Routes" on page 100.

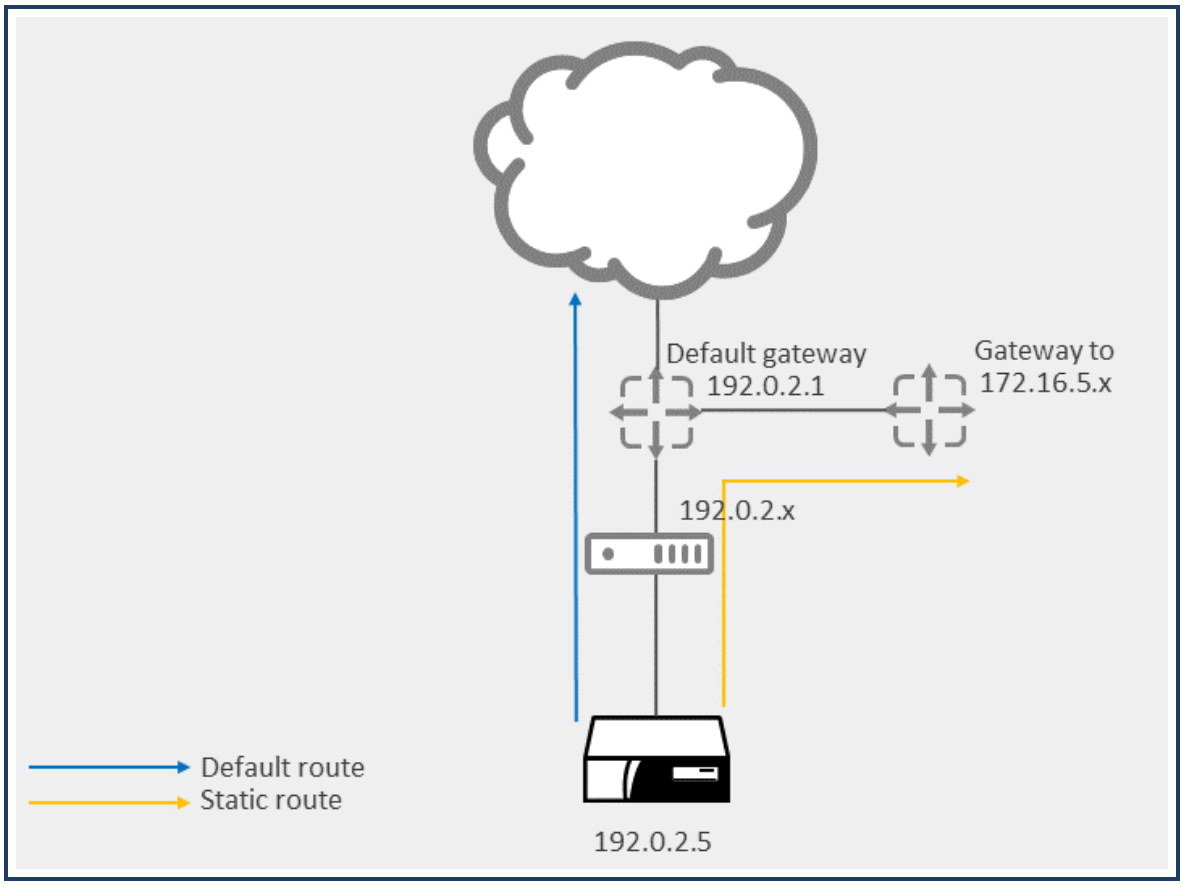

The appliance can be configured to distribute traffic through multiple default gateways, as explained in "Distribute Traffic Through Multiple Default Gateways" on page 91.

See Also:

- "Distribute Traffic Through Multiple Default Gateways" on page 91
- "Routing in Transparent Deployments" on page 93

IP Forwarding

IP forwarding is a special type of transparent proxy. The appliance is configured to act as a gateway and is configured so that if a packet is addressed to the appliance adapter, but not its IP address, the packet is forwarded toward the final destination. If IP forwarding is disabled, the packet is rejected as being misaddressed.

By default, IP forwarding is disabled to maintain a secure network. To enable IP forwarding, switch the **Enable** toggle.

> **Caution:** When IP forwarding is enabled, be aware that all appliance ports are open and all the traffic coming through them is not subjected to policy, with the exception of the ports that have explicitly defined in **Configuration > Services > Proxy Services**.

How Routes are Determined

Typically, the appliance uses the routing table to determine which interface to send the outbound packets to. When a packet is received, the appliance does a routing lookup to see if it can determine the correct route, either by using a static route or, if one is not defined, by sending it over the default route; however, the routing lookup might be bypassed depending on how the appliance is deployed and configured.

See Also

- "About Static Routes" on page 100

Routing Domains

Routing domains allow you to route traffic for unique networks through the same appliance, where each network has its own gateway and DNS server.

The following illustrates typical network flow through an appliance with Routing Domains configured:

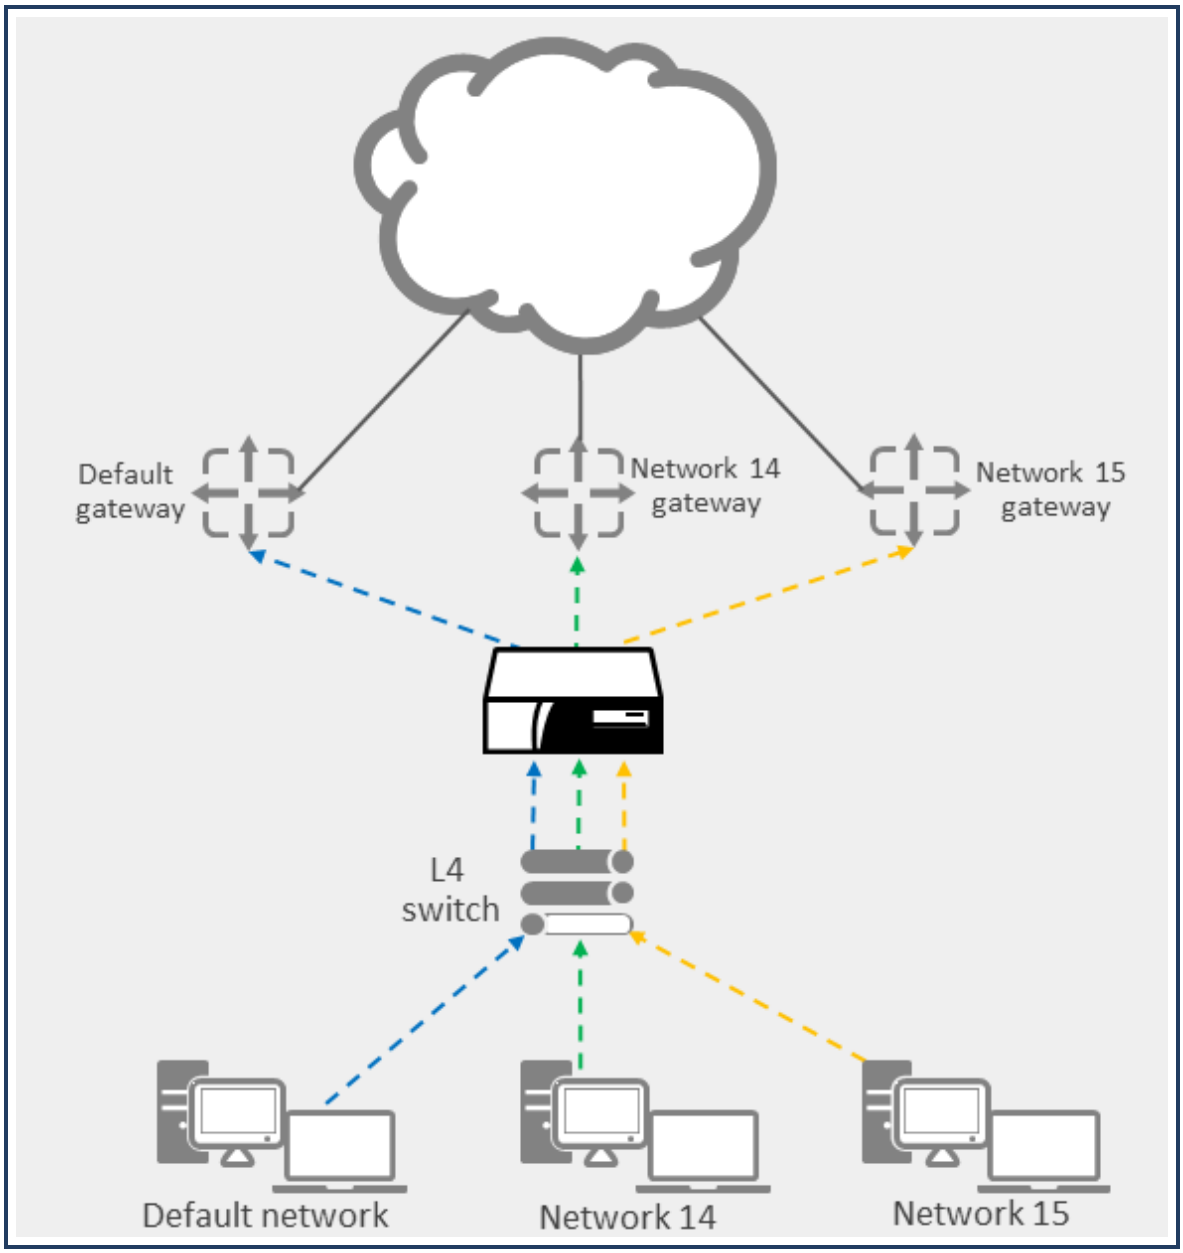

In large network architectures, it is common for multiple logical networks to traverse the same set of physical network devices. Often, segregation of these networks is required to:

- Increase visibility
- Reduce maintenance cost
- Minimize security risk

Routing domains provides this segregation by partitioning network interfaces into disjoint groups that only allow traffic to be forwarded to one of the other interfaces in the same group. Traffic cannot be forwarded between interfaces in different routing domains. Thus, network traffic is effectively segregated and can never cross routing domains.

Each routing domain object includes its own routing table that enforces layer 3 segregation as follows:

- Each routing table is associated with one or more logical interfaces.
- IP traffic that arrives on these interfaces is subject to routing and forwarding decisions defined by the routing table.
- To enforce network segregation, traffic never crosses routing domains.

Interfaces that are not assigned to any routing domain are automatically added to the default routing domain, which is subject to the configuration specified in the Admin Console in **Configuration > Network > Routing**.

See Also

- "Implement a Routing Domain" on page 103
- *Creating Multiple Logical Networks on a Single Appliance with Routing Domains* document at MySymantec

Distribute Traffic Through Multiple Default Gateways

You can distribute traffic originating at the ProxySG appliance through multiple default gateways and fine-tune how the traffic is distributed. This feature works with any routing protocol.

By using multiple gateways, an administrator can assign a number of available gateways into a preference group and configure the load distribution to the gateways within the group. Multiple preference groups are supported. In a mixed IPv4/IPv6 environment, you must define two default gateways: one for IPv4 and one for IPv6.

The specified gateway applies to all network adapters in the system.

> **Note:** Load balancing through multiple gateways is independent from the per interface load balancing the appliance automatically does when more than one network interface is installed.

Which default gateway the ProxySG appliance uses at a given time is determined by the preference group configuration assigned by the administrator. An appliance can have from 1 to 10 preference groups. A group can contain multiple gateways or only a single gateway.

Each gateway within a group can be assigned a relative weight value from 1 to 100. The weight determines how much bandwidth a gateway is given relative to the other gateways in the same group. For example, in a group with two gateways, assigning both gateways the same weight, whether 1 or 100, results in the same traffic distribution pattern. Alternatively, assigning one gateway a value of 10 and the other gateway a value of 20 results in the appliance sending approximately twice the traffic to the gateway with a weight value of 20. If there is only one gateway, it automatically has a weight of 100.

All gateways in the lowest preference group are considered to be active until one of them becomes unreachable and is dropped from the active gateway list. Any remaining gateways within the group continue to be used. If all gateways in the lowest preference group become unreachable, the gateways in the next lowest preference group become the active gateways (unless a gateway in a lower preference group becomes reachable again).

Configure Multiple Gateways for Load Balancing

When a gateway goes down, the networking code detects the unreachable gateway in 20 seconds, and the switchover takes place immediately if a secondary gateway is configured.

To create multiple gateways:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In the **Routing Config** section, enter or edit a routing domain name in **Domain Name**.
3. In the **IP Gateways** section:

a. Select **Add Gateway** to add a gateway.

b. Select an existing gateway in the list or select Edit under Actions.

4. Enter or edit the following settings:

Setting or field	Description
Group	Select a preference group.
Weight	Which default gateway the ProxySG appliance uses at a given time is determined by the assigned preference group configuration. Note that: ■ An appliance can have from 1 to 10 preference groups. ■ A group can contain multiple gateways or only a single gateway. ■ Each gateway within a group can be assigned a relative weight value from 1 to 100.
Gateway	The gateway IP address (IPv4 or IPv6).

Repeat the previous steps to define all gateways as needed.

5. (Optional) Add a new routing domain in **Routing Config** to set gateways for additional routing domains. See "Implement a Routing Domain" on page 103 for information.

> **Note:** The appliance uses the default route configuration for all appliance-initiated traffic requests.

Routing in Transparent Deployments

This section describes the mechanisms the appliance uses to route packets.

Outbound Routing

By default, the appliance sends outbound traffic to the default gateway unless one of the following is used (in order of precedence):

- The Trust Destination MAC feature, which is used when the appliance is in transparent bridging mode (unless certain other conditions are true—see "About Trust Destination MAC" below.
- A static route, if one is defined. For more information, see "About Static Routes" on page 100.
- The outbound RTS feature. For more information, see "About Return-to-Sender (RTS)" on the next page.
- An interface route, if the device is on the same subnet as the appliance. The appliance automatically adds an interface route to the routing table for hosts on the same subnet as the appliance interface. The interface route maps the subnet to the interface. The appliance can then do an ARP lookup for those hosts and send the packets directly to the client's MAC address.

Inbound Routing

By default, the appliance sends inbound traffic to the default gateway unless one of the following is used (in order of precedence):

- A static route, if one is defined. For more information, see "About Static Routes" on page 100.
- The inbound RTS feature. Inbound RTS is enabled by default. For more information, see "About Return-to-Sender (RTS)" on the next page.
- An interface route, if the device is on the same subnet as the appliance.

About Trust Destination MAC

When the appliance is in transparent bridging mode (in-path), it "trusts" the destination MAC address of the first client SYN packet and does not consult its routing table. The appliance notes the destination MAC address and outgoing interface specified in the frame and passes that information to the software process initiating the server connection, thus avoiding a routing lookup on the appliance. This feature is called Trust Destination MAC. It is enabled by default when the appliance is in transparent bridging mode and cannot be disabled.

Trust Destination MAC eliminates the need to create static routes and circumvents any routing issues encountered when the information in a packet is not sufficient for the appliance to make a routing decision.

> **Note:** Trust Destination MAC uses only the first client SYN packet to determine the MAC address and outgoing interface and continues to use this information even if the destination MAC address is not responding. To work around this limitation, enable outbound RTS, as described in "About Return-to-Sender (RTS)" below.

Override Trust Destination MAC

Unlike RTS, non-default static routes cannot override Trust Destination MAC. However, Trust Destination MAC behavior can be overridden if any of the following conditions are true:

- The result of the DNS lookup does not match the destination IP address. If Trust Destination IP is enabled, the DNS lookup is bypassed, and the IP address should always match.
- A policy rule does one of the following:
  - Specifies a forwarding host or SOCKS gateway
  - Rewrites the server URL in a way that causes the server connection to be forwarded to a different host

About Return-to-Sender (RTS)

As stated previously, the appliance does a routing lookup to see if it can determine the correct route for a packet, either by using a static route or, if one is not defined, by sending it over the default route. However, using the default route is sometimes suboptimal. For example, if the appliance satisfies a request and sends the client response traffic over the default route, the gateway router simply returns the traffic to the client router on the LAN side of the appliance. This causes unnecessary traffic between the switch and gateway, before the packet is finally received by the client router.

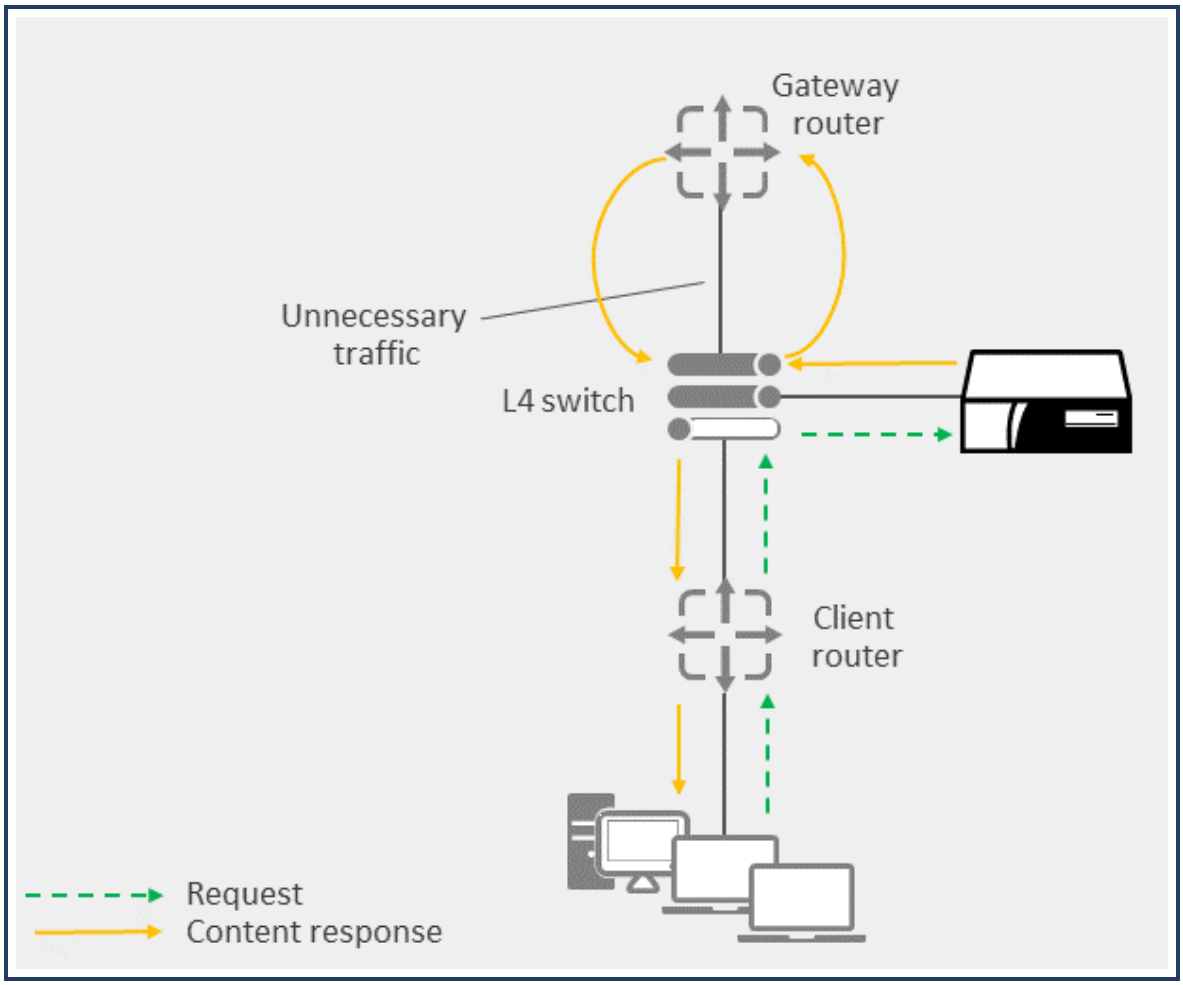

To specify a direct route, a static route must be created, which requires the administrator to update the routing table every time a new route is needed.

The Return-to-Sender (RTS) option eliminates the need to create static routes by configuring the appliance to send response packets back to the same interface that received the request packet, entirely bypassing any routing lookup on the appliance. Essentially, the appliance stores the source Ethernet MAC address that the client's packet came from and sends all responses to that address.

The RTS interface mapping is updated each time a packet is received. For example, if there are two gateways and both of them send packets to the appliance, the packets are sent back to the last MAC address and interface that received the packet.

> **Note:** Non-default static routes override RTS settings.

RTS can be configured in two ways, inbound or outbound. These two options can be enabled at the same time.

Inbound RTS

Inbound RTS affects connections initiated to the appliance by clients and is enabled by default. Inbound RTS configures the appliance to send SYN-ACK packets to the same interface that the SYN packet arrived on. All subsequent TCP/IP response packets are also sent to the same interface that received the request packet.

RTS inbound applies only to clients who are on a different subnet than the appliance. If clients are on the same subnet, interface routes are used.

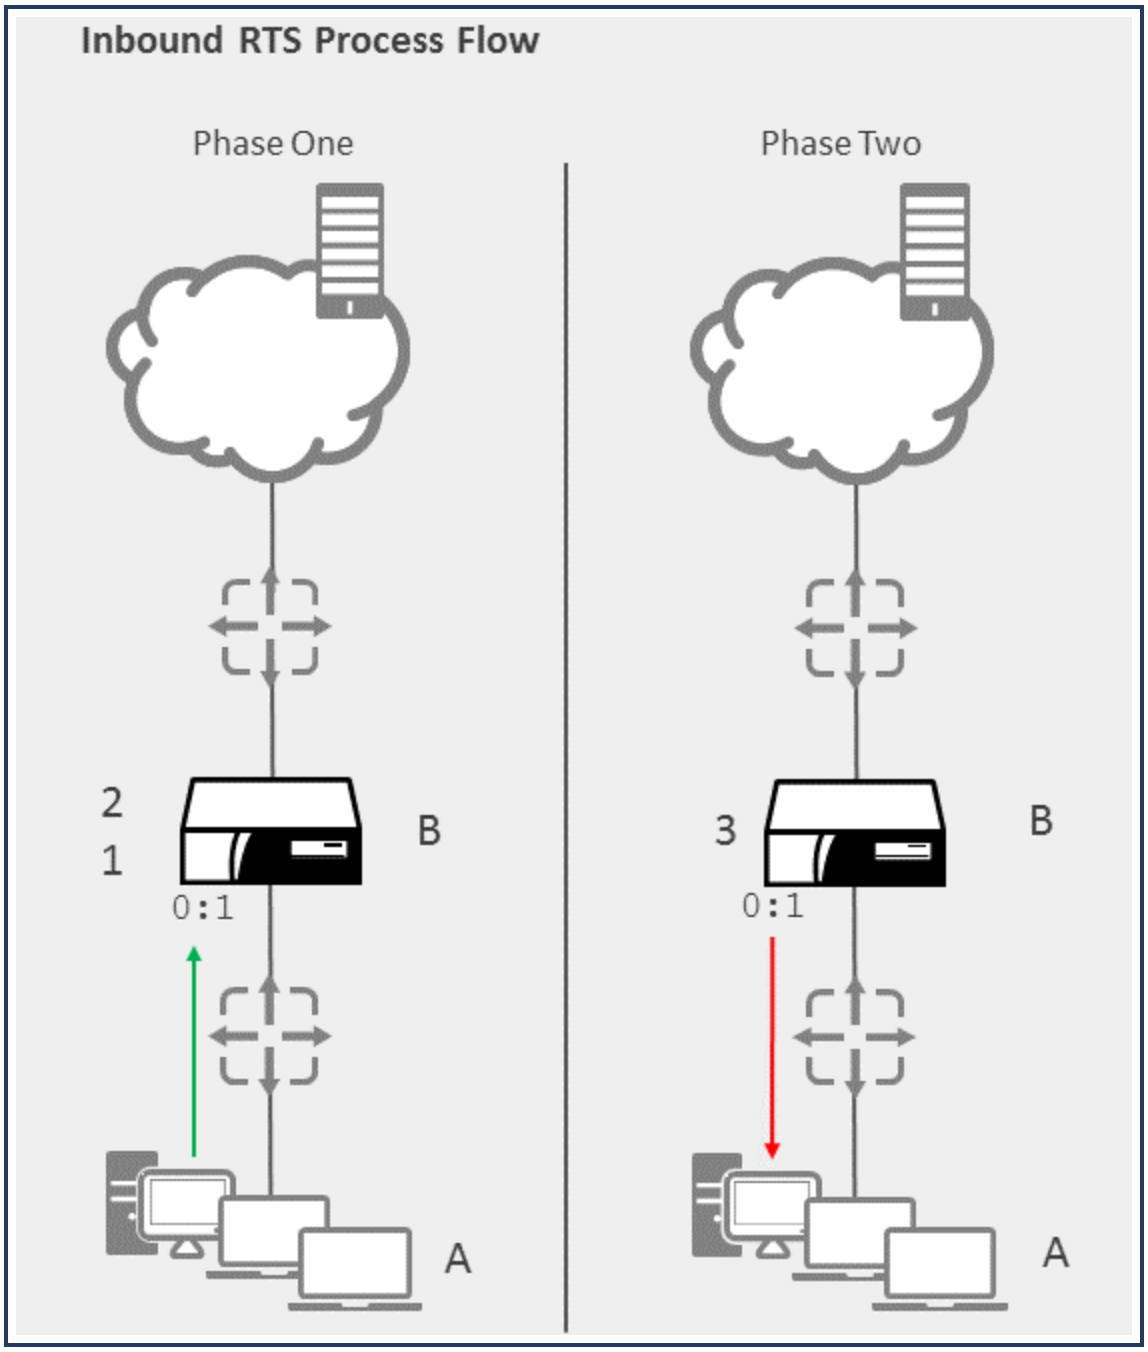

Inbound RTS Process Flow

1. Client A sends SYN to Server C across the WAN. The SYN is intercepted on the appliance B interface 0:1.
2. The appliance maps Client A to interface 0:1. All packets to Client A will now go to interface 0:1.
3. When a packet arrives for Client A, the appliance checks its interface mapping and sends the packet to the client over interface 0:1.

Outbound RTS

Outbound RTS affects connections initiated by the appliance to OCSes and causes the appliance to send ACK and subsequent packets to the same interface that the SYN-ACK packet arrived on. Outbound RTS requires a route to the client, either through a default gateway or static route.

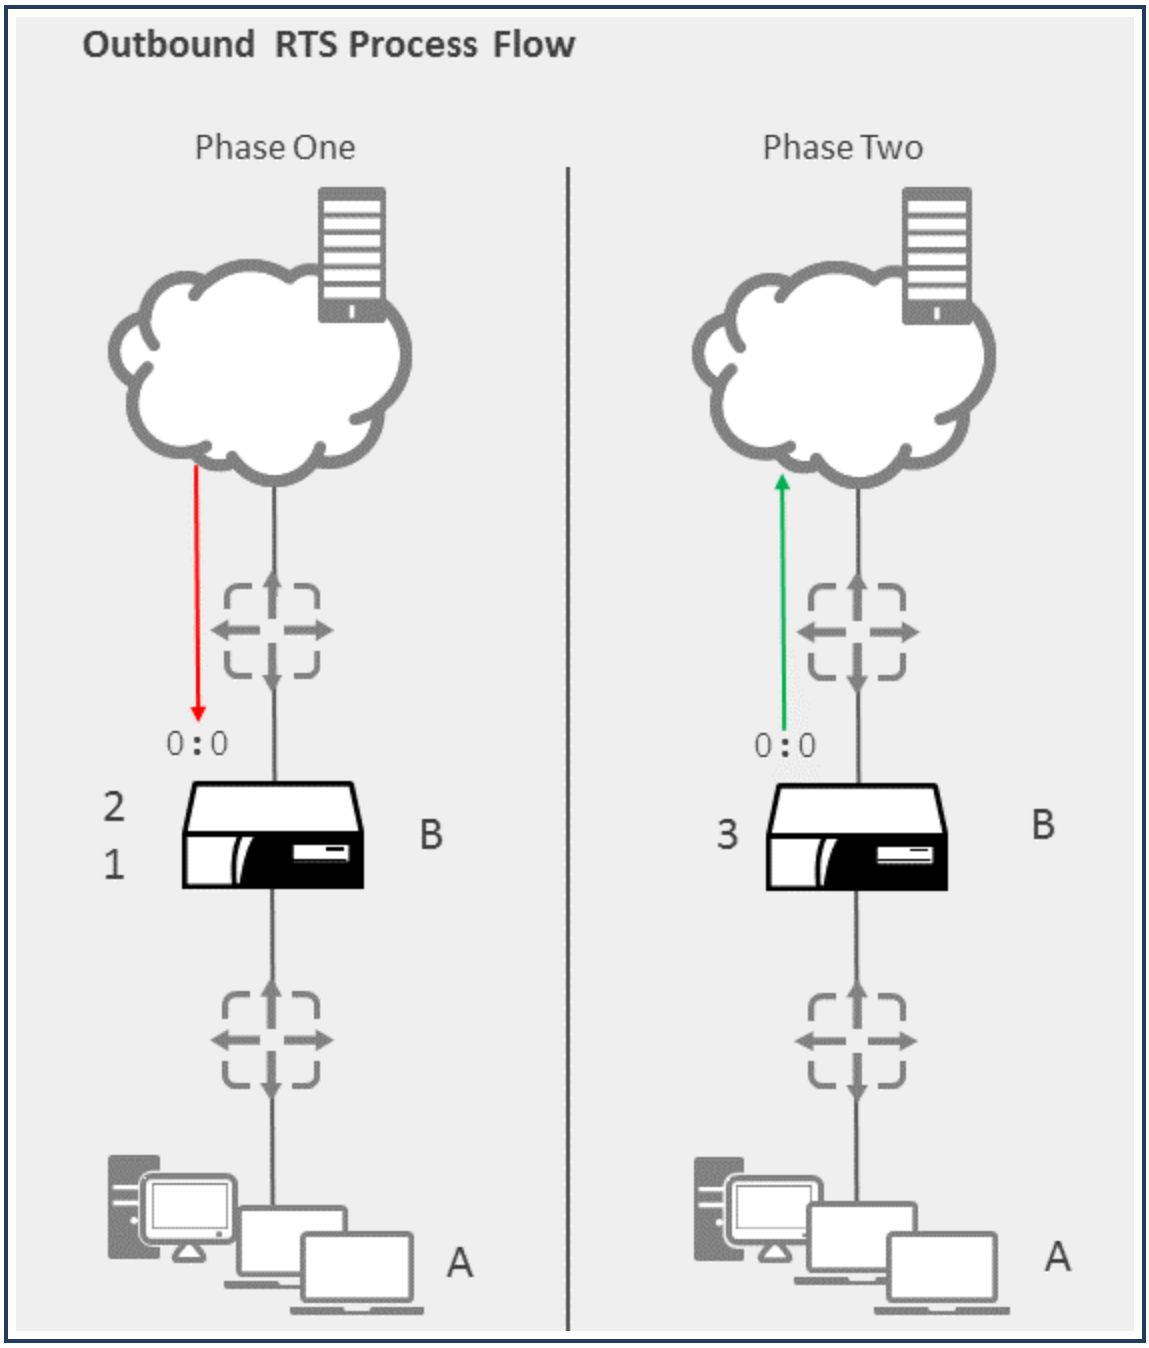

Outbound RTS Process Flow

1. Server C sends SYN-ACK to appliance B. The SYN-ACK is received on appliance interface 0:0.
2. The appliance maps Server B to interface 0:0. All packets to Server B will now go to interface 0:0.
3. When a packet arrives for Server B, the appliance checks its interface mapping and sends the packet out of interface 0:0.

Configure RTS

To configure RTS, use the following CLI commands:

```
# (config) return-to-sender inbound {disable | enable}
```

Enable or disable return-to-sender for inbound sessions.

```
# (config) return-to-sender outbound {disable | enable}
```

Enable or disable return-to-sender for outbound sessions.

For more information, refer to the *Command Line Interface Reference*.

DNS Verification

In transparent deployments, the appliance verifies the destination IP addresses provided by the client. This is known as L2/L3 transparency.

> **Note:** The Trust Destination IP option overrides DNS verification. This option is recommended for acceleration deployments only. For more information about this option, see "Trust the Destination IP Address Provided by the Client" on page 184

For hostname-less protocols such as CIFS and FTP, the IP address can always be trusted. For other protocols, such as HTTP, RTSP, and MMS, which have a hostname that must be resolved, verification can be an issue. URL rewrites that modify the hostname also can cause verification to fail.

L2/L3 transparency is not supported in explicit proxy deployments, or if the destination IP addresses cannot be verified by the appliance. In these cases, you must configure static routes to hosts that are only accessible through gateways other than the default gateway.

Transparent ADN connections that are handed off to an application proxy (HTTP or MAPI, for example) can utilize L2/L3 transparency. Also, transparent ADN connections that are tunneled but not handed off can utilize the functionality.

About Static Routes

Static routes define alternate gateways that the appliance can send packets to. A static route is a manually-configured route that specifies a destination network or device and the specific router that should be used to reach it.

> **Note:** For transparent bridge deployments, Trust Destination MAC overrides any static routes.

How Static Routes are Defined

You define static routes in a routing table. The routing table is a text file containing a list of static routes made up of destination IP addresses (IPv4 or IPv6), subnet masks (for IPv4) or prefix lengths (for IPv6), and gateway IP addresses (IPv4 or IPv6). The table is broken out into sections for each routing domain; the topmost entries are in the default routing domain. You are limited to 10,000 entries in the static routes table.

The following is a sample routing table:

```
10.25.36.0   255.255.255.0   10.25.36.1
10.25.37.0   255.255.255.0   10.25.37.1
2001::/64               fe80::2%0:1

[ROUTING_DOMAIN1]
10.73.37.0   255.255.255.0    10.73.37.1
[ROUTING_DOMAIN2]
10.73.38.0    255.255.255.0   10.73.38.1
```

Note that a routing table can contain a combination of IPv4 and IPv6 entries, but the gateway for each destination must be on the appropriate network type. For example, an IPv6 destination must use an IPv6 gateway.

When a routing table is installed, all requested URLs are compared to the list and routed based on the best match.

View the Current Routing Table Configuration

To view the current routing table configuration:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **Routing Table**, select **Source**.

   The routing table configuration page opens in a new browser window.

Install a Routing Table

You can install the routing table using the following methods:

- Text editor, which allows you to enter settings (or copy and paste the contents of an existing file) directly onto the appliance.
- Local file, to which the appliance can browse and install.
- Remote URL, where you put an already-created file on an FTP or HTTP server to be downloaded to the appliance.

  Use the `# (config) static-routes path url` command to set the path and the `# load static-route-table` command to load the new routing table.
- CLI command `# (config routing_domain_name) inline static-route-table EOF` to install the static routing table on the appliance.

For more information on the CLI commands, refer to the *Command Line Interface Reference*.

To install a routing table:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **Routing Table**, choose a method to install the routing table:

Method	Instructions
Text editor	1. Select **Install from text editor**. 2. In the Load Edit and Install from text editor dialog, enter the contents manually or copy and paste the contents from another file. 3. Select **Install**.
Local file	1. Select **... > Local File**. 2. In the Load Edit and Install from local file dialog, browse to the directory of the file on the local system. 3. Select **Install**.
Remote URL	1. Select **... > Remote URL**. 2. In the Load Edit and Install from remote URL dialog, enter the fully-qualified URL, including the filename, where the file is located. 3. Select **Download**.

> **Note:** If you use URL host rewrite functionality in your policies, mismatches can occur between the client-provided IP address and the resolved, rewritten hostname. In these cases, a routing lookup is performed and an interface route, static route, or default route is used.

Advanced Static Route Statistics

To view advanced static route statistics:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **Routing Table**, select **Static Routing Table**.

   The TCP/IP Static Routes page opens in a new browser window.

> **Caution:** Proceed with caution. Symantec recommends that you refer to Symantec Support for assistance with this URL. Support might direct you to this link to provide additional information during troubleshooting.

Implement a Routing Domain

The Admin Console shows the following details about all configured routing domains and interfaces. Select **Configuration > Network > Routing** and go to the specified section.

Section	Setting or field	Description
Routing Domains	**Name**	Name of the routing domain.
	Referenced	Indicates whether or not the routing domain is referenced elsewhere in the configuration (such as in **Configuration > Network >DNS**). Before deleting a routing domain, remove any references to it.
	Gateways	All configured interfaces and aggregate interfaces on the appliance. For information on gateways, see "Distribute Traffic Through Multiple Default Gateways" on page 91.
Interface to Routing Domain Mapping	See "Assign an Interface to a Routing Domain" on the next page for details.	

The appliance is configured with a default routing domain which cannot be deleted. Interfaces that are not assigned to any routing domain are automatically added to the default routing domain. See "Assign an Interface to a Routing Domain" on the next page to assign an interface to a different routing domain.

To implement a routing domain:

1. Add or edit an adapter for the network.

   Determine the interfaces will be used for each routing domain (ingress and egress) and define them in the Admin Console (**Configuration > Network > Adapters**). You can use either physical or VLAN interfaces, based on your network configuration:

   - If adding the new logical network to physical ports, connect your network to those ports on the appliance. Then, add a VLAN to the appropriate interface (see step 5 in Configure a Network Adapter). Leave the **VLAN ID** field blank, and define an IP address congruent with the new network.
   - If your new logical network is based on VLAN tagging on a preexisting physical connection, add a VLAN to the physical interface used for the connection (see step 5 in Configure a Network Adapter). Configure the VLAN ID and IP address as appropriate for the new network.

2. Configure the routing domain for the new network. Select **Configuration > Network > Domains**. In Routing Domains:

   a. Select **Add Domain** to add a routing domain.

   b. Select an existing routing domain in the list or select Edit under Actions.

3. Add a gateway for the new network. In the Properties section:
   a. In **Domain Name**, enter a routing domain name.
   b. To add or edit IP gateways, follow instructions in "Configure Multiple Gateways for Load Balancing" on page 91.

      Add as many gateways as needed.
4. Set a unique DNS server for the new network.

   If you don't define a DNS server for the new network, the routing domain uses the primary/alternate DNS server configuration.

   > **Note:** The appliance uses the primary/alternate DNS configuration for all proxy-initiated lookups, regardless of additional configuration.

   See "Configure DNS Forwarding Groups" on page 117 for instructions.

Assign an Interface to a Routing Domain

You can assign an interface to a routing domain or change the current routing domain assignment.

The Admin Console provides an overview of the network interface configuration. Select **Configuration > Network > Routing** and go to **Interface to Routing Domain Mapping** for the following configuration information:

Setting or field	Description
Interface	Interface in the form *adapter:interface* or (if an aggregate interface) aggr:*number*.
VLAN ID	VLAN identifier, if applicable. Otherwise, states "Physical Interface".
IP Address	All VLANs (both IPv4 and IPv6) on the interface.
Routing Domain	Routing domain to which the interface belongs. If the interface doesn't belong to a domain, value is "None". To assign the interface to a different routing domain, select the domain from the list.

Configure Routing Information Protocol

RIP is designed to select the fastest route to a destination. Default routes added by RIP are treated the same as the static default routes; that is, the default route load-balancing schemes apply to the default routes from RIP. The appliance's RIP implementation also supports advertising default gateways.

View the Current RIP Configuration

To view the current RIP configuration:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **RIP**, select **Sources**.

   The RIP configuration page opens in a new browser window.

Install RIP Settings

No RIP configuration file is shipped with the appliance; to configure RIP, you create a RIP installation file and install it using one of the following methods:

- Text editor, which allows you to enter settings (or copy and paste the contents of an existing file) directly onto the appliance.
- Local file, to which the appliance can browse and install.
- Remote URL, where you put an already-created file on an FTP or HTTP server to be downloaded to the appliance.

  Use the `# (config) rip` *`url`* command to set the path and the `# load rip-settings` command to load the RIP settings.
- CLI command `# inline rip-settings`*`EOF`* to install the RIP configuration on the appliance.

For more information on the CLI commands, refer to the *Command Line Interface Reference*.

To install RIP settings:

> **Note:** Do not enable RIP until after you have installed RIP settings.

1. Create the RIP configuration file. See "RIP Configuration File Commands" on page 108 for details.
2. In the Admin Console, select **Configuration > Network > Routing**.
3. In **RIP**, choose a method to install RIP settings:

Method	Instructions
Text editor	1. Select **Install from text editor**. 2. In the Load Edit and Install from text editor dialog, enter the contents manually or copy and paste the contents from another file. 3. Select **Install**.
Local file	1. Select **... > Local File**. 2. In the Load Edit and Install from local file dialog, browse to the directory of the file on the local system. 3. Select **Install**.
Remote URL	1. Select **... > Remote URL**. 2. In the Load Edit and Install from remote URL dialog, enter the fully-qualified URL, including the filename, where the file is located. 3. Select **Download**.

Enable RIP

RIP is disabled by default. After you have installed RIP settings, you can enable RIP.

To enable RIP:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **RIP**, switch the **Enable RIP (Routing Information Protocol)** toggle.

View the Current RIP Commands

To view the current RIP commands:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **RIP**, select **RIP Settings**.

   The RIP Parameters page opens in a new browser window. See "RIP Configuration File Commands" on page 108 for information on supported commands.

Advanced RIP Statistics

To view advanced RIP statistics:

1. In the Admin Console, select **Configuration > Network > Routing**.
2. In **RIP**, select **RIP Routes**.

The RIP Information page opens in a new browser window.

> **Caution:** Proceed with caution. Symantec recommends that you refer to Symantec Support for assistance with this URL. Support might direct you to this link to provide additional information during troubleshooting.

Create Advertising Default Gateway Routes

Default routes advertisements are treated the same as the static default routes; that is, the default route load balancing schemes also apply to the default routes from RIP.

By default, RIP ignores the default routes advertisement. You can change the default from disable to enable and set the preference group and weight through the CLI only.

To create advertising default gateway routes:

1. Log in to the CLI mode and enter `# (config)` mode.
2. Issue the following commands:

   ```
   # (config) rip default-route enable
   # (config) rip default-route group group_number
   # (config) rip default-route weight weight_number
   ```

   where *group_number* defaults to 1 and *weight_number* defaults to 100, the same as the static default route set by the `# (config) ip-default-gateway` command.
3. (Optional) To view the default advertising routes, issue the `# (config) show rip default-route` command. The CLI shows output similar to the following:

   ```
   # (config) show rip default-route
   RIP default route settings:
   Enabled: Yes
   Preference group: 3
   Weight: 30
   ```

RIP Configuration File Commands

You can put any of the commands below into a RIP configuration text file. You cannot edit a RIP file through the command line, but you can overwrite a RIP file using the command `# inline rip-settings`*EOF*. After the file is complete, put it on an HTTP or FTP server accessible to the appliance and download it.

To view current RIP commands, see "View the Current RIP Commands" on page 106.

> **Note:** RIP parameters are accepted in the order that they are entered. If a RIP parameter is added, it is appended to the default RIP parameters. If a subsequent parameter conflicts with a previous parameter, the most recent one is used.

net

`net` *Nname*`[/`*mask*`]` `gateway` *Gname* `metric` *Value* `{passive | active | external}`

Parameter	Description
Nname	Name of the destination network. It can be a symbolic network name, or an Internet address specified in dot notation.
/mask	Optional number between 1 and 32 indicating the netmask associated with *Nname*.
Gname	Name or address of the gateway to which RIP responses should be forwarded.
Value	The hop count to the destination host or network. A `net Nname/32` specification is equivalent to the `host Hname` command.
`passive | active | external`	Specifies whether the gateway is treated as passive or active, or whether the gateway is external to the scope of the RIP protocol.

host

`host` *Hname* `gateway` *Gname* `metric` *Value* `{passive | active | external}`

Parameter	Description
Hname	Name of the destination network. It can be a symbolic network name, or an Internet address specified in dot notation.
Gname	Name or address of the gateway to which RIP responses should be forwarded. It can be a symbolic network name, or an Internet address specified in dot notation.
Value	The hop count to the destination host or network. A `net Nname/32` specification is equivalent to the `host Hname` command.
`passive | active | external`	Specifies whether the gateway is treated as passive or active, or whether the gateway is external to the scope of the RIP protocol.

RIP Parameters

Lines that do not start with `net` or `host` commands must consist of one or more of the following parameter settings, separated by commas or blank spaces:

Parameter	Description
`bcast_rdisc`	Specifies that Router Discovery packets should be broadcast instead of multicast.
`if=[0|1|2|3]`	Specifies that the other parameters on the line apply to the interface numbered 0,1,2, or 3 in SGOS terms.
`no_ag`	Turns off aggregation of subnets in RIPv1 and RIPv2 responses.
`no_rdisc`	Disables the Internet Router Discovery Protocol. This parameter is set by default.
`no_rdisc_adv`	Disables the transmission of Router Discovery Advertisements.
`no_rip`	Disables all RIP processing on the specified interface.
`no_ripv1_in`	Causes RIPv1 received responses to be ignored.
`no_ripv2_in`	Causes RIPv2 received responses to be ignored.
`no_solicit`	Disables the transmission of Router Discovery Solicitations.
`no_super_ag`	Turns off aggregation of networks into supernets in RIPv2 responses.
`passive`	Marks the interface to not be advertised in updates sent through other interfaces, and turns off all RIP and router discovery through the interface.
`passwd=XXX`	Specifies an RIPv2 password included on all RIPv2 responses sent and checked on all RIPv2 responses received. The password must not contain any blanks, tab characters, commas or '#' characters.
`rdisc_adv`	Specifies that Router Discovery Advertisements should be sent, even on point-to-point links, which by default only listen to Router Discovery messages.
`rdisc_pref=N`	Sets the preference in Router Discovery Advertisements to the integer N.
`rdisc_ interval=N`	Sets the nominal interval with which Router Discovery Advertisements are transmitted to N seconds and their lifetime to 3*N.
`redirect_ok`	Causes RIP to allow ICMP Redirect messages when the system is acting as a router and forwarding packets. Otherwise, ICMP Redirect messages are overridden.
`ripv2_out`	Turns off RIPv1 output and causes RIPv2 advertisements to be multicast when possible. ripv2 Is equivalent to no_ripv1_in and no_ripv1_out. This parameter is set by default.
`send_solicit`	Specifies that Router Discovery solicitations should be sent, even on point-topoint links, which by default only listen to Router Discovery messages.
`trust_ gateway=rname`	Causes RIP packets from that router and other routers named in other trust_gateway keywords to be accept, and packets from other routers to be ignored.

ProxySG-Specific RIP Parameters

The following RIP parameters are unique to ProxySG configurations.

Parameter	Description
`advertise_ routes`	`-s` option: Supplying this option forces routers to supply routing information whether it is acting as an Internetwork router or not. This is the default if multiple network interfaces are present or if a point-to-point link is in use. `-g` option: This flag is used on Internetwork routers to offer a route to the 'default' destination. This is typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers.
`supply_ routing_ info`	`-h` option: `Suppress_extra_host_routes advertise_host_route` `-m` option: `Advertise_host_route` on multi-homed hosts `-A` option: Ignore_authentication `//`
`no_supply_ routing_ info`	`-q` option: Opposite of `-s`.
`no_rip_out`	Disables the transmission of all RIP packets. This setting is the default.
`no_ripv1_ out`	Disables the transmission of RIPv1 packets.
`no_ripv2_ out`	Disables the transmission of RIPv2 packets.
`rdisc`	Enables the transmission of Router Discovery Advertisements.
`rip_out`	Enables the transmission of RIPv1 packets.
`ripv1`	Causes RIPv1 packets to be sent.
`ripv1_in`	Causes RIPv1 received responses to be handled.
`ripv1_out`	Enables the transmission of RIPv1 packets.

Using Passwords with RIP

The first password specified for an interface is used for output. All passwords pertaining to an interface are accepted on input. Consider the following settings:

```
if=0 passwd=aaa
if=1 passwd=bbb
passwd=ccc
```

Interface `0` accepts passwords `aaa` and `ccc`, and transmits using password `aaa`.

Interface `1` accepts passwords `bbb` and `ccc`, and transmits using password `bbb`.

Other interfaces accept and transmit the password `ccc`.

Configure DNS

This section describes various configuration tasks associated with DNS services. During initial configuration of the appliance, you configured the IP address of a single primary DNS server. You can add one or more alternate DNS servers, and also define custom DNS service groups.

DNS comprises a hierarchical set of DNS servers. For each domain or sub-domain, one or more authoritative DNS servers publish information about that domain and the name servers of any domains that are under it.

There are two types of DNS queries:

- Non-recursive, which means that a DNS server can provide a partial answer or return an error to the client.
- Recursive, which means that the DNS server either fully answers the query or returns an error to the client.

By default, recursive DNS is disabled on the appliance.

See Also:

- "Recursive and Non-Recursive DNS" on page 114

DNS Forwarding Groups

DNS servers are configured in groups. Customers with split DNS server configuration (for example, environments that maintain private internal DNS servers and external DNS servers) might choose to add servers to an alternate DNS server group as well as to the primary DNS server group. In addition, you can create custom DNS server groups.

On the appliance, internal DNS servers are placed in the primary group, while external DNS servers (with the Internet information) populate the alternate group.

If routing domains are configured on the appliance, you can associate DNS server groups with a routing domain. See "Implement a Routing Domain" on page 103 for details.

> **Note:** The appliance always uses the primary and alternate DNS groups for requests initiated by the appliance.

See Also:

- "Configure DNS Forwarding Groups" on page 117

Resolve Hostnames Using Name Imputing Suffixes

The ProxySG appliance queries the original hostname before checking imputing suffixes unless there is no period in the hostname. If there is no period in the hostname, imputing is applied first.

The appliance uses name imputing to resolve hostnames based on a partial name specification (DNS name imputing suffix). When the appliance submits a hostname to the DNS server, the DNS server resolves the hostname to an IP address.

The appliance then tries each entry in the name-imputing suffixes list until the name is resolved or it reaches the end of the list. If by the end of the list the name is not resolved, the appliance returns a DNS failure. For example, if the name-imputing list contains the entries example.com and com, and a user submits the URL http://www.eedept, the appliance resolves the host names in the following order.

www.eedept
www.eedept.example.com
www.eedept.com

See Also:

- "Configure DNS Name Imputing Suffixes" on page 121

Negative DNS Response Caching

By default, the appliance caches negative DNS responses sent by a DNS server. You can configure the appliance to set the TTL value for a negative DNS response to be cached. You can also disable negative DNS response caching. This functionality is only available through the CLI. You can't configure DNS negative caching through the Admin Console.

The appliance supports caching of both type A and type PTR DNS negative responses.

> **Note:** The appliance generates more DNS requests when negative caching is disabled.

To configure negative caching TTL values, issue the CLI command:

```
# (config) dns negative-cache-ttl-override seconds
```

where *seconds* is an integer between 0 and 600. Setting the TTL value to 0 seconds disables negative DNS caching; setting the TTL setting to a non-zero value overrides the TTL value from the DNS response.

To restore negative caching defaults, issue the CLI command:

```
# (config) dns no negative-cache-ttl-override
```

Recursive and Non-Recursive DNS

There are two types of DNS queries:

- Non-recursive, which means that a DNS server can provide a partial answer or return an error to the client
- Recursive, which means that the DNS server either fully answers the query or returns an error to the client

Non-Recursive DNS

If you have defined more than one DNS server, the ProxySG appliance uses the following logic to determine which servers are used to resolve a DNS host ame and when to return an error to the client.

> **Note:** Servers are always contacted in the specified order in a group list.

- The ProxySG appliance first checks all the DNS groups for a domain match, using domain-suffix matching to match a request to a group.
    - If there is a match, the servers in the matched group are queried until a response is received; no other DNS groups are queried.
    - If there is no match, the appliance selects the Primary DNS group.
- The appliance sends requests to DNS servers in the Primary DNS server group in the order in which they appear in the list. If a response is received from one of the servers in the Primary group, no attempts are made to contact any other Primary DNS servers.
- If none of the servers in the Primary group resolve the host name, the appliance sends requests to the servers in the Alternate DNS server group. (If no Alternate servers have been defined, an error is returned to the client.)
    - If a response is received from a server in the Alternate group list, there are no further queries to the Alternate group.
    - If a server in the Alternate DNS server group is unable to resolve the host name, an error is returned to the client, and no attempt is made to contact any other DNS servers.

      The Alternate DNS server is not used as a failover DNS server. It is only used when DNS resolution of the Primary DNS server returns a name error. If the query to each server in the Primary list times out, no alternate DNS server is contacted.
- If the appliance receives a referral (authoritative server information), DNS recursion takes over if it is enabled.

  If the appliance receives a negative DNS response (a response with an error code set to name error), it caches that negative response. See "Negative DNS Response Caching" on the previous page.

Recursive DNS

If you have enabled recursive DNS, the appliance uses the following logic to determine how to resolve a DNS host name and when to return an error to the client.

If the DNS server response does not contain an A record with an IP address but instead contains authoritative server information (a referral), the appliance follows all referrals until it receives an answer. If the appliance follows more than eight referrals, it assumes there is a recursion loop, aborts the request, and sends an error to the client.

When to Enable Recursive DNS

If you have a DNS server that cannot resolve all host names, it might return a list of authoritative DNS servers instead of a DNS A record that contains an IP address. To avoid this situation, configure the appliance to recursively query authoritative DNS servers.

To enable recursive DNS:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Settings**, switch the **Enable DNS Recursion** toggle.

Extensions for DNS

The DNS proxy supports EDNS, which allows DNS requesters to receive DNS UDP messages longer than the default 512 bytes. For detailed information on EDNS, refer to the following RFCs:

- RFC 2671: https://www.ietf.org/rfc/rfc2671.txt
- RFC 6891: https://tools.ietf.org/html/rfc6891

EDNS request messages consist of the following sections:

- constant section: Includes an acceptable DNS response size. The DNS proxy uses this value to send long responses.
- variable section: The DNS proxy does not parse this section and transparently forwards it to upstream DNS server instead.

When EDNS is enabled on the appliance, the DNS proxy:

- accepts EDNS queries
- allocates DNS RESPONSE buffers with respect to the original queries
- forwards EDNS variable sections to the upstream DNS server
- saves long DNS server responses in the DNS cache
- creates long responses to EDNS clients

To enable EDNS:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Settings**, switch the **Enable EDNS** toggle.

Configure DNS Forwarding Groups

The following rules apply to DNS forwarding groups:

- You can add servers to the primary and alternate groups, but you cannot change the domain or add additional domains; these groups are defined at initial configuration.
- The primary and alternate DNS groups can't be deleted.
- A custom DNS group must include at least one DNS server address before adding a domain. In addition, you must delete all domains associated with the last server in the DNS group before you can delete the server.

DNS Health Checks

Each time you add a DNS server to a group, the appliance automatically creates a DNS health check for that server IP address and uses a default configuration for the health check. For example, if you add a DNS server to a primary or alternate DNS group, the created health check has a default hostname of symantec.com. If you add a DNS server to a custom group, the longest domain name is used as the default hostname for the health check.

After you add DNS servers to a group, Symantec recommends that you check the DNS server health check configurations and edit them as required.

DNS Forwarding Group Details

In the Admin Console, select **Configuration > Network > DNS**. The Forwarding Groups section lists all configured DNS server groups on the appliance and displays information about each one:

Setting or field	Description
Name	Name of the DNS server group.
Servers	The DNS server addresses belonging to the specified group. Both IPv4 and IPv6 addresses are supported.
Domains	The DNS server domains belonging to the specified group. Does not apply to the primary and alternate groups.
Routing Domain	Routing domains associated with the DNS server group.

Configure the Primary or Alternate Group

When you first set up the appliance, you configured a primary DNS server. If your deployment makes use of more than one DNS server, you can them to the primary or alternate server group; you can also delete DNS servers from the primary group, but you cannot delete the group or change the domain or add additional domains–the group is defined at initial configuration.

If you are using the appliance in a mixed IPv4/IPv6 environment, you should configure both IPv4 and IPv6 DNS servers.

To configure the primary/alternate group:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Forwarding Groups**:
   - Select **Add Group** to add a group.
   - Select an existing group in the list or select Edit under Actions.

     The Editing group page displays group settings.
3. (Optional) Beside **Routing Domain**, select a routing domain to associate with the group. See "Implement a Routing Domain" on page 103 for details.
4. Configure DNS server addresses in the group:
   - To add addresses:
     - Select **Add > Add Server** to add entries one at a time.
     - Select **Add > Paste Multiple** to add multiple entries at a time. In the Paste Multiple dialog, enter each address on a separate line. The addresses are added in the specified order.

       If needed later, you can change the order of DNS servers using the CLI. See "Change Order of DNS Servers in a List" on page 120.
   - To edit an existing address, type directly in the **Server** field.

Define a Custom DNS Group

Custom groups enable you to specify servers and domains for specific company needs (such as resolving internal or external hostnames) depending on how you have set up your primary and alternate DNS groups.

DNS Group Requirements

- Valid DNS entry formats are www.example.com, example.com.
- You can create a maximum of 8 custom groups , and each custom group can contain a maximum of four DNS servers and eight domains.
- Groups don't accept wild cards, such as *.example.com.
- Groups don't partially match domain names, such as *.example.com and .example.com. exam.com doesn't match queries for www.example.com.
- DNS record requirements have been relaxed, as discussed in RFC 2181. Review sections 10 and 11 for more information.

To define a custom DNS group:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Forwarding Groups**:
   - Select **Add Group** to add a group.
   - Select an existing group in the list or select Edit under Actions.

     The Editing group page displays group settings.
3. Specify the following settings as needed:

Setting or field	Description
Group Name	Name of the DNS server group.
Routing Domain	(Optional) Routing domain associated with the DNS server group. See "Implement a Routing Domain" on page 103 for details.

4. Configure DNS server addresses in the group:
   - To add addresses:
     - Select **Add > Add Server** to add entries one at a time.
     - Select **Add > Paste Multiple** to add multiple entries at a time. In the Paste Multiple dialog, enter each address on a separate line. The addresses are added in the specified order.

       If needed later, you can change the order of DNS servers using the CLI. See "Change Order of DNS Servers in a List" on the next page.
   - To edit an existing address, type directly in the **Server** field.
5. Configure DNS server domains in the group:
   - To add domains:
     - Select **Add > Add Server** to add entries one at a time.
     - Select **Add > Paste Multiple** to add multiple entries at a time. In the Paste Multiple dialog, enter each domain on a separate line. The domains are added in the specified order.
   - To edit an existing domain, type directly in the **Domain** field.

> **Note:** If a domain becomes defunct, you can delete it from a DNS group. In addition, you must delete all domains associated with the last server in any DNS group before you can delete the server.

Change Order of DNS Servers in a List

Using the CLI, you can promote DNS servers in the list for any DNS forwarding group.

To promote or demote DNS servers in a list:

```
# (config dns forwarding) edit group_alias
```

This changes the prompt to:

```
# (config dns forwarding group_alias)
# (config dns forwarding group_alias) { demote | promote} server_ip_address number
```

where the DNS server specified by *server_ip_address* is demoted or promoted *number* places in the list, and *number* is a positive number.

Configure DNS Name Imputing Suffixes

Using name imputing suffixes is particularly useful for a company's internal domains. For example, it enables you to simply enter:

webServer

rather than:

webServer.inOurInternalDomain.ForOurCompany.com

Also, this resolves any problem with external root servers being unable to resolve names that are internal only. The ProxySG appliance supports up to 30 name imputing suffixes.

To configure DNS name imputing suffixes:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Imputing**:
   - To add an entry, select **Add DNS name imputing suffix** and enter the suffx in the **Name Imputing Suffix** field that becomes available.
   - To edit an existing entry in the list, type directly in the **Name Imputing Suffix** field.

Change the Order of DNS Name Imputing Suffixes

The appliance uses imputing suffixes according to the list order. You can organize the list of suffixes so the preferred suffix displays at the top of the list.

> **Note:** This functionality is only available through the Admin Console. You can't configure it using the CLI.

To change the order of DNS name imputing suffixes:

1. In the Admin Console, select **Configuration > Network > DNS**.
2. In **Imputing**:
   - To promote an entry, select the up arrow under Actions.
   - To demote an entry, select the down arrow under Actions.

About WCCP

The WCCP protocol allows certain Cisco routers and switches to transparently redirect traffic to a cache engine such as a ProxySG appliance. This traffic redirection helps to improve response time and optimize network resource usage. The appliance can be configured to participate in a WCCP scheme, in which WCCP-capable switches or routers collaborate with appliances to form one or more groups that service requests from clients.

> **Note:** WCCP is supported for the default routing domain only. You cannot configure WCCP for multiple routing domains. See "Implement a Routing Domain" on page 103 for information on routing domains.

WCCP on the ProxySG Appliance

In virtually in-path deployments, when the ProxySG appliance is not in the physical path of clients and servers, a WCCP-capable router is used to redirect traffic to the appliance for transparent proxy services.

In a transparent proxy deployment the client is not aware that it is interacting with an intermediate proxy and not the OCS. The process works as follows:

1. The client sends a packet addressed for the OCS.
2. The WCCP-enabled router redirects the packet to the appliance.
3. The appliance determines what to do with it based on the transparent proxy services that have been configured for the traffic type. If it cannot service the request locally (for example by returning a page from its local cache), it sends a request to the specified OCS on behalf of the client.

The OCS response is routed (or redirected depending on the configuration) back to the appliance.

The appliance then forwards the response back to the client. To implement this transparent redirection scheme, one or more appliances and one or more routers/switches must form a service group.

The appliance offers VLAN Support for WCCP and allows you to redirect traffic from the router over physical or virtual interfaces. If you configure multiple virtual interfaces between the appliance and the WCCP-capable router, you can segregate WAN and LAN traffic on the same physical interface by enabling a VLAN trunk between the appliances. By default, VLAN trunking is enabled on the appliance. For information on VLANs, see "Configure a VLAN" on page 66.

See Also:

- "Requirements for Configuring WCCP" on page 126
- "Configure WCCP" on page 127

Service Groups

A service group unites one or more routers/switches with one or more appliances in a transparent redirection scheme governed by a common set of rules. The service group members agree on these rules by announcing their specific capabilities and configuration to each other in WCCP protocol packets. When creating a service group on the appliance, you define the following:

Home Router Address

In order to establish and maintain a service group, the appliances and routers must be able to communicate. To specify the addresses of all the routers in a service group, you must choose one of the following methods:

- Unicast: Each ProxySG appliance must be explicitly configured with the IP address of every router in the service group. You will need to reconfigure each appliance whenever you add or remove a router from the group.
- Multicast: The routers and ProxySG appliances in the service group communicate using a single IP address in the range of 224.0.0.0 to 239.255.255.255. To configure this, each appliance and each router in the group must be configured with the multicast IP address. If the WCCP routers and/or appliances are more than one hop apart, IP multicast routing must also be enabled on the intervening routers.

Service Group Authentication

If you are using WCCP v2, you can secure a service group by configuring an MD5 authentication between the ProxySG appliances and the routers in the group.

To configure authentication, you must define the same password on all routers and all appliances in the service group. When authentication is enabled, an appliance is not allowed to join the service group unless it knows the password.

Packet Forward and Return Methods

The packet forward and return method for a service group defines how the router forwards packets to the ProxySG appliance as well as how the appliance returns packets that it does not intercept because of the policy or services configured on it, back to the router.

Symantec recommends that all service groups configured on a router use the same forwarding and return methods.

The appliance supports the following forward/return methods:

- GRE Forwarding/GRE Return: With GRE forwarding, the router encapsulates the intercepted packet in an additional IP and GRE header that shows the router address as the source IP address and the address of the appliance as the destination IP address. When the appliance receives the packet, it strips the outside header and then determines how to process the request, either forwarding the request on to the OCS or servicing it locally.

  When returning the redirected packet, the appliance encapsulates the packet with an IP and GRE header that bears the IP address of the appliance as the source and the router IP address as the destination.

- L2 Forwarding/L2 Return: With Layer 2 (L2) forwarding the router rewrites the destination MAC address of the

intercepted packet to the MAC address of the appliance to which it is redirecting the packet. This method is faster than GRE forwarding because the forwarding is done at the hardware level and doesn't require encapsulating and decapsulating the packet at Layer 3. However, to use L2 forwarding, the appliance and the routers in the service group must all be on the same L2 broadcast domain (that is, there cannot be more than one hop between them).

When returning the redirected packet, the appliance rewrites the destination MAC address to that of the router.

To determine whether L2 forwarding is supported on your hardware platform, refer to your Cisco documentation. For a list of the Cisco platforms on which Symantec has tested L2 forwarding with the appliance, refer to the *WCCP Reference Guide*.

- L2 Forwarding/GRE Return: With L2 forwarding the router rewrites the destination MAC address of the intercepted packet to the MAC address of the appliance to which it is redirecting the packet.

  When returning the redirected packet, the appliance encapsulates the packet with an IP and GRE header that bears the IP address of the appliance as the source and the router IP address as the destination.

> **Note:** The appliance does not support GRE forwarding and L2 packet return. If you configure this combination, the appliance will generate a capability mismatch error. To view the errors and warnings, use the `# show wccp status` CLI command.

Router Affinity

By default, the ProxySG appliance uses the configured return method to return bypassed traffic to the router that redirected it and uses regular routing table lookups to determine the next hop for intercepted traffic. With router affinity, the appliance also uses the configured return method to return intercepted clientand/ or server-bound traffic to the WCCP router that redirected it, bypassing the routing table lookup. This is a useful feature if you have routing policies that may prevent your client- and/or server-bound traffic from reaching its destination and simplifies the appliance configuration process by eliminating the need to replicate these policies on the appliance. It is also useful in configurations where you have multiple home routers or where your WCCP router is multiple hops away from the appliance because it ensures that the traffic is always returned to the same WCCP router that redirected it.

Keep in mind, however, that enabling this feature unnecessarily when using GRE return does add additional CPU overhead on the router due to the need to decapsulate the GRE packets. In addition, the appliance and the router use a reduced MTU for GRE packets, which reduces the amount of data that can be transferred per packet.

Assignment Types

For every service group, you must configure the way the router determines the appliance to which to redirect a given packet, by setting an assignment type on the appliance. When the service group is formed, the appliance with the lowest IP address automatically becomes the designated cache (and if there is only one appliance in the service group, it is automatically the designated cache). The designated cache is responsible for communicating the assignment settings to the router, that is which appliance should be assigned a particular packet.

The appliance supports two assignment types:

- Hash Assignment (Default): With hash assignment, the designated cache assigns each appliance in the service group a portion of a 256-bucket hash table and communicates the assignment to the routers in the group. When the router receives a packet for redirection, it runs the hashing algorithm against one or more of the fields in the packet header to determine the hash value. It then compares the value to the hash assignment table to see which appliance is assigned to the corresponding bucket and then forwards the packet to that appliance. When you configure the service group on the appliances, you specify which field(s)—destination IP address, destination port, source IP address, and/or source port—should be used to calculate the hash value.

  In some cases, since all of the packets are hashed using the same fields and algorithm, it is possible that one of the caches in the group can become overloaded. For example, if you have a large proportion of traffic that is directed to the same server and you are using the destination IP address to run the hashing function, it is possible that the bulk of the traffic will be redirected to the same appliance. Therefore, you can configure an alternate field or group of fields to use to run the hashing algorithm. The router will then use this alternate hashing algorithm if the number of GRE packets or MAC addresses (depending on the forwarding method you're using) redirected to a given appliance exceeds a certain number.

  For details on configuring a hash-weight value to adjust the proportion of the hash table that gets assigned to an appliance, see "WCCP Load Balancing " below below.

- Mask Assignment: With mask assignment, each router in the service group has a table of masks and values that it uses to distribute traffic across the appliances in the service group. When the router receives a packet, it performs a bitwise AND operation between the mask value and the field of the packet header that is designated in the appliance mask assignment configuration. It then compares the result against its list of values for each mask; each value is assigned to a specific appliance in the service group.

WCCP Load Balancing

Each ProxySG appliance in the service group is assigned roughly an even percentage of the load by default, regardless of assignment type. If you would like to adjust or balance the load across multiple appliances, you can assign a weight value to each appliance in the group. ProxySG appliances with higher weight values receive a larger portion of the redirected traffic.

For example, suppose you have assigned the following weight values: ProxySG appliance1=100, ProxySG appliance2=100, and ProxySG appliance3=50 respectively. The total weight value is 250, and so ProxySG appliance1 and ProxySG appliance2 will each receive 2/5 of the traffic (100/250) and ProxySG appliance3 will receive 1/5 of the traffic (50/250).

If a ProxySG appliance becomes unavailable, the load will automatically be redistributed across the remaining ProxySG appliances in the service group.

See Also:

- "Requirements for Configuring WCCP" on the next page
- "Configure Service Groups" on page 130

Requirements for Configuring WCCP

Before you configure WCCP on the ProxySG appliance, complete the following tasks:

- Plan your service groups by determining:
  - Which routers and which appliances will work together in the redirection scheme.
  - The WCCP capabilities that your router/switch supports. Refer to the documentation that came with your router for the specifics on your router/switch.
  - What traffic you want to redirect. Do you want to redirect all traffic, or just a specific protocol or specific ports? Do you want to exclude certain hosts or traffic from redirection?
  - What forwarding and return method you plan to use and make sure that all the routers in the service group support the chosen method(s).
  - If you want to enable router affinity so that the appliance uses the chosen return method to return intercepted server- and/or client-bound traffic to the originating WCCP router as well as bypassed traffic.
  - How the router will assign a specific redirected packet to an appliance. Make sure the router(s) in the service group support the assignment method you plan to use. If there is more than one appliance in the service group, decide whether you want to distribute traffic equally, or if you want to assign varying weights.
- Configure the routers. For information on the feature sets and the capabilities of your router and for instructions on how to configure WCCP on the router, refer to the router documentation. For sample router WCCP configurations, refer to the *WCCP Reference Guide*.

Configure WCCP

Make sure you have met requirements before proceeding with WCCP configuration.

View the Current WCCP Configuration

To view the current WCCP configuration:

1. In the Admin Console, select **Configuration > Network > WCCP**.
2. In **WCCP Settings**, select **Show Source**.

   The console displays the WCCP settings source.

WCCP Configuration Example

The following is an example of an enabled WCCP configuration:

```
wccp enable
wccp version 2
service-group 9
  forwarding-type L2
  returning-type GRE
  router-affinity both
  assignment-type mask
  mask-scheme source-port
  priority 1
  protocol 6
  service-flags ports-defined
  ports 80 21 1755 554 0 0 0 0
  interface 0:0
  home-router 10.16.18.2
end
```

For descriptions of the values in the configuration file, refer to the *WCCP Reference Guide*.

Install WCCP Settings

You must create a WCCP configuration file on the appliance that contains the WCCP settings specific to the appliance. When installed, these configuration settings enable the appliance to collaborate with the WCCP-capable router or switch. To configure WCCP, you create a WCCP installation file and install it using one of the following methods:

- Text editor, which allows you to enter settings (or copy and paste the contents of an existing file) directly onto the appliance.
- Local file, to which the appliance can browse and install.

- Remote URL, where you put an already-created file on an FTP or HTTP server to be downloaded to the appliance.

  Use the `# (config) wccp`*`url`* command to set the path and the `# load wccp-settings` command to load the WCCP settings.

- CLI command `# inline wccp-settings`*`EOF`* to install the wccp configuration on the appliance.

For more information on the CLI commands, refer to the *Command Line Interface Reference*.

To install WCCP settings:

> **Note:** Do not enable WCCP until after you have installed WCCP settings.

1. Create the WCCP configuration file.
2. In the Admin Console, select **Configuration > Network > WCCP**.
3. In **WCCP Settings**, choose a method to install WCCP settings:

Method	Instructions
Text editor	1. Select **Install from text editor**. 2. In the Load Edit and Install from text editor dialog, enter the contents manually or copy and paste the contents from another file. 3. Select **Install**.
Local file	1. Select **... > Local File**. 2. In the Load Edit and Install from local file dialog, browse to the directory of the file on the local system. 3. Select **Install**.
Remote URL	1. Select **... > Remote URL**. 2. In the Load Edit and Install from remote URL dialog, enter the fully-qualified URL, including the filename, where the file is located. 3. Select **Download**.

Enable WCCP

WCCP is disabled by default. After you have installed WCCP settings, you can enable WCCP.

To enable WCCP:

1. In the Admin Console, select **Configuration > Network > WCCP**.
2. In **WCCP**, switch the **Enable WCCP** toggle.

Note: To exclude a ProxySG appliance from receiving traffic or from participating in any of the services groups configured on it, you can disable WCCP on the appliance. Disabling WCCP does not delete the WCCP configuration settings; it makes them out-of-service until WCCP is re-enabled on the appliance.

Configure Service Groups

Make sure you have met requirements before proceeding with service group configuration.

In the Admin Console (**Configuration > Network > WCCP**), the **Web Cache Service Groups** section lists all configured service groups on the appliance and displays basic information about each one:

Setting or field	Description
Group Number	Unique identifier (from 0 to 255).
Home Router	Home router address.
Ports	Ports to redirect on: All, Source, or Destination.

To configure WCCP service groups:

1. In the Admin Console, select **Configuration > Network > WCCP**.
2. In **Web Cache Service Groups**:
   - To add a service group, select **Add Aggregate**.
   - To edit an existing service group, select it in the list.

     The Add/Edit Web Cache Service Group page opens.
3. Enter or edit the following settings as needed:

Section	Setting or field	Description
Number, Priority	**Service Group Number**	Unique identifier (from 0 to 255).
	Priority	(Optional) Select a service group priority. When multiple service groups that are redirecting the same traffic (such as HTTP on port 80) are assigned to a common router interface, the priority defines the order in which the router evaluates the service groups.
Password	**Use Password**	Switch the **Use password** toggle to configure MD5 authentication for added security. The password can include 1 to 8 characters. When authentication is enabled on the router, the ProxySG appliance must be configured with the same password to join the service group.

Section	Setting or field	Description
Interfaces	**Interface**	Physical or virtual interface to which the service group is applied. Virtual interfaces are depicted as *adapter:interface.vlan_id*, for example, 0:0.3. Select **Add** to apply the service group to more interfaces. To remove interfaces, delete them from the list.
	Weight	(Optional) Specify a weight value (from 1 to 255) to determine the proportion of traffic that the router redirects to this interface. Use this field only if you are redirecting traffic in the same service group to multiple interfaces or to multiple ProxySG appliances and you want to allocate the percentage of traffic redirected to each appliance and/or interface in the service group.
Redirect	**Redirect On**	Ports to redirect on: **All**, **Source**, or **Destination**.

Section	Setting or field	Description
Forwarding, Returning	**Forwarding Type**	How the router and the ProxySG appliance handle packet forwarding. Select GRE or L2 forwarding. For the GRE forwarding method, the ProxySG appliance only supports GRE return. For a description of these options, see "Packet Forward and Return Methods" on page 123.
	Returning Type	How the router and the ProxySG appliance handle packet return. (Only applicable if you select L2 forwarding) Select GRE or L2 returning. For a description of these options, see "Packet Forward and Return Methods" on page 123.
	Router Affinity	(Optional) To ensure that intercepted traffic is always routed through the WCCP router that redirected it, select a Router affinity value: ■ **Client** indicates that the appliance will return client-bound traffic to the originating WCCP router using the configured Returning Type. ■ **Server** indicates that the appliance will return server-bound traffic to the originating WCCP router using the configured Returning Type. ■ **Both** indicates that the appliance will return both client- and server-bound traffic to the originating WCCP router using the configured Returning Type. ■ **None** (the default) indicates that the appliance will use regular routing table lookups rather than the configured Returning Type to route the client- and server-bound traffic that it intercepts.
	Home Router	Home router address. Specify individual unicast or a single multicast address for the router(s) in the service group: ■ To use multicast addressing, select **Multicast Home Router** and enter the **Group Address**, Optionally, enter a **Multicast TTL** value. ■ To use unicast addresses, select **Individual Home Router Addresses**. For each router in the service group, select Add and enter a **Home Router** address. The home router address that you use for a service group on the appliance should be consistent with the IP address (virtual or physical) over which the appliance communicates with the router.
	Group Address	Group address for multicast addressing. Specify if **Multicast Home Router** is selected for Home Router.
	Multicast TTL	How long a multicast packet is propagated. Default value is 1. Specify if **Multicast Home Router** is selected for Home Router.

Section	Setting or field	Description
Assignment Type	**Hash**	Instructs the router to use hash to distribute redirected traffic, using packet header information. You can select a different assignment method for each service group configured on the same appliance. To use hash: 1. Select **Hash**. 2. For **Primary Hash**, select one or more packet header fields. 3. (Optional) For **Alternate Hash**, select one or more header fields. The alternate hashing function is used to distribute traffic when a particular appliance exceeds a given number of redirected packets.
	Mask	Instructs the router to use maskto distribute redirected traffic, using packet header information. You can select a different assignment method for each service group configured on the same appliance. To use mask: 1. Select **Mask**. 2. For **Mask Scheme**, select the packet header field to use to run the mask function. 3. Enter a **Value** in either decimal or, when prefixed by 0x, a hexadecimal value. The default value for this field is 0x3f. The following Cisco Web page describes the mask value in detail: https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/white_paper_c11-629052.html

Configure a Forwarding System

You can configure forwarding hosts to which client requests can be redirected. Those hosts can be servers or proxies. Rules to redirect requests are set up in policy. High-level steps to configure forwarding are:

- Create the forwarding hosts and lists, including parameters such as protocol agent and port.

  A forwarding host is an appliance configured as a proxy to which certain traffic is redirected for the purpose of leveraging object caching to minimize trips to the Internet and over WAN links. After hosts are created, you can add them to forwarding groups.

- Set load balancing and host affinity values.

See Also:

- "Configure Forwarding Hosts " on page 142
- "Configure Forwarding Lists" on page 143

General Settings

Configure general forwarding settings that apply to all connections. Unlike global defaults for hosts and lists, these settings apply per connection and can't be overridden in individual host or list configurations.

To configure general settings:

1. In the Admin Console, select **Configuration > Networking > Forwarding Hosts**.
2. In Settings, enter or edit the following as needed:

Setting or field	Description
If no forwarding host available	Specify how connections behave if no forwarding is available, **Connect directly (fail open)** or **Deny the request (fail closed)**. Fail open is insecure; fail closed is the default setting. This setting can be overridden by policy.
Use forwarding for administrative downloads	Specify whether forwarding is applied to requests generated for administrative reasons on the system, such as downloading policy files or new system images. If enabled (default setting), you can control the forwarding in policy as needed. This option also affects the use of SOCKS gateways.
Integrated host timeout	An integrated host is an OCS that has been added to the health check list through the `integrate_new_hosts()` CPL property. The host ages out after being idle for the specified time. Default is 60 minutes.
Host affinity timeout	The amount of time a user's IP address, SSL ID, or cookie remains valid after its most recent use. The default is 30 minutes, meaning that the IP address, SSL ID or cookie must be used once every 30 minutes to restart the timeout period.

About Load Balancing

Load balancing distributes forwarding traffic among multiple IP addresses to achieve optimal resource utilization, to maximize throughput, and to minimize response time. Typically, you use load balancing to distribute requests to more than one ProxySG appliance, although you can also distribute requests to multiple IP addresses on a single appliance–or a combination of the two.

ProxySG load balancing methods include round robin–which selects the next system in the list–or least connections–which selects the system with the fewest number of connections. You can configure load balancing in any of the following ways:

- For hosts: If a host is DNS-resolved to multiple IP addresses, then that host's load-balancing method (round robin, least connections, or none) is applied to those IP addresses. The method is either explicitly set for that host or taken from the configurable global default settings.
- For lists: Load balancing for groups works exactly the same as load balancing for hosts with multiple IP addresses except there are two additional load balancing methods for groups:
    - URL hash: Requests are hashed based on the request URL.
    - Domain hash: Requests are hashed based on the domain name in the request.

About Host Affinity

Host affinity is the attempt to direct multiple connections by a single user to the same group member. Host affinity causes the user's connections to return to the same server until the configurable host affinity timeout period is exceeded.

For example, suppose a Web site with a shopping cart has several load-balanced Web servers, but only one Web server has the session data for a given user's shopping cart transaction. If a connection is sent to a different Web server that has no data about the user's session, the user has to start over. ProxySG host affinity helps make sure each request goes to its proper destination; however, the proxy does not interact with the session or with session data.

Host affinity allows you to use the following options:

- Client IP address: Use the client IP address to determine which group member was last used. When the same client IP sends another request, the host makes the connection to that group member.
- Accelerator cookie: Insert a cookie in the response to the client. When the client makes future requests, the cookie data is used to determine which group member the client last used. The host makes the connection to that group member.
- SSL session ID: For HTTPS, extract the SSL session ID name from the connection information. The host uses the session ID in place of a cookie or client IP address to determine which group member was last used. The host makes the connection to that group member.

See Also:

- "Configure Forwarding Hosts " on page 142
- "Configure Forwarding Lists" on page 143
- "Set Global Defaults for Load Balancing and Host Affinity" on page 145
- "Set the Forwarding Default Sequence" on page 147

About the Forwarding System

Forwarding creates a proxy hierarchy, which consists of a set of proxies (including ProxySG appliances that are configured as proxies (see "About Proxy Services " on page 150). Appliances close to the origin server perform object caching for server content and distribute the content to the object caches of other proxies that are farther away from the origin server. If

forwarding is set up in an organized manner, the load involved with object caching is distributed throughout the proxy hierarchy, which avoids sending any piece of content across any given WAN link more than once.

Forwarding redirects content requests to IP addresses other than those specified in the requesting URL. Forwarding affects only the IP address of the upstream device to which a request is sent; forwarding does not affect the URL in the request.

The ProxySG forwarding system consists of forwarding, upstream SOCKS gateways, load balancing, host affinity, and health checks. The forwarding system determines the upstream address where a request is sent, and is tied in with all the protocol proxies.

> **Note:** The ProxySG forwarding system directly supports the forwarding of HTTP, HTTPS, FTP, MMS, RTSP, Telnet, and TCP tunnels.

Example of a Forwarding System

This section discusses an example of using forwarding to minimize traffic over WAN links and to the Internet by leveraging object caching on proxies in the forwarding system.

> **Note:** The ProxySG forwarding system directly supports the forwarding of HTTP, HTTPS, FTP, MMS, RTSP, Telnet, and TCP tunnels.

The example discussed in this section uses the following logical proxy hierarchy:

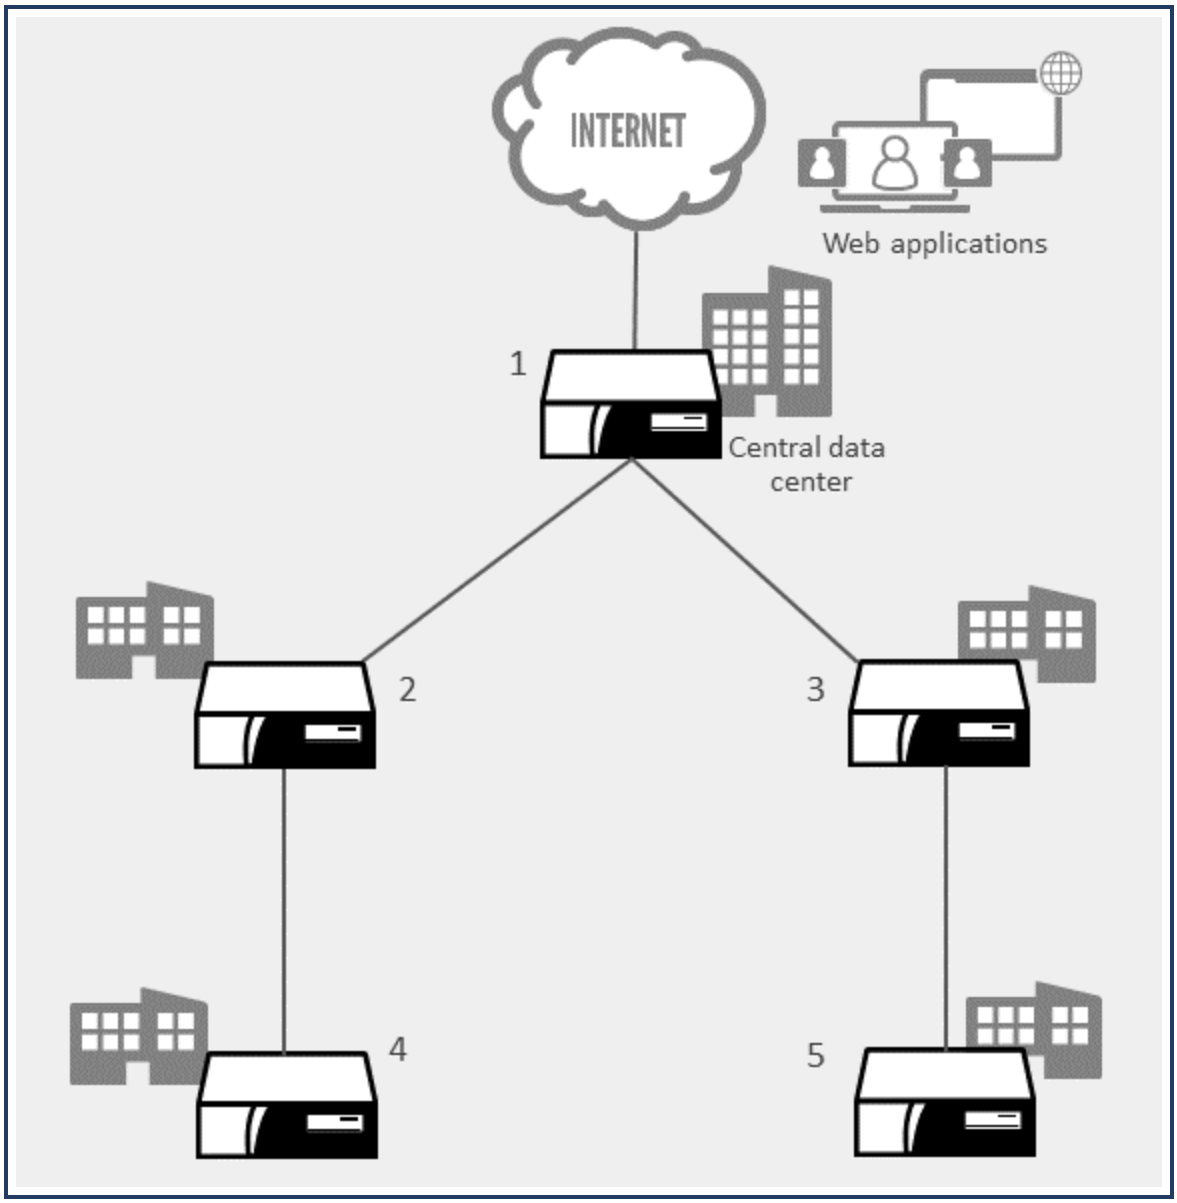

There are five ProxySG appliances configured as proxies: one in the central data center and one apiece in four branch offices or sites.

Appliance 1, located in the central data center, provides Internet access for the entire system. Appliance 4 uses Appliance 2 as its forwarding host and Appliance 2 uses Appliance 1 as its forwarding host. Similarly, Appliance 5 uses Appliance 3 as its forwarding host and Appliance 3 uses Appliance 1 as its forwarding host. This means that, for example, any piece of content in Appliance 1's object cache can be distributed to Appliance 2 or Appliance 3's object cache without having to send the content over the Internet.

The following figure shows a more detailed view of the example network.

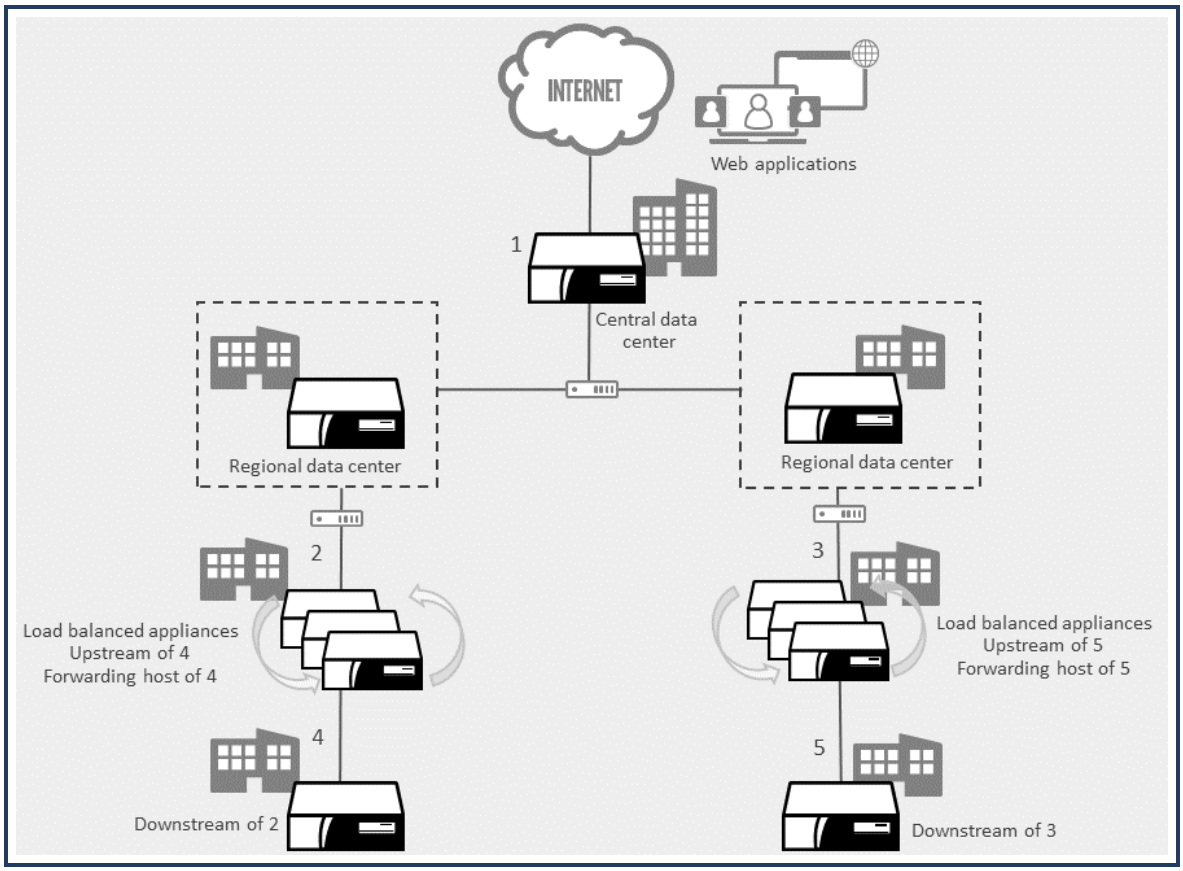

- Appliance 1 (located in the central data center) acts as a gateway to the Internet; in other words, all Internet access goes through Appliance 1. Two regional data centers accept requests from four branch offices or sites, each with appliances configured as a proxies.
- Appliance 1 is the gateway to the Internet, so it is upstream of all other Appliances shown above.
- Load-balanced Appliance 2 and Appliance 3 are configured to use Appliance 1 as their forwarding host, so they are downstream of Appliance 1.
- Appliance 4 is configured to use the load-balanced group of Appliances labeled 2 as its forwarding host, so Appliance 4 is downstream of both Appliance 2 and Appliance 1.
- Appliance 5 is downstream of Appliance 3 and Appliance 1.

Another way of stating this, using Appliance 4 as an example, is that any request to the Internet goes through Appliance 2 and Appliance 1 instead of going directly to the host specified in the URL of the request.

How the Example Uses Object Caching

The following example shows how forwarding and object caching work together to minimize traffic over the example network's WAN links and to the Internet.

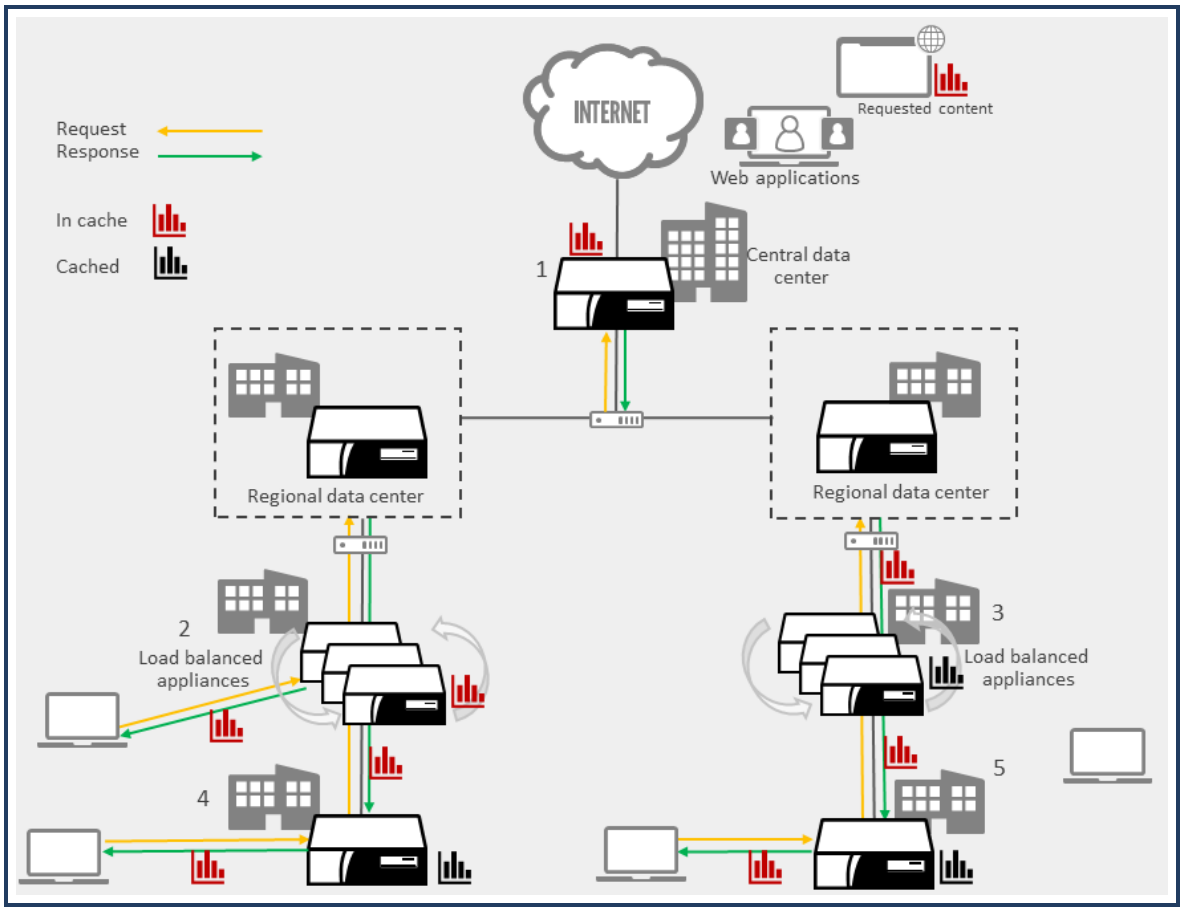

In this example:

- A user connected to Appliance 4 requests content located on a Web server in the Internet. The content—which might be a spreadsheet or multimedia—is in the object cache of load-balanced Appliance 2, and therefore is retrieved from the object cache. Neither the WAN links nor the origin server are used to retrieve the content. The content is then cached on Appliance 4's object cache so the next time a user requests the same content, it is retrieved from Appliance 4's object cache.
- If a user connected to Appliance 5 requests the same content—and the content is in neither Appliance 5's nor Appliance 3's object cache—load-balanced Appliance 3 gets the content from Appliance 1 and object caches it. Subsequently, Appliance 5 gets the content from Appliance 3 and object caches it.
- Because the content is in Appliance 1's object cache, the content is not retrieved from the origin server. In this scenario, only the WAN links are used; the Internet link is not used to retrieve the content.

Note: In the figures above, each appliance is assumed to use one IP address for forwarding. You could achieve similar results using load balancing if you configure a DNS host name as a forwarding host and used DNS load balancing to forward requests to more than one appliance. For more information, see "About Load Balancing " on page 141.

Configure a Forwarding System

You can configure forwarding hosts to which client requests can be redirected. Those hosts can be servers or proxies. Rules to redirect requests are set up in policy. High-level steps to configure forwarding are:

- Create the forwarding hosts and lists, including parameters such as protocol agent and port.

  A forwarding host is an appliance configured as a proxy to which certain traffic is redirected for the purpose of leveraging object caching to minimize trips to the Internet and over WAN links. After hosts are created, you can add them to forwarding groups.

- Set load balancing and host affinity values.

See Also:

- "Configure Forwarding Hosts " on page 142
- "Configure Forwarding Lists" on page 143

General Settings

Configure general forwarding settings that apply to all connections. Unlike global defaults for hosts and lists, these settings apply per connection and can't be overridden in individual host or list configurations.

To configure general settings:

1. In the Admin Console, select **Configuration > Networking > Forwarding Hosts**.
2. In Settings, enter or edit the following as needed:

Setting or field	Description
If no forwarding host available	Specify how connections behave if no forwarding is available, **Connect directly (fail open)** or **Deny the request (fail closed)**. Fail open is insecure; fail closed is the default setting. This setting can be overridden by policy.
Use forwarding for administrative downloads	Specify whether forwarding is applied to requests generated for administrative reasons on the system, such as downloading policy files or new system images. If enabled (default setting), you can control the forwarding in policy as needed. This option also affects the use of SOCKS gateways.
Integrated host timeout	An integrated host is an OCS that has been added to the health check list through the `integrate_new_hosts()` CPL property. The host ages out after being idle for the specified time. Default is 60 minutes.

Setting or field	Description
Host affinity timeout	The amount of time a user's IP address, SSL ID, or cookie remains valid after its most recent use. The default is 30 minutes, meaning that the IP address, SSL ID or cookie must be used once every 30 minutes to restart the timeout period.

About Load Balancing

Load balancing distributes forwarding traffic among multiple IP addresses to achieve optimal resource utilization, to maximize throughput, and to minimize response time. Typically, you use load balancing to distribute requests to more than one ProxySG appliance, although you can also distribute requests to multiple IP addresses on a single appliance—or a combination of the two.

ProxySG load balancing methods include round robin—which selects the next system in the list—or least connections—which selects the system with the fewest number of connections. You can configure load balancing in any of the following ways:

- For hosts: If a host is DNS-resolved to multiple IP addresses, then that host's load-balancing method (round robin, least connections, or none) is applied to those IP addresses. The method is either explicitly set for that host or taken from the configurable global default settings.
- For lists: Load balancing for groups works exactly the same as load balancing for hosts with multiple IP addresses except there are two additional load balancing methods for groups:
    - URL hash: Requests are hashed based on the request URL.
    - Domain hash: Requests are hashed based on the domain name in the request.

About Host Affinity

Host affinity is the attempt to direct multiple connections by a single user to the same group member. Host affinity causes the user's connections to return to the same server until the configurable host affinity timeout period is exceeded.

For example, suppose a Web site with a shopping cart has several load-balanced Web servers, but only one Web server has the session data for a given user's shopping cart transaction. If a connection is sent to a different Web server that has no data about the user's session, the user has to start over. ProxySG host affinity helps make sure each request goes to its proper destination; however, the proxy does not interact with the session or with session data.

Host affinity allows you to use the following options:

- Client IP address: Use the client IP address to determine which group member was last used. When the same client IP sends another request, the host makes the connection to that group member.
- Accelerator cookie: Insert a cookie in the response to the client. When the client makes future requests, the cookie data is used to determine which group member the client last used. The host makes the connection to that group member.
- SSL session ID: For HTTPS, extract the SSL session ID name from the connection information. The host uses the session ID in place of a cookie or client IP address to determine which group member was last used. The host makes the connection to that group member.

See Also:

- "Configure Forwarding Hosts " below
- "Configure Forwarding Lists" on the facing page
- "Set Global Defaults for Load Balancing and Host Affinity" on page 145
- "Set the Forwarding Default Sequence" on page 147

Configure Forwarding Hosts

To configure forwarding hosts:

1. In the Admin Console, select **Configuration > Forwarding > Forwarding Hosts**.
2. In Forwarding Hosts:
   - Select **Add Host** to add a new host.
   - Select an existing host to edit it.
3. On the Add New /Edit Host page, enter or edit the following as needed:

Section	Setting or field	Description
Properties	**Alias**	Unique name for the forwarding host. **Note:** Because the forwarding host alias is used in policy, the alias cannot be a CPL keyword, such as `no`, `default`, or `forward`.
	Host	Host's fully qualified domain name (FQDN) or IPv4/IPv6 address.
	Type	Select the forwarding host type: - **Server**: Use for reverse proxy deployments. This choice means the appliance the relative path for URLs in the HTTP header because the next hop is a web server, not a proxy server. HTTPS, TCP tunnels, and Telnet can be forwarded to a server only; they cannot be forwarded to a proxy. - **Proxy**: Use for forward proxy deployments.
	Ports	Select the protocol(s) to forward. When **Server** is selected, more protocols are available. When selected, a protocol's default port is (except for TCP, which has no default). Change the port if needed.

Section	Setting or field	Description
Load balancing and host affinity	**Use balancing for Forwarding Hosts**	Specify the load balancing method to override the global setting, or accept the global default: ■ **Round Robin**: Forward the request to the next host in the sequence. ■ **Least Connections**: Forward the request to the host that currently has the least number of connections.
	Host affinity methods	Select host affinity methods to override the global setting, or accept the global default.See "Set Global Defaults for Load Balancing and Host Affinity" on page 145.

If a setting says **Override**, it is set to override the global default and shows the current value. If a setting says **Default**, it inherits the global value that's displayed.

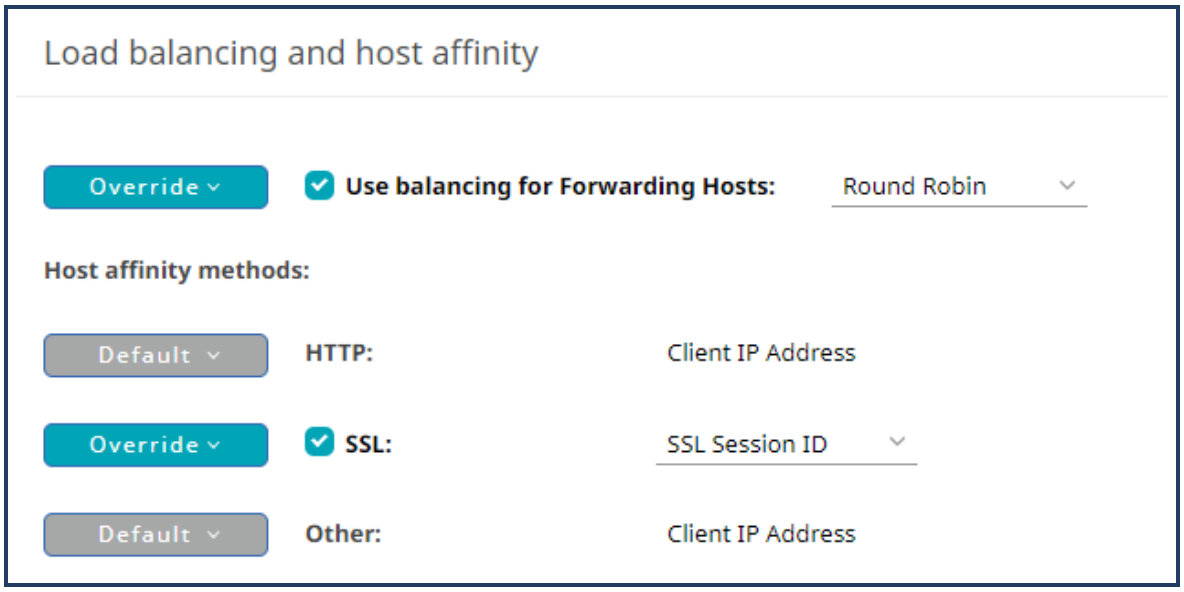

See Also:

- "Set Global Defaults for Load Balancing and Host Affinity" on page 145

Configure Forwarding Lists

To configure forwarding lists:

1. In the Admin Console, select **Configuration > Forwarding > Forwarding Hosts**.
2. In Forwarding Lists:

- Select **Add List** to add a new list.
- Select an existing list to edit it.

3. On the Add New List page, enter an **Alias** for the forwarding list.

   If editing a list, the alias can't be changed.

4. The Selected Forwarding Hosts section shows hosts belonging to the list. Select one of the following:

   - **Select Hosts**: Adds hosts to the list. On the Select Forwarding Hosts dialog, select hosts or clear selections. The **Show Selected** toggle filters the view to show only selected hosts.
   - **Remove**: Removes the host from the list.

5. Review load balancing and host affinity settings.

Setting or field	Description
Use balancing for Forwarding Lists	Specify the load balancing method to override the global setting, or accept the global default: ■ **Round Robin**: Forward the request to the next list in the sequence. ■ **Least Connections**: Forward the request to the list that currently has the least number of connections. ■ **Domain Hash**: Hash requests based on the domain name in the request. ■ **URL Hash**: Hash requests based on the URL in the request.
Host affinity methods	Select host affinity methods to override the global setting, or accept the global default. See "Set Global Defaults for Load Balancing and Host Affinity" on the facing page.

If a setting says **Override**, it is set to override the global default and shows the current value. If a setting says **Default**, it inherits the global value that's displayed.

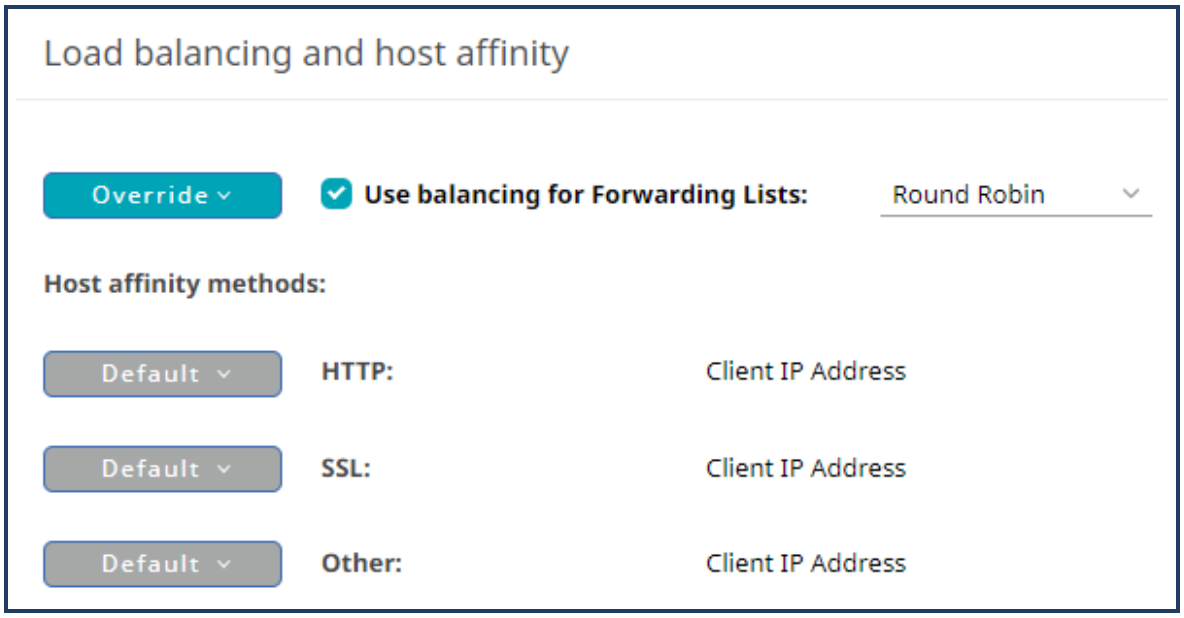

See Also:

- "Set Global Defaults for Load Balancing and Host Affinity" below

Set Global Defaults for Load Balancing and Host Affinity

The global defaults apply to all forwarding hosts and lists that are configured to use the **Default** setting. If hosts and lists have different settings specified, they override the global defaults.

To configure global defaults:

1. In the Admin Console, select **Configuration > Networking > Forwarding Hosts**.
2. In Global Defaults, enter or edit the following as needed:

Setting or field	Description
Use balancing for Forwarding Hosts	Specify the load-balancing method for all forwarding hosts except those whose configuration specifically override the global settings. ▪ **Round Robin**: Forward the request to the next host in the sequence. ▪ **Least Connections**: Forward the request to the host that currently has the least number of connections. Clear the **Use balancing for Forwarding Hosts** setting to specify no global default.

Setting or field	Description
Use balancing for Forwarding Lists	Specify the load-balancing method for all forwarding lists except those whose configuration specifically override the global settings. ■ **Round Robin**: Forward the request to the next list in the sequence. ■ **Least Connections**: Forward the request to the list that currently has the least number of connections. ■ **Domain Hash**: Hash requests based on the domain name in the request. ■ **URL Hash**: Hash requests based on the URL in the request. Clear the **Use balancing for Forwarding Lists** setting to specify no global default.
Host affinity methods	Select host affinity methods. See "Host Affinity Methods" below.

Host Affinity Methods

Symantec highly recommends that if you enable load balancing, you also enable host affinity. By default, if you use load balancing, each connection is treated independently. The connection is made to whichever member of the load-balancing group the load-balancing algorithm selects.

If host affinity is configured, the system checks host affinity first to see if the request comes from a known client. If this is a first connection, the load-balancing algorithm selects the group member to make the connection.

Host affinity records the result of the load balancing and uses it if that client connects again. Host affinity does not make a connection to a host that health checks report is down; instead, if host affinity breaks, the load-balancing algorithm selects a group member that is healthy and re-establishes affinity on that working group member.

Host affinity methods are discussed in the following table.

Setting	Description	HTTP	SSL	Other (TCP Tunnel or Telnet)
Global Default	Uses the default setting for all forwarding hosts on the system.	x	x	x
None	Disables host affinity.	x	x	x
Client IP Address	Uses the client IP address to determine which forwarding group member was last used.	x	x	x
Accelerator Cookie	Inserts a cookie into the response to the client.	x	x	
SSL Session ID	Used in place of a cookie or client IP address. Extracts the SSL session ID name from the connection information.		x	

Set the Forwarding Default Sequence

The default sequence is the forwarding sequence used when there is no matching forwarding rule in policy. The following is an example of forwarding policy:

```
<Forward>
  url.domain=symantec.com forward(FWGrp2, FWGrp1)
```

In the example, requests that match the URL domain symantec.com are sent to a forwarding group named FWGrp2 unless all of the members in FWGrp2 are down, in which case requests are sent to FWGrp1. Health checks are performed continually to minimize the possibility that requests are sent to a forwarding host or group that is known to be down.

The default sequence (and any sequence specified in policy) works by allowing healthy hosts to take over for an unhealthy host or group (one that is failing its DNS resolution or its health check). If more than one member is in the sequence, the sequence specifies the order of failover, with the second host or group taking over for the first one, the third taking over for the second, and so on.

If all of the hosts in the sequence are down, the request either fails open or fails closed (that is, the connection is denied). Symantec recommends you set this behavior in policy as follows:

```
forward.fail_open(yes|no)
```

You can also configure it using global defaults.

> **Note:** To create forwarding policy (including sequences), use the VPM or CPL. For information on using VPM, refer to the *Web Visual Policy Manager Reference*; for information on using CPL, refer to the *Content Policy Language Reference*.

To configure the default sequence:

1. In the Admin Console, select **Configuration > Networking > Forwarding Hosts**.

   The Default Sequence section shows the current default sequence, if one is configured.

2. Modify the default sequence as needed:

   - To add configured forwarding hosts or forwarding lists to the default sequence, select **Select Hosts, Lists**. On the dialog that opens, select hosts or lists.
   - To remove hosts or lists from the current default sequence, select **Remove** in the list of entries.
   - To change the order of entries in the sequence, use the Move Up or Move Down arrows in the entries.

   Save the changes to the default sequence.

> **Note:** Any host or list in the default sequence is considered in use by policy. As a result, if you try to delete a host or list while it is in the default sequence, you receive an error message. You must remove the host/list from the sequence before removing it from the forwarding configuration.

Manage ProxySG Services

You can configure two types of services on the appliance: proxy services and management services. Refer to the following topics to learn about proxy services and services groups and their roles in intercepting traffic:

- "About Proxy Services " on the next page
  - "Change the State of a Proxy Service" on page 154
  - "Configure a Proxy Service" on page 156
    - "Configure an HTTPS Reverse Proxy" on page 161
    - "Configure Listeners" on page 169
  - "Manage Service Groups" on page 170
  - "Exempt Requests From Specific Clients" on page 171
  - "Static Bypass and Dynamic Bypass" on page 174
  - "Troubleshooting: Bypass All Proxy Services or Restrict Interception" on page 177
  - "Reference: Proxy Services, Proxy Configurations, and Policy" on page 179
- General Proxy Settings
  - "Global Settings for Proxy Services" on page 182
- "Manage the SSL Proxy" on page 190
  - "Intercept HTTPS" on page 194
  - "Configure SSL Proxy" on page 195
  - "Configure SSL Rules through Policy" on page 200

About Proxy Services

> **Note:** By default, the Admin Console opens the **Configuration > Services > Proxy Services** page with **Proxy Services** selected for Interception. In normal operations, you should not need to change this setting. Set Interception to **Force Bypass** or **Restricted Intercept List** for troubleshooting only. See "Troubleshooting: Bypass All Proxy Services or Restrict Interception" on page 177 for details.

In Symantec terminology, proxy service defines:

- The combinations of IP addresses and ports that the proxy matches against.
- Whether to intercept or bypass matched traffic; if intercepted, which proxy to use to process the traffic.
  - When a service is set to **Intercept**, the ProxySG appliance listens on the port for traffic and upon detection, terminates the connection, performs an action (such as a policy check), and initiates a new connection to the traffic destination.
  - When a service is set to **Bypass**, the traffic pass through the appliance. By default, services are set to **Bypass**.
- A collection of attributes that control what type of processing the appliance performs on the intercepted traffic.

See Also:

- "Change the State of a Proxy Service" on page 154

About Proxy Service Groups

The ProxySG appliance groups services into predefined service groups based on the type of traffic that service carries. Service groups enable you to:

- Quickly locate a specific service and view its attributes.
- Create a custom service group and add custom services or existing services to that group.

Predefined Service Groups and Services

The following table lists all service groups and their associated services. This list applies to new installations or the result of restoring the appliance to factory defaults after the ab upgraded from a lower version. Upon upgrading to the current version, the Admin Console retains existing services, service group names, and policies.

For a complete list of supported proxy services and listeners, see "Reference: Proxy Services, Proxy Configurations, and Policy" on page 179.

Services Group	Description	Predefined Service Types or Examples
Standard	The most commonly intercepted services.	■ HTTP/HTTPS–external (transparent and explicit) and internal ■ Endpoint Mapper (for MAPI protocol–Microsoft Exchange) ■ CIFS (file sharing) ■ Streaming (MMS, RTSP) ■ FTP ■ DNS ■ SOCKS
Bypass Recommended	Services that contain encrypted data and therefore recommended to not be ADN-optimized; also includes other interactive services.	■ Cisco VPN ■ Symantec ADN ■ Symantec management ■ Oracle over SSL ■ Other encrypted services
Tunnel Recommended	Services that employ the TCP Tunnel proxy to provide basic application-independent acceleration.	Citrix, IMAP, LDAP, and various other common business applications.
Default	See "Default Listener" on the next page.	

To move a service to another group, edit the service. See "Configure a Proxy Service" on page 156.

See Also:

- "Configure an HTTPS Reverse Proxy" on page 161
- "Manage Service Groups" on page 170

> **Note:** The HTTPS Reverse Proxy service is available but not created by default. For information about configuring the service, see "Configure an HTTPS Reverse Proxy" on page 161.

About Proxy Service Listeners

A proxy service listener specifies where a ProxySG appliance service listens for traffic. The listener comprises four attributes:

- Source address: Most of the time, this attribute is set to all source addresses, which means any IPv4 or IPv6 address that originates the request. You can also specify specific IP addresses and subnets. For example, you want to exclude a network segment, so you specify a subnet and set to Bypass.
- Destination address: All addresses, which means any IPv4 or IPv6 destination.
    - Transparent—Acts on connections without awareness from the client or server. Only connections to IPv4 or IPv6 destination addresses that do not belong to the appliance are intercepted. This setting requires a bridge, such as that available in the appliance; a Layer-4 switch, or a WCCP compliant router. You can also transparently redirect requests through an appliance by setting the workstation's gateway to the appliance IP address.
    - Explicit—Requires web browser and service configuration. It sends requests explicitly to a proxy instead of to the origin content servers. Only destination addresses that match one of the IPv4 or IPv6 addresses on the appliance are intercepted.
    - Destination IP address or subnet/prefix length—This listener type ensures that only destination addresses matching the IPv4/IPv6 address or subnet/prefix length are intercepted.
- Port: A specific port or port range. All default appliance services configured to their industry-standard ports. For example, the explicit HTTP service is configured to listen on ports 80 and 8080.
- Action: The aforementioned action to take on traffic detected by this service: Intercept or Bypass.

For a complete list of supported proxy services and listeners, see "Reference: Proxy Services, Proxy Configurations, and Policy" on page 179.

Default Listener

The **Default** listener detects any traffic that does not match any other listeners on any of the services.

Multiple Listeners

A listener identifies network traffic based on a source IP address or range, destination IP address or range, or both. Multiple listeners can be defined for a proxy service or console service. Each service has a set of default actions to apply to the traffic identified by the listeners it owns.

The destination IP address of a connection can match multiple proxy service listeners. Multiple matches are resolved using the most-specific match algorithm used by routing devices. A listener is more specific if it has a larger Destination IP subnet prefix. For example, the subnet 10.0.0.0/24 is more specific than 10.0.0.0/16, which is more specific than 10.0.0.0/8.

When a new connection is established, the ProxySG appliance first finds the most specific listener destination IP. If a match is found, and the destination port also matches, the connection is then handled by that listener. If the destination port of the listener with the most specific destination IP does not match, the next mostspecific destination IP is found; this process continues until either a complete match is found or no more matching addresses are found. If a destination IP address is not specified, the closest matching explicit proxy service listener has priority over a subnet match. In that instance, the explicit proxy service listener handles the connection instead of the subnet listener. Explicit port 80 listeners with a destination host IP identical to the appliance have priority over other explicit listeners.

For example, assume the following services were defined as given in the following table.

Proxy Service Name and Type	Listener Source IP Address	Listener Destination IP Address	Listener Port Range
New York Data Center, HTTP proxy	192.168.20.22	10.167.10.0/24	80
New York CRM, HTTP proxy	All	10.167.10.2	80
HTTP Service, HTTP proxy	All	transparent	80

An HTTP connection initiated to server 10.167.10.2 could match any of the three listeners in the above table. The most specific match algorithm finds that a listener in the New York CRM service is the most specific and since the destination port of the connection and the listener match, the connection is handled by this service. The advantage of the most specific match algorithm becomes evident when at some later point another server is added in the New York Data Center subnet. If that server needs to be handled by a different service from the New York Data Center service, a new service with a listener specific to the new server would be added. The administrator does not need to be concerned about rule order in order to intercept traffic to this particular server using the new, most specific service listener.

See Also:

- "Configure Listeners" on page 169

Change the State of a Proxy Service

A proxy service can be in one of two states:

- Bypass: Traffic for this service passes through the ProxySG appliance without receiving an optimization or policy checking (as applicable).
- Intercept: The appliance intercepts traffic for this service and applies optimization or policy checks (as applicable).

The state of a service depends on whether the appliance was upgraded:

- Upon upgrade from a previous release, supported services remain in their original service groups and retain their bypass/intercept states.
- On a new installation or a re-initialization, all services are set to Bypass unless otherwise specified during the installation/re-initialization process.

You cannot change the state of entire predefined group; you must set each service required for your deployment to Intercept.

Changing the state of a service to Intercept is only the first step in configuring a protocol proxy. To achieve your corporate deployment goals, you must also configure the proxy settings and define policy, both of which determine how the appliance processes the intercepted traffic. These aspects are discussed in each proxy section later in this guide.

For more conceptual information about services, see "Manage ProxySG Services" on page 149.

To change the state of a service:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. In **Proxy Services**, select the arrow beside a service name. The row expands to display more information about the selected service.
3. For each configured listener, set the Action to **Intercept** or **Bypass**.

   In the following example, both listener ports for the Explicit HTTP service are set to Intercept.

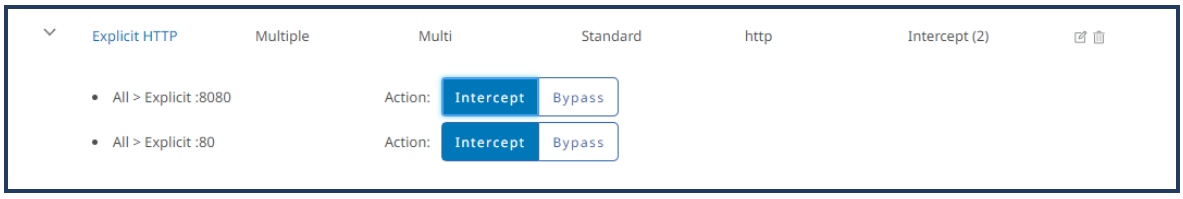

You can also change the Intercept/Bypass action when you edit a service. See "Configure a Proxy Service" on page 156.

Tip: Setting a service to Intercept is one step in controlling specific traffic types. There are other options for the services themselves, plus proxy configurations and policy definitions. You can also create custom services and service groups.

See Also:

- "Configure a Proxy Service" on the next page
- "Manage Service Groups" on page 170
- "Reference: Proxy Services, Proxy Configurations, and Policy" on page 179
- *Content Policy Language Reference*

Configure a Proxy Service

You can configure settings for a predefined proxy service, an imported proxy service, or a new proxy service that you created for a custom application.

- Predefined services: Services that were defined on the system upon installation.
- Imported services: Import a proxy service if you delete a default service and want to re-add it. If you import an existing service, you are prompted to confirm the replacement of a service. Existing service settings are overwritten with the default settings.

  > **Note:** After an SGOS upgrade, you must import any services added to the service library in order to use them.

- Custom services: You can create custom proxy service groups and populate them with custom services or move default services to them. For example, this ProxySG appliance serves a specific purpose and you want a custom group that contains only those services.

Before proceeding, make sure that you understand the goal of your deployment, how the application proxy operates, and the IP addresses (source and/or destination) and ports to intercept. Some proxy services, such as DNS, are simple–composed only of IP addresses and ports. Others, such as HTTP, have more attributes to consider. For a high-level description of these options, see "About Proxy Service Attributes" on page 158.

To configure a proxy service:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. In **Proxy Services:**
   - To add a service: Select **New Service**.
   - To edit a service: Select the existing service.
   - To import a service: Select **Import Service**.

   The Add New/Edit Proxy Service page displays service and listener settings.

3. Enter or edit the following settings as needed:

Setting or field	Description
Name	(Creating new or editing existing services) Name for the proxy service.
Service to Import	(Importing services) Service to import from library.
Service Group	Select a predefined service group or a custom group. Predefined services are not anchored to their default groups. You can move a service to any other predefined or custom group. **Note:** The entire service is moved; you can't move individual listeners.
Proxy Settings	The selected proxy type determines what, if any, proxy-specific fields are available: ■ CIFS: See "ADN Optimizations" on the next page and "Early Intercept" on page 159. ■ DNS: No other settings are needed. ■ Endpoint Mapper: See "ADN Optimizations" on the next page. ■ FTP: See "ADN Optimizations" on the next page. ■ FTPS: See "ADN Optimizations" on the next page. ■ HTTP: See "ADN Optimizations" on the next page, "Authenticate-401" on page 159, and "Protocol Detection" on page 160. ■ HTTPS reverse proxy: See "ADN Optimizations" on the next page and "Configure an HTTPS Reverse Proxy" on page 161. ■ MMS: No other settings are needed. ■ RTMP: No other settings are needed. ■ RTSP: No other settings are needed. ■ SOCKS: See "Protocol Detection" on page 160. ■ SSL: See "ADN Optimizations" on the next page and "Protocol Detection" on page 160. ■ TCP-tunnel: See "ADN Optimizations" on the next page, "Early Intercept" on page 159, and "Protocol Detection" on page 160. ■ Telnet: No other settings are needed.

4. To add or edit listeners, see "Configure Listeners" on page 169.

About Proxy Service Attributes

In addition to the listener information, each service contains one or more settings that affect how the appliance proxies the traffic. The following sections provide an overview of those settings. The proxy configuration topics provide more information about these attributes.

ADN Optimizations

Applies to the HTTP, HTTPS, CIFS, Endpoint Mapper, FTP, FTPS, SSL, and TCP Tunnel proxies in supported deployments.

Controls whether ADN optimizations—byte caching and/or compression—are enabled on a specific service. Note that enabling these ADN optimizations does not guarantee accelerated connections. It depends on ADN routing (for explicit deployments) and network configuration (for transparent deployments).

Byte caching is an optimization that replaces byte sequences in traffic flows with reference tokens. The byte sequences and the token are stored in a byte cache on a pair of ProxySG appliances (for example, one at the branch, the other at the data center). When a matching byte sequence is requested or saved, the ProxySG appliance transmits the token instead of the byte sequence.

GZIP compression removes extraneous/predictable information from traffic before it is transmitted. The information is decompressed at the destination's ProxySG appliance.

The ADN settings are:

- **Enable ADN**: This setting does not guarantee acceleration for this service—it also depends on ADN routing (for explicit deployments) or networksetup (for transparent deployments).
- **Enable byte caching**: This acceleration technique replaces byte sequences in traffic flows with reference tokens and stores them in a byte cache on a pair of ProxySG appliances at each end of the WAN. When a matching byte sequence is requested again, the ProxySG appliance transmits a token instead of the byte sequence.
- **Enable compression**: Uses a variety of algorithms to remove extraneous/ predictable information from the traffic before it is transmitted. The information is reconstituted at the destination based on the same algorithms.

  To get the maximum benefit of ADN, both byte caching and compression should be enabled. In cases where byte caching may be causing issues for an ADN deployment, you can turn off the Enable byte caching option and just use compression (or vice versa). If you know the traffic for this proxy is already compressed or encrypted, you can conserve resources by clearing the Enable byte caching and Enable compression options.
- **Enable thin client processing**: Applies special treatment to application traffic from thin client applications (such as RDP, VNC, and Citrix). This processing improves responsiveness of thin client actions. For example, end-users will notice that the desktop displays significantly faster. In addition, thin client data is not retained in the byte cache as long as other types of data because this data is are more temporal in nature; the byte cache, therefore, can be used more efficiently for other types of traffic that can better leverage it.

This option is available for TCP Tunnel proxies only, and is only available when ADN is enabled and byte caching and/or compression is enabled. Retention priority and thin client processing are mutually exclusive settings; you cannot enable both options for a service.

For thin client processing to be most effective, you must deactivate the thin client's software-based encryption and compression.

- **Retention priority**: You can control how long data is stored in the byte cache dictionary by assigning a retention priority to a particular service. If you want to keep certain types of data in the dictionary for as long as possible, set a high retention priority for the service. Or for data that isn't likely to get much benefit from byte caching, you can set a low retention priority for the related service. Most services are set to normal priority by default. This option is available only if byte caching is enabled for the service.

  You can use this option to preserve the most relevant content in the byte cache in the face of continually incoming, competing byte cache data. For example, when an application is being used for backup, you may want to set the retention priority to high so that competing traffic doesn't evict the backup data. However, if an application is being used for data replication, you may want to set the service's retention priority to low as the data most likely will only be hit in the next short duration.

Authenticate-401

Available on the Explicit HTTP and External HTTP services.

When this option is selected, all transparent and explicit requests received on the port always use transparent authentication (cookie or IP, depending on the policy configuration).

If you have deployed Authentication in the way recommended by Symantec– where only the ProxySG appliance nearest the user performs the authentication tasks–configuring Authenticate-401 is not necessary. However, multiple explicitly-configured appliances in a proxy chain are all attempting to perform authentication tasks can cause issues with browsers. By forcing one of the proxies (recommended: the one furthest away from the client) to use 401-style authentication instead of the standard proxy 407-style authentication, the browser can better handle the multiple authentication challenges.

Early Intercept

Applies to the CIFS and TCP Tunnel services.

Opening a TCP connection involves a three-way handshake involving packets: the client contacts the server, the server acknowledges the client, and the client acknowledges the server.

- With early intercept, the appliance returns a server acknowledgment back to the client and waits for the client acknowledgment, which completes the TCP three-way handshake, before the appliance connects upstream to the server. Furthermore, proxies that support object caching (such as HTTP), the appliance serves from the cache–a server connection is not necessary.
- With delayed intercept, the appliance attempts to connect upstream immediately after receiving the client's initial connection request, but waits to return the server acknowledgment until determining whether or not the upstream connection succeeds. This provides greater transparency, as the client receives either an RST or no response, which mirrors what is sent from a server when connections fail.

For every proxy listener except CIFS and TCP Tunnel services, early intercept is hard-coded to enabled.

- For CIFS, the listener is hard-coded as delayed intercept because of a specific issue with the way clients attempt to connect to ports 139 and 445 simultaneously. Without a full transparency in our response to the TCP three-way handshakes, client connections might break.
- For TCP Tunnel, you have the option to select either (disabled by default). For the TCP Tunnel service, the Early Intercept option is selectable and disabled by default. When this option is disabled, the proxy delays responding to the client until after it has attempted to contact the server. For maximum transparency, disable this option. If reduced latency is more important, enable it.

Protocol Detection

Applies to the HTTP, HTTPS, SOCKS, and TCP Tunnel services.

Protocol detection identifies HTTP, SOCKS CONNECT requests, and TCP tunnels. You can enable protocol detection on the aforementioned services or implement it using policy. Policy can further be used to negate protocol detection for SSL requests. Defining a policy for protocol detection enhances granularity by matching on a richer set of conditions rather than the specific service; policy always overrides manual settings.

If protocol detection is enabled, the appliance inspects the first bytes sent from the client and determines if a corresponding application proxy is available to hand off the connection. For example, an HTTP request identified on a TCP tunnel has full HTTP policy applied to it, rather than just simple TCP tunnel policy. In particular, this means that:

- The request arrives as a client protocol HTTP rather than a TCP Tunnel.
- The URL used while evaluating policy is an http:// URL of the tunneled HTTP request, not a tcp:// URL to which the tunnel was connecting.
- Forwarding policy is applied based on the new HTTP request; therefore, the selected forwarding host selected support HTTP. A forwarding host of type TCP cannot handle the request, which forces the request to be blocked.

Enabling protocol detection helps accelerate the flow of traffic. However, the TCP session must be fully established with the client before either the application proxy or the TCP tunnel proxy contacts the origin server. In some cases, like in the active-mode FTP data connections, enabling protocol detection might cause a delay in setting up the connection.

To avoid this connection delay, either use a protocol specific proxy, such as the FTP proxy, or disable protocol detection.

If protocol detection is disabled, traffic flows over a TCP tunnel without acceleration provided by a protocol-specific proxy. The setting is disabled by default.

See Also:

- "Configure an HTTPS Reverse Proxy" on the facing page
- "Configure Listeners" on page 169

Configure an HTTPS Reverse Proxy

The Symantec HTTPS Reverse Proxy implementation:

- Combines hardware-based SSL acceleration with full caching functionality.
- Establishes and services incoming SSL sessions.
- Provides TLS v1.2, TLS v1.1, TLSv1, SSL v3.0, and SSL v2.0 support. In version 7.2 and later, supports TLS 1.3.
- Supports IPv6 connections.
- (In version 7.1 and later) Supports the HTTP/2 protocol (TLS must be enabled).

> **Note:** Reverse Proxy and Web Application Firewall solution guides that provide additional solutions, workflows, and configuration options can be found at: https://support.symantec.com/en_US/article.DOC10942.html

Prerequisite Tasks

Before creating an HTTP reverse proxy service, you must:

- Create or import a keyring (**Configuration > SSL > Keyrings and Keylists**).
- (If necessary) Create a CSR that can be sent to a CA. After the CSR has been signed and the certificate has been created, import the certificate to the keyring you created or imported.

Alternatively:

- Create a certificate for internal use and associate it with the keyring.
- (Optional, if using server certificates from CAs) Import Certificate Revocation Lists (CRLs) so the ProxySG appliance can verify that certificates are still valid. When these steps are complete, you can configure the HTTPS reverse proxy service.

Creating an HTTPS Reverse Proxy Service

> **Note:** One common scenario in using HTTPS reverse proxy, which connects the client to the ProxySG appliance, is in conjunction with HTTPS origination, which is used to connect to the OCS. For more information on this option, see "Configure HTTP or HTTPS Origination to the OCS" on page 165.

The ProxySG appliance does not create an HTTPS Reverse Proxy service by default. (The ProxySG appliance has an HTTPS proxy service configured on port 443.)

To create an HTTPS reverse proxy service:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. In **Proxy Services:**
   - To add a service: Select **New Service**.
   - To edit a service: Select the existing service.
   - To import a service: Select **Import Service**.

   The Add New/Edit Proxy Service page displays service and listener settings.

3. Enter or edit the following settings as needed:

Setting or field	Description
Name	(Creating new or editing existing services) Name for the proxy service.
Service to Import	(Importing services) Service to import from library.
Service Group	Select a predefined service group or a custom group. Predefined services are not anchored to their default groups. You can move a service to any other predefined or custom group. **Note:** The entire service is moved; you can't move individual listeners.
Proxy Settings	■ **Keyring**: Select a keyring. The system ships with a default keyring that is reusable for each HTTPS service. The list also includes keylists configured on the appliance. **Note:** The configuration-passwords-key keyring that shipped with the ProxySG appliance does not contain a certificate. The appliance-key keyring does conta in a certificate if you have Internet connectivity, but it cannot be used for purpos es other than appliance authentication. ■ **CCL**: Select a CA certificate list. ■ **SSL Protocols**: Select the version(s) to use for this service from the list. ■ **Verify Client**: Select to enable mutual SSL authentication. See "About Mutual SSL Authentication " below for information. ■ **Forward Client Cert**: (Available if **Verify Client** is selected) Select this option to put the extracted client certificate information into the Client-Cert header that is included in the request when it is forwarded to the origin content server. The header contains the certificate serial number, subject, validity dates, and issuer (all as name=value pairs). The actual certificate itself is not forwarded. ■ ADN options: See "ADN Optimizations" on page 158.

4. To add or edit listeners, see "Configure Listeners" on page 169.

About Mutual SSL Authentication

During an SSL handshake, the client and server negotiate the mode of operation, the type of authentication required by both parties, the cryptographic and hashing algorithms to use for providing confidentiality and integrity, and the compression algorithm to use for the session.

SSL authentication can use the following modes of operation:

- Typical SSL authentication: This mode provides confidentiality and integrity of the data sent between the client and the server, and requires the server to authenticate to the client using an X.509 certificate.
- Mutual SSL authentication: In this mode, the server authenticates to the client using an X.509 certificate and the client must authenticate to the server with a separate X.509 certificate.

  When a Common Access Card (CAC) is used, the certificate identifies the user who owns the CAC. For information on CAC authentication, refer to the *Common Access Card Solutions Guide*.

In mutual SSL authentication, an SSL connection between a client and a server is established only if the client and server validate each other's identity during the SSL handshake. Both the server and the client must have their own valid certificate and the associated private key in order to authenticate.

> **Note:** TLS is supported based on the server and client in use. For brevity, this section refers only to SSL; however, SSL can be used interchangeably with TLS.

Typical SSL Authentication

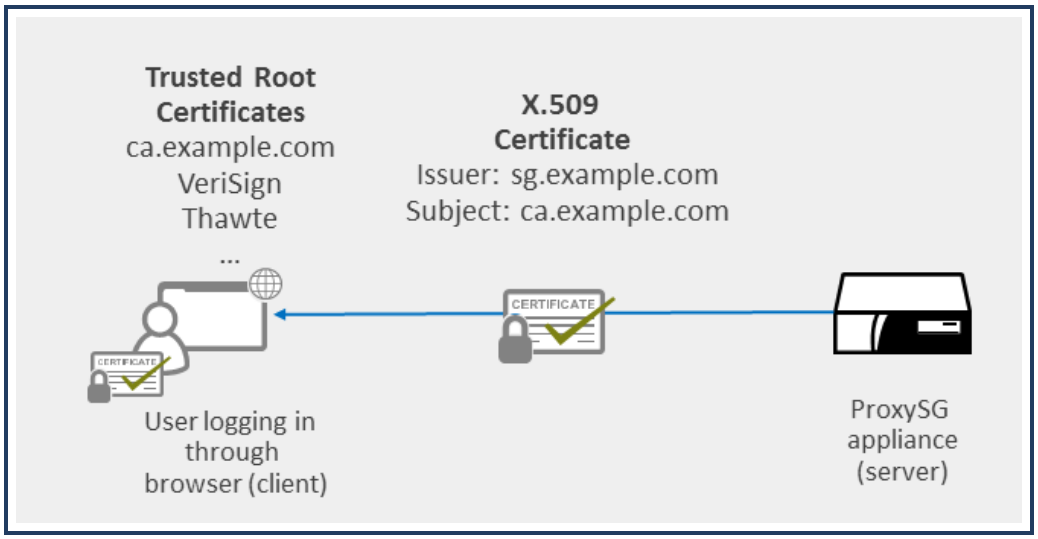

In this scenario, the user logs in to the ProxySG appliance (server) using a browser (client). During this process, the client (browser) validates the server (ProxySG appliance) certificate. This includes the following checks:

- The certificate subject must match the server's hostname.
- The certificate must be issued by a CA listed in the browser's Trusted Root Certificate store.
- The client confirms that the server has the certificate's private key by challenging the server to sign random data. The client validates the signature using the server's certificate.

Mutual SSL Authentication

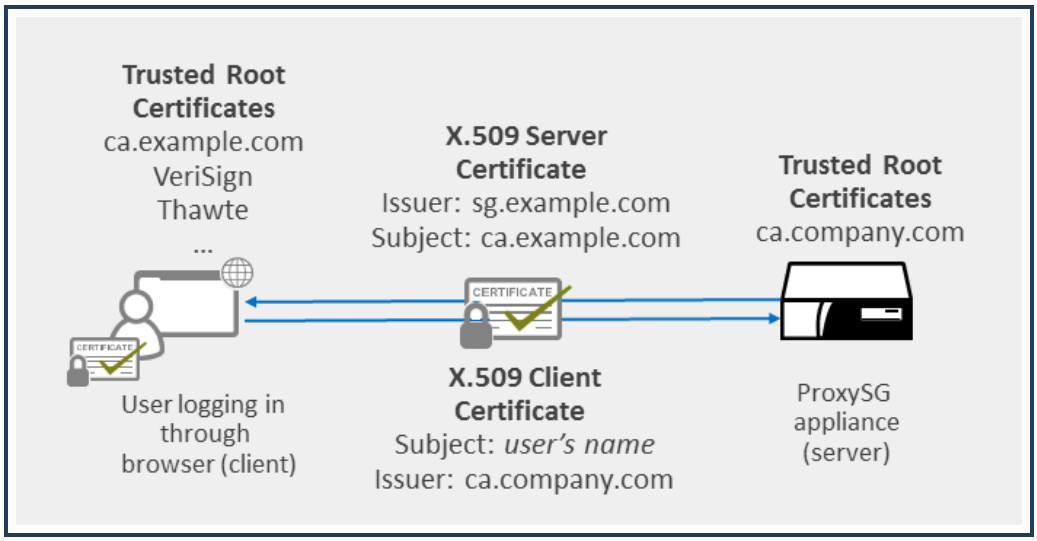

In this scenario, the user logs in to the ProxySG appliance using mutual SSL authentication. During this process:

1. The client (browser) validates the server (ProxySG appliance) certificate. This includes the following checks:
   - The certificate subject must match the server's hostname.
   - The certificate must be issued by a CA listed in the browser's Trusted Root Certificate store.
   - The client confirms that the server has the certificate's private key by challenging the server to sign random data. The client validates the signature using the server's certificate.
2. The server (ProxySG appliance) validates the client certificate that the browser presents. This includes the following checks:
   - The certificate must be issued by a CA in the CCL for the ProxySG appliance service that is performing the validation.
   - The server confirms that the client has the certificate's private key by challenging the client to sign random data. The server validates the signature using the client's certificate.
   - The certificate must be valid; it must have a valid signature and not be expired.
   - (If using a CRL) The certificate must not have been revoked.

Configure HTTP or HTTPS Origination to the OCS

In previous procedures, you configured HTTPS Reverse Proxy to the ProxySG appliance. In two common termination scenarios, you must also configure HTTPS origination to the OCS.

The first two scenarios are used to provide a secure connection between the proxy and server, if, for example, the proxy is in a branch office and is not co-located with the server.

Scenario 1: HTTPS Reverse Proxy with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG appliance	ProxySG appliance > HTTPS > OCS
Steps: ■ Configure a keyring. ■ Configure the SSL client. ■ Configure the HTTPS service.	Steps (optional): ■ Add a forwarding host. ■ Set an HTTPS port. ■ Enable server certificate verification.

Scenario 2: HTTP Termination with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG appliance	ProxySG appliance > HTTPS > OCS
Steps: ■ Client is explicitly proxied.	Steps: ■ Server URL rewrite. OR ■ Add a forwarding host. ■ Set an HTTPS port. ■ (Optional) Enable server certificate verification.

Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *Content Policy Language Reference*.

To configure HTTPS origination, use the following CLI command:

```
# (config forwarding) create host_alias host_name https[=port_number] server ssl-verify-server=yes
```

where:

- *host_alias*: Specifies the alias name of the OCS.
- *host_name*: Specifies the host name or IPv4/IPv6 address of the OCS, such as www.symantec.com.
- *port_number*: Specifies the port number on which the OCS is listening.

- **server**: Specifies to use the relative path for URLs in the HTTP header because the next hop is a Web server, not a proxy server. Proxy is the default.
- **ssl-verify-server**: Specifies whether the upstream server certificate should be verified. You can only enable this command if the upstream host is a server, not a proxy.

The next scenario is useful when the appliance is deployed as a reverse proxy. This scenario is used when it is not necessary for a secure connection between the proxy and server.

Scenario 3: HTTP Reverse Proxy with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG appliance	ProxySG appliance > HTTPS > OCS
Steps: ■ Configure a keyring. ■ Configure the SSL client. ■ Configure the HTTPS service.	Steps: ■ Server URL rewrite. OR ■ Add a forwarding host. ■ Set an HTTP port.

Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *Content Policy Language Reference*.

To configure HTTP origination, use the following CLI command:

```
# (config forwarding) create host_alias host_name http[=port_number] server
```

where:

- *host_alias*: Specifies the alias name of the OCS.
- *host_name*: Specifies the host name or IPv4/IPv6 address of the OCS, such as www.symantec.com.
- *port_number*: Specifies the port number on the OCS in which HTTP is listening.
- **server**: Specifies to use the relative path for URLs in the HTTP header because the next hop is a web server, not a proxy server. Proxy is the default.

Create Policy for HTTP and HTTPS Origination

Forwarding hosts must be already created on the appliance before forwarding policy can be created.

To create policy using CPL:

```
<Forward>
  url.host=host_name forward(host_alias)
```

To create policy using the VPM:

1. In the VPM, add a Forwarding layer.
2. Set the Destination to be the URL of the OCS.
3. Set the Action to forward to the forwarding host and configure parameters to control forwarding behavior.

For specific instructions for the legacy VPM or the web VPM, refer to the respective documentation.

Configure Listeners

You can configure listeners for a proxy service. For details, see "About Proxy Service Listeners" on page 151

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. In **Proxy Services**, select the existing custom service.

   The Edit Proxy Service page displays service and listener settings.
3. In Listeners:
   - To add a listener: Select **New**.
   - To edit a listener: Select the existing listener.

     The Add/Edit Listener dialog opens.
4. Enter or edit the following settings as needed:

Setting or field	Description
Source Address	The client addresses to which the service applies. The most common selection is **All**, which means the service applies to requests from any client (IPv4 and IPv6). You can also restrict this listener to a specific IP address (IPv4 or IPv6) or user subnet (for IPv4) or prefix length (for IPv6).
Destination Address	The correct selection might depend on network configuration. See "About Proxy Services " on page 150.
Port Range	A single port number or a port range on which this application protocol broadcasts. For a port ranges, type a hyphen between the start and end ports, such as: 8080-8085
Action	The fefault action for the service: - **Bypass** configures the service to ignore any traffic matching this listener. - **Intercept** configures the service to intercept and proxy the associated traffic.

   Create or edit as many listeners as needed.

See Also:

- "Configure a Proxy Service" on page 156

Manage Service Groups

The ProxySG appliance groups services into predefined service groups based on the type of traffic that service carries.

To manage service groups:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. In **Proxy Services**, select **Manage Groups**.

   The Manage Service Groups dialog displays the configured service groups, including the predefined groups Standard, Bypass Recommended, and Tunnel Recommended. The **Service Count** column shows the number of services in each group.
3. Do the following as needed:
   - To add a new service group, select **New** and specify a name for the group.
   - To rename a service group, select it and type the name.
   - To remove a service group, delete it. A service group must be empty before it is deleted. If a service group contains services, the delete icon is unavailable.

Exempt Requests From Specific Clients

The bypass list contains IP addresses/subnet masks of client and server workstations. Used only in a transparent proxy environment, the bypass list allows the appliance to skip processing requests sent from specific clients to specific servers. The list allows traffic between protocol non-compliant clients and servers to pass through the appliance without a disruption in service.

> **Note:** This prevents the appliance from enforcing any policy on these requests and disables any caching of the corresponding responses. Because bypass entries bypass Symantec policy, use bypass sparingly and only for specific situations.

Create static and dynamic bypass entries prevent the appliance from intercepting requests from specified systems. You can add static bypass entries in the Admin Console, but dynamic bypass entries can be configured through CLI or policy.

Configure Static Bypass

To configure static bypass

1. Select **Configuration > Services > Proxy Services**.
2. Under Static Bypass:
   - To add a new entry, select **New**.
   - To edit an existing entry, select Edit under **Actions**.
3. Specify the following information:

Field or Value	Description
Client/Prefix	IPv4 address with subnet mask or IPv6 address with prefix length.
Server/Prefix	IPv4 address with subnet mask or IPv6 address with prefix length.
Comment	(Optional) Comment that indicates why you are creating the static bypass rule for the specific source/destination combination. This is useful if another administrator needs to tune the settings later.

Configure Dynamic Bypass

Dynamic bypass is disabled by default. Enabling and fine-tuning dynamic bypass is a two-step process:

- Set the desired dynamic bypass timeout and threshold parameters.
- Use policy (recommended) or the CLI to enable dynamic bypass and set the types of errors that cause dynamic bypass to add an entry to the bypass list.

Dynamic bypass keeps its own (dynamic) list of which connections to bypass, where connections are identified by both source and destination. Dynamic bypass can be based on any combination of policy triggers. In addition, some global settings can be used to selectively enable dynamic bypass based on specific HTTP response codes. After an entry exists in the dynamic bypass table for a specific source/destination IP pair, all connections from that source IP to that destination IP are bypassed in the same way as connections that match against the static bypass list.

For a configured period of time, further requests for the error-causing URLs are sent immediately to the OCS, bypassing the appliance. The amount of time a dynamic bypass entry stays in the list and the types of errors that cause the appliance to add a site to the list, as well as several other settings, are configurable from the CLI.

After the dynamic bypass timeout for a client and server IP address entry ends, the appliance removes the entry from the bypass list. On the next client request for the client and server IP address, the appliance attempts to contact the OCS. If the OCS still returns an error, the entry is again added to the local bypass list for the configured dynamic bypass timeout. If the entry does not return an error, entries are again added to the dynamic list and not the local list.

Note: Dynamic bypass entries are lost when the appliance is restarted. No policy enforcement occurs on client requests that match entries in the dynamic or static bypass list. If a site that requires forwarding policy to reach its destination is entered into the bypass list, the site is inaccessible.

Caution: Dynamic bypass policy can automatically compile a list of response URLs that return various types of errors. Because bypass entries bypass Symantec policy, the feature should be used sparingly and only for specific situations.

Adding Dynamic Bypass Parameters to the Local Bypass List

The first step in configuring dynamic bypass is to set the server-threshold, max-entries, or timeout values in the CLI.

Note: This step is optional because the appliance uses default configurations if you do not specify them. Use the default values unless you have specific reasons for changing them. Contact Symantec Support for detailed advice on customizing these settings.

- The server-threshold value defines the maximum number of client entries before the appliance consolidates client-server pair entries into a single server entry that then applies to all clients connecting to that server. The range is 1 to 256. The default is 16. When a consolidation occurs, the lifetime of the consolidated entry is set to the value of timeout.

- The max-entries defines the maximum number of total dynamic bypass entries. The range is 100 to 50,000. The default value is 10,000. When the number of entries exceeds the max-entries value, the oldest entry is replaced by the newest entry.
- The timeout value defines the number of minutes a dynamic bypass entry can remain unreferenced before it is deleted from the bypass list. The range is 1 to 86400. The default value is 60.

Enabling Dynamic Bypass and Specifying Triggers

Enabling dynamic bypass and specifying the types of errors that causes a URL to be added to the local bypass list are done with the CLI.

Using policy to enable dynamic bypass and specify trigger events preferable to using the CLI, because the CLI supports a limited set of HTTP responses. For information about available CLI triggers, refer to the *Content Policy Language Reference*. For information about using policy to configure dynamic bypass, refer to the *Web Visual Policy Manager Reference*.

Bypassing Connection and Receiving Errors

In addition to setting HTTP code triggers, you can enable connection and receive errors for dynamic bypass.

If connect-error is enabled, any connection failure to the OCS, including timeouts, inserts the OCS destination IP address into the dynamic bypass list.

If receive-error is enabled, when the cache does not receive an HTTP response on a successful TCP connection to the OCS, the OCS destination IP address is inserted into the dynamic bypass list. Server timeouts can also trigger receive error.

The default timeout value is 180 seconds, which can be changed.

CLI Syntax to Enable Dynamic Bypass and Trigger Events

To enter configuration mode for the service:

```
# (config) proxy-services
# (config proxy-services) dynamic-bypass
```

The following subcommands are available:

```
# (config dynamic-bypass) {enable | disable}
# (config dynamic-bypass) max-entries number
# (config dynamic-bypass) server-threshold number
# (config dynamic-bypass) [no] trigger {all | connect-error | non-http | receive-error | 400 | 401 |
403 | 405 | 406 | 500 | 502 | 503 | 504}
# (config dynamic-bypass) timeout minutes
# (config dynamic-bypass) clear
# (config dynamic-bypass) view
```

Static Bypass and Dynamic Bypass

Create static and dynamic bypass entries prevent the appliance from intercepting requests from specified systems. You can add static bypass entries in the Admin Console, but dynamic bypass entries can be configured through CLI or policy.

Configure Static Bypass

To configure static bypass

1. Select **Configuration > Services > Proxy Services**.
2. For Interception, select **Restricted Intercept**.
3. In **Static Bypass**:
   - To add a new entry, select **New**.
   - To edit an existing entry, select Edit under **Actions**.
4. Specify the following information:

Field or Value	Description
Client/Prefix	IPv4 address with subnet mask or IPv6 address with prefix length.
Server/Prefix	IPv4 address with subnet mask or IPv6 address with prefix length.
Comment	(Optional) Comment that indicates why you are creating the static bypass rule for the specific source/destination combination. This is useful if another administrator needs to tune the settings later.

Configure Dynamic Bypass

Dynamic bypass is disabled by default. Enabling and fine-tuning dynamic bypass is a two-step process:

- Set the desired dynamic bypass timeout and threshold parameters.
- Use policy (recommended) or the CLI to enable dynamic bypass and set the types of errors that cause dynamic bypass to add an entry to the bypass list.

Learn more

Dynamic bypass keeps its own (dynamic) list of which connections to bypass, where connections are identified by both source and destination. Dynamic bypass can be based on any combination of policy triggers. In addition, some global settings can be used to selectively enable dynamic bypass based on specific HTTP response codes. After an entry exists in the dynamic bypass table for a specific source/destination IP pair, all connections from that source IP to that destination IP are bypassed in the same way as connections that match against the static bypass list.

For a configured period of time, further requests for the error-causing URLs are sent immediately to the OCS, bypassing the appliance. The amount of time a dynamic bypass entry stays in the list and the types of errors that cause the appliance to add a site to the list, as well as several other settings, are configurable from the CLI.

After the dynamic bypass timeout for a client and server IP address entry ends, the appliance removes the entry from the bypass list. On the next client request for the client and server IP address, the appliance attempts to contact the OCS. If the OCS still returns an error, the entry is again added to the local bypass list for the configured dynamic bypass timeout. If the entry does not return an error, entries are again added to the dynamic list and not the local list.

Note: Dynamic bypass entries are lost when the appliance is restarted. No policy enforcement occurs on client requests that match entries in the dynamic or static bypass list. If a site that requires forwarding policy to reach its destination is entered into the bypass list, the site is inaccessible.

Caution: Dynamic bypass policy can automatically compile a list of response URLs that return various types of errors. Because bypass entries bypass Symantec policy, the feature should be used sparingly and only for specific situations.

Adding Dynamic Bypass Parameters to the Local Bypass List

The first step in configuring dynamic bypass is to set the server-threshold, max-entries, or timeout values in the CLI.

Note: This step is optional because the appliance uses default configurations if you do not specify them. Use the default values unless you have specific reasons for changing them. Contact Symantec Support for detailed advice on customizing these settings.

- The server-threshold value defines the maximum number of client entries before the appliance consolidates client-server pair entries into a single server entry that then applies to all clients connecting to that server. The range is 1 to 256. The default is 16. When a consolidation occurs, the lifetime of the consolidated entry is set to the value of timeout.
- The max-entries defines the maximum number of total dynamic bypass entries. The range is 100 to 50,000. The default value is 10,000. When the number of entries exceeds the max-entries value, the oldest entry is replaced by the newest entry.
- The timeout value defines the number of minutes a dynamic bypass entry can remain unreferenced before it is deleted from the bypass list. The range is 1 to 86400. The default value is 60.

Enabling Dynamic Bypass and Specifying Triggers

Enabling dynamic bypass and specifying the types of errors that causes a URL to be added to the local bypass list are done with the CLI.

Using policy to enable dynamic bypass and specify trigger events preferable to using the CLI, because the CLI supports a limited set of HTTP responses. For information about available CLI triggers, refer to the *Content Policy Language Reference*. For information about using policy to configure dynamic bypass, refer to the *Web Visual Policy Manager Reference*.

Bypassing Connection and Receiving Errors

In addition to setting HTTP code triggers, you can enable connection and receive errors for dynamic bypass.

If connect-error is enabled, any connection failure to the OCS, including timeouts, inserts the OCS destination IP address into the dynamic bypass list.

If receive-error is enabled, when the cache does not receive an HTTP response on a successful TCP connection to the OCS, the OCS destination IP address is inserted into the dynamic bypass list. Server timeouts can also trigger receive error.

The default timeout value is 180 seconds, which can be changed.

CLI Syntax to Enable Dynamic Bypass and Trigger Events

To enter configuration mode for the service:

```
# (config) proxy-services
# (config proxy-services) dynamic-bypass
```

The following subcommands are available:

```
# (config dynamic-bypass) {enable | disable}
# (config dynamic-bypass) max-entries number
# (config dynamic-bypass) server-threshold number
# (config dynamic-bypass) [no] trigger {all | connect-error | non-http | receive-error | 400 | 401 |
403 | 405 | 406 | 500 | 502 | 503 | 504}
# (config dynamic-bypass) timeout minutes
# (config dynamic-bypass) clear
# (config dynamic-bypass) view
```

Troubleshooting: Bypass All Proxy Services or Restrict Interception

Refer to the following topics to troubleshoot issues with traffic interception.

Bypass All Proxy Services

As an interim solution, you can bypass all proxy services while troubleshooting problems. When this feature is enabled, transparent proxy connections are bypassed and explicit proxy connections are rejected.

To bypass all proxy services:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. For Interception, select **Force Bypass**.

> **Tip:** When you are done troubleshooting, select **Configuration > Services > Proxy Services** and select **Proxy Services** for Interception. Save the setting to disable **Force Bypass** and resume normal operation of proxy services.

Restrict Interception From Clients or To Servers

By default, all clients and servers evaluate the entries in Proxy Services where the decision is made to intercept or bypass a connection. To restrict or reduce the clients and servers that can be intercepted by proxy services, create restricted intercept lists. A restricted intercept list is useful in a rollout, before entering full production–you only want to intercept a subset of the clients. After the appliance is in full production mode, you can disable the restricted intercept list.

A restricted intercept list is also useful when troubleshooting an issue because you can reduce the set of systems that are intercepted.

> **Note:** Restricted intercepts lists are only applicable to transparent connections.
>
> if an entry exists in both the static bypass list and restricted intercept list, the static bypass entry overrides the restricted intercept entry. See "Static Bypass and Dynamic Bypass" on page 174.

To restrict interception from clients or to servers:

1. In the Admin Console, select **Configuration > Services > Proxy Services**.
2. For Interception, select **Restricted Intercept List**.

3. In **Restricted Interceptions**:
   a. To add entries, select **Add**.
   b. To edit entries, select the edit icon.

      The Add/Edit Intercept List Entry dialog appears.

4. Enter or edit the following settings as needed, selecting **Apply** to save each entry in the list:

Area	Setting or field	Description
Client Address	**All**	Restrict interception from all clients.
	Specific Host or Subnet	Restrict interception from specific clients. Enter an IPv4 or IPv6 address in the **IP Address** field and enter the subnet mask (for IPv4 addresses) or prefix length (IPv6) in the **Prefix Length/Subnet Mask** field.
Server Address	**All**	Restrict interception to all servers.
	Specific Host or Subnet	Restrict interception to specific servers. Enter an IPv4 or IPv6 address in the **IP Address** field and enter the subnet mask (for IPv4 addresses) or prefix length (IPv6) in the **Prefix Length/Subnet Mask** field.

Reference: Proxy Services, Proxy Configurations, and Policy

Refer to the following descriptions of proxies on the ProxySG appliance, proxy configurations, and policy. Access log fields are also included.

Proxy Types

Descriptions of proxies on the ProxySG appliance.

Proxy	Capabilities and Benefits
CIFS	■ Optimizes/accelerates file sharing across the WAN to users in branch offices.
DNS	■ Speeds up domain name resolution by looking up domain names in the appliance's DNS cache. If the name isn't found in the cache, the appliance forwards the request to the configured DNS server list. ■ Ability to rewrite DNS requests and responses.
Flash	■ Live streaming: The appliance fetches the live Flash stream once from the OCS and serves it to all users behind the appliance. ■ Video-on-demand: As Flash clients stream pre-recorded content from the OCS through the appliance, the content is cached on the appliance. After content gets cached on the appliance, subsequent requests for the cached portions are served from the appliance; uncached portions are fetched from the OCS.
FTP	■ Controls, secures, and accelerates file transfer requests. ■ Caches FTP objects.
HTTP	■ Controls, secures, and accelerates Web traffic. ■ Caches copies of frequently requested web pages and objects.
HTTPS Reverse Proxy	■ A proxy positioned in front of an HTTPS server that answers secure web requests from clients (using the appliance's local cache when possible). ■ Accelerates secure web requests, improving the response time to clients. ■ Because the Reverse Proxy is processing the requests, it allows the HTTPS server to handle a heavier traffic load.
MAPI	■ Accelerates the following Outlook processes: sending/ receiving e-mail, accessing message folders, changing calendar elements.
MMS	■ Monitors, controls, limits, or blocks streaming media traffic that uses Microsoft's proprietary streaming protocol. ■ Reduces stutter and improves the quality of streaming media. ■ Logs streaming connections.

Proxy	Capabilities and Benefits
RTSP	■ Monitors, controls, limits, or blocks streaming media traffic that uses the Internet standard RTSP protocol. ■ Reduces stutter and improves the quality of streaming media. ■ Logs streaming connections.
Shell	■ A proxy that allows a client to connect to other destinations via Telnet, after the client has created an authenticated Telnet connection to the appliance ■ Monitors, controls, limits, or blocks outbound Telnet connections. ■ Enforces access control to a group of users and destinations via policy. ■ Logs all connections.
SOCKS	■ A proxy that allows a client to connect to other destination servers/ports in a SOCKS tunnel, after the client's connection to the SOCKS proxy is authenticated. ■ Monitors, controls, limits, or blocks outbound client connections requested using the SOCKS protocol. ■ Through policy, enforces access control to a group of users and destinations. ■ SOCKS traffic can be passed to other proxies (such as HTTP) for acceleration. ■ Logs all connections.
SSL	■ Allows authentication, virus scanning and URL filtering of encrypted HTTPS content. ■ Accelerates performance of HTTPS content, using HTTP caching. ■ Validates server certificates presented by various secure websites at the gateway.
TCP-Tunnel	■ A tunnel for any TCP based protocol for which a more specific proxy is not available. ■ Compresses and accelerates tunneled traffic.

Service/Policy Matrix

The table below provides a list of the pre-defined proxy services and listeners that the proxy can accelerate and interpret.

Service Name	Proxy	Destination IP Address	Port Range
CIFS	CIFS	Transparent	445, 139
DNS	DNS	All	53
Endpoint Mapper	Endpoint Mapper	All	135
Explicit HTTP	HTTP	Explicit	8080, 80
External HTTP	HTTP	Transparent	80
FTP	FTP	All	21

Service Name	Proxy	Destination IP Address	Port Range
HTTPS	SSL	All	443
Internal HTTP	TCP-Tunnel	192.168.0.0/16 10.0.0.0/8 172.16.0.0/16 169.254.0.0/16 192.0.2.0/24	80
MMS	MMS	All	1755
MS Terminal Services	TCP-Tunnel	Transparent	3389
SOCKS	SOCKS	Explicit	1080

Access Log Fields

The following access log fields apply to service name and service group name.

- Name of the service used to intercept this connection:

  `x-service-name` (ELFF token)

  `service.name=` (CPL token)

- Service group name:

  `x-service-group` (ELFF token)

  `service.group=` (CPL token)

See "Configure an Access Log Format" on page 415 and "About Access Log Formats" on page 402 for detailed information about creating and editing log formats.

Global Settings for Proxy Services

The following optional settings apply to all proxy services when configured:

- "Ensure Application Availability (Tunnel-on-Protocol Error)" below
- "Use the Client IP Address for Server Connections" on the facing page
- "Improve Performance by Not Performing a DNS Lookup" on page 184
- "Manage Licensed User Connection Limits (Appliance to Server)" on page 187
- "Optimize HTTP Traffic using CachePulse" on page 188

Ensure Application Availability (Tunnel-on-Protocol Error)

This section discusses configuring the Tunnel-on-Protocol Error setting for HTTP and HTTPS traffic.

HTTP Proxy

In many networks, business-critical applications send traffic over port 80—the default HTTP port—because it is used as a generic route through the firewall. However, the ProxySG appliance HTTP proxy encounters problems when it receives non-HTTP requests from clients or browsers. The client receives an exception page and the connection closes. The following deployment operations create this situation:

- The client request from an application or browser is not HTTP.
- The request is HTTP but also contains components that are not HTTP.
- The request contains an unexpected formatting error in a line or header.

The appliance provides an option that enables the HTTP proxy to tunnel the connection when it receives non-HTTP traffic or broken HTTP request. This allows application traffic to continue and employee production to continue. The transactions remain labeled as HTTP; therefore, the access logs and the Traffic Mix and Active Sessions statistics display TCP_TUNNELED to indicate when a connection passed through the HTTP proxy. The HTTP proxy cannot apply security policies; however, benefits provided by ADN configurations might occur. The TCP Tunnel on Error option is viable with the following deployments:

- Applies only to HTTP traffic; HTTPS is not supported in either forward or reverse proxy modes.
- Applies only to errors in requests from the client browser or application to the appliance. Any issues that arise from server responses are not accommodated by this feature.

SSL Proxy

For the SSL proxy, the Tunnel on Protocol Error option applies when non-SSL traffic arrives at the SSL port (443 by default).A common scenario that causes this is having peer-to-peer applications (such viz, Skype, BitTorrent, and Gnutella) configured to enable port 443 for peer-to-peer traffic without SSL set as the transport protocol. An appliance transparently intercepting all 443

traffic cannot process these connections, rendering the application unusable. With an explicit proxy deployment, SSL errors during the initial handshake causes the same issue. The following example illustrates this:

- The appliance is configured to have an explicit HTTP service on port 8080.
- The HTTP service is configured with detect protocol enabled, which hands off SSL traffic to the SSL proxy from an HTTP CONNECT request. Detect Protocol is set to OFF by default.

> **Note:** The same applies to an explicit SOCKS proxy deployment with protocol detection enabled or an explicit TCP listener.

Considerations for Forwarding

Enabling the TCP Tunnel on Error option might cause issues if the appliance has forwarding rules that direct traffic to upstream proxies or other devices:

- Forwarding hosts are not viewed as HTTP proxies (even if they are). The initial HTTP proxy connects with a TCP tunnel to the forwarding host. If the appliance has a policy to forward and tunnels on error, the forwarding rule might not match if the forwarding rule has a condition based on information that is not present–any HTTP conditions, such as:
  - Request method
  - Request URL
  - Request headers
- In the case of tunnel on error with explicit proxy, HTTP must match a forwarding host for the connection of a successful TCP tunnel to occur. If no forwarding host matches, HTTP will not tunnel on error.

To ensure application availability by enabling TCP tunnel on HTTP protocol errors:

1. In the Admin Console, select **Configuration > Services > General Proxy Settings**.
2. In Settings, switch the **TCP tunnel request when a protocol error is detected** toggle.

Related Policy

The VPM provides the **Client Certificate Requested** object (**SSL Intercept > Service**).Use this policy in conjunction with an `ssl.intercept(no)` action, or a Do Not Intercept SSL action in the VPM, to minimize traffic disruption when the SSL proxy intercepts secure traffic where the OCS requests a client certificate. When Tunnel on Error is enabled, the first detection of a client certificate request from an OCS causes the connection to fail. The appliance adds the details for that exchange to an internal list of connections for which SSL interception should be negated. Subsequent requests function as expected.

Use the Client IP Address for Server Connections

This section discusses configuring the ProxySG appliance to use the IP address of the client to connect to destination servers rather than use the appliance address.

Reflect the Client Source IP when Connecting to Servers

By default, the ProxySG appliance uses its own IP address as the source IP address for requests (when connecting to servers). If Reflect Client IP is enabled, the appliance uses the client IP address for all requests. Enabling this option is not an arbitrary decision; it depends on the deployment and role of the appliance. For example, if this ProxySG is acting as a branch peer in an ADN deployment, enable client IP address reflection. This provides maximum visibility for network usage statistics and enables user-based access control to network resources.

> **Note:** The Reflect Client IP option is only supported in transparent deployments.

You can globally enable the Reflect Client IP option for all services that will be intercepted. To apply Reflect Client IP option to only a few services, first enable this option globally and then create policy to disable the Reflect Client IP option for the exceptions. Or, disable the option globally and create policy to enable it.

To reflect client IP:

1. In the Admin Console, select **Configuration > Services > General Proxy Settings**.
2. In **Settings**, switch the **Reflect client's source IP when connecting to servers** toggle.

> **Note:** If you enable Reflect Client IP and want the appliance to preserve persistent client connections, you must also add policy. In the VPM, in the Web Access Layer, use the **Support Persistent Client Requests** static action object, or use the following CPL:
>
> ```
> <proxy>
>     http.client.persistence(preserve)
> ```

Improve Performance by Not Performing a DNS Lookup

This section describes how to improve performance by configuring the appliance to trust the destination IP address provided by the client.

Trust the Destination IP Address Provided by the Client

If, in your environment, a client sometimes provides a destination IP address that the ProxySG appliance cannot identify, you have the option to configure the appliance to not perform a DNS lookup and allow that IP address. This can improve performance, but potentially presents a security issue.

You can configure the appliance to trust a client-provided destination IP address in transparent proxy deployments where:

- DNS configuration on the client is correct, but is not correct on the appliance.
- The client obtains the destination IP address using Windows Internet Name Service (WINS) for NetBIOS name resolution.
- DNS imputing on the appliance is not configured correctly. On the appliance, you can configure a list of suffixes to help with DNS resolution. In the event that the host name is not found, these suffixes are appended to the host name provided by the client. For information on DNS imputing, see "Resolve Hostnames Using Name Imputing Suffixes" on page 112.

In each of the cases above, the appliance cannot obtain the destination IP address to serve client requests. When you enable the appliance to trust a client-provided destination IP address, the appliance uses the IP address provided by the client and does not perform a DNS lookup.

In the following example, no DNS lookup occurs; the transactions goes straight to the OCS.

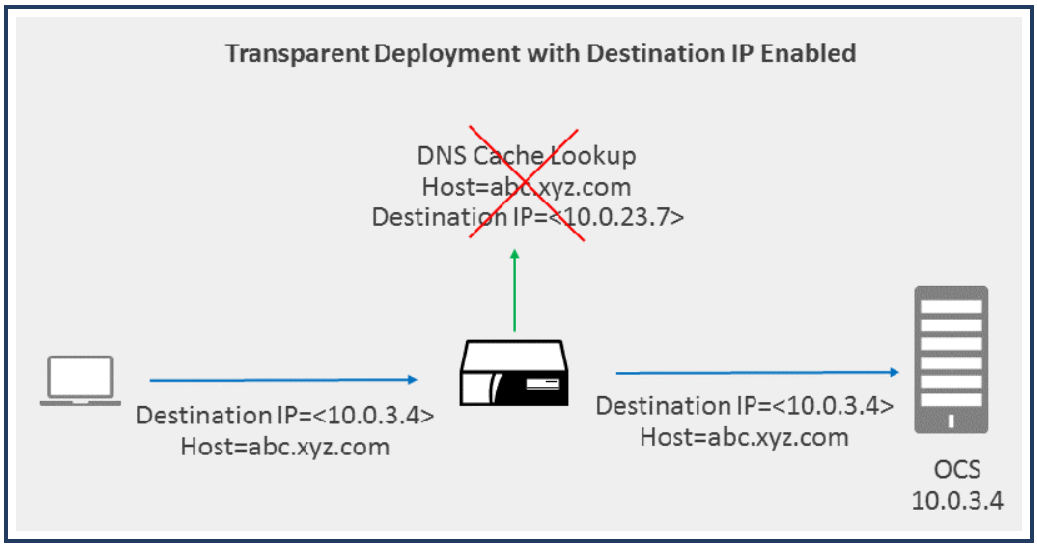

In the following example, the appliance initiates a DNS lookup and initiates a new connection to the server.

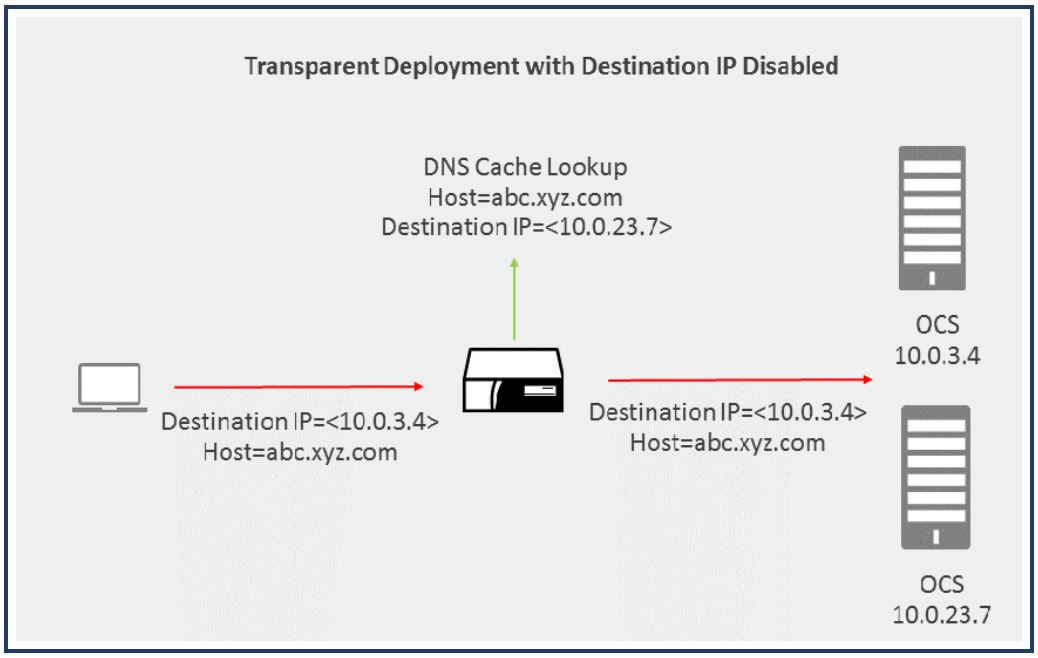

The appliance cannot trust the client-provided destination IP address in the following situations if the appliance:

- receives the client requests in an explicit proxy deployment.
- has a forwarding rule configured for the request.
- has a SOCKS gateway rule configured for the request.
- has policy that rewrites the server URL.

A transproxy deployment is one where a client is configured to contact an appliance explicitly, and a new appliance is deployed between the client and its explicit proxy. The new appliance, now transparently intercepts the traffic between the client and its explicit proxy. In a transproxy deployment, the destination IP address used by the client does not match the host header in the HTTP request, since the client is configured to use the explicit proxy. The path that the client request takes in a transproxy deployment depends on whether or not Trust Destination IP is enabled on the transparently deployed appliance.

- When Trust Destination IP is enabled on the transparent appliance, the transparent proxy trusts the destination IP included in the request and forwards the request to the explicit proxy which is serviced either from cache or from the OCS.
- When Trust Destination IP is disabled on the transparent appliance, the transparent proxy performs a DNS resolution on the host header in the request. The request is then completed based on the configured policy– forwarding rules, SOCKS gateway policy, and server URL rewrite policy.

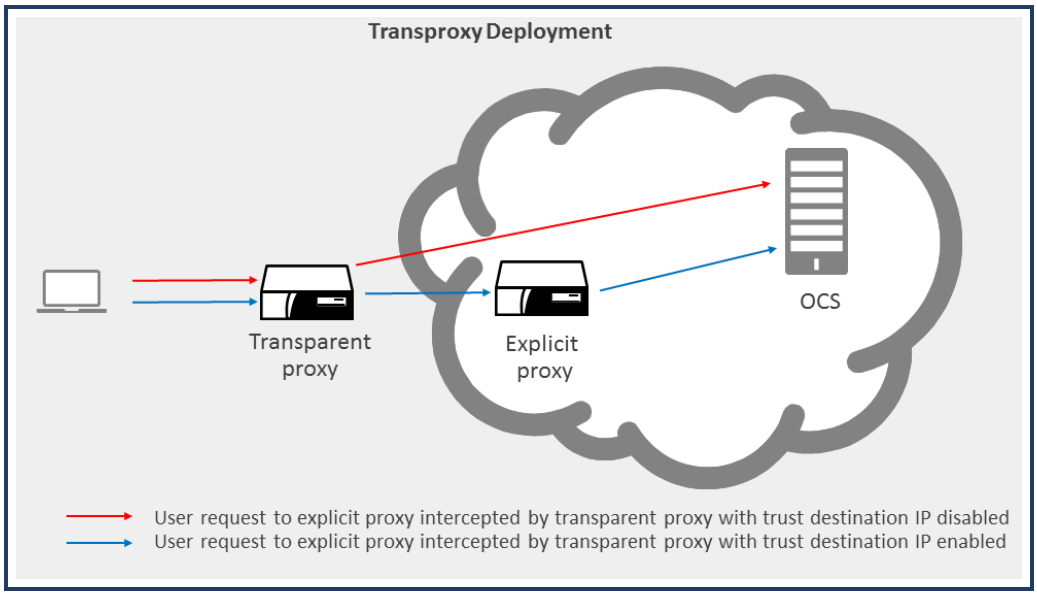

Note: If a client gives the destination address of a blocked site but the host name of a non-blocked site, with Trust Destination IP enabled, the appliance connects to the destination address. This might allow clients to bypass the configured security policy for your environment.

About the Default Settings

During the ProxySG initial configuration tasks, the administrator determined the default Trust Destination IP setting. In most deployments, the role of the appliance determines the setting:

- Acceleration role: enabled.
- Most other proxy deployments: disabled for tighter security.

You can change these defaults through the Admin Console, the CLI, or through policy. If you use policy, however, be aware that it overrides the setting in the in Admin Console. For information about using the `trust_destination_ip(yes|no)` CPL property, refer to the *Content Policy Language Reference*.

To trust the destination IP address provided by the client:

1. In the Admin Console, select **Configuration > Services > General Proxy Settings**.
2. In Settings, switch the **Trust client-provided destination IP when connecting to servers** toggle.

Manage Licensed User Connection Limits (Appliance to Server)

This section describes ProxySG appliance how to enable license-enforced user limits, describes how to monitor user numbers, and describes how to configure the ProxySG appliance to behave when a limit is breached.

> **Note:** If your platform and license support unlimited user connections, you do not have to configure user limit notifications because the thresholds cannot be exceeded. See User Limits per Model.

Modify User Limits Notifications

You can set and monitor user limit thresholds of the model license. A threshold breach triggers a notification and/or event log entry. Frequent breaches indicate that constant user connections to this particular ProxySG model are exceeding the optimal design.

To specify what happens when overflow connections occur:

1. In the Admin Console, select **Configuration > Services > General Proxy Settings**.
2. In **Settings**, select an option for User Overflow Action:

Setting	Description
Do not enforce licensed user limit	(Default behavior) Unlimited user connections are possible. If the limit is exceeded, the appliance health changes to CRITICAL. This option is not available on the ProxySG VA because licensed user limits are always enforced.
Bypass connections from users over license limit	Any transaction from a user whose connection exceeds the licensed limit is not susceptible to policy checks or any other ProxySG benefit, such as acceleration. This option provides the best user experience (with the caveat of potentially slower performance), but presents a Web security risk. This is the default option for the ProxySG VA.
Queue connections from users over license limit	Any transaction from a user whose connection exceeds the licensed limit must wait (in order) for an available ProxySG connection. This option provides the lowest user experience (and users might become frustrated and, perceiving a hang, might attempt request refreshes), but preserves Web security policies.

Optimize HTTP Traffic using CachePulse

> **Note:** Before you can use CachePulse, you must have a valid license for the feature. Refer to your Symantec point of contact for more information. If you do not have a valid license, the Admin Console might display Health Monitoring errors. The event log might also contain error messages about the subscription.

To optimize HTTP traffic using CachePulse:

1. In the Admin Console, select **Configuration > Services > General Proxy Settings**.
2. In **CachePulse**, switch the **Enable** toggle to enable the service.

   When you save the changes, the appliance attempts to download the database. The License and Download Status field shows download progress and information about the license and database.

Download the Database

You can download the CachePulse database at any time if the feature is enabled. If the initial download failed, and you resolved the issue that caused the failure, you can select **Download Now** to download database updates.

Manage the SSL Proxy

HTTPS traffic poses a major security risk to enterprises. Because the SSL content is encrypted, it cannot be monitored by normal means. This enables users to bring in viruses, access forbidden sites, or leak confidential business information over the HTTPS connection on port 443.

The SSL proxy intercepts, decrypts and re-encrypts HTTPS traffic (in explicit and transparent modes) so that security measures such as authentication, virus scanning, and URL filtering, and performance enhancements such as HTTP caching can be applied to HTTPS content. Additionally, the SSL proxy validates server certificates presented by various HTTPS sites at the gateway and offers information about the HTTPS traffic in the access log.

The SSL proxy tunnels all HTTPS traffic by default unless there is an exception, such as a certificate error or a policy denial. In such cases, the SSL proxy intercepts the SSL connection and sends an error page to the user. The SSL proxy also enables interception of HTTPS traffic for monitoring purposes.

The SSL proxy can perform the following operations while tunneling HTTPS traffic:

- Validate server certificates, including revocation checks using CRLs.
- Check various SSL parameters such as cipher and version.
- Log useful information about the HTTPS connection.

When the SSL proxy is used to intercept HTTPS traffic, it can also:

- Cache HTTPS content.
- Apply HTTP-based authentication mechanism.
- Send decrypted data to a configured ICAP scanning appliance for virus scanning and URL filtering. Apply granular policy (such as validating mime type and filename extension).

IPv6 Support

The SSL proxy is able to communicate using either IPv4 or IPv6, either explicitly or transparently.

In addition, for any service that uses the SSL proxy, you can create listeners that bypass or intercept connections for IPv6 sources or destinations.

Validating the Server Certificate

The SSL proxy can perform the following checks on server certificates:

- Verification of issuer signature.
- Verification of certificate dates.
- Host names in server certificates are important because the SSL proxy can identify aWeb site just by looking at the

server certificate if the host name is in the certificate. Most content-filtering HTTPS sites follow the guideline of putting the name of the site as the common name in the server's certificate.

Comparison of host name in the URL and certificate (intercepted connections only).

- Verification of revocation status.

  To mimic the overrides supported by browsers, the SSL proxy can be configured to ignore failures for the verification of issuer signatures and certificate dates and comparison of the host name in the URL and the certificate.

Checking CRLs

An additional check on the server certificate is done through CRLs. CRLs show which certificates are no longer valid; the CRLs are created and maintained by Certificate Signing Authorities that issued the original certificates.

Only CRLs that are issued by a trusted issuer can be used by the appliance. The CRL issuer certificate must exist as CA certificate on the appliance before the CRL can be imported.

The appliance allows:

- One local CRL per certificate issuing authority.
- An import of a CRL that is expired; a warning is displayed in the log.
- An import of a CRL that is effective in the future; a warning is displayed in the log.

See Also:

- "Configure Server Certificate Validation" on page 201

Working with SSL Traffic

The STunnel (SSL interception and tunnel) configuration intercepts all SSL traffic, handing HTTPS traffic off to the HTTPS forward proxy for compression and acceleration. STunnel decrypted traffic may be tapped and read by a third-party application such as Wireshark or Snort.

Recommendations for intercepting traffic include:

- Intercept non-HTTPS traffic for acceleration.
- Intercept any SSL traffic for tap, when you don’t know the application protocol over SSL.
- The HTTPS information in the next section, "Determining What HTTPS Traffic to Intercept" below.

Determining What HTTPS Traffic to Intercept

The SSL proxy tunnels HTTPS traffic by default; it does not intercept HTTPS traffic.

Many existing policy conditions, such as destination IP address and port number, can be used to decide which HTTPS connections to intercept.

Additionally, the SSL proxy allows the host name in the server certificate to be used to make the decision to intercept or tunnel the traffic. The server certificate host name can be used as is to make intercept decisions for individual sites, or it can be categorized using any of the various URL databases supported by Symantec.

Categorization of server certificate host names can help place the intercept decision for various sites into a single policy rule.

Recommendations for intercepting traffic include:

- Intercept Intranet traffic.
- Intercept suspicious Internet sites, particularly those that are categorized as none in the server certificate.

Recommendations for traffic to not intercept includes sensitive information, such as personal financial information.

Managing Decrypted Traffic

After the HTTPS connection is intercepted, you can do:

- Anti-virus scanning over ICAP
- URL filtering
- Filtering based on the server certificate host name

Caching

HTTPS applications that require browsers to present client certificates to secure Web servers do not work if you are intercepting traffic. To address this, you can create a policy rule to prevent the interception of such applications, or add client certificates to the appliance, and write policy to present the correct certificate.

If you configure the appliance to intercept HTTPS traffic, be aware that local privacy laws might require you to notify the user about interception or obtain consent prior to interception. You can option to use the **Notify User** object to notify users after interception or you can use consent certificates to obtain consent before interception. The **Notify User** is the easiest option; however, the appliance must decrypt the first request from the user before it can issue an HTML notification page.

Using the SSL Proxy with ADN Optimization

The SSL proxy itself can be used as a split proxy, which requires two SSL proxies, one at the branch and one at the core, working together. A split proxy can be configured (see below) to implement functionality that is not possible in a standalone proxy.

In this configuration, the SSL proxy supports ADN optimization on WAN networks, and SSL traffic performance can be increased through the byte caching capability offered. The branch proxy, which makes the decisions, is configured with both ADN optimization and SSL proxy functionality.

The Concentrator proxy (a ProxySG appliance that provides access to data center resources) does not require any configuration related to the SSL proxy. It only requires the necessary ADN configuration for applying byte caching capabilities to intercepted SSL content.

No special configuration is required to the SSL proxy.

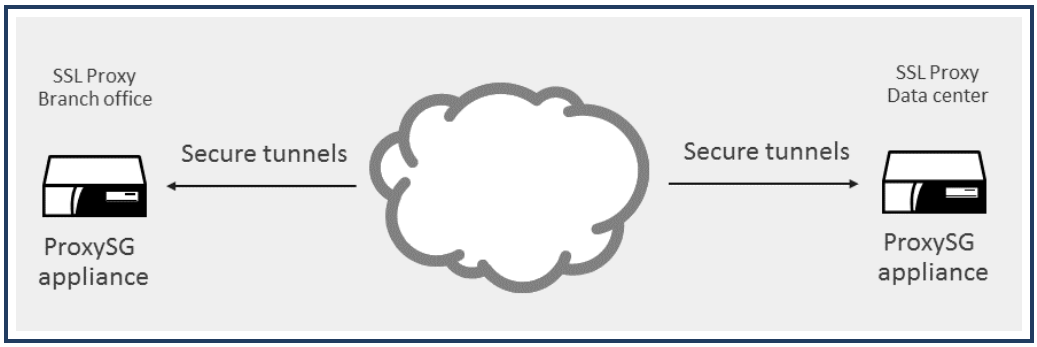

Notes on Pipelining

- Pipelining configuration for HTTP is ignored for HTTPS requests intercepted by the SSL proxy. When the SSL proxy intercepts an HTTPS request, and the response is an HTML page with embedded images, the embedded images are not pre-fetched by the appliance.
- If the appliance and the origin content server cannot agree on a common cipher suite for intercepted connections, the connection is aborted.

> **Note:** Pipelining configuration for HTTP is ignored for HTTPS requests intercepted by the SSL proxy. When the SSL proxy intercepts an HTTPS request, and the response is an HTML page with embedded images, the embedded images are not pre-fetched by the appliance.

- Server-Gated Cryptography and step-up certificates are treated just as regular certificates; special extensions present in these certificates are not be copied into the emulated certificate. Clients relying on SGC/step-up certificates continue using weaker ciphers between the client and the appliance when the SSL proxy intercepts the traffic.

See Also:

- "Intercept HTTPS" on the next page
- "Configure SSL Proxy" on page 195
- "Configure SSL Rules through Policy" on page 200

Intercept HTTPS

Intercepting HTTPS traffic (by decrypting SSL connections at the appliance) allows you to apply security measures like virus scanning and URL filtering.

> **Note:** An SSL license is required before you can make use of the SSL proxy for interception. Verify the license at the **Administration > Licensing** page in the Admin Console.

Configuration to intercept HTTPS traffic requires the following tasks:

- Determine whether you are using transparent or explicit mode.
- Create an SSL service or HTTP/SOCKS services with protocol detection enabled, depending on whether you are using transparent or explicit mode. The Detect Protocol setting is disabled by default. For more information on creating an SSL service, skip to "Configure the SSL Proxy in Transparent Proxy Mode" on the facing page.
- Create or import an issuer keyring, which is used to sign emulated server certificates to clients on the fly, allowing the SSL proxy to examine SSL content. For more information on creating an issuer keyring, see "Specify an Issuer Keyring and CCL for SSL Interception" on page 196.
- (Optional) Use the **Notify User** object or client consent certificates to notify users that their requests are being intercepted and monitored. Whether this is required depends on local privacy laws. The appliance has to decrypt the first request from the user to issue an HTML notification page. If this is not desirable, use client consent certificates instead. For more information on configuring the Notify User policy, refer to the *Web Visual Policy Manager Reference*. For information on managing client consent certificates, see "Configure Client Consent Certificates" on page 202.
- Download CA certificates to desktops to avoid a security warning from the client browsers when the appliance is intercepting HTTPS traffic. For information, see "Download an Issuer Certificate" on page 197.
- Using policy (VPM or CPL), create rules to intercept SSL traffic and to control validation of server certificates. By default, such traffic is tunneled and not intercepted. You must create suitable policy before intercepting SSL traffic. For more information on using policy to intercept SSL traffic, see "Configure SSL Rules through Policy" on page 200.
- (Version 7.1 and later) Enabled Policy Services. Refer to the *SGOS Administration Guide* for details.
- Configure a content filtering service. For more information, see "Configure Data Services" on page 323.
- Configure access logging. For more information on configuring access logging, see "General Logging Configuration" on page 384.
- Customize exception pages: For more information, see "About Policy Exception Pages" on page 27. To customize exception pages (in case of server certificate verification failure), refer to "Advanced Policy Tasks" in the *Web Visual Policy Manager Reference*.

Configure SSL Proxy

Refer to the appropriate steps to configure SSL proxy in your deployment:

- "Configure the SSL Proxy in Transparent Proxy Mode" below
- "Configure the SSL Proxy in Explicit Proxy Mode" below

Configure the SSL Proxy in Transparent Proxy Mode

You can configure the SSL proxy service in the Admin Console. This procedure applies if you are running the SSL proxy in transparent mode.

If you are using the SSL proxy in explicit mode, you might need an HTTP proxy or a SOCKS proxy. For information on configuring an SSL proxy in explicit mode, see "Configure the SSL Proxy in Explicit Proxy Mode" below.

You can use a TCP Tunnel service in transparent mode to get the same functionality. A TCP tunnel service is useful when you have a combination of SSL and non-SSL traffic going over port 443 and you do not want to break the non-SSL traffic. The SSL service requires that all requests to its port be SSL.

To configure an SSL service in transparent proxy mode, add an SSL proxy service with the following **Proxy Settings**:

- **Enable ADN**: Select this option to configure this service to use ADN. Enabling ADN does not guarantee the connections are accelerated by ADN. The actual enable decision is determined by ADN routing (for explicit deployment) or network setup (for transparent deployment).
- **Optimize Bandwidth**: Selected by default if you enabled WAN optimization during initial configuration. Clear the option if you are not configuring WAN optimization.

See "Configure a Proxy Service" on page 156 for instructions.

Continue with "Specify an Issuer Keyring and CCL for SSL Interception" on the next page.

Configure the SSL Proxy in Explicit Proxy Mode

The SSL proxy can be used in explicit mode in conjunction with the HTTP Proxy or SOCKS Proxy. You must create an HTTP Proxy service or a SOCKS Proxy service and use it as the explicit proxy from desktop browsers. You must also ensure that **Detect Protocol** attribute is enabled for these services. See "Configure a Proxy Service" on page 156 for instructions.

When requests for HTTPS content are sent to either a SOCKS proxy or an HTTP proxy, the proxies can detect the use of the SSL protocol on such connections and enable SSL proxy functionality.

Continue with "Specify an Issuer Keyring and CCL for SSL Interception" on the next page.

Specify an Issuer Keyring and CCL for SSL Interception

The SSL proxy can emulate server certificates; that is, present a certificate that appears to come from the origin content server. In actuality, Symantec has emulated the certificate and signed it using the issuer keyring. By default only the subjectName and the expiration date from the server certificate are copied to the new certificate sent to the client. The appliance supports the following key sizes when it emulates server certificates:

- DSA and ECDSA certificates: Key size up to 2048 bits
- RSA certificates: Key size up to 4096 bits

> **Note:** Only keyrings with both a certificate and a keypair can be used as issuer keyrings.

You can also change the CCLs that contain the CAs to be trusted during client and server certificate validation. The appliance can verify DSA, RSA, and ECDSA signed client and server CA certificates. The defaults are adequate for the majority of situations.

To specify an issuer keyring and CCLs:

1. In the Admin Console, select **Configuration > Services > SSL Proxy Settings**.
2. In **General Settings**, specify the following settings:

Setting or field	Description
Issuer Keyring	Select the keyring to use as the issuer keyring. Any keyring with both a certificate and a keypair can be used. If configured, HSM keyrings and keylists are also listed here.
CCL for Client Certificates	Specify which CAs are trusted when the SSL proxy validates client certificates. By default, all CA certificates are selected.
CCL for Server Certificates	Specify which CAs are trusted when the SSL proxy validates server certificates. The CCL for server certificates is relevant even when SSL proxy is tunneling SSL traffic. The default is browser-trusted. It is possible to set the CCL to validate client or server certificates from within policy on a per-request basis. For details, see the `client.certificate.validate.ccl()` and `server.certificate.validate.ccl()` properties in the *Content Policy Language Reference*.

To configure policy, see "Configure SSL Rules through Policy" on page 200.

Configure OCSP Stapling and CRLs for Emulated Certificates

OCSP stapling (RFC 6066) and CRLs allow you to obtain the revocation status of an X.509 digital certificate. The ProxySG appliance supports configuring these options for emulated certificates that the SSL forward proxy delivers.

> **Note:** You might want to configure both OCSP stapling and a CRL if some clients only support one of the methods.

To enable or disable OCSP stapling, use the following command:

```
# (config ssl) proxy ocsp-stapling {enable | disable}
```

OCSP stapling is enabled by default.

To configure the appliance to provide a CRL for the emulated certificates:

1. In the Admin Console, select **Configuration > Services > SSL Proxy Settings**.
2. In **General Settings**, switch the **Enable CRL on emulated certificates** toggle.
3. In the **CRL distribution point host name** field, enter the host name that is used to reach the ProxySG appliance for CRL downloads.

Download an Issuer Certificate

When the SSL proxy intercepts an SSL connection, it presents an emulated server certificate to the client browser. The client browser issues a security pop-up to the end-user because the browser does not trust the issuer used by the appliance. This pop-up does not occur if the issuer certificate used by SSL proxy is imported as a trusted root in the client browser's certificate store.

The appliance makes all configured certificates available for download via Admin Console. You can ask end users to download the issuer certificate through a browser and install it as a trusted CA in their browser of choice. This eliminates the certificate pop-up for emulated certificates.

To download an issuer certificate:

1. In the Admin Console, select **Dashboard > Advanced URLs**.
2. Select **SSL**.
3. Select **Download a Certificate as a CA Certificate** to display a list of certificates on the system. Continue with downloading the certificate.

Warn Users When Accessing Websites with Untrusted Certificates

The **Preserve untrusted certificate issuer** setting allows the appliance to present the browser with a certificate that is signed by its untrusted issuer keyring. The browser displays certificate information to the user, and lets the user accept the security risk of an untrusted certificate and proceed to the website.

The default-untrusted keyring has been added to the appliance to use with the **Preserve untrusted certificate issuer** setting. The default-untrusted keyring should not be added to any trusted CA lists.

> **Note:** This only applies to SSL forward proxy transactions with HTTPS interception enabled.

To display a warning to users about untrusted certificates on website, complete the following tasks:

- "Present Untrusted Certificates to a Browser" below
- "Set the Behavior when Encountering Untrusted Certificates" below

Present Untrusted Certificates to a Browser

Configure the appliance to act as a certificate authority and present a certificate signed by a specific keyring for all traffic. The default is the default-untrusted keyring.

1. In the Admin Console, select **Configuration > Services > SSL Proxy Settings**.
2. In **Preserve Untrusted Issuer Settings**, specify the following settings:

Setting or field	Description
Preserve untrusted certificate issuer	Enable to have the appliance act as a CA and present the browser with an untrusted certificate.
Untrusted Issuer Keyring	Select a keyring to use for signing untrusted server certificates presented by the appliance.

Set the Behavior when Encountering Untrusted Certificates

In the web VPM or CPL, define what the appliance should do for specific traffic if the user tries to access a website with an untrusted certificate.

Define Behavior in the VPM

Override the Admin Console settings for specific traffic, to specify whether the users should be prompted when a certificate that has not been signed by a trusted CA is encountered.

In the SSL Intercept Layer, add one of the following actions:

- **Do not Preserve Untrusted Issuer**

  If an OCS presents a certificate to the appliance that is not signed by a trusted CA, the appliance either sends an error message to the browser, or ignores the error and processes the request, based on the configuration of the Server Certificate Validation object.

- **Preserve Untrusted Issuer**

  If an OCS presents a certificate to the appliance that is not signed by a trusted CA, the appliance acts as a CA and presents the browser with an untrusted certificate. A warning message is displayed to the user, and they can decide to ignore the warning and visit the website or cancel the request.

- **Use Default Setting for Preserve Untrusted Issuer**

  The Preserve untrusted certificate issuer configuration setting in the Admin Console is used to determine whether or not untrusted certificate issuer should be preserved for a connection. This is the default behavior.

Define Behavior in CPL

Include the following syntax in policy to specify the behavior of the appliance when users encounter a website with an untrusted certificate:

```
ssl.forward_proxy.preserve_untrusted(auto|yes|no)
```

where:

- `auto`: Uses the Preserve untrusted certificate issuer configuration setting in the Admin Console to determine whether untrusted certificate issuer should be preserved for a connection. This is the default.
- `yes`: Preserves untrusted certificate issuer is enabled for the connection.
- `no`: Preserves untrusted certificate issuer is disabled for the connection.

  For example, to use the enable using the preserve untrusted certificate issuer, use the following syntax:

  ```
  <ssl-intercept>
   ssl.forward_proxy.preserve_untrusted(yes)
  ```

Configure SSL Rules through Policy

SSL interception and access rules, including server certificate validation, are configured through policy—either the VPM or CPL. In the VPM, use the SSL Intercept Layer to configure SSL interception; use the SSL Access Layer to control other aspects of SSL communication such as server certificate validation and SSL versions.

The policy examples in this section are for in-path deployments of appliances:

- "Configure SSL Interception " below
- "Configure Server Certificate Validation" on the facing page

Configure SSL Interception

The SSL Intercept layer in the web VPM allows you to set intercept options:

- "Intercept HTTPS " below
- "Intercept HTTPS Requests to Specific Sites" on the facing page
- "Categorize Host Names in Server Certificates" on the facing page
- "Configure Client Consent Certificates" on page 202

For detailed instructions on using the VPM, refer to the *Web Visual Policy Manager Reference*.

Intercept HTTPS

1. In the web VPM, add an SSL Intercept layer (select **Add Layer > SSL Intercept**).
2. Add an **Enable SSL Interception** action object. The **Enable HTTPS interception** option is selected by default.
3. Select the following settings and configure as needed

Setting or field	Description
Issuer Keyring	If you selected an issuer keyring previously, the keyring is displayed. If you did not select an issuer keyring previously, the default keyring is displayed. You can change the keyring that is used as the issuer keyring.
Hostname	Host name in the emulated certificate.
Splash Text	The splash text is added to the emulated certificate as a certificate extension. You are limited to a maximum of 200 characters.
Splash URL	The splash URL is added to the emulated certificate as a certificate extension.

4. Configure STunnel policy.

    a. Enable STunnel Interception: Establish a policy where configured STunnel services (such as POP3S and SMTPS) are terminated and accelerated.

    b. Enable SSL interception with automatic protocol detection: In addition to STunnel interception as described above, discovered HTTPS is handed off to the HTTPS proxy. Otherwise, SSL traffic continues in STunnel mode.

Intercept HTTPS Requests to Specific Sites

1. In the web VPM, add an SSL Intercept layer (select **Add Layer > SSL Intercept**).
2. Add an **Server Certificate** destination object.
3. Select one of the following settings:

Setting or field	Description
Hostname	Host name of the server whose traffic you want to intercept. Enter text for the host name and specify how to match: Exact Match, Contains, At Beginning, At End, Domain, or Regex.
Subject	Subject field in the server's certificate. After you Enter text for the subject and specify how to match: Exact Match, Contains, At Beginning, At End, Domain, or Regex.

Categorize Host Names in Server Certificates

1. In the web VPM, add an SSL Intercept layer (select **Add Layer > SSL Intercept**).
2. Add an **Server Certificate Category** destination object.
3. Add categories as needed to the selected list.

Configure Server Certificate Validation

For detailed instructions on using VPM, refer to the *Web Visual Policy Manager Reference*. For a list of the conditions, properties, and actions that can be used in the SSL Access Layer, refer to the *Content Policy Language Reference*.

> **Note:** The CPL property `server.certificate.validate()`, if set, overrides the `ssl-verify-server` command for either HTTP or for forwarding hosts.

To configure server certificate validation:

1. In the web VPM, add an SSL Access layer (select **Add Layer > SSL Access**).
2. Add a **Set Server Certificate Validation** action object.
3. Select **Enable server certificate validation** and configure the following as needed

Setting or field	Description
Behavior when server certificate validation is enabled	Select one of: ■ **Ignore hostname mismatch** ■ **Ignore expiration** ■ **Ignore untrusted issuer** These options mimic the overrides supported by most browsers.
Whether or not to check certificate revocation status	Select one of: ■ **Use only local certificate revocation check**: Uses the CRL configured on the appliance to perform the revocation check for a server certificate. ■ **Do not check certificate revocation**: Does not check the revocation status of the server certificate; however it still carries out the other certificate validation checks.

Configure Client Consent Certificates

The SSL proxy, in forward proxy deployments, can specify whether a client (typically a browser) certificate is required. These certificates are used for user consent, not for user authentication. Whether they are needed depends upon local privacy laws.

With client consent certificates, each user is issued a pair of certificates with the corresponding private keys. Both certificates have a meaningful user-readable string in the common name field. One certificate has a string that indicates grant of consent something like: “Yes, I agree to SSL interception”. The other certificate has a common name indicating denial of consent, something like: “No, I do not agree to SSL interception”.

Policy is installed on the appliance to look for these common names and to allow or deny actions. For example, when the string “Yes, I agree to SSL interception” is seen in the client certificate common name, the connection is allowed; otherwise, it is denied.

To configure client consent certificates:

1. Install the issuer of the client consent certificates as a CA certificate.
2. In the web VPM, add an SSL Access layer (select **Add Layer > SSL Access**).
3. Add a **Require Client Certificate** action object (it is a static object).
4. Add a **Client Certificate** source object to common names.

Configure SSL and Device Authentication

Refer to the following to configure SSL and device authentication:

- "PKI Concepts" on the next page
- "Manage SSL Keyrings and Keylists" on page 207
  - "Configure a Keyring" on page 209
    - "Associate a Certificate with the Keyring" on page 210
  - "Configure a Keylist" on page 212
  - "Import Keydata" on page 214
- "Manage CA Certificates" on page 216
  - "Manage CA Certificate Lists" on page 218
- "About SSL Device Profiles" on page 221
  - "Configure SSL Device Profiles for Device Authentication" on page 222
- "Integrate with an HSM Appliance" on page 224
  - "Manage HSMs" on page 226
  - "Manage HSM Keyrings" on page 227
  - "Manage HSM Keylists" on page 228
  - "Configure HSM Policy" on page 228

PKI Concepts

X.509 certificates are a cryptographic standard for PKI that specifies standard formats for public key certificates. Several RFCs and books exist on the public-key cryptographic system. This discussion of the elements of public-key cryptography system is relevant to their implementation in SGOS. Symantec uses certificates for various applications, including:

- authenticating the identity of a server
- authenticating the ProxySG appliance
- securing an intranet
- encrypting data

The following sections describe the concepts of PKI you must understand in order to use certificate authentication on the ProxySG appliance. The concepts included are the following:

Public Keys and Private Keys

In public-key cryptography systems, the intended recipient of encrypted data generates a private/public keypair, and publishes the public key, keeping the private key secret. The sender encrypts the data with the recipient's public key, and sends the encrypted data to the recipient. The recipient uses the corresponding private key to decrypt the data.

For two-way encrypted communication, the endpoints can exchange public keys, or one endpoint can choose a symmetric encryption key, encrypt it with the other endpoint's public key, and send it.

Certificates

Certificates are encrypted files that contain a public/private keypair. They can be used to verify the identity of a server, a website or to encrypt files. The SGOS software uses:

- SSL certificates
- CA certificates
- External certificates
- Certificate chains

You can also use wildcard certificates during HTTPS termination. Microsoft's implementation of wildcard certificates is as described in RFC 2595, allowing an * (asterisk) in the leftmost-element of the server's common name only.

For information on wildcards supported by Internet Explorer, refer to article 258858 at the Microsoft Knowledge Base. Any SSL certificate can contain a common name with wildcard characters.

SSL Certificates

SSL certificates are used to authenticate the identity of a server or a client. A certificate is confirmation of the association between an identity (expressed as a string of characters) and a public key. If a party can prove they hold the corresponding private key, you can conclude that the party is who the certificate says it is. The certificate contains other information, such as what functions the certificate can be used for, and the time rang it is valid.

The association between a public key and a particular client or server is done by generating a certificate signing request using the server's or client's public key. A CA verifies the identity of the server or client and generates a signed certificate. The resulting certificate can then be offered by the server to clients (or from clients to servers) who can recognize the CA's signature. Such use of certificates issued by CAs has become the primary infrastructure for authentication of communications over the Internet.

CA Certificates

CA certificates are certificates that belong to certificate authorities. ProxySG appliances use CA certificates to verify certificates presented by a client or a server during secure communication. There can be multiple levels of CA certificates, with a client or server certificate being signed by an intermediate CA which has a certificate signed by a parent CA, and so on, ending in a trusted root CA that must be installed on the appliance for the validation to succeed. Browsers offer a certificate if the server is configured to ask for one and an appropriate certificate is available to the browser.

The appliance trusts all root CA certificates trusted by Internet Explorer and Firefox. The list is updated periodically to be in sync with the latest versions of IE and Firefox. You can review these certificates using the Management Console or the CLI.

You can delete some of these pre-installed certificates if you decide you don't want to trust them. You can also add your own root and intermediate CA certificates for your own internal certificate authorities.

External Certificates

An external certificate is any X.509 certificate for which the appliance does not have the private key. The certificate can be used to encrypt data, such as access logs, with a public key so that it can only be decrypted by someone who has the corresponding private key. See "Encrypt the Access Log " on page 388 for information about encrypting access logs.

Using Keyrings and SSL Certificates

Keyrings are virtual containers. Each keyring holds a public/private key pair and a customized key length. You can associate certificates, certificate chains or certificate signing requests with keyrings. In general, SSL certificates involve three parties:

- The subject of the certificate.
- The CA, which signs the certificate, attesting to the binding between the public key in the certificate and the subject.
- The relying party, which is the entity that trusts the CA and relies on the certificate to authenticate the subject.

Keyrings and certificates are used in:

- Encrypting data
- Digitally signing access logs
- Authenticating end users
- Authenticating an appliance

You must perform these steps using a secure connection such as HTTPS, SSH, or a serial console:

- Create a keyring. A default keyring is shipped with the system and is used for accessing the Management Console, although you can use others. You can also use the default keyring for other purposes. You can create other keyrings for each SSL service. (See "Configure a Keyring" on page 209.)

  > **Note:** You can also import keyrings. For information on importing keyrings, see "Configure a Keyring" on page 209.

- (Optional) Create CSRs to be sent to CAs. Refer to the *SGOS Administration Guide*.
- Import X.509 certificates issued by trusted CA authorities for external use and associate them with the keyring. (See "Managing SSL Certificates" on page 1283.) Alternatively, create certificates and associate them with the keyring. You can also associate a certificate chain with a keyring.
- (Optional, if using SSL certificates from CAs) Import CRLs so the appliance can verify that certificates are still valid.

Manage SSL Keyrings and Keylists

The appliance ships with several keyrings already created:

- default: The default keyring contains a certificate and an automatically generated keyring and a self signed certificate which can be used for accessing the appliance through HTTPS.
- configuration-passwords-key: The configuration-passwords-key keyring contains a keypair but does not contain a certificate. This keyring is used to encrypt passwords in the show config command and should not be used for other purposes.
- appliance-key: The appliance-key keyring contains an internally-generated keypair. If the appliance is authenticated (has obtained a certificate from the Symantec CA appliance-certificate server), that certificate is associated with this keyring, which is used to authenticate the device.
- passive-attack-protection-only-key: The passive-attack-protection-only-key keyring allows data to be encrypted, but with no endpoint authentication. Although the traffic cannot be sniffed, it can be intercepted with a man-in-the-middle attack. The passive-attack-protection-only-key keyring is NOT considered secure; therefore, it should not be used on production networks.

If an OCS requires a client certificate and no keyring is associated with the ProxySG SSL client, the HTTPS connections fails.

To view the properties of the keyrings, in the Admin Console, select **Configuration > SSL > Keyrings and Keylists**. In Keyrings, refer to the following details:

Field or Setting	Description
Name	Unique name of the keyring. If a pre-defined keyring, one of the names listed above.
Referenced	Indicates whether the keyring is associated with a client certificate. Refer to the *SGOS Administration Guide* for details on how to do this in the Management Console.
Show Private Key	Indicates whether **Allow the key to be viewed and exported** is enabled. See "Configure a Keyring" on page 209 for details.
CSR	Indicates whether a certificate signing request was created for the keyring.
Certificate Expiry	If the keyring is associated with a client certificate, shows an expiry date. If the keyring is not associated with a client certificate, shows "No certificate".

View Certificate Details

If a keyring is associated with a client certificate, you can select **View** in the Actions column to see the certificate details and PEM file data. To associate a client certificate with a keyring, see "Associate a Certificate with the Keyring" on page 210.

Group Related Keyrings into a Keylist

To easily reference client certificate keyrings in policy, use keylists to group them together. For example, it is often useful to group certificates into keylists bundled by:

- all client certificates for a specific web address
- all client certificates for a group of users
- all client certificates for a specific user

All keyrings in the keylist must have the same extractor, but each certificate must have a unique value for the extractor. The evaluation of the keylist extractor string must be unique across the client certificates in the keylist, otherwise changes being applied to the keylist will fail with an error.

To view the properties of the keylists, in the Admin Console, select **Configuration > SSL > Keyrings and Keylists**. In Keylists, refer to the following details:

Field or Setting	Description
Name	Unique name of the keylist.
Keyring Count	Number of keyrings in the keylist.
Extractor	Extractor used for the keyrings in the keylist. For more information, see "Configure a Keylist" on page 212. All keyrings in the keylist must have the same extractor, but each certificate must have a unique value for the extractor. The evaluation of the keylist extractor string must be unique across the client certificates in the keylist; otherwise, changes being applied to the keylist will fail with an error.

To create or edit a keylist, see "Configure a Keylist" on page 212.

Configure a Keyring

Note: Full configuration requires working in the Management Console. Refer to the *SGOS Administration Guide*.

Note: To import keyring data from a file, see "Import Keydata" on page 214.

To configure a keyring:

1. In the Admin Console, select **Configuration > SSL > Keyrings & Keylists**.
2. In **Keyrings**, select an existing keyring to edit it, or select **Add Keyring > Keyring** to add a new keyring.
3. On the Add/Edit Keyring page, specify the following settings:

Setting or field	Description
Keyring Name	Unique name for the keyring. Spaces in keyring names are not supported. Including a space can cause unexpected errors while using the keyrings.
Private Key Visibility	If **Allow the key to be viewed and exported** is enabled, the private keys can be viewed and exported.

Setting or field	Description
Private Key	Create a new private key or import an existing one. To create a new key: 1. Select **Generate new**. 2. In **Private Key Length**, specify a key length from 384 to 4096 bits. If needed, determine the maximum key length allowed for export. To import and existing key: 1. Select **Import**. 2. Copy the key and select **Paste From Clipboard**. The key is pasted into the field as obfuscated text. 3. (If needed) Select **Show in plain text** to display the key in plain text. 4. (If applicable) If the private key that is being imported is encrypted with a password, enter the password in the **Private key password** field. **Note:** The only way to retrieve a keyring's private key from the appliance is through the CLI.

See Also:

- "Associate a Certificate with the Keyring" below
- "Import Keydata" on page 214

Associate a Certificate with the Keyring

You can associate a client certificate with the keyring using the following methods:

- Create a self-signed certificate on the appliance. See "Create a Self-Signed Certificate" below.
- Obtain a CA-signed certificate by creating a CSR. See "Obtain a CA-Signed Certificate" on the facing page.

Create a Self-Signed Certificate

The appliance ships with a self-signed certificate, which is associated with the default keyring. Only one certificate can be associated with a keyring. If you have multiple uses, use a different keyring and associated certificate for each one. Self-signed certificates are generally meant for intranet use, not Internet.

1. In the Admin Console, select **Configuration > SSL > Keyrings & Keylists**.
2. Select an existing keyring or select **Add Keyring** to add a new keyring.

3. On the Add New/Edit Keyring page, in **Certificate**, do one of the following:
   - Select **Create**. In the Create Certificate dialog, enter all applicable information. See "Certificate Fields" below.
   - Select **Paste from Clipboard** and paste the certificate you created off-box.

   After you create the certificate:
   - In **Keyrings** (select **Configuration > SSL > Keyrings & Keylists**), the **Certificate Expiry** column shows an expiration date and the **Actions** column has a **View** icon.
   - You can copy the certificate from the Certificate field by selecting **Copy to clipboard**.

     The certificate consists of data between `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.

Obtain a CA-Signed Certificate

CSRs are used to obtain a certificate signed by a CA. You can create CSRs off-box. If you have already done so, be ready to copy the CSR while performing the following steps.

To create a CSR:

1. In the Admin Console, select **Configuration > SSL > Keyrings & Keylists**.
2. Select an existing keyring or select **Add Keyring** to add a new keyring.
3. On the Add New/Edit Keyring page, in **Certificate Signing Request**, do one of the following:
   - Select **Create**. In the Create Certificate Signing Request dialog, enter all applicable information. See "Certificate Fields" below.
   - Select **Paste from Clipboard** and paste the CSR you created off-box.

   After you create the CSR:
   - You must submit it to a CA in the format the CA requires. You can copy the CSR from the Certificate Signing Request field by selecting **Copy to clipboard**.

     The CSR consists of data between `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` statements.
   - In **Keyrings** (select **Configuration > SSL > Keyrings & Keylists**), the **CSR** column indicates that a CSR exists for the keyring.

Certificate Fields

Most field limits are counted in terms of bytes rather than characters. The number of non-ASCII characters a field can accommodate will be less than the size limit because non-ASCII characters can occupy more than one byte, depending on the encoding. The only exception is the Challenge field, which is counted in terms of characters.

Field	Description
State/Province	State or province where the appliance is located.
Country Code	Two-character ISO code of the country.
City/Locality	City or locality where the appliance is located.
Organization	Organization name.
Unit	Name of the group that is managing the appliance.
Common Name	Name containing an URL with client access to the specific OCS.
Challenge	A 4-20 character alphanumeric challenge.
E-mail	E-mail address for the organization.
Company	Company name.
Digest Type	Signing hash used to create the certificate; default is SHA256.

Configure a Keylist

> **Note:** To import keylist data from a file, see "Import Keydata" on page 214.

To configure a keylist:

1. In the Admin Console, select **Configuration > SSL > Keyrings and Keylists**. In Keylists, select **Add Keylist**.
2. On the Add/Edit Keylist page, under Properties, specify the following settings:

Setting or field	Description
Keylist Name	Unique name for the keylist.

Setting or field	Description
Extractor	String that identifies which certificate field attribute value to extract to determine a policy match. Enter the string in `$(field.attribute)` format. For example, to extract the value of the CN attribute from the Subject field of the certificate, you would enter `$(subject.CN)`. Alternatively, select values from the Field, Attribute, and Group Name drop down lists to build an extractor string, and click Add to extractor. The new extractor string is appended to any existing text in the **Extractor** field. The Group Name drop down list only appears for IssuerAltName and SubjectAltName fields. The **Extractor** field can have a maximum of 255 characters. The extractor supports substitutions from all attributes of Subject, Issuer, SubjectAltName, IssuerAltName, SerialNumber, and ServerName certificate fields. The default extractor value is $(Subject.CN); many other subject attributes are recognized, among them OU, O, L, ST, C, and DC. Field indexes can be used in substitutions on a group name or attribute; for example $(SubjectAltName.DNS.1).

3. Under Selected Keyrings, specify the following settings:

Setting or field	Description
Default Keyring	(Applicable to SNI in reverse proxy mode) Select an option for Default keyring to ensure that a keyring is used if the client has not implemented SNI or sends incompatible SNI information. **Note:** You can't remove the default keyring from a keylist. Select a different keyring as default or disable **Default Keyring** before attempting to remove the keyring from the keylist.
Keyrings selection	On the list of available keyrings, select keyrings to add to the list. To remove keyrings, clear the selection. Use the **Show Selected Only** toggle to show only the keyrings in the keylist or all keyrings.

If errors occur while saving the keylist, possible causes might be:

- The included keyring does not contain the specified extractor pattern or substitution variable.
- Multiple keyrings have the same value for the specified extractor. The extracted value in the keyring allows the policy action object to find the appropriate keyring certificate to use. Only one keyring can be utilized by each policy transaction; thus, the extractor string evaluation must be unique across the certificates in the keylist. A keyring whose extractor value matches the extractor value of any existing keyring in the keylist will not be added to the keylist. For example, if the extractor $(Subject.DC) is selected, and all keyrings have the same value in the certificate for that extractor, the policy would not be able to determine which keyring to select.

See Also

- "Configure a Keyring" on page 209
- "Associate a Certificate with the Keyring" on page 210
- "Import Keydata" below

Import Keydata

You can import keyrings and keylists from a file on a server that the appliance can access. A keydata file can consist of keyrings, keylists, or both. Make sure that the file is correctly formatted; see "Create a Keydata File" below.

> **Note:** If you are importing a keyring and one or more certificates onto an appliance, first import the keyring, followed by the related certificates. The certificates contain the public key from the keyring, and the keyring and certificates are related.

To import a keyring or keylist:

1. In the Admin Console, select **Configuration > SSL > Keyrings and Keylists**. In Keyrings or Keylists, select **Import**.
2. Specify the following on the Import Keyrings, Keylists dialog:
   - **URL**: Enter the URL for the file that contains the keyrings/keylists information.
   - **Passphrase**: (If applicable) If the keypair that is being imported has been encrypted with a passphrase, enter it.
3. Select **Import** to import the data.

Create a Keydata File

Bundle multiple keyrings and keylists into a single keydata file for simple importing into the ProxySG appliance. The keydata file does not need to include both keyring and keylist information.

> **Note:** All keyrings or keylists being imported must have the same passphrase for the import to be successful.

To create a keydata file:

1. Open a new text file.
2. Add keyring information to the keydata file in the following format:

```
#keyring: <keyring_id>
#visibility: {show | show-director | no-show}
<Private_Key>
<Certificate>
<CSR>
```

where:

- *keyring_id*: Name of the keyring.
- visibility: How the keyring is displayed in the show configuration output. Options include:
  - show: Private keys associated with keyrings created with this attribute can be displayed in the CLI or included as part of a profile or overlay pushed by Director.
  - show-director: Private keys associated with keyrings created with this attribute are part of the show configuration output if the CLI connection is secure (SSH/RSA) and the command is issued from Director.
  - no-show: Private keys associated with keyrings created with this attribute are not displayed in the show configuration output and cannot be part of a Director profile. The no-show option is provided as additional security for environments where the keys will never be used outside of the particular ProxySG.
- *Private_Key*, *Certificate*, and *CSR* – Paste the contents of the key, certificate or CSR into the text file, including the ---Begin and ---End tags.

In the following example, the private key and certificate has been truncated.

```
#keyring:Keyring1
#visibility:no-show
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQE...KvBgDmSIw6dTXxAT/mMUHGRd7cRew==
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIDdjCCAl4CCQC...TjUwxwboMEyL60z/tixM=
-----END CERTIFICATE-----
#keyring:Keyring2
```

3. Add keylist information to the file in the following format:

```
#keylist: <keylist_name>
#extractor: <extractor>
<keyring_id>
<keyring_id>
...
```

where:

- *keylist_name*: Name of the keylist.
- *extractor*: String that identifies which certificate field attribute value to extract to determine a policy match, using the `$(field.attribute)` syntax. Substitutions from all attributes of Subject, Issuer, SubjectAltName, IssuerAltName, and SerialNumber certificate fields are supported.
- *keyring_id*: Keyrings to include in the keylist. The keyrings may be included in the keydata file, or may be keyrings that already exist on the ProxySG appliance.

See the example:

```
#keylist:mylist
#extractor: $(Subject.CN)
Keyring1
Keyring2
```

4. Save the file as .txt on a web server that the ProxySG appliance can access.

Manage CA Certificates

The appliance is pre-installed with browser-trusted root CA certificates; the list is updated periodically to be in sync with the latest browser versions.

You can also import non-standard third-party CA certificates into the appliance CA certificate store, including root and intermediate CA certificates. CA certificates added to the CA certificate store are available for use by the CA certificate lists (CCL) for validating the security of connections.

View Installed CA Certificates

To view a CA certificate, select **View Certificate**.

To search for specific certificate(s), enter text in the **Name** field filter.

Use the options at the bottom of the section to search for more certificates:

- **Certificates per page**: Select a number to decrease or increase the number of certificates displayed per page.
- **Back and forward arrows**: Use the arrows to go back or forward one page, or to go to the first or last page of certificates.

Import a CA Certificate

Before proceeding, copy the file contents of the certificate to import to clipboard.

To import a CA certificate:

1. In the Admin Console, select the **Configuration > SSL > CA Certificates**.
2. In Certificates, select **Import**.

3. In the Import CA Certificate dialog, enter a unique name for the certificate. Do not include spaces in the name.
4. Select **Paste from Clipboard**. Certificate contents should appear in the dialog.
5. Select **OK** to save changes.

About CA Certificate Lists

A CA certificate list (CCL), which contains some of the CA certificates available on the appliance, allows the administrator to control the set of CA certificates trusted for a particular set of SSL connections. A CCL contains a subset of the available CA certificates on the appliance, and restricts trust to those certificates. The CCL referenced by the profile or service configuration is used when an SSL connection is established to that service or using that profile. Three CCLs are created by default on the appliance:

- appliance-ccl: This CCL is used for authenticating connections among devices manufactured by Symantec. By default it contains the Symantec ABRCA root certificate (ABRCA_root). This list is used by default in the bluecoat-appliance-certificate SSL device profile. This CCL can be edited but not deleted.
- browser-trusted: This CCL includes most of the well-known CAs trusted by common browsers. This CCL can be edited but not deleted. You can manually add CAs to this list. In addition, the appliance automatically retrieves an updated browser-trusted CCL from Symantec every seven days. The browser-trusted CCL is used by default during certificate verification by the SSL client and by the default SSL device profile.
- image-validation: This CCL is used to validate signed SGOS images.

You can customize the CCLs available on the appliance to ensure that the appliance has the CA certificates it needs to handle HTTPS requests. You can create your own CA certificate lists or modify the default CCLs by adding or removing trusted CAs.

See Also:

- "Manage CA Certificate Lists" on the next page

Manage the Certificate Cache

The appliance automatically stores unrecognized intermediate CA certificates that are included with validated CA certificate chains whenever an SSL connection is established. These intermediate CA certificates are stored within a separate cache on the appliance and are used to validate SSL connections when an incomplete certificate chain is encountered. For security purposes, OCSP and CRL validation checks are performed to confirm the safety of the certificate chain. As an additional layer of security, the intermediate CA certificates in the chain must end with a trusted root certificate from the CCL (CA certificate list) that is associated with the connection. If a compatible certificate is not found, the connection is considered insecure and the user will be given a security warning.

Turn off Intermediate Certificate Caching

Turning off caching automatically clears the existing cache of intermediate CA certificates and prevents any validated intermediate certificates from being added to the cache. To turn off intermediate certificate caching:

1. In the Admin Console, select the **Configuration > SSL > CA Certificates**.
2. In Cached Intermediate Certificates, select the **Enable Caching** toggle.
3. On the dialog that appears, select **Disable** to confirm your choice.

View Cached Intermediate Certificates

You can view information about the CA certificates, which conform to, at a minimum, the standards established within the PKI ITU-T X.509 standard.

To view the details of a specific cached intermediate certificate:

1. In the Admin Console, select the **Configuration > SSL > CA Certificates**.
2. In Cached Intermediate Certificates, select **View Certificate**.

   The appliance does not allow automatic retrieval of issuing certificates for Intermediate certificates that include an AIA (Authority Information Access) entry.

   - General: Includes top level information about a digital certificate, including the DN (distinguished name) identifying the owner and issuer, the dates when the certificate is valid, and the public key fingerprints using MD5 and SHA-1 cryptographic hash functions.
   - Details: Includes certificate field information as defined in the ITU-T X.509 public key certificate standard.
   - PEM (Privacy-enhanced mail): Displays the certificate contents in a Base64 encoded format.

Clear Cached Intermediate Certificates

Clearing the CA certificate cache removes all stored intermediate CA certificates.

1. In the Admin Console, select the **Configuration > SSL > CA Certificates**.
2. In Cached Intermediate Certificates, select **Clear Cache**.
3. On the dialog that appears, select **Clear** to confirm your choice.

> **Note:** The appliance retains the list of cached intermediate CA certificates even after the appliance is shut down and restarted. The only way to delete the cache is to manually clear or turn off certificate caching.

Manage CA Certificate Lists

Because the list of trusted CAs changes over time, you may want to update your CCLs to ensure that they contain the most up-to-date list of CA certificates. You can manually edit the default appliance-ccl and browser-trusted CCLs as well as any

custom-produced CCL. The bluecoat-services and image-validation CCLs are read-only and cannot be modified; however, you can still view the contents.

Keep in mind that if you plan to add a CA to a CCL, you must first import the corresponding CA certificate as described in "Import a CA Certificate" on page 216.

For the browser-trusted CCL, you also have the option to configure the appliance to download an updated browser-trusted list of CAs on demand or automatically on a schedule. This smart download compares the existing browser-trusted list on the appliance to the new list and only adds those CA certificates that are new since the last update. Any manual changes that you have made to the file are preserved.

To manage a CCL:

1. In the Admin Console, select the **Configuration > SSL > CA Certificates**.
2. In CA Certificate Lists:
   - Select **Add List** to add a CCL.
   - Select an existing CCL to edit it.
3. On the Add CA Certificate List page, enter a **Certificate List Name**.

   If editing a CCL, the name can't be changed.
4. The Selected Certificates section shows CA certificates belonging to the list. Select one of the following:
   - **Select Certificates**: Adds CA certificates to the list. On the Select Certificates dialog, select certificates or clear selections. The **Show Selected Only** toggle filters the view to show only selected certificates.
   - **View**: Displays the certificate contents.
   - **Remove**: Removes the certificate from the list.

Configuring Download of CCL Updates from Symantec

By default, the appliance will automatically download and install a package containing the updated CA Certificates and CCL updates—called a trust package—from Symantec every seven days. This trust package contains any updates to the browser-trusted and image-validation CCLs and their associated CA certificates since the last update, based on the timestamp at the time the trust package was created. Note that any manual changes you have made to the CCLs and CA certificates will be preserved.

You can customize the CA download list updates as described in the following sections.

Change the Download Location

The downloadable CA list—called a trust package—is hosted at the following URL:

http://appliance.bluecoat.com/sgos/trust_package.bctp

By default, the appliance is configured to download the trust package directly from this URL. As an alternative you can set up your own download site on premise. To do this, you must download the trust package from the URL to your download server and then configure the download path on the appliances in your network.

After you determine the download location, you must configure the appliance to point to the location using the following command:

```
#(config) security trust-package download-path <URL>
```

For example, to configure the appliance to download the trust package from a bluecoat folder on your download.acme.com server, you would enter the following command:

```
#(config) security trust-package download-path http://downloads.acme.com/bluecoat/trust_package.bctp
```

Configure Automatic Updates

By default, the appliance automatically downloads and installs the latest trust package every seven days by default. You can disable automatic updates or modify the update interval as follows:

To disable automatic updates:

If you prefer to manually download and install the trust package, you can download automatic updates as follows:

```
#(config) security trust-package auto-update disable
```

To change the update interval:

```
#(config) security trust-package auto-update interval <days>
```

where <*days*> is the number of days between updates. This value can be from 1 to 30 inclusive. For example, to set the auto-update interval to 10 days, you would enter the following command:

```
#(config) security trust-package auto-update interval 10
```

If you previously disabled automatic updates, you can re-enable them using the following command:

```
#(config) security trust-package auto-update enable
```

> **Note:** The appliance can only download and install a trust_package.bctp trust package created by Symantec.

Note that if you previously modified the automatic update interval, your settings will be preserved.

Load the Trust Package

If you want to manually download and install the trust package—either because you have disabled automatic updates or you want to force an update before the next automatic update—enter the following command:

```
# load trust-package
Downloading from "http://appliance.bluecoat.com/sgos/trust_package.bctp"
```

```
The trust package has been successfully downloaded.
trust package successfully installed
```

Verify Trust Package Downloads

Use the following command to view the status of the trust package downloads:

```
# show security trust-package
Download url: http://appliance.bluecoat.com/sgos/trust_package.bctp
Auto-update: enabled Auto-update interval: 7 days
Previous (success) install via manual
Creation time: Saturday October 1 2011 00:26:43 UTC
CA Certificate List changes:
browser-trusted: CAs - 3 added, 4 deleted, 0 modified
image-validation install: Tuesday October 11 2011 00:26:27 UTC
Download log:
Downloaded at: Tuesday October 11 2011 00:26:27 Success
Downloaded from: http://appliance.bluecoat.com/sgos/trust_package.bctp
```

About SSL Device Profiles

An SSL device profile contains the following information required for device authentication:

- The name of the keyring that contains the private key and certificate this device uses to authenticate itself. The default keyring is `appliance-key`.
- The name of the CA Certificate List (CCL) that contains the names of certificates of CAs trusted by this profile. If another device offers a valid certificate signed by an authority in this list, the certificate is accepted. The default is appliance-ccl.
- Verification of the peer certificate (enabled by default).

  When the appliance is participating in device authentication as an SSL client, the peer certificate verification option controls whether the server certificate is validated against the CCL. If verification is disabled, the CCL is ignored.

  When the appliance is participating in device authentication as an SSL server, the peer certificate verification option controls whether to require a client certificate. If verification is disabled, no client certificate is obtained during the SSL handshake.
- The format of device ID authorization data extracted from the certificate. The default is `$(subject.CN)`.
- SSL cipher settings. The default is SHA256.

For an overview of these device profile properties, in the Admin Console, select **Configuration > SSL > Device Profiles**. In Device Profiles, refer to the details for each installed device profile.

Each Symantec appliance has a built-in profile called `bluecoat-appliance-certificate` that can be used for device-to-device authentication. This profile cannot be deleted or edited. If you cannot use the built-in profile (for example, if you are using your own appliance certificates) you must create a different profile that references the keyring containing your certificate.

If you create your own profile, it must contain the same kind of information that is contained in the Symantec profile. To create your own profile, see "Configure SSL Device Profiles for Device Authentication" below.

> **Note:** If you do not want to use peer verification, you can use the built-in `passive-attack-detection-only` profile instead of the `bluecoat-appliance-certificate` profile. This profile uses a self-signed certificate and disables the verify-peer option, so that no authentication is done on the endpoints of the connection. The traffic is encrypted, but is vulnerable to active attacks. This profile can be used only when there is no threat of an active man-in-the-middle attack. Like the `bluecoat-appliance certificate` profile, the `passive-attack-detection-only` profile cannot be edited or deleted.

See Also:

- "Configure SSL Device Profiles for Device Authentication" below

Configure SSL Device Profiles for Device Authentication

You can use the built-in `bluecoat-appliance-certificate` profile, which can't be deleted or edited, for device authentication. If the built-in profile does not meet your needs, you can create a new profile to meet your specific requirements.

> **Note:** Full configuration requires working in the Management Console. Refer to the *SGOS Administration Guide*.

To configure an SSL device profile:

1. In the Admin Console, select the **Configuration > SSL > Device Profiles**.
2. In Device Profiles, select an existing (non-read-only) device profile to edit it, or select **Add Profile** to add a new device profile.
3. On the Add/Edit Profile page, specify the following settings:

Setting or field	Description
Device Profile Name	Unique name for the device profile. The only valid characters are alphanumeric, the underscore, and hyphen, and the first character must be a letter.

Setting or field	Description
Keyring	If the OCS requires a client certificate, select the keyring used to negotiate with the OCS through an encrypted connection. Only keyrings with certificates can be associated with the SSL client. **Note:** You must create a new keyring for device authentication if you do not use the appliance-key keyring. The other keyrings shipped with the appliance are dedicated to other purposes. For information on creating a new keyring, see "Configure a Keyring" on page 209.
CCL	Select the CA Certificate List (CCL) you want to use so that the appliance knows which CA certificates to use when validating the trust of any received certificates.
Verify Peer	When verification enabled, the appliance performs a set of checks to ensure any received certificates are both trusted (as determined by the CA certificates contained within the specified CCL) and valid (that is, not expired, hostname matches, etc.). Enable this setting to aid in troubleshooting. If peer verification is not enabled, the ProxySG appliance doesn't actupon any verification failures it finds when checking the received certificate.
Device ID Extractor	How device ID information is extracted from a presented certificate. The string contains references to the attributes of the subject or issuer in the form $(subject.attr[.n]) or $(issuer.attr[.n]), where attr is the short-form name of the attribute and n is the ordinal instance of that attribute, counting from 1 when the subject is in LDAP (RFC 2253) order. If n is omitted, it is assumed to be 1. The default is $(subject.CN); many other subject attributes are recognized, among them OU, O, L, ST, C, and DC.

Select Cipher Suites

The Selected Ciphers displays cipher suites selected by default when you create the device profile.

To configure the ciphers list, select:

- **Select Ciphers** to add ciphers to the list. On the Select Ciphers dialog that opens, add ciphers and save your changes. For improved security, select only ciphers with High strength.
- **Remove** to remove a cipher from the list.

For more information on cipher suites supported on the appliance, refer to https://www.symantec.com/docs/TECH247556.

See Also:

- "About SSL Device Profiles" on page 221

Integrate with an HSM Appliance

A Hardware Security Module (HSM) provides additional security for storing cryptographic keys and certificates, which is required in some highly regulated industries. The appliance is able to use a network-attached HSM appliance to store resigning CA keys, and to perform digital signature operations. The appliance exchanges signing requests and responses with the attached HSM appliance, over mutually authenticated HTTPS requests. The appliance sends certificate data to the HSM.

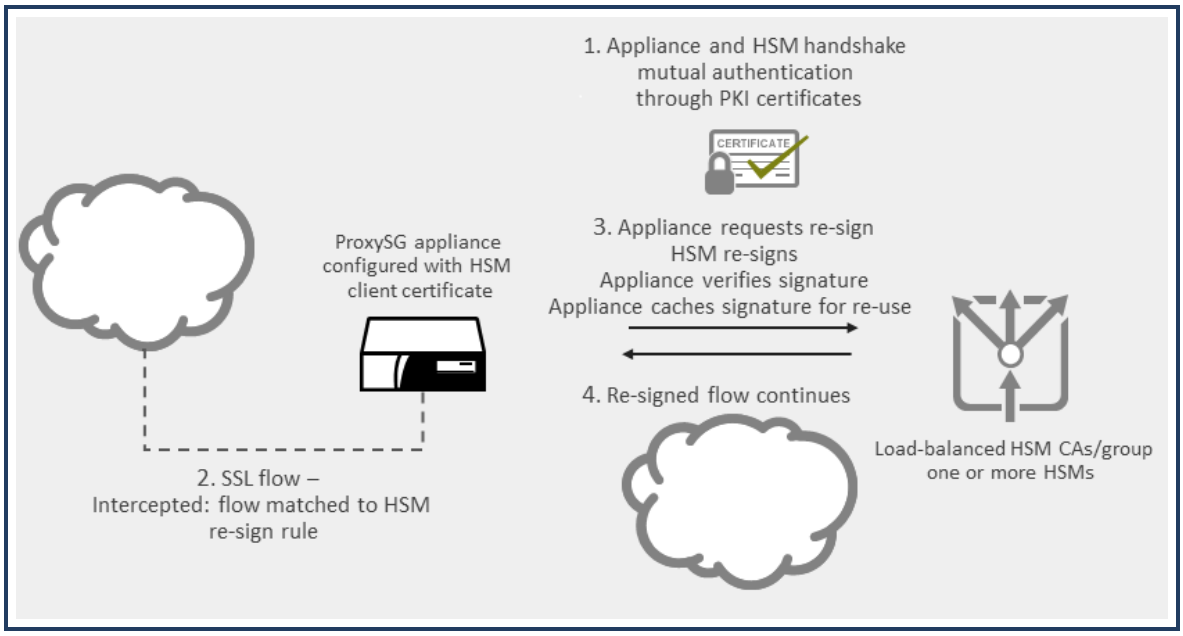

The appliance can work with multiple HSM appliances, and multiple appliances can work with the same HSM. In the event that a policy rule using an HSM to sign cannot work due to lack of response from the HSM, the attempt is logged, and the appliance responds with an exception. In addition to the resigning certificates, a mutually authenticated connection (communication pipeline) must be set up by verified certificates.

Configure a Thales Luna HSM

The Thales HSM must be configured separately. Additionally, Symantec provides an agent to install on the Thales Java HSM, which will be used to interact with Symantec appliances. A certificate to authorize the agent is included. The Symantec HSM Agent operates on top of a secure session. It communicates to the external Symantec entity (ProxySG appliance), and is used remotely.

In order for the appliance to trust the HSM, you must import the server certificate for the HSM, and put it in to a CA Certificate List. In the Admin Console, select **Configuration > SSL > CA Certificates**, and import the certificate. Name the certificate and paste the .PEM data in to the appropriate field. For further information, see "Manage CA Certificates" on page 216.

An HSM requires a linked device profile (in the Admin Console, select **SSL > Device Profiles**). See "Manage HSMs" on page 226.

See Also:

- *Symantec HSM Agent for the Thales Luna 7 HSM* documentation at Tech Docs

View Added HSMs

The HSM section (**Configuration >SSL > HSM**) lists all added HSMs and displays information about each one:

Setting or field	Description
Name	Name of the HSM.
Referenced	Whether or not the HSM is referenced in ProxySG policy.
Hostname/IP Address	Configured hostname or IP address of the HSM.
Port	Configured port of the HSM.
Device Profile	Linked device profile.

See Also:

- "Manage HSMs" on the next page

View Added HSM Keyrings

The HSM Keyrings section (**Configuration >SSL > HSM**) lists all added HSM keyrings and displays information about each one:

Setting or field	Description
Name	Name of the HSM keyring.
Referenced	Whether or not the keyring is referenced in ProxySG policy.
HSM	Name of the HSM.
Key Label	Name of the private key in HSM configuration.
Certificate Expiry	HSM certificate expiration date.
View Certificate (under Actions)	View certificate details and PEM file data. On the Certificate dialog, select **PEM** to copy the certificate.

See Also:

- "Manage HSM Keyrings" on page 227

View Added HSM Keylists

The HSM Keylists section (**Configuration >SSL > HSM**) lists all added HSM keylists and displays information about each one:

Setting or field	Description
Name	Name of the HSM keylist.

Setting or field	Description
Referenced	Whether or not the keylist is referenced in ProxySG policy.
HSM Keyring Count	Number of HSM keyrings in the list.

See Also:

- "Manage HSM Keylists" on page 228

Manage HSMs

> **Note:** A supported HSM must be configured separately. To configure a Thales Luna HSM, refer to HSM documentation at Tech Docs.

To add or edit an HSM:

1. In the Admin Console, select **Configuration > SSL > HSM**.

   In the HSM section, select **Add HSM** to add an HSM or select an existing HSM to edit it.

2. On the Add/Edit HSM page, specify the following settings:

Setting or field	Description
Name	Unique name for the HSM.
Hostname or IP address	Hostname or IP address of the HSM.
Port	Port of the HSM.
Device Profile	Select a device profile. To have an appropriate device profile in the list, in the Admin Console, select **SSL > Device Profiles**. Create a FIPS-compliant or non-compliant profile as required, and then enter the HSM credentials. For more information, see "Configure SSL Device Profiles for Device Authentication" on page 222.

Manage HSM Keyrings

The HSM keyring is used for certificate emulation when intercepting SSL with the HSM. The keyring must contain a certificate for which the HSM has the private key.

Adding an HSM keyring follows the same steps as adding any SSL keyring. Configured HSM keyrings are also available in **Configuration > Services > SSL Proxy Settings > General Settings**.

> **Tip:** Before creating the keyring, make sure that you can access the HSM's certificate file. You will copy and paste it in the keyring configuration.

> **Note:** You can't edit or remove a keyring if it is referenced in policy. Remove the references in policy before attempting to make changes to existing keyrings. See "Configure HSM Policy" on the next page for details.

To add or edit an HSM keyring:

1. In the Admin Console, select **Configuration > SSL > HSM**.

   In the HSM Keyings section, select **Add HSM Keyring** to add a keyring or select an existing keyring to edit it.

2. On the Add/Edit HSM Keyring page, specify the following settings:

Setting or field	Description
Name	Unique name for the HSM keyring.
HSM	Select an HSM from the list.
Key label	Name associated with the private key created on the HSM.

3. To enter or update the certificate, copy the contents of the PEM file including `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----`.

   Then, in the Certificate section, select **Paste From Clipboard**.

See Also:

- "Configure HSM Policy" on the next page

Manage HSM Keylists

Keylists, rather than individual keyrings, can be referenced in policy. When a keylist is used, the SSL connections are load balanced, either within one HSM or across an HSM cluster.

Adding an HSM keylist follows the same steps as adding any SSL keylist. Configured HSM keylists are also available in **Configuration > Services > SSL Proxy Settings > General Settings**.

To add or edit an HSM keylist:

1. In the Admin Console, select **Configuration > SSL > HSM**.

   In the HSM Keylists section, select **Add HSM Keylist** to add a keylist or select an existing keylist to edit it.

2. On the Add/Edit HSM Keylist page, in the Properties section, specify the name of the keylist.

3. In the Select HSM Keyrings section, select **Select HSM Keyrings** to add or remove configured keyrings. The Select HSM Keyrings dialog opens.

   The **Select HSM Keyrings** option does not appear if no keyrings are configured yet.

4. On the Select HSM Keyrings dialog:

   - Select keyrings to add them to the keylist.
   - Clear keyring selections to remove them from the list.

   If **Show Selected only** is enabled, only keyrings that have been added to the keylist are displayed. If **Show Selected only** is disabled, all keyrings available are displayed.

See Also:

- "Configure HSM Policy" below

Configure HSM Policy

Use policy to direct the SSL proxy to use an HSM keyring or keylist to sign an emulated certificate from an intercepted authenticated SSL connection. The following procedure uses the Web Visual Policy Manager: (VPM)

1. In the ProxySG management console, launch the web VPM:

   a. Select **Configuration > Policy > Visual Policy Manager**.

   b. Select **Launch Web VPM**. The web VPM launches in a new browser tab.

2. Select **Add Layer**. In the dialog, select **SSL Intercept**.

3. In the layer you added, add a new Action object:

a. Under action, select **New > Enable SSL Interception**.

b. The **Issuer Keyring** list displays configured HSM keyrings and keylists. Select the keyring or keylist to use for HSM signatures.

c. Apply your changes.

4. Select **Apply Policy** to install the HSM policy.

Authenticate Administrators and Users

Refer to the following topics to learn about authenticating ProxySG administrators and users:

- " About Appliance Security" on page 232
  - "Limit Access to the Appliance" on page 237
    - "Moderate Security: Restrict Access Using the ACL" on page 239
    - "Maximum Security: Administrative Authentication and Authorization Policy" on page 240
- "Control User Access with Identity-based Access Controls" on page 243
  - Authentication with Active Directory Using IWA
    - Prepare for a Kerberos Deployment
    - Configure an IWA Authentication Realm
    - Configure IWA Servers
    - Configure General IWA Properties
    - "Join a Windows Domain" on page 258
  - "Authentication with Active Directory Using LDAP" on page 260
    - "Configure an LDAP Authentication Realm " on page 261
    - "Configure LDAP Realm Details" on page 264
    - "Configure LDAP Servers " on page 267
    - "Define LDAP Authorization Properties" on page 268
  - "RADIUS Realm Authentication and Authorization" on page 272
    - "Configure a RADIUS Authentication Realm " on page 273
    - "Configure RADIUS Realm Details" on page 274
    - "Configure RADIUS Servers " on page 276

- "SAML Realm Authentication and Authorization" on page 277
  - "SAML Authentication Process" on page 280
  - "Overview: SAML Realm Configuration" on page 282
    - "Prepare the Appliance" on page 285
    - "Configure a SAML Realm " on page 286
    - "Configure SAML Attributes" on page 287
    - "Configure SAML Realm Details" on page 289
    - "Configure SAML Authorization" on page 292
  - "Considerations" on page 294

- "Authenticate Outbound SSH Client Connections" on page 295
  - "Manage Host Keys for Outbound SSH Connections" on page 295
  - "Manage SSH Ciphers for Outbound Connections" on page 299
  - "Manage SSH HMACs for Outbound Connections" on page 301
- "Manage the SSH Console" on page 303
  - "Manage the Host Key Pairs" on page 304
  - "Manage SSH Client Keys for Inbound Connections" on page 306
  - "Manage SSH Ciphers for Inbound Connections" on page 307
  - "Manage SSH HMACs for Inbound Connections" on page 309

About Appliance Security

You can limit access to the appliance by configuring Admin Console and CLI console account settings. You can also specify read-only or read-write access to admins as required.

Configure Login Parameters

You can change the console username and password, the console realm name which displays when you log in to the appliance, and the auto-logout time. The default value is 900 seconds. The Admi Console requires a valid administrator username and password to have full read-write access; you do not need to enter a privileged-mode password as you do when using the CLI. A privileged-mode password, however, must already be set.

> **Note:** To prevent unauthorized access to the ProxySG appliance, only give the console username and password to those who administer the system.

Change Admin Account Credentials

During the initial configuration of your ProxySG appliance, a console administrator username and password was created. This is a special account that can always be used to administer the appliance from either the Admin Console or the CLI. You can change the username and the password of this administrator account.

Edit the username of the administrator who is authorized to view and revise console properties. Only one console account exists on the ProxySG appliance. If you change the console account username, that username overwrites the existing console account username. The console account username can be changed to anything that is not null and contains no more than 64 characters.

After changing the username and password, you can further restrict access by creating an access control list or by creating a policy file containing `<Admin>` layer rules. See "Limit Access to the Appliance" on page 237.

> **Note:** Changing the console account's username or password causes the Admin Console to refresh, requiring you to log in again using the new credentials. Each parameter must be changed and individually refreshed. You cannot change both parameters at the same time.

To change the username:

1. In the Admin Console, select **Configuration > Authentication > Console Access**.
2. In **Console Account**, enter a new name in the **Username** field.

   After you save changes, if you refresh the browser page, the console prompts you to re-enter your credentials. Enter the new username,

To change the password:

> **Note:** The console password and privileged-mode password were defined during initial configuration of the system. The console password can be changed at any time. The privileged-mode, or enabled-mode, password can only be changed through the CLI or the serial console.

1. In the Admin Console, select **Configuration > Authentication > Console Access**.
2. In **Console Account**, select **Change Password**.
3. In the Change Password dialog, enter the current password. Enter and confirm the new password.

   The passwords are obscured by asterisks. You can select the eye icon to display the passwords.

   The password must be from 1 to 64 characters long.

After you save changes, the console prompts you to re-enter your credentials. Enter the new password.

Change the Appliance Realm Name

When you have multiple ProxySG appliances in your network, you can configure a console realm name to identify the appliance that you are accessing.

When you log in to the Admin Console, using a browser, the browser's pop-up dialog displays. This dialog identifies the ProxySG appliance that is requesting the username and password.

If configured, the realm name displays on the pop-up dialog. The default realm name is usually the IP address of the ProxySG appliance. You can, however, change the display string to reflect your description of the appliance.

To change the appliance realm name:

1. In the Admin Console, select **Configuration > Authentication > Console Access**.
2. In **Console Account**, enter a new realm name in the **Console realm name** field.

The next time you log in to the Admin Console, the new realm name displays on the browser's pop-up dialog.

Change the ProxySG Appliance Timeout

The timeout is the length of time an Admin Console or CLI session persists before you are logged out. The default timeout for these options is as follows:

- Enforce Web auto-logout–15 minutes
- Enforce CLI auto-logout–5 minutes

To change the timeout:

1. In the Admin Console, select **Configuration > Authentication > Console Access**.
2. In **Console Account**:
   - To enable an Admin Console timeout, switch the **Enforce Web auto-logout** toggle and specify the number of minutes in the **Web auto-logout (minutes)** field.
   - To enable a CLI timeout, switch the **Enforce CLI auto-logout** toggle and specify the number of minutes in the **CLI auto-logout (minutes)** field.

Control Access to the Appliance

- Restricting physical access to the system and by requiring a PIN to access the front panel.
- Restricting the IP addresses that are permitted to connect to the ProxySG CLI.
- Requiring a password to secure the setup console.

For more information, see "Limit Access to the Appliance" on page 237.

Refer to the following to determine how to give other users read-only or read-write access to the ProxySG appliance.

> **Note:** If Telnet Console access is configured, Telnet can be used to manage the appliance with behavior similar to SSH with password authentication. SSL configuration is not allowed through Telnet, but is permissible through SSH. Behavior in the following sections that applies to SSH with password authentication also applies to Telnet. Use of Telnet is not recommended because it is not a secure protocol

- ### Limit access to the console account (minimum security)

  The console account username and password are evaluated when the ProxySG appliance is accessed from the Admin Console through a browser and from the CLI through SSH with password authentication. The Enable (privileged mode) password is evaluated when the console account is used through SSH with password authentication and when the CLI is accessed through the serial console and through SSH with RSA authentication. The simplest way to give access to others is sharing this basic console account information, but it is the least secure and is not recommended.

  To give read-only access to the CLI, do not give out the Enable (privileged mode) password.

- ### Create an ACL (moderate security)

  Using the ACL allows you to further restrict use of the console account and SSH with RSA authentication to workstations identified by their IP address and subnet mask. When the ACL is enforced, the console account can only be used by workstations defined in the console ACL. Also, SSH with RSA authentication connections are only valid from workstations specified in the console ACL (provided it is enabled).

  After setting the console account username, password, and Enable (privileged-mode) password, use the CLI or the Admin Console to create a console ACL. See "Moderate Security: Restrict Access Using the ACL" on page 239.

- Per-user RSA public key authentication (moderate security)

  Each administrator's public keys are stored on the appliance. When connecting through SSH, the administrator logs in with no password exchange. Authentication occurs by verifying knowledge of the corresponding private key. This is secure because the passwords never go over the network.

  This is a less flexible option than CPL because you cannot control level of access with policy, but it is a better choice than sharing the console credentials.

- Use CPL to control access (maximum security)

  CPL allows you to control administrative access to the ProxySG appliance through policy. If the credentials supplied are not the console account username and password, policy is evaluated when the appliance is accessed through SSH with password authentication or the Admin Console.

  Policy is never evaluated on direct serial console connections or SSH connections using RSA authentication.

  - Using the CLI or the Admin Console, create an authentication realm to be used for authorizing administrative access. For administrative access, the realm must support BASIC credentials—for example, LDAP, RADIUS, Local, or IWA with BASIC credentials enabled.
  - Using the VPM, or by adding CPL rules to the Local or Central policy file, specify policy rules that meet both criteria:
    - Require administrators to log in using credentials from the previously-created administrative realm:
    - Specify the conditions under which administrators are either denied all access, given read-only access, or given read-write access.

    Authorization can be based on IP address, group membership, time of day, and many other conditions. For more information, refer to the *Web Visual Policy Manager Reference*.
  - To prevent anyone from using the console credentials to manage the ProxySG appliance, set the console ACL to deny all access (unless you plan to use SSH with RSA authentication). For more information, see "Moderate Security: Restrict Access Using the ACL" on page 239. You can also restrict access to a single IP address that can be used as the emergency recovery workstation.

  For details on using policy to limit access, see "Maximum Security: Administrative Authentication and Authorization Policy" on page 240.

The following chart details the various ways administrators can access the ProxySG console and the authentication and authorization methods that apply to each.

Security Measures Available	Serial Console	SSH with Password Authentication	SSH with RSA Authentication	Admin Console
Username and password evaluated (console-level credentials)		X		X

Security Measures Available	Serial Console	SSH with Password Authentication	SSH with RSA Authentication	Admin Console
Console Access List evaluated		X (if console credentials are offered)	X	X (if console credentials are offered)
CPL `<Admin>` Layer evaluated		X [1]		X [2]
Enable password required to enter privileged mode [2]	X	X	X	
CLI line-vty timeout command applies.	X	X	X	
Admin Console Login/Logout				X

[1] When using SSH (with a password) and credentials other than the console account, the enable password is actually the same as the login password. The privileged mode password set during configuration is used only in the serial console, SSH with RSA authentication, or when logging in with the console account.

[2] In this case, user credentials are evaluated against the policy before executing each CLI command. If you log in using the console account, user credentials are not evaluated against the policy.

Limit Access to the Appliance

Limit access to the appliance using the following methods:

- (If applicable) "Specify a PIN for the Front Panel" below
- "Limit Workstation Access" below
- "Secure the Serial Port" below

Specify a PIN for the Front Panel

On systems that have a front-panel display, you can create a four-digit PIN to protect the system from unauthorized use. The PIN is hashed and stored. You can only create a PIN from the command line.

To create a front panel PIN, after initial configuration is complete, type

```
# (config) security front-panel-pin PIN
```

where *PIN* is a four-digit number.

To clear the front-panel PIN, enter:

```
# (config) security front-panel-pin 0000
```

Limit Workstation Access

During initial configuration, you have the option of preventing workstations with unauthorized IP addresses from accessing the ProxySG appliance for administrative purposes. This covers all access methods - Telnet, SNMP, HTTP, HTTPS and SSH. If this option is not enabled, all workstations are allowed to access the appliance administration points. You can also add allowed workstations later to the ACL.

For more information on limiting workstation access, see "Moderate Security: Restrict Access Using the ACL" on page 239.

Secure the Serial Port

If you choose to secure the serial port, you must provide a setup console password that is required to access the Setup Console in the future.

After the secure serial port is enabled:

- The setup console password is required to access the setup console.
- An authentication challenge (username and password) is issued to access the CLI through the serial port.

To recover from a lost setup console password, you can:

- Use the front-panel display to either disable the secure serial port or enter a new setup console password.
- Use the `restore-defaults factory-defaults` CLI command to delete all system settings.
- (If the appliance has a reset button) Use the reset button to delete all system settings. Otherwise, reset the appliance to its factory settings by holding down the left arrow key on the front panel for five seconds. The appliance will be reinitialized.

To reconfigure the appliance or secure the serial port, refer to the hardware guides for your appliance.

Moderate Security: Restrict Access Using the ACL

The ProxySG appliance allows you to limit access to the Admin Console and CLI through the console ACL. An ACL, once set up, is enforced only when console credentials are used to access either the CLI or the Admin Console, or when an SSH with RSA authentication connection is attempted. The following procedure specifies an ACL that lists the IP addresses permitted access.

To restrict Admin Console Access using the ACL:

To create an ACL:

1. In the Admin Console, select **Configuration > Authentication > Console Access**.
2. (Optional) In **Console Access**, modify the list of IP addresses in the ACL. Select Add and specify the following:

Field or Value	Description
Source Address	Static IP address. .
Prefix Length (Subnet Mask)	Subnet mask. To restrict access to an individual workstation, enter 255.255.255.255

Select **Apply** after each entry.

> **Caution:** Before you enforce the ACL, verify the IP address for the workstation you are using is included in the list. If you forget, or you find that you mistyped the IP address, you must correct the problem using the serial console.

3. To enable ACL, switch the **Enforce ACL for built-in administrators** toggle.

See Also:

- "Maximum Security: Administrative Authentication and Authorization Policy" on the next page

Maximum Security: Administrative Authentication and Authorization Policy

The ProxySG appliance permits you to define a rule-based administrative access policy. This policy is enforced when accessing:

- the Admin Console through HTTP or HTTPS
- the CLI through SSH when using password authentication
- the CLI through telnet
- the CLI through the serial port if the secure serial port is enabled

These policy rules can be specified either by using the VPM or by editing the Local policy file. Using policy rules, you can deny access, allow access without providing credentials, or require administrators to identify themselves by entering a username and password. If access is allowed, you can specify whether read-only or read-write access is given. You can make this policy contingent on IP address, time of day, group membership (if credentials were required), and many other conditions.

Serial-console access is not controlled by policy rules. For maximum security to the serial console, physical access must be limited.

Define Administrator Authentication and Authorization Policies

Administrative authentication uses policy, (either VPM or CPL in the local policy file) to authenticate administrative users to the appliance. This is done with two layers in policy: one to define the realm that is used to authenticate users (**Admin Authentication** layer) and the other to define security rights for authenticated users or groups (**Admin Access** layer).

> **Note:** If you choose a realm that relies on an external server and that server is unavailable, the appliance will not be able to authenticate against that realm. For best security, Symantec recommends the following authentication realms for administrative authentication to the appliance.

- IWA-BCAAA (with TLS -- not SSL) with basic credentials
- Local
- .509 certificate based (including certificate realms; refer to the *Common Access Card Solutions Guide* for information)
- LDAP with TLS (not SSL)
- IWA-Direct with basic credentials RADIUS

The following realms can be configured for administrative authentication, but pass administrative credentials in clear text. These realms should not be used for administrative authentication:

- Windows SSO
- Novell SSO
- IWA-BCAAA without SSL or TLS
- LDAP without SSL or TLS

The following realms do not support administrative authentication:

- IWA-BCAAA/IWA-Direct realms that do not accept basic credentials
- SiteMinder
- COREid
- SAML (Policy Substitution)
- XML

> **Note:** Other authentication realms can be used, but will result in administrative credentials being sent in clear text.

Configure Administrative Authentication with a Local Realm

> **Note:** Some of the following steps require working in the Management Console.

The process to provide read-only access for administrators includes the following steps:

- Create a local authentication realm.
- Create a list that includes usernames and passwords for members whom you wish to provide read-only access in the Admin Console.
- Connect the list to the local realm.
- Create policy to enforce read-only access to members included in the list.

To configure administrative authentication with a local realm:

1. Create a local realm. Refer to the *SGOS Administration Guide*.
2. Using the CLI, create a list of users who need read-only access. The list must include a username and password for each user.

a. Enter configuration mode in the CLI; this example creates a list called Read_Access:

```
# (config) security local-user-list create Read_Access
```

b. Edit the list to add user(s) and to create usernames and passwords. See the following example:

```
# (config) security local-user-list edit Read_Access
# (config) user create Bob_Kent
# (config) user edit Bob_Kent
# (config) password 12345
```

3. In the Management Console, connect the user list (created in step 2) to the local realm (created in step 1). Refer to the *SGOS Administration Guide*.
4. Use the for creating policy to enforce read-only access to the users in your list:

a. In the web VPM, create an **Admin Authentication** layer (or add a new rule in an existing layer). This layer determines the authentication realm that will be used to authenticate users who access the appliance Admin Console.

b. Add an **Authentication** action object.

c. In the Authentication object, select the local realm you created in step 1.

d. Add an **Admin Access** layer.

e. Add a **User** source object.

f. In the User object, enter the name of the user for whom you want to provide read-only access.

g. Add an **Allow Read-only Access** action object.

After you save all changes and apply policy, the user can log into the Admin Console with read-only access.

> **Note:** Use the previous steps, with an **Allow Read/Write Access** action object, to define user access with read-write access.

Control User Access with Identity-based Access Controls

Enterprise-wide security begins with security on the ProxySG appliance, and continues with controlling user access to the Intranet and Internet. SSH and HTTPS are the recommended (and default) methods for managing access to the appliance. SSL is the recommended protocol for communication between the and a realm's off-box authentication server.

You can control access to the appliance in several ways: by limiting physical access to the system, by using passwords, restricting the use of console account, through per-user RSA public key authentication, and through CPL. How secure the system needs to be depends upon the environment. See " About Appliance Security" on page 232 for details.

By using every possible method (physically limiting access, limiting workstation IP addresses, and using passwords), the appliance is very secure. After the appliance is secure, you can limit access to the Internet and intranet. It is possible to control access to the network without using authentication. You only need to use authentication if you want to use identity-based access controls.

View Configured Realms

The appliance provides a flexible authentication architecture that supports multiple services with multiple backend servers (for example, LDAP directory servers together with NT domains with no trust relationship) within each authentication scheme with the introduction of the realm.

A realm authenticates and authorizes users for access to ProxySG services using either explicit proxy or transparent proxy mode, discussed in "About Proxy Services " on page 150.

Multiple authentication realms can be used on a single appliance. Multiple realms are essential if the enterprise is a managed provider or the company has merged with or acquired another company. Even for companies using only one protocol, multiple realms might be necessary, such as the case of a company using an LDAP server with multiple authentication boundaries. You can use realm sequencing to search the multiple realms all at once.

To display the list of realms on the appliance, in the Admin Console, select **Configuration > Authentication > Realms and Domains**. The Realms section displays realms configured on the appliance and indicates the Type/Protocol of each one.

Realm configuration includes:

- Unique realm name.
- Authentication service. This release of the Admin Console supports configuration of IWA realms. Refer to Authentication with Active Directory Using IWA for details. To configure other realms, use the Management Console and refer to the *SGOS Administration Guide*.
- External server configuration: Back-end server configuration information, such as host, port, and other relevant information based on the selected service.
- Authentication schema: The definition used to authenticate users.
- Authorization schema: The definition used to authorize users for membership in defined groups and check for attributes that trigger evaluation against any defined policy rules.

See Also:

- Configure an IWA Authentication Realm

Integration with the Windows Domain

The following configurations require that you join your ProxySG appliance to your Windows Domain:

- To accelerate encrypted MAPI traffic, the appliance at the branch office must join the same domain as the Exchange server.
- To configure an IWA Direct realm, you must first join the appliance to your Windows domain. If you want to authenticate users in different AD domains that do not have trust relationships, you must join the appliance to each domain.

To integrate the ProxySG appliance into one or more Windows domains:

1. Synchronize the appliance and DC clocks.

   The appliance can't join a Windows domain unless its internal clock is in sync with the DC. To ensure that the clocks are synchronized with the DC clock, specify the same NTP servers for the DC as the NTP source server. To configure the appliance NTP servers, select **Configuration > General > Time**. See Manage NTP Servers .

2. Join the appliance to the Windows domain. See "Join a Windows Domain" on page 258.

Authentication with Active Directory Using IWA

Integrated Windows Authentication (IWA) is an authentication scheme that allows you to authenticate and authorize users against your Windows AD. One of the main benefits of IWA is that it can provide a single sign-on experience for your users. When configured properly, the user agent or browser will automatically provide the users' domain credentials to the appliance when challenged without prompting the end users.

Another benefit of IWA is that it provides authorization without any additional configuration because it automatically returns group membership information for the user as part of the authentication response. The ProxySG appliance can then use this group membership information to enforce its authorization policies.

Symantec supports two methods to integrate the appliance with Active Directory using IWA:

- IWA BCAAA: Connect via an authentication agent running on a Windows server in your domain.

  For instructions, refer to the *BCAAA Service Requirements* document posted on MySymantec:

  https://support.symantec.com/content/unifiedweb/en_US/ Documentation.html?prodRefKey=1145522

- IWA Direct: Connect the appliance directly to your AD domains on Windows Server 2008, 2012, or 2016. Refer to the following sections.

> **Note:** You can configure IWA Direct authentication with both IPv4 and IPv6 network addresses.

You can authenticate administrative users to the appliance with IWA realms. IWA Direct, though less secure than IWA BCAAA, is compatible with administrative authentication. For maximum security, Symantec recommends using an IWA BCAAA configuration, where the communication between the appliance and BCAAA is secured with TLS.

About IWA Challenge Protocols

When configured for IWA, the ProxySG appliance determines which of the following protocols to use to obtain Windows domain login credentials each time it receives a client request that requires authentication:

- Kerberos: This is the most secure protocol because it establishes mutual authentication between the client and the server using an encrypted shared key. This protocol requires additional configuration and the appliance will silently downgrade to NTLM if Kerberos is not set up properly or if the client cannot do Kerberos. For more information, see "Prepare for a Kerberos Deployment" on page 247.
- NTLM: Uses an encrypted challenge/response that includes a hash of the password. NTLM requires two trips between the workstation and the appliance, and one trip between the appliance and the DC. It therefore puts more load on the network than Kerberos, which only requires one trip between the workstation and the appliance, and doesn't require a trip between the appliance and the DC.
- Basic: Prompts the user for a username and password to authenticate the user against the Windows Active Directory.

When the appliance receives a request that requires authentication, it consults the IWA configuration settings you have defined to determine what type of challenge to return to the client. It will try to use the strongest authentication protocol that is configured and, if the browser cannot use that protocol or if it is not configured properly, the appliance will downgrade to the next authentication protocol. For example, if you configure the IWA realm to allow Kerberos and NTLM authentication, but the user agent/browser does not support Kerberos, the appliance will automatically downgrade to NTLM. IWA authentication realms (with basic credentials) can be used to authenticate administrative users (read only and read/write) to the Admin Console. To ensure that credentials are not sent in clear text, configure the IWA realm to use TLS to secure the communication with the BCAAA server, or in the case of IWA direct, secure the communication from the appliance to the domain.

About IWA Failover

The way IWA failover works depends on your deployment:

- "IWA Direct Failover" on the next page
- "IWA BCAAA Failover " on the next page

IWA Direct Failover

For IWA Direct, the realm is considered "healthy" if the appliance is able to establish a connection to the Windows domain to which it is a member. As with any other device in the Windows domain, the appliance will establish a connection with the closest Windows DC upon successful domain login. If the DC to which the appliance is connected goes down, the appliance will send an LDAP ping to locate and connect to the next closest DC.

Because communication between the appliance and the Windows Active Directory relies on DNS, you must make sure that the appliance is configured to use more than one DNS server to ensure proper failover. This will ensure that the appliance will still be able to communicate with AD, should the primary DNS server go offline. For instructions, see Configure the Primary or Alternate Group.

IWA BCAAA Failover

For IWA BCAAA, the realm is considered "healthy" (and therefore won't fail over) if the appliance is able to establish a connection to the BCAAA service. This means that the appliance is able to complete the TCP handshake with BCAAA on port 16101 (or whichever port the BCAAA service is configured to use), and the appliance has been able to send BCAAA its "login" message. There are several different failover scenarios in a BCAAA deployment:

Each time an appliance connects to BCAAA, the BCAAA service (bcaaa.exe) spawns a new BCAAA process (such as bcaaa-130.exe). If the BCAAA process crashes, the TCP connection with the corresponding appliance will be reset and the appliance will attempt to reconnect to the BCAAA service. Other appliances that are connected to other instances of the BCAAA process will be unaffected.

- If the BCAAA service (bcaaa.exe) crashes or is stopped, but the Windows system on which it is running remains available, any appliance that is already connected to the BCAAA process will have their connections reset. The appliances will not be able to reconnect to BCAAA because the service is no longer running, and will instead fail over to the secondary BCAAA server.
- If the Windows server on which BCAAA is running crashes or becomes unavailable, it cannot reset the TCP connection. In this case, BCAAA must wait for the appliance's TCP connection to the Windows server to time out. This can take a couple of minutes, and won't occur until the appliance attempts to send a new authentication request.
- If the BCAAA server loses its connection to the Windows DC, it will automatically fail over to a different DC. Keep in mind that BCAAA cannot detect when Windows fails to connect to any DCs in a particular domain. In this case all authentication requests will fail, but because the connection between the BCAAA service and the appliance is still considered healthy, the appliance will not fail over to the secondary BCAAA service.

In addition, authentication requests can be slowed significantly if BCAAA is querying a slow DC. However, this will not cause the appliance to fail over to the secondary BCAAA server. By default, BCAAA will query whichever DC is chosen at boot time by the server it is installed on, and it only changes if the DC goes down or the server reboots. You can see and/or modify what DC the BCAAA server is communicating with using the nltest.exe utility, which is part of the Windows Support Tools.

To see which DC the BCAAA server is communicating with:

```
nltest /sc_query:internal.domain.com
```

To switch to a different DC:

```
nltest /sc_reset:internal.domain.com\new_dc_name
```

Prepare for a Kerberos Deployment

Kerberos is the recommended authentication protocol for IWA because it is more secure than NTLM or Basic and it puts the least load on your network. To ensure that IWA uses the Kerberos protocol rather then downgrading to NTLM, you just need to make sure that authentication requests are directed to the Kerberos service associated with the appliance. The way you do this depends on how your IWA realm is connecting to the Active Directory as follows:

- "Enable Kerberos in an IWA Direct Deployment" below
- "Enable Kerberos in a BCAAA Deployment " below

Enable Kerberos in an IWA Direct Deployment

In an IWA Direct realm, Kerberos configuration is minimal because the appliance has its own machine account in Active Directory and it uses its account password to decrypt service tickets from clients. Therefore, there is no need for you to create a privileged Active Directory account or generate an SPN for the appliance as is required with an IWA BCAAA realm.

To ensure that IWA uses the Kerberos protocol rather then downgrading to NTLM, you just need to make sure that authentication requests are directed to the DNS name of the appliance's Active Directory machine account name as follows:

1. Create a DNS "A" record for the appliance that resolves to the DNS name of the appliance's Active Directory machine account name. For example, if you have an appliance named 1 with IP address 1.2.3.4 in the blue9 Active Directory domain at acme.com, you would create the following DNS record:

   ```
   Proxy1.blue9.acme.com Host(A) 1.2.3.4
   ```

2. Ensure that client requests are directed to the DNS name for the appliance's Active Directory machine account:
   - Explicit deployments: Configure the client browser explicit proxy settings to point to this DNS name.
   - Transparent deployments: Set the **Virtual URL** in the general realm settings to this DNS name. In addition, make sure that the DNS name for the appliance's Active Directory domain is either included in the workstation's list of imputing DNS suffixes or explicitly specified as part of Internet Explorer's local intranet zone. For example, if your AD domain DNS name is blue9.acme.com, then you would add *.blue9.acme.com to IE's local intranet zone. See "Configure IWA Realm Details" on page 250.

Enable Kerberos in a BCAAA Deployment

For the BCAAA service to participate in an IWA Kerberos authentication exchange, it must share a secret with the Kerberos server and have registered an appropriate SPN. To prepare for a BCAAA Kerberos deployment:

1. Create a DNS "A" record for the appliance that resolves to the appliance's FQDN. Keep in mind that the DNS name you choose must not match the Active Directory machine account name for the appliance. For example, rather than using

the machine name, you might create a DNS entry for the appliance using a name such as bcaaaUser1.

Supposing the appliance is in the acme.com domain and has an IP address of 1.2.3.4, you would create the following DNS record:

bcaaaUser1.acme.com Host(A) 1.2.3.4

After you create the DNS mapping, make sure you can ping the appliance using the FQDN.

2. Create a domain user account for the BCAAA service in the Windows AD.
3. Install BCAAA. *Refer to the BCAAA Service Requirements* document for installation instructions. Go to MySymantec:

   https://support.symantec.com/content/unifiedweb/en_US/Documentation.html?prodRefKey=1145522
4. Configure the BCAAA Windows service on the system where you just installed BCAAA to log on using the domain user account you created for it in step 2 rather than using the local system account.
5. In the Local Security Policy of the server on which BCAAA is running, modify the user rights assignment for the BCAAA domain user to have the following rights:
   - Full access to the directory where you installed BCAAA
   - Act as part of the operating system (not required for BCAAA 6.0)
   - Log on as a service
6. Register the Kerberos SPN for the appliance:
   a. Log in to the DC using an account with administrative access and open a command prompt.
   b. Enter the following case-sensitive command:

      ```
      setspn -A HTTP/<FQDN_of_Proxy> <AD_Account_Name>
      ```

      where:
      - *<FQDN_of_Proxy>* is the FQDN of the appliance as specified in the browser's explicit proxy configuration (explicit deployments) or in the **Virtual URL** setting in the IWA realm configuration (transparent deployments)
      - *<AD_Account_Name>* is the name of the BCAAA domain service account, for example:

        ```
        setspn -A HTTP/bcaaaUser1.acme.com AcmeDomain\BCAAAuser
        ```

> **Note:** Do not assign the same SPN to multiple AD accounts or the browser will fall back to NTLM without providing any warning or explanation. To list all SPNs that are currently registered on an account, use the `setspn -L <AD_Account_Name>` command. If you find a duplicate, remove the extraneous SPN using the `setspn -D <SPN>` command.

Configure an IWA Authentication Realm

You can configure an IWA realm in the Admin Console. See Authentication with Active Directory Using IWA for details.

> **Note:** All other realms are configurable in the Management Console. Refer to the *SGOS Administration Guide*.

Before you can create an IWA realm, you must integrate with the Windows domain. The way you do this depends on how you plan to connect to your Active Directory:

- Direct: The appliance will communicate directly with your DCs to obtain authentication information. Before you can use this option, you must join the appliance to the Windows domains that contain your users as described in "Join a Windows Domain" on page 258.

  Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAAA compatibility. The *BCAAA Service Requirements* document is posted at MySymantec:

  https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/bcaaa-service-requirements.html

  > **Note:** In the **Realms** list (**Configuration > Authentication > Realms and Domains**), an IWA Direct realm says "IWA On-box" under Type/Protocol.

- BCAAA: The appliance will contact the BCAAA server when it needs to authenticate a user. To use this option, you must first install BCAAA on a dedicated server in your Windows domain and configure it to communicate with both the DC and with the appliance as an authentication agent. Use this option if you do not want to allow the appliance to join your Windows domain. For more information, refer to the *BCAAA Service Requirements* document.

  > **Note:** In the **Realms** list (**Configuration > Authentication > Realms and Domains**), an IWA BCAAA realm says "IWA" under Type/Protocol.

To configure an IWA authentication realm:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > IWA** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add New /Edit IWA Realm page, enter or edit the following as needed:

Setting or field	Description
Realm name	Unique name for the realm.
Active Directory connection	Select the type of Active Directory Connection you are using and then provide the appropriate configuration information as follows: ■ **Direct**: (Available if the appliance is joined to at least one Windows domain) Connect directly to the Windows Domain to obtain authentication information. ■ **BCAAA**: Connect via the BCAAA server. See below for additional settings.
Primary Server	(BCAAA only) Hostname or IP address of the server where you installed BCAAA. If you plan to secure communication between the appliance and BCAAA, use a host name rather than an IP address. The DNS server that the appliance is configured to use must be able to resolve the hostname.
Port	(BCAAA only) If you configured BCAAA to use a port other than the default (16101), specify it.
Advanced Settings	Select **Show** to specify custom settings: ■ "Configure IWA Servers " on page 252 ■ "Configure IWA Realm Details" below
Test Configuration	(Only available when editing) Select to test the configuration. When prompted, enter the username and password of a user in the Windows domain and select **OK**. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm properly. If the test succeeds, the message also displays a list of any groups of interest (that is, groups that are referenced in policy) to which the user belongs.

Configure IWA Realm Details

You can configure the behavior of the authentication transaction, such as timeout and refresh intervals and cookie usage. You also configure the Virtual URL for transparent authentication requests.

To configure IWA realm details:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > IWA** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add/Edit IWA Realm page, in Advanced Settings, select **Show**.
4. In the sections that appear, enter or edit the following as needed:

Task	Setting or field	Description
Set display name	**Display name**	(Optional) By default, the appliance displays the authentication realm name when prompting the user for authentication credentials.
Configure refresh times for Basic credentials	**Use the same refresh time for all**	Select to use the same time for credential refresh and surrogate refresh.
	Credential refresh time	Specifies the amount of time the appliance will cache Basic credentials (username and password) and use these cached credentials to authenticate the user rather than sending another request to the authentication server. By default, basic credentials are good for 900 seconds (15 minutes).
	Surrogate refresh time	(Unavailable if set to use the same refresh time for all) After the appliance successfully authenticates a client, it caches the client's IP address or a cookie (depending on the authentication mode that is in use) in its surrogate cache. If it receives subsequent requests from the same client during the surrogate refresh time, it uses the IP address or cookie in its cache to authenticate the user instead of sending a request to the authentication server. By default, the surrogate credential is good for 900 seconds (15 minutes).
Configure timeouts	**Inactivity timeout**	When a client request is successfully authenticated, the appliance establishes an active session with the client and as long as that session stays active, the appliance will not attempt to reauthenticate requests from that client. This setting specifies how long the client session can be inactive before the appliance terminates the session; subsequent requests from that client will require authentication. By default, the client can be inactive for 900 seconds (15 minutes). **Note:** If the **Challenge user after logout** option is selected, the appliance will automatically challenge the client for credentials when the session becomes inactive. If you are using a challenge method that prompts the user for credentials, you might want to clear this option.
	Rejected credentials timeout	Specifies whether to cache failed authentication attempts (bad password, expired account, disabled account, old password, or server down). If the client attempts to connect again during the rejected credentials time, the appliance will automatically reject the request for the specified period of time. Enter a value from 1 second (the default) to 10 seconds. To disable this option, enter 0.

Task	Setting or field	Description
Configure cookie settings to specify how the appliance uses cookie surrogates Applicable if you plan to use an authentication mode that uses cookie surrogates.	**Use persistent cookies**	By default, this option is not selected, which means that the appliance will use session cookies when creating a cookie surrogate for a client. Session cookies are only valid during the current browser session and are deleted when the user closes the browser. The appliance must re-authenticate the client each time the user starts a new browser session. If you select this option, the appliance will use persistent cookies instead of session cookies. Persistent cookies are stored on the client system and are therefore not deleted at the end of the browser session. When using persistent cookies, the appliance will only need to re-authenticate a client when the cookie in its surrogate credential database expires.
	Verify the IP address in the cookie	By default, this option is selected, which means that the appliance will only accept a cookie from a client if the client IP matches the IP address in the surrogate cookie. To enable the appliance to accept cookies from IP addresses that do not match the address in the cookie–for example, if you use DHCP–clear this option.
Set virtual URL	**Virtual URL**	(Transparent proxy only) Specify the URL to which to redirect client requests that require authentication in the Virtual URL field. For best results, the virtual URL you specify must: ■ Contain a simple hostname that does not contain any dots (for example, use http://myproxy rather than http://myproxy.acme.com. ■ Resolve to the IP address of the appliance. To accomplish this, you must add an "A" record to your internal DNS server that associates the Virtual URL with the IP address of the appliance. ■ (IWA Direct Kerberos only) If you're using Kerberos in a non-load balancing IWA Direct realm, the Virtual URL must be the DNS name of the appliance in the Active Directory domain. Typically this will be the DNS name of the Active Directory domain prefixed with the appliance machine account name. For example, proxy.blue9.local. If you do not use the Active Directory DNS name of the appliance as the Virtual URL, all authentication transactions will be downgraded to NTLM.
Configure re-authentication after timeout	**Challenge user after logout**	(Optional) Select to prompt the client for authentication credentials whenever the inactivity timeout expires.

Configure IWA Servers

You can configure the connection between the appliance and the authentication server (either directly or via BCAAA), as well as failover servers for BCAAA. Also, specify the type of credentials to accept from the browser/user agent.

The way you set up the configuration depends on whether you connecting directly to the DC or you are using BCAAA to connect to the domain:

- "Configure Custom Settings for an IWA BCAAA Realm" below
- "Configure Custom Settings for an IWA Direct Realm" on the next page

Configure Custom Settings for an IWA BCAAA Realm

If you plan to use a BCAAA server to act as an intermediary between your appliance and your Active Directory, you can configure and verify the authentication settings for connecting to the Windows domain using BCAAA.

To configure custom settings for an IWA BCAAA realm:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > IWA Realm** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add/Edit IWA Realm page, in Advanced Settings, select **Show**.
4. In **Servers**, enter or edit the following as needed:

Task	Setting or field	Description
Configure failover	**Alternative Server Host**	(Optional) Hostname or IP address of the secondary server where you installed BCAAA. If you plan to secure communication between the appliance and BCAAA, use a host name rather than an IP address. The DNS server that the appliance is configured to use must be able to resolve the hostname.
	Port	(Optional) If you configured the secondary BCAAA server to use a port other than the default (16101), specify it.

Task	Setting or field	Description
Specify credentials to use	**Client Authentication Mode**	Type of credentials to accept from the browser/user agent. By default, all credential types are allowed and the appliance will try to use Kerberos (the default authentication method for Windows clients), but will automatically downgrade to a different challenge type depending on the browser/user agent capabilities. ■ **Allow Basic credentials**: Prompts the user for a username and password to authenticate the user against the Windows Active Directory. **Note:** Basic credentials can't be disabled in the IWA realm if the IWA realm is part of a sequence realm but is not the first realm in the sequence with try IWA authentication only once enabled. ■ **Allow NTLM credentials**: Uses an encrypted challenge/response that includes a hash of the password. ■ **Allow Kerberos credentials**: Uses a ticket containing an encrypted session key in place of a user name and password. This is the most secure method because it establishes mutual authentication between the client and the server using an encrypted shared key. However, if you select this option, NTLM is automatically selected as well; in the event that the browser/user agent and/or the appliances are not configured properly for Kerberos, the appliance will automatically downgrade to NTLM. To use Kerberos, you must complete some additional configuration tasks. See "Enable Kerberos in a BCAAA Deployment " on page 247. **Note:** Forms authentication modes can't be used with an IWA realm that allows only NTLM/Kerberos credentials. If a form mode is in use and the authentication realm is an IWA realm, you will receive a configuration error.
Configure SSL options	**Enable SSL**	Select to enable SSL.
	Select Domain	The SSL device profile that the realm uses to make an SSL connection to the BCAAA server. Refer to the *SGOS Administration Guide* for details on using SSL device profiles.
Configure timeout	**Timeout request after**	Amount of time the appliance waits for an authentication response before timing out. Default value is 60 seconds.

Configure Custom Settings for an IWA Direct Realm

If you have joined your appliance to your Windows domain and created a realm for connecting to your Active Directory directly, you can configure and verify the authentication settings for connecting to Windows domain directly.

To configure custom settings for an IWA Direct realm:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > IWA** to add a new realm.
   - Select an existing realm to edit it.
3. (If adding a new realm) On the Add New IWA Realm page, in **Realm Config**, select **Custom Configuration**.
4. On the Add New /Edit IWA Realm page, in **Servers**, enter or edit the following as needed:

Task	Setting or field	Description
Specify credentials to use	**Client Authentication Mode**	Type of credentials to accept from the browser/user agent. By default, all credential types are allowed and the appliance will try to use Kerberos (the default authentication method for Windows clients), but will automatically downgrade to a different challenge type depending on the browser/user agent capabilities. ■ **Allow Basic credentials**: Prompts the user for a username and password to authenticate the user against the Windows Active Directory. Because the username and password are sent in plaintext, it is important to enable SSL if you allow Basic. **Note:** Basic credentials can't be disabled in the IWA realm if the IWA realm is part of a sequence realm but is not the first realm in the sequence with try IWA authentication only once enabled. ■ **Allow NTLM credentials**: Uses an encrypted challenge/response that includes a hash of the password. Because the plaintext username and password are not sent across the wire, this method is more secure than Basic authentication. ■ **Allow Kerberos credentials**: Uses a ticket containing an encrypted session key in place of a user name and password. This is the most secure method because it establishes mutual authentication between the client and the server using an encrypted shared key. However, if you select this option, NTLM is automatically selected as well; in the event that the browser/ user agent and/or the appliance are not configured properly for Kerberos, the appliance will automatically downgrade to NTLM. To use Kerberos, you must complete some additional configuration tasks. See "Enable Kerberos in an IWA Direct Deployment" on page 247. **Note:** Forms authentication modes can't be used with an IWA realm that allows only NTLM/Kerberos credentials. If a form mode is in use and the authentication realm is an IWA realm, you will receive a configuration error.
Configure timeout	**Timeout request after**	Amount of time the appliance waits for an authentication response before timing out. Default value is 60 seconds.

Task	Setting or field	Description
Verify the configuration	**Test Configuration**	Select to test the configuration. When prompted, enter the username and password of a user in the Windows domain and select **OK**. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm properly. If the test succeeds, the message also displays a list of any groups of interest (that is, groups that are referenced in policy) to which the user belongs.
Configure load balancing	**Set Credentials**	(Optional) Select if you are sharing an SPN across multiple appliances in a load balancing configuration. In the Enter Kerberos Credentials dialog, specify the **Username** and **Password** for the AD account.
	Clear Credentials	(If needed) Select to disable load balancing.

Join a Windows Domain

After you have synchronized the appliance's internal clock with the DC, you can join the appliance to one or more Windows domains.

To edit a domain, see "Edit a Windows Domain (NTLM Authentication Only)" on the facing page.

To join a Windows domain:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Domains**, beside Hostname, specify the hostname to use:
   - (Recommended) Select **Default** to use the default hostname.
   - Select **Custom** to enter a hostname.

   > **Note:** Unless you have a specific need to use a particular hostname (for example, to ensure correct DNS lookup), Symantec recommends that you use the default hostname to guarantee that each appliance's hostname is unique. In addition, you must use unique hostnames for multiple appliances joined to the same domain.

3. To add a domain, select **Add Domain**.

   On the Add Domain page, specify the following:

Section	Setting or field	Description
Domain Name	**Domain Name Alias**	Domain name alias
	DNS Domain Name	DNS name for the Windows Active Directory domain. This is not the fully qualified domain name of the appliance. **Note:** The appliance must be able to resolve the DNS domain name you supply for the Active Directory domain or the appliance will not be able to join the domain. If DNS resolution fails, check your DNS configuration.

Section	Setting or field	Description
Windows Domain Credentials and Status	**Username** and **Password**	User name and password for the primary Windows domain. You can either enter the plain user name (for example, administrator) or use the username@dnsname format (administrator@acme.com). This account must have rights for joining the domain.
	Join	Select to join the appliance to the domain. When joined, the domain State says "Joined" in the **Domains** list. Otherwise, the domain State says "Not joined". If the domain is used in an IWA realm, the State says "Joined & used". This option is not available until all of the previous settings are specified.

Edit a Windows Domain (NTLM Authentication Only)

You can optionally configure a MaxConcurrentApi value for each of the joined Windows domains on the appliance. The default value is 2 (same as the default for a Windows member server). The minimum value is 2 and the maximum value is 150.

You can also specify a preferred Schannel DC and alternate Schannel DC for each domain. If the preferred Schannel DC is available, the appliance will always connect to it, even if it sees another DC that appears to be faster. That serves two purposes:

- You need only increase Maximum number of concurrent Schannel connections on the preferred and alternate DCs, rather than on every DC in the domain.
- The preferred and alternate DCs can be read-only. Some customers are willing to deploy a read-only DC that is dedicated just to handling authentication requests for an appliance, whereas they would not be willing to deploy a regular “writable” DC for that purpose.

If the appliance can't connect to either of its preferred or alternate DCs, it connects to the fastest DC available. The appliance periodically checks to see if the preferred or alternate DC comes back online, and reconnects to it if it does.

To edit a Windows domain:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **NTLM Domain Controller**, specify the following:

Setting or field	Description
Preferred domain controller	Preferred DC.
Alternate domain controller:	Alternate DC, used if the preferred one is not available.

Setting or field	Description
Max number of Schannel connections	Maximum number of Schannel connections. The range is 2-150. In order for the maximum number of concurrent connections to take effect, you must enter the same number in the registry for the DC(s). The registry setting on the DC is MaxConcurrentAPI. If you change the MaxConcurrent API setting, you must restart the NetLogon service on the DC or reboot the DC after changing the MaxConcurrent API setting.

Notes on Preferred and Alternate DCs

If the appliance cannot connect to either of its preferred or alternate DCs, it connects to the fastest DC available. The appliance periodically checks to see if the preferred or alternate DC comes back online, and reconnects to it if it does.

The preferred and alternate DCs can be read-only. If you use a read-only DC, you must replicate user passwords to that domain controller. If the domain controller doesn't have a copy of the user's password, it must forward the request to a writable domain controller that has a copy, which will diminish performance. Consult Microsoft documentation to determine how to do this in your environment:

https://technet.microsoft.com/en-us/library/cc732801(v=ws.10).aspx

Authentication with Active Directory Using LDAP

Lightweight Directory Access Protocol (LDAP) is a client protocol used to access information stored in an LDAP-compatible directory service. It is the vehicle by which LDAP-enabled applications speak to one another. As a shared protocol, LDAP integrates compatible applications in your network to a single authentication interface. Any additions or changes made to information in the directory are available to authorized users, directory-enabled applications, devices, and ProxySG appliances. This central control gives administrators simplified application management.

LDAP authentication realms can be used to authenticate administrative users (read only and read/write) to the management console. To ensure that credentials are not sent in clear text, configure the LDAP realm to use TLS to secure the communication with the LDAP server.

> **Note:** To ensure that only TLS is used to communicate with the LDAP Server, check Enable SSL in the LDAP Server configuration page and edit the SSL device profile configured on the LDAP Server configuration page in the Admin Console.

About the Symantec LDAP Solution

The appliance uses existing directory-based authentication by passing log in requests to the directory service. By keeping authentication centralized on your directory, a security administrator will always know who is accessing network resources and can easily define user/group-based policies to control access through the appliance.

Symantec supports both LDAP v2 and LDAP v3, but recommends LDAP v3 because it supports additional authentication mechanisms. In LDAP v3, servers can return referrals to other servers back to the client, allowing the client to follow those referrals if desired.

Supported Directory Services

LDAP group-based authentication for the appliance can be configured to support any LDAP-compliant directory including:

- Microsoft Active Directory Server
- Novell NDS/eDirectory Server
- Netscape/Sun iPlanet Directory Server

How to Implement LDAP Authentication

Configuring the SGOS for LDAP authentication involves the following steps:

1. Create an LDAP realm on the appliance.
2. Configure LDAP properties on the appliance.
    a. Configure LDAP server settings
    b. Define LDAP Base Distinguished Names
    c. Define Authorization and Group information
    d. Define objectclass attributes on an LDAP entry
    e. Configure general LDAP realm settings
3. Create policies on the appliance.

See Also:

- "Configure an LDAP Authentication Realm " below

Configure an LDAP Authentication Realm

You can configure an LDAP realm in the Admin Console. See "Authentication with Active Directory Using LDAP" on the previous page for details.

An LDAP authentication realm authenticates and authorizes users to access services using either explicit proxy or transparent proxy mode. These realms integrate third-party vendors, such as LDAP, Windows, and Novell, with the Symantec operating system.

To configure an LDAP authentication realm:

Note: Realm creation requires knowledge of LDAP server type, server host information, and attribute type. This section describes realm configuration options and how to set up and add additional realms. For more information, see "Authentication with Active Directory Using LDAP" on page 260.

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > LDAP Realm** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add New /Edit LDAP Realm page, enter or edit the following as needed:

Section	Setting or field	Description
Properties	**Realm name**	Unique name for the realm.
	Server type	Specify the server type: ■ Microsoft Active Directory Server ■ Novell NDS/eDirectory Server ■ Netscape/Sun iPlanet Directory Server
	Primary server	Hostname or IP address of the primary server.
	Port	If you configured the primary server to use a port other than the default (389), specify it. If you enable SSL, change the port to an SSL listening port, such as port 636.
SSL	**Enable SSL**	(For LDAP v3 only) Select to enable SSL.
	Device Profile	Select the SSL device profile that this realm uses to make an SSL connection to a remote system.

Section	Setting or field	Description
LDAP Distinguished Name (DN)	**User attribute type**	This field displays the default user attribute type for the type of LDAP server you specified when creating the realm: ■ Microsoft Active Directory Servers: `sAMAccountName=` ■ Novell NDS/eDirectory Server/Other: `cn=` ■ Netscape/iPlanet Directory Server: `uid=` If you entered information correctly when creating the realm, you do not need to change the User attribute type in this step. If you do need to change or edit the entry, do so directly in the field. Add as many Base DNs as required for the realm. To specify entries for the Base DNs field, select **Add DN** and enter the Base DN(s). To search all of Sample_Company, enter `o` values. To search the manufacturing organizations, rather than starting at the top, enter `ou` and `o` values. You can add, edit, and delete Base DNs for an appliance to search. The appliance searches multiple DNs in the order listed, starting at the top and working down. To promote or demote a DN, select **Move up** or **Move down** in the Actions column.
Advanced Settings	**Advanced settings**	Select **Show** to specify custom settings: ■ "Configure LDAP Realm Details" on the next page ■ "Configure LDAP Servers " on page 267 ■ search ■ group ■ object classes
Verify the configuration	**Test Configuration**	Select to test the configuration. When prompted, enter the username and password of a user in the Windows domain and select **OK**. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm correctly. If the test succeeds, the message also displays a list of any groups of interest (that is, groups that are referenced in policy) to which the user belongs. You can also look up LDAP users and groups from the CLI using the lookup-user and lookup-group commands. Refer to the *Command Line Interface Reference* for details.

Configure LDAP Realm Details

You can configure the behavior of the authentication transaction, such as timeout and refresh intervals and cookie usage. You also configure the Virtual URL and group comparison method.

To configure LDAP realm details:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > LDAP** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add/Edit LDAP Realm page, in Advanced Settings, select **Show**.
4. In the sections that appear, enter or edit the following as needed:

Task	Setting or field	Description
Set display name	**Display name**	(Optional) By default, the appliance displays the authentication realm name when prompting the user for authentication credentials.

Task	Setting or field	Description
Configure refresh times for Basic credentials	**Use the same refresh time for all**	Select to use the same time for credential refresh and surrogate refresh.
	Credential refresh time	Specifies the amount of time the appliance will cache Basic credentials (username and password) and use these cached credentials to authenticate the user rather than sending another request to the authentication server. By default, basic credentials are good for 900 seconds (15 minutes). Before the refresh time expires, the appliance will authenticate the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.
	Surrogate refresh time	Specifies a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here. Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the appliance authenticates the transaction. After the refresh time expires, the appliance verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials. The main goal of this feature is to verify that the user-agent still has the appropriate credentials.
	Authorization refresh time	Specifies how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Task	Setting or field	Description
Configure timeouts	**Inactivity timeout**	When a client request is successfully authenticated, the appliance establishes an active session with the client and as long as that session stays active, the appliance will not attempt to reauthenticate requests from that client. This setting specifies how long the client session can be inactive before the appliance terminates the session; subsequent requests from that client will require authentication. By default, the client can be inactive for 900 seconds (15 minutes). **Note:** If the **Challenge user after logout** option is selected, the appliance will automatically challenge the client for credentials when the session becomes inactive. If you are using a challenge method that prompts the user for credentials, you might want to clear this option.
	Rejected credentials timeout	If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds. This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication service. The original failed authentication result is returned for the new request. All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down. To disable caching for failed authentication attempts, set to 0.
Configure cookie settings to specify how the appliance uses cookie surrogates Applicable if you plan to use an authentication mode that uses cookie surrogates.	**Use persistent cookies**	By default, this option is not selected, which means that the appliance will use session cookies when creating a cookie surrogate for a client. Session cookies are only valid during the current browser session and are deleted when the user closes the browser. The appliance must re-authenticate the client each time the user starts a new browser session. If you select this option, the appliance will use persistent cookies instead of session cookies. Persistent cookies are stored on the client system and are therefore not deleted at the end of the browser session. When using persistent cookies, the appliance will only need to re-authenticate a client when the cookie in its surrogate credential database expires.
	Verify the IP address in the cookie	By default, this option is selected, which means that the appliance will only accept a cookie from a client if the client IP matches the IP address in the surrogate cookie. To enable the appliance to accept cookies from IP addresses that do not match the address in the cookie–for example, if you use DHCP–clear this option.

Task	Setting or field	Description
Set virtual URL	**Virtual URL**	(Transparent proxy only) Specify the URL to which to redirect client requests that require authentication in the Virtual URL field. For best results, the virtual URL you specify must: ■ Contain a simple hostname that does not contain any dots (for example, use http://myproxy rather than http://myproxy.acme.com. ■ Resolve to the IP address of the appliance. To accomplish this, you must add an "A" record to your internal DNS server that associates the Virtual URL with the IP address of the appliance.
Configure re-authentication after timeout	**Challenge user after logout**	(Optional) Select to prompt the client for authentication credentials whenever the inactivity timeout expires.
Group comparison search method	**Group comparison**	Select the search method: ■ Local: The local method performs compare operations on the appliance after retrieving the appropriate entries. Because the compares are performed locally, this method typically reduces load on the LDAP server. ■ Server: The server method queries the LDAP server for each compare operation. If there are a large number of compares to perform, it can result in significant server load. There is a minute possibility that local compares can produce differing results from server compares. If you suspect erroneous compare results, set to server.

Configure LDAP Servers

After an LDAP authentication realm is created, you must set LDAP realm properties according to your directory type.

To configure LDAP servers:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > LDAP** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add/Edit LDAP Realm page, in Advanced Settings, select **Show**.
4. In **Servers**, enter or edit the following as needed:

Setting or field	Description
LDAP protocol version	Select LDAP v3 or v2.
Follow referrals	(LDAVP v3 only) Allow the client to follow referrals to other servers. (This feature is not available with LDAP v2.) The default is Disabled.
Alternate server	Specify the host and port for the alternate LDAP server.
Timeout request after	Amount of time the appliance waits for an authentication response before timing out. Default value is 60 seconds.
Case sensitive	Select if the LDAP server is configured to expect case-sensitive usernames and passwords.

Define LDAP Authorization Properties

After creating an LDAP realm, providing at least the required fields of the LDAP server for that realm, and defining base DNs for the realm, you must define authorization properties for each LDAP realm you created.

Search Properties

You can enable the following types of LDAP searches:

- Anonymous searches, which allows a user to perform an LDAP search without entering a distinguished name.

  The LDAP directory attributes available for an anonymous client are typically a subset of those available when a valid user distinguished name and password have been used as search credentials.

- Authenticated searches, which require a search user DN to function properly.

  For Microsoft Active Directory, you must use the full name and not the login name.

Group Properties

Authorization decisions are completely handled by policy. The groups that the appliance looks up and queries are derived from the groups specified in policy in `group=`, `attribute=`, `ldap.attribute=`, and `has_attribute=-` conditions. If you do not have any of those conditions, then Symantec does not look up any groups or attributes to make policy decisions based on authorization.

`Objectclass` Attribute Values

The *`objectclass`* attributes on an LDAP object define the type of object an entry is. For example, a user entry might have an attribute value of person while a group entry might have an objectclass attribute value of group.

The *`objectclass`* attribute values defined on a particular entry can differ among LDAP servers. The attribute values are not DNs of any kind.

Currently, the attribute values are used by Symantec during a VPM browse of an LDAP server. If an administrator wants to browse the groups in a particular realm, the appliance searches the LDAP server for objects that have objectclass attribute values matching those in the group list and in the container list. The list of objectclass attribute values in the container list is needed so that containers that contain groups can be fetched and expanded correctly.

Specify LDAP Authorization Properties

Define the search, group, and *objectclass* properties:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > LDAP Realm** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add New /Edit LDAP Realm page, enter or edit the following as needed:

Section	Setting or field	Description
Search	**Search**	Specify the type of search: ■ **Only authorized** To enforce user authentication before binding to the LDAP service, specify a user distinguished name in the **Search User DN** field. This username can identify a single user or a user object that acts as a proxy for multiple users (a pool of administrators, for example). A search user distinguished name can be up to 512 characters long. You can set a **Password** up to 64 alphanumeric characters long. **Note:** You might want to create a separate user (such as Symantec, for example) instead of using an Administrator distinguished name and password. ■ **Anonymous search allowed** For example, with Netscape/iPlanet Directory Server, when anonymous access is allowed, no username or password is required by the LDAP client to retrieve information. Some directories require a valid user to be able to perform an LDAP search; they do not allow anonymous bind. (Active Directory is one such example.) For these directories, you must specify a valid fully-qualified distinguished username and the password that permits directory access privileges. (For example, cn=user1,cn=users,dc=symantec,dc=com is a possible fully-qualified distinguished name.)
	Dereference aliases	Specify when to search for a specific object rather than search for the object's alias: ■ **Always** (default) ■ **Searching** ■ **Finding** ■ **Never**

Section	Setting or field	Description
Groups	**Membership type** and **Membership attribute**	The appliance enters defaults for the following LDAP directories: ■ Microsoft Active Directory Membership type: user Membership attribute type: memberOf ■ Netscape/Sun iPlanet Membership attribute type: unique Member Membership type: group ■ Novell NDS eDirectory Membership type: group Membership attribute type: member ■ Other Membership type: user Membership attribute type: member
	Username type to lookup	(If membership type is **Group**) Select either **FQDN** or **Relative**. FQDN indicates that the lookup is done only on the user object. If membership type is **User**, FQDN is selected by default.
	Nested groups	Select if the LDAP server you use does not natively support group membership tests of nested groups. When a group of interest referenced within policy is part of a loop, User Authorization results in Access Denied(policy_denied). For example, a loop forms if the group member Testgroup has the nested group member Testgroup2, which in turn has the aforementioned Testgroup as a nested member. When loops are removed from an LDAP server, the Nested Groups Support option must be disabled and then re-enabled for the appliance to re-fetch the correct group structure.Nested group attribute: For other, ad and nds, the default attribute is member. For iPlanet, the attribute is uniqueMember.

Section	Setting or field	Description
	Group constraint filter	(Available if **Nested groups** is enabled) Enter a search limiting clause to reduce the number of groups returned for an LDAP search. This feature is generally used only when the user wishes to limit the scope of a comparison due to a very large number of groups. Constraints must be valid LDAP search filters and are ANDed to the search filter when performing a group search. Example 1: If you enter (cn=p*) into the Group constraint filter field, only groups starting with the letter P are returned. Example 2: If you enter (cn=proxy) into the Group constraint filter field, only the proxy group is returned. **Note:** The Group constraint filter functions only for local comparisons. To enable local group comparisons, see "Configure LDAP Realm Details" on page 264.
Object Classes	**Object type**	Select **container**, **group**, or **user** to display the defined object classes for the type. You can edit an object class by updating its name and applying your changes.
	Add Object class	Select to add a new class to the selected type.

RADIUS Realm Authentication and Authorization

RADIUS is often the protocol of choice for ISPs or enterprises with very large numbers of users. RADIUS is designed to handle these large numbers through centralized user administration that eases the repetitive tasks of adding and deleting users and their authentication information. RADIUS also inherently provides some protection against sniffing.

Some RADIUS servers support one-time passwords. One-time passwords are passwords that become invalid as soon as they are used. The passwords are often generated by a token or program, although pre-printed lists are also used. Using one-time passwords ensures that the password cannot be used in a replay attack.

The ProxySG appliance's one-time password support works with products such as Secure Computing SafeWord synchronous and asynchronous tokens and RSA SecurID tokens.

The appliance supports RADIUS servers that use challenge/response as part of the authentication process. SafeWord asynchronous tokens use challenge/ response to provide authentication. SecurID tokens use challenge/response to initialize or change PINs.

The challenge is displayed as the realm information in the authentication dialog; Symantec recommends that you use form authentication if you create a challenge/response realm, particularly if you use SecurID tokens.

If you set an authentication mode that uses forms, the system detects what type of question is being asked. If it is a yes/no question, it displays the query form with a yes and no button. If it is a new PIN question, the system displays a form with entry fields for the new PIN.

Using policy, you can fine-tune RADIUS realms based on RADIUS attributes. If you use the Symantec attribute, groups are supported within a RADIUS realm. RADIUS authentication is compatible with administrative authentication for the appliance management console.

See Also:

- "Configure a RADIUS Authentication Realm " below

Configure a RADIUS Authentication Realm

To configure a RADIUS authentication realm:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > RADIUS** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add New/Edit IWA Realm page, enter or edit the following as needed:

Setting or field	Description
Realm name	Unique name for the realm. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name must start with a letter.
Primary Server	Specify the IP address for the primary RADIUS server.
Port	Specify the port for the primary RADIUS server. The default port is 1812.
Secret	Specify the RADIUS secret. RADIUS secrets can be up to 64 characters long and are always case-sensitive.
Advanced Settings	Select **Show** to specify custom settings: ■ "Configure RADIUS Servers " on page 276 ■ "Configure RADIUS Realm Details" on the next page

Setting or field	Description
Test Configuration	(Only available when editing) Select to test the configuration. When prompted, enter the username and password of a user in the Windows domain and select **OK**. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm properly. If the test succeeds, the message also displays a list of any groups of interest (that is, groups that are referenced in policy) to which the user belongs.

Configure RADIUS Realm Details

Once you have created the RADIUS realm, you can change the primary host, port, and secret of the RADIUS server for that realm.

To configure RADIUS realm details:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > IWA** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add/Edit IWA Realm page, in Advanced Settings, select **Show**.
4. In the sections that appear, enter or edit the following as needed:

Task	Setting or field	Description
Set display name	**Display name**	(Optional) By default, the appliance displays the authentication realm name when prompting the user for authentication credentials.

Task	Setting or field	Description
Configure refresh times for Basic credentials	**Use the same refresh time for all**	Select to use the same time for credential refresh and surrogate refresh.
	Credential refresh time	Specifies the amount of time the appliance will cache Basic credentials (username and password) and use these cached credentials to authenticate the user rather than sending another request to the authentication server. By default, basic credentials are good for 900 seconds (15 minutes). Before the refresh time expires, the appliance authenticates the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.
	Surrogate refresh time	(Unavailable if set to use the same refresh time for all) After the appliance successfully authenticates a client, it caches the client's IP address or a cookie (depending on the authentication mode that is in use) in its surrogate cache. If it receives subsequent requests from the same client during the surrogate refresh time, it uses the IP address or cookie in its cache to authenticate the user instead of sending a request to the authentication server. By default, the surrogate credential is good for 900 seconds (15 minutes).
Configure timeouts	**Inactivity timeout**	When a client request is successfully authenticated, the appliance establishes an active session with the client and as long as that session stays active, the appliance will not attempt to reauthenticate requests from that client. This setting specifies how long the client session can be inactive before the appliance terminates the session; subsequent requests from that client will require authentication. By default, the client can be inactive for 900 seconds (15 minutes). **Note:** If the **Challenge user after logout** option is selected, the appliance will automatically challenge the client for credentials when the session becomes inactive. If you are using a challenge method that prompts the user for credentials, you might want to clear this option.
	Rejected credentials timeout	Specifies whether to cache failed authentication attempts (bad password, expired account, disabled account, old password, or server down). If the client attempts to connect again during the rejected credentials time, the appliance will automatically reject the request for the specified period of time. Enter a value from 1 second (the default) to 10 seconds. To disable this option, enter 0.

Task	Setting or field	Description
Configure cookie settings to specify how the appliance uses cookie surrogates Applicable if you plan to use an authentication mode that uses cookie surrogates.	**Use persistent cookies**	By default, this option is not selected, which means that the appliance will use session cookies when creating a cookie surrogate for a client. Session cookies are only valid during the current browser session and are deleted when the user closes the browser. The appliance must re-authenticate the client each time the user starts a new browser session. If you select this option, the appliance will use persistent cookies instead of session cookies. Persistent cookies are stored on the client system and are therefore not deleted at the end of the browser session. When using persistent cookies, the appliance will only need to re-authenticate a client when the cookie in its surrogate credential database expires.
	Verify the IP address in the cookie	By default, this option is selected, which means that the appliance will only accept a cookie from a client if the client IP matches the IP address in the surrogate cookie. To enable the appliance to accept cookies from IP addresses that do not match the address in the cookie–for example, if you use DHCP–clear this option.
Set virtual URL	**Virtual URL**	(Transparent proxy only) Specify the URL to which to redirect client requests that require authentication in the Virtual URL field. For best results, the virtual URL you specify must: ■ Contain a simple hostname that does not contain any dots (for example, use http://myproxy rather than http://myproxy.acme.com. ■ Resolve to the IP address of the appliance. To accomplish this, you must add an "A" record to your internal DNS server that associates the Virtual URL with the IP address of the appliance.
Configure re-authentication after timeout	**Challenge user after logout**	(Optional) Select to prompt the client for authentication credentials whenever the inactivity timeout expires.

Configure RADIUS Servers

After creating a RADIUS authentication realm, you can change the primary host, port, and secret of the RADIUS server for the realm.

To configure RADIUS servers:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > RADIUS** to add a new realm.
   - Select an existing realm to edit it.

3. On the Add/Edit RADIUS Realm page, in Advanced Settings, select **Show**.
4. In **Servers**, enter or edit the following as needed:

Setting or field	Description
Alternate server	Specify the host and port for the alternate RADIUS server.
Secret	Specify the RADIUS secret. RADIUS secrets can be up to 64 characters long and are always case-sensitive.
Encode credentials with character set	Select the character set used for encoding credentials; the RADIUS server needs the same character set. A character set is a Multipurpose Internet Mail Extension (MIME) charset name. Any of the standard charset names for encodings commonly supported by Web browsers can be used. The default is Unicode:UTF8.
Timeout request after	Enter the amount of time the appliance waits for an authentication response before timing out. Default value is 10 seconds.
Retry	Specify the number of attempts you want to permit before marking a server offline. The client maintains an average response time from the server; the retry interval is initially twice the average. If that retry packet fails, then the next packet waits twice as long again. This increases until it reaches the timeout value. The default number of retries is 10.
One-time passwords	Enable if you created a challenge/response realm.
Case sensitive	Select if the RADIUS server is configured to expect case-sensitive usernames and passwords.

SAML Realm Authentication and Authorization

SAML 2.0 was developed by the OASIS Security Services Technical Committee.

It is an industry standard for retrieving authorization and identity information in XML documents to facilitate single sign-on (SSO) applications or services on the Internet. In SAML authentication, the exchange of information is performed by the following entities:

- Identity providers (IDPs), which are identity stores. For example, an IDP may have a back-end directory of users. The IDP authenticates the users. Supported IDPs are listed in "Requirements for SAML Authentication" on page 279.
- Service providers (SPs), which provide access to applications or services to users. It is the entity that creates an authenticated session for the user.

With SAML authentication, the ProxySG appliance acts as the service provider. SAML realms are not compatible with administrative authentication to the appliance management console.

> **Note:** This document assumes that you are familiar with SAML concepts and practices.

About SAML

The following sections provide conceptual information you must understand before configuring SAML:

Federation and Metadata

The entities (IDP and SP) must federate before authentication can occur. During federation, configuration data is exchanged in metadata files. Each entity publishes information about itself in these files and publishes them to a specific location, for example, on the internet or a network drive. When the entities share metadata, they establish and agree on the parameters that they will use for authentication requests and responses. They also share information such as:

- Entity IDs, which entities use to identify themselves to each other. For example, the Entity ID tells the IDP if an authentication request comes from a federated relying party.
- The SSO POST endpoint and SSO redirect endpoint to which entities send assertions during authentication. (See "Assertions" below for more information.)
- Each entity's public key certificate, which is used to validate assertions.
- Whether each entity requires encryption. If one of the entities requires encrypted assertions, it will publish a separate encryption certificate.

> **Note:** The appliance may consume encrypted assertions from the IDP, but it does not encrypt authentication requests that it sends to the IDP.

Assertions

The ProxySG appliance and the IDP exchange data in XML documents called assertions. After a user is authenticated, the IDP sends an authentication assertion and the appliance establishes an authenticated session with the appropriate authorization for the user.

The appliance processes SAML authentication responses from the IDP; these responses may contain assertion attributes that describe the authenticated user.

For example, <saml:Attribute Name="mail"> is an assertion attribute that contains the user's email address inside the <saml:AttributeValue> element.

You can configure the appliance to use assertion attributes in authorization decisions. For more information on attributes, see "Configure SAML Attributes" on page 287.

Profiles and Bindings

A profile contains information about how SAML supports a defined use case. For example, theWeb Browser SSO Profile enables single sign-on authentication for resources on the internet.

SAML 2.0 includes protocol-specific bindings, which describe how SAML data is exchanged over those protocols. SAML authentication supports the following for the Web Browser SSO Profile:

- The HTTP POST binding for authentication responses
- The HTTP POST binding and the redirect binding for authentication requests

Requirements for SAML Authentication

Setting up a SAML realm for the ProxySG appliance requires one of the following IDPs:

- Microsoft® Active Directory Federation Services (AD FS) 2.0

> **Note:** ADFS 1.0 ships with Windows Server 2008. If you want to use the SAML realm with AD FS, you must download AD FS 2.0 from the Microsoft website and install it.

- CA SiteMinder® Federation Partnership R12
- Oracle® Identity Federation 11g
- Shibboleth 2.3.5

Checklist: Preparing the IDP

Before you set up a SAML realm, make sure that you have done the following for your IDP:

- Installed and configured the administration software.
- Set up the identity store for authentication.
- Identified the default user attribute to be passed in SAML assertions, for example, the User Principal Name attribute in LDAP.
- Identified any additional attributes that you want to be passed in assertions, for example, the memberOf attribute, which identifies the groups of which a user is a direct member in LDAP.
- Determined the location (URL) of the IDP's metadata file. This is needed to complete the steps in "Configure a SAML Realm " on page 286. If you import metadata, the realm uses its preconfigured settings. See "SAML Realm Authentication and Authorization" on page 277 for instructions on locating the metadata file for the IDP.

> **Note:** To import SiteMinder and Oracle metadata, use the `#(config saml <realm_name>)inline idp-metadata <XML>` CLI command. To avoid errors, import metadata through the CLI instead of entering the information manually in the Management Console.

SAML Authentication Process

After you have defined and configured a SAML realm, and both entities have federated and exchanged metadata, authentication can occur.

The following example is an overview of what happens when a user goes to a website that requires authentication.

Step 1: Initial user request

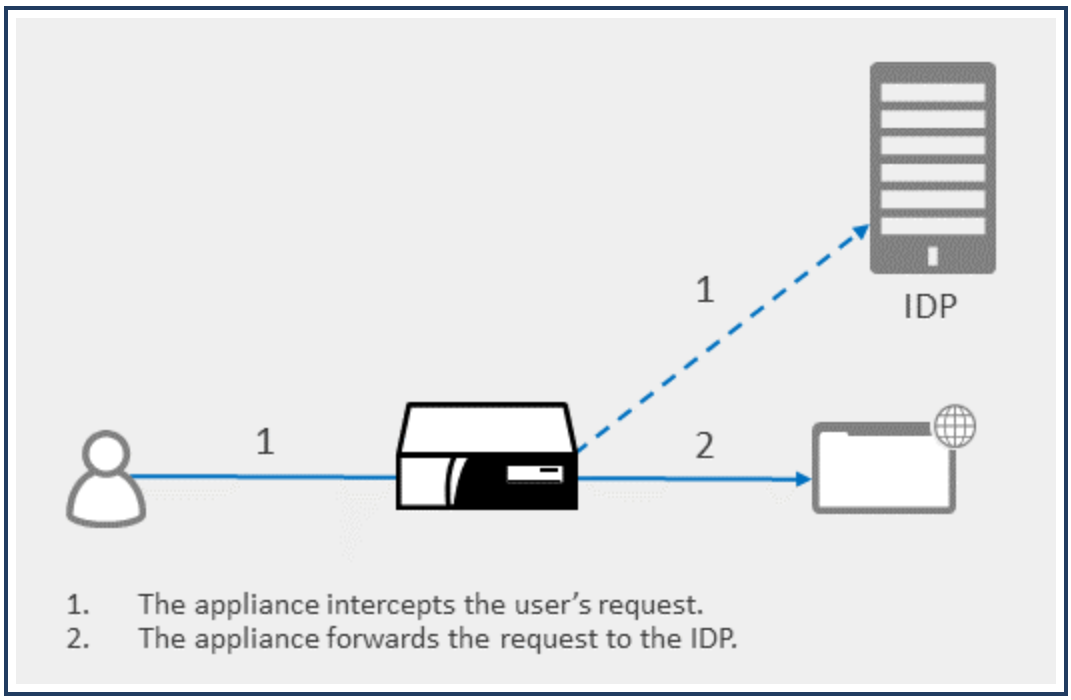

The appliance intercepts the user's request and redirects the web browser to the IDP. The redirect URL includes the SAML authentication request that should be submitted to the IDP's SSO service. If the **Disable Client Redirects** check box is checked, the appliance does not redirect the client to the IDP.

Step 2: Authentication request and response

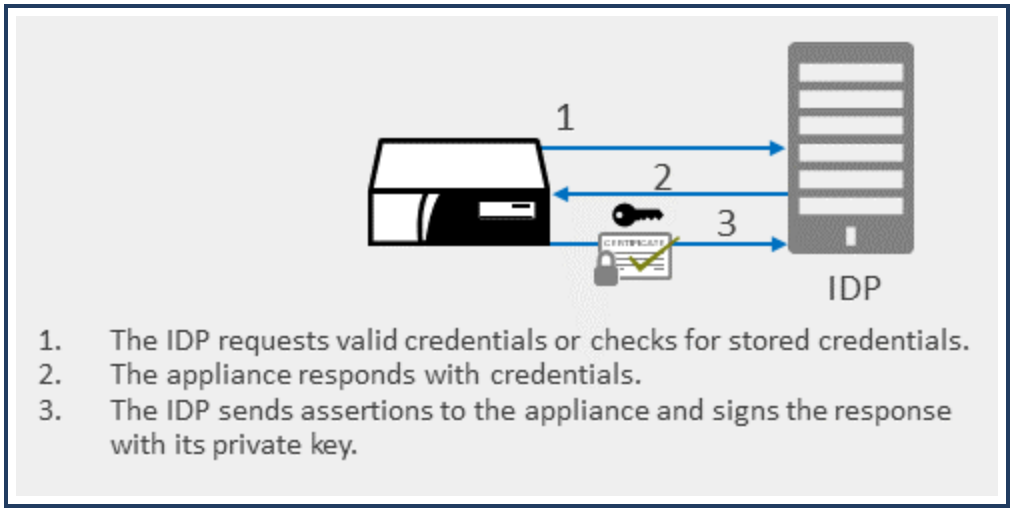

The IDP asks the appliance for the user's credentials, for example by asking for valid login credentials or checking for valid session cookies for stored credentials. If the appliance responds with valid credentials, the IDP:

- Signs an authentication response with its private signing key. If the IDP has been configured to send encrypted assertions, the IDP encrypts the assertion before sending it to the appliance.
- Sends the authentication response to the appliance, which contains the user's username (however, the appliance is not aware of the user's credentials).

Step 3: Assertion decryption and validation

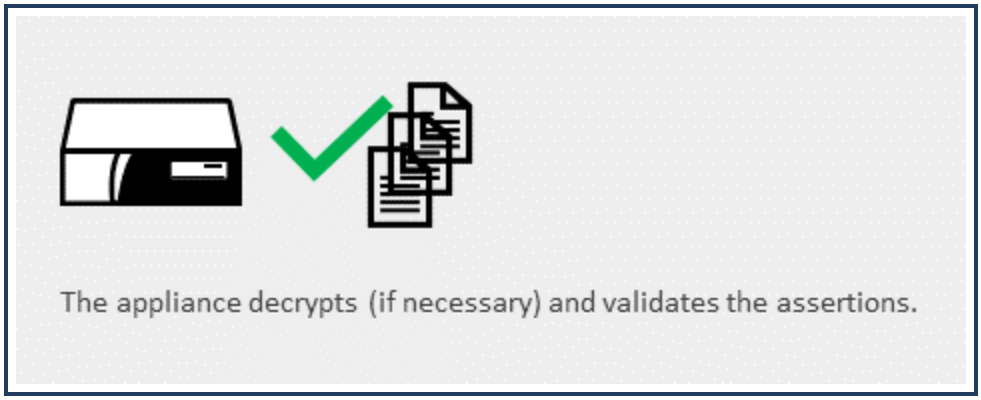

If the assertion is encrypted, the appliance decrypts it. The appliance rejects any unsigned assertions.

The appliance validates the assertion and then retrieves the user's name and group memberships (as specified in assertion attributes) from the assertion using the private key.

Step 4: User request validation

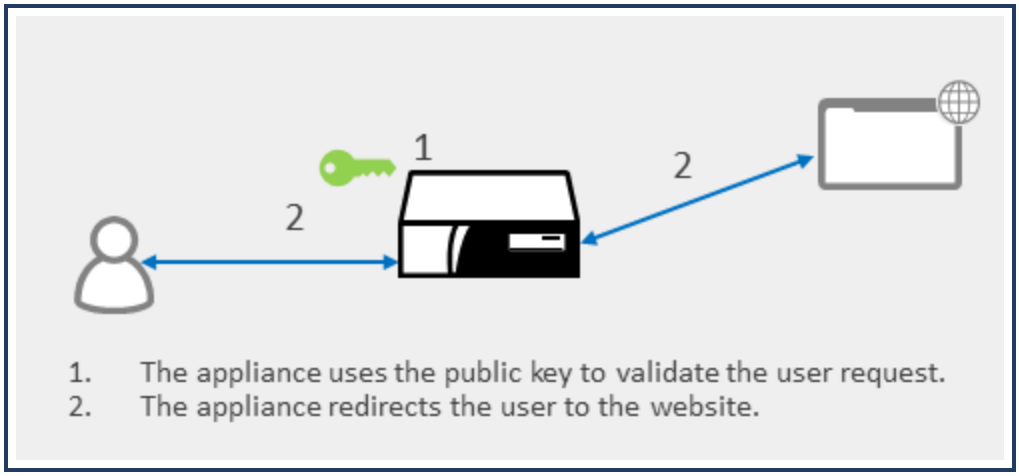

The appliance validates the user request using the corresponding public key, which is embedded in the IDP's signing certificate. Then, the appliance redirects the user to the website and creates an authenticated session for the user.

Overview: SAML Realm Configuration

Perform the following steps to set up a SAML realm.

Task #	Task	Reference(s)
1	Save a local copy of the IDP's configuration data.	"Export the IDP Metadata File" on the facing page
2	Prepare the appliance for SAML authentication: ■ Import certificates to the CCL. ■ Set up an HTTPS reverse proxy service. **Note:** This is required only if the SAML realm is using an HTTPS POST endpoint. ■ (If needed) Define assertion attributes. ■ Configure the SAML realm.	"Prepare the Appliance" on page 285 "Configure a SAML Realm " on page 286 "Configure SAML Attributes" on page 287 "Configure SAML Realm Details" on page 289
3	(Optional) To authorize users through one or more SAML realms, specify the criteria to use when searching for users.	"Configure SAML Authorization" on page 292

Task #	Task	Reference(s)
4	Configure your specific IDP.	Refer to the appropriate section in the *SGOS Administration Guide*. ■ "Configure AD FS" ■ "Configure SiteMinder" ■ "Configure Oracle" ■ "Configure Shibboleth"
5	Include the SAML realm in your policy.	Refer to "Add the SAML Realm to Policy" in the *SGOS Administration Guide*.

See Also:

- *SGOS Administration Guide*

Export the IDP Metadata File

You can export the IDP metadata file and import it via local file or remote URL to the SAML realm configuration.

To export the IDP metadata file, log in to the IDP's administration software.

Exporting IDP metadata entails saving the XML document to disk. It is important to save the metdata file without opening it in a browser first. Browsers do not necessarily support XML file structure and may change the XML tags.

If you use SiteMinder, Oracle, or Shibboleth, you will need to copy and paste the metadata file contents to the CLI using the `inline idp-metadata` command.

Because XML files are text-based, it is best to use a text editor such as Notepad to open the file to copy its contents.

> **Note:** To ensure that the SAML realm is configured correctly, import metadata instead of entering the information manually. If there are issues with realm configuration, the Authentication debug log shows the following error: `The SAML realm configuration is invalid`.

Export Metadata from AD FS

To export metadata from AD FS:

1. Log in to the AD FS MMC.
2. Select **Endpoints** and look under Metadata for the URL beside the Federation Metadata type.

3. Copy the URL and paste it into a browser address bar.
4. Save the XML file to a location that the appliance can access.

Export Metadata from SiteMinder

Before you can export metadata, make sure that you have created a SAML 2.0 IDP. This assumes that you have already created the IDP (entity) in SiteMinder.

To export metadata from SiteMinder:

1. Log in to the CA Federation Manager.
2. Select **Federation > Entities**.
3. Beside the entity you created, select **Action > Export Metadata**.
4. In the Partnership Name field, enter a name to identify the partnership between the appliance and SiteMinder. You will refer to this partnership name later, when you configure the partnership in SiteMinder.
5. Select **Export**. SiteMinder generates the metadata document.
6. Save the XML file to a location that the appliance can access.

Export Metadata from Oracle

To export metadata from Oracle:

1. Log in to the Oracle Enterprise Manager.
2. In the navigation tree on the left, select **Identity and Access > OIF**.
3. On the main page, select **Oracle Identity Federation > Administration > Security and Trust**.
4. Select the **Provider Metadata** tab.
5. In the Generate Metadata section, select **Identity Provider** from the Provider Type menu.
6. Select **SAML 2.0** from the Protocol menu.
7. Select **Generate**. OIF generates the metadata document.
8. Save the XML file to a location that the appliance can access.

Export Metadata from Shibboleth

To export data from Shibboleth:

1. On the server where Shibboleth is installed, browse to the <shibboleth>/metadata folder.
2. Copy the idp-metadata.xml file.
3. Save the XML file to a location that the appliance can access.

See Also:

- "Configure a SAML Realm " on the next page

Prepare the Appliance

Complete the following steps to prepare the appliance for SAML authentication.

Configure the CCL

The appliance CCL must contain at least a root certification authority (CA) certificate, but depending on other considerations, you may require more certificates. Refer to the following list to determine which certificates you must import to the CCL.

- Root CA certificate: Required; add the certificate for the root CA that issued the IDP's signing certificate to the CCL.
- IDP's signing certificate: Required if self-signed; if the IDP's signing certificate is self-signed, add it to the CCL. Certificates signed by the CA are included in SAML assertions.
- Intermediate CA certificate: Optional; you must import intermediate CA certificates to the appliance, but it is not necessary to add them to the CCL. For instructions on importing certificates to the appliance, see "Import Keydata" on page 214.

> **Note:** In explicit deployments, if you do not add the certificate for the CA that issued the IDP's certificate to the appliance's CCL, HTTPS connections to the IDP fail.

Create an HTTPS Reverse Proxy Service

You should create an HTTPS reverse proxy service only if the SAML realm uses an HTTPS, rather than HTTP, POST endpoint.

Determine whether to use an HTTP or HTTPS POST endpoint:

- If you use AD FS, the POST endpoint must use HTTPS. If you use SiteMinder or Oracle, the POST endpoint can use either HTTPS or HTTP.
- For better security, use an HTTPS POST endpoint.

> **Note:** Regardless of which IDP you use, use an HTTPS POST for best security.

Create an HTTPS reverse proxy service to act as the SAML realm's HTTPS POST endpoint. The IDP redirects browsers to this service when it creates assertions. SSL connections require a certificate and a private key. Browsers must trust the certificate that the HTTPS reverse proxy service uses, and the certificate's Subject value must match the **Virtual URL**

configured in the SAML realm (see "Configure a SAML Realm " below). If the names do not match, SSL hostname mismatch errors occur.

To create the HTTPS reverse proxy service, see "Configure an HTTPS Reverse Proxy" on page 161.

See Also:

- "Configure a SAML Realm " below

Configure a SAML Realm

Before proceeding, make sure you have saved a local copy of the IDP's configuration data (see "SAML Realm Authentication and Authorization" on page 277) and prepared the ProxySG appliance (see "Prepare the Appliance" on the previous page).

To configure a SAML authentication realm:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > SAML** to add a new realm.
   - Select an existing realm to edit it.
3. On the Add New/Edit SAML Realm page, enter or edit the following as needed:

Setting or field	Description
Realm name	Unique name for the realm.
Virtual URL	Specify the hostname for the SAML endpoint; in other words, point to the HTTPS reverse proxy listener you set up. Enter the host and port in format https://<*hostname_or_IP_address*>:<*port_number*>. The hostname must match the name of the SSL certificate for the HTTPS reverse proxy service. See "Create an HTTPS Reverse Proxy Service" on the previous page.

Setting or field	Description
Federated IDP Import IDP metadata or configure the IDP manually	You can import IDP metadata from a file saved locally or on a remote server. Complete "Export the IDP Metadata File" on page 283 before proceeding. To import the IDP metadata file, save the current SAML realm first. Then, edit the realm and select **Import metadata from text**. ■ If you exported the IDP metadata file to the local computer, select **Local File** and browse to the file ■ If you exported the IDP metadata file to an FTP or HTTP server, select **Remote URL** and enter the URL. **Note:** Importing metadata will overwrite any manually-entered IDP values. To configure the IDP manually, specify the following settings: ■ **Entity ID**: SAML entity ID. ■ **POST URL**: SSO POST endpoint. ■ **Redirect URL**: SSO redirect endpoint. ■ **CCL**: Select the trusted CCL created in "Configure the CCL" on page 285. ■ **Allowed Signing ECL**: (Optional) Select a configured external certificate list. To create external certificate lists, use the Management Console. Refer to "Creating an External Certificate List" in the *SGOS Administration Guide*.
Attributes and **Attributes Definitions**	See "Configure SAML Attributes" below.
Advanced Settings	Select **Show** to specify realm details and authorization settings: ■ "Configure SAML Realm Details" on page 289 ■ "Configure SAML Authorization" on page 292

Configure SAML Attributes

The appliance maps policy conditions to assertion attribute values. If you require more attributes than the ones included in SAML assertions, you can define them in the SAML realm.

You can forward SAML assertion attributes via custom headers to back-end or front-end servers using the **SAML Attribute** VPM object. Refer to the *Visual Policy Manager Reference* for instructions on using this object in policy.

To configure SAML attributes:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > SAML** to add a new realm.
   - Select an existing realm to edit it.

   The console displays the Add New/Edit SAML Realm page.
3. In the Attribute Definitions section, select **Add Attribute**. The console displays the Add Attribute dialog.
4. Enter or modify your own attribute definitions:
   - **Attribute Name**: The name of the attribute as it appears in the appliance and IDP configuration, and when referring to the attribute in the `attribute.<name>=` policy condition. The name must be unique.
   - **Attribute Data Type**: Select **String (Ignore Case)** or **Case (Exact Case)**. The appliance uses this setting to match assertion attribute values with policy conditions.
   - **SAML Name**: The name of the attribute as it will appear in assertions from the IDP, in the Name= XML attribute of the <Attribute> element. For example, an assertion might include the line <saml:Attribute Name="*mail*"> where *mail* is the SAML attribute name.

   Select **Add** or **Update** to save the attribute definition.
5. In the Attributes section, specify the attribute definition to use for the following attributes:
   - **User Attribute**: Attribute containing the relative username. If you do not specify the attribute, the appliance uses the SAML Name ID value for the username.
   - **Fullname Attribute**:Attribute containing the full username.
   - **Group Attribute**: Name of the group membership attribute. Values of this attribute match the `group=` policy condition.
   - **Email Address Attribute**: Attribute that represents the user's email address. For more information, refer to KB article 169260.

See Also:

- "Configure a SAML Realm " on page 286

SAML Attribute Substitutions

You can forward SAML attributes, including the user name, in SAML 2.0 assertions through HTTP headers to back-end servers by creating substitution policy. The SAML substitution uses the following CPL syntax:

`$(saml.attribute.<name>)`

To forward SAML attributes using substitution policy:

1. In the SAML realm configured on the appliance, define the attribute you want to forward.

   > **Note:** If you remove an attribute that is referenced in policy, the appliance issues a warning the next time that policy is compiled.

2. In the Web VPM, add a **Web Access** layer.
3. In the Source column, add a **SAML Attribute** object that specifies the attribute to forward.

   > **Note:** If the attribute you specify does not exist in the SAML realm, the appliance issues a warning when it compiles policy.

4. In the Action column, add a **Control Request Header** or **Control Response Header** object. Select **Set Value** and enter the substitution:

   `$(saml.attribute.<name>)`

   where <*name*> is the name of the attribute to forward.

5. Install the policy.
6. After you install the substitution policy, the appliance evaluates the configured policy rules and modifies the HTTP request/response headers with the additional x-header fields corresponding to attribute names with the attribute values. When the appliance compiles policy, it validates the field attribute <*name*> against configured SAML attributes.

   > **Note:** You can alternatively add the substitution policy through the **`#inline policy`** command or CPL. For information on those methods, refer to the *Command Line Interface Reference* and the *Content Policy Language Reference*, respectively.

Configure SAML Realm Details

To configure SAML realm details:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > SAML** to add a new realm.
   - Select an existing realm to edit it.

3. On the Add/Edit SAML Realm page, in Advanced Settings, select **Show**.
4. In the sections that appear, enter or edit the following as needed:

Task	Setting or field	Description
Set display name	**Display Name**	(Optional) By default, the appliance displays the authentication realm name when prompting the user for authentication credentials.
Configure refresh times for Basic credentials	**Use the Same Refresh Time For All**	Select to use the same time for credential refresh and surrogate refresh (in seconds). Otherwise, specify different refresh times for credentials and surrogates.
Configure timeouts	**Inactivity Timeout**	When a client request is successfully authenticated, the appliance establishes an active session with the client and as long as that session stays active, the appliance will not attempt to reauthenticate requests from that client. This setting specifies how long the client session can be inactive before the appliance terminates the session; subsequent requests from that client will require authentication. By default, the client can be inactive for 900 seconds (15 minutes). **Note:** If this option is selected, the appliance wil l automatically challenge the client for credentia ls when the session becomes inactive. If you ar e using a challenge method that prompts the us er for credentials, you might want to clear this o ption.
	Rejected Credentials Timeout	Specifies whether to cache failed authentication attempts (bad password, expired account, disabled account, old password, or server down). If the client attempts to connect again during the rejected credentials time, the appliance will automatically reject the request for the specified period of time. Enter a value from 1 second (the default) to 10 seconds. To disable this option, enter 0.

Task	Setting or field	Description
Configure cookie settings to specify how the appliance uses cookie surrogates Applicable if you plan to use an authentication mode that uses cookie surrogates.	**Use Persistent Cookies**	Select to use persistent cookies.
	Verify the IP Address in the Cookie	Select to enable verification of cookies' IP addresses.
Configure re-authentication after timeout	**Challenge User After Logout**	(Optional) Select to enable challenging after logout. For example, if this setting is enabled and a user logs out of a web site, the user must enter credentials again the next time they access the web site.
Define limits for assertions' timestamps. Assertions with timestamps that fall outside of these limits are invalid.	**Not Before**	(Optional) Specify an interval before the current time. Assertions stamped before this interval are invalid. In the Not before field, specify the number of seconds. The default value is 60.
	Not After	(Optional) Specify an interval after the current time. Assertions stamped after this interval are invalid. In the Not after field, specify the number of seconds. The default value is 60.
Manage encrypted assertions	**Encryption Keyring**	(Optional) Select the keyring to use for decrypting encrypted assertions.
	Require Encryption	(Optional) Select to require that assertions from the IDP be encrypted. If selected, the appliance rejects unencrypted assertions. **Note:** As long as the encryption keyring is configured, the appliance attempts to decrypt encrypted assertions whether or not the Require encryption check box is marked.

Task	Setting or field	Description
Configure redirect behavior	**Disable Client Redirects**	(Optional) If the client cannot be redirected to, or communicate directly with, the IDP–for example, if a firewall exists between users and the IDP–select this option to disable the appliance's ability to redirect the browser to the IDP for authentication. When you disable client redirects, the appliance handles all communication with the IDP and OCS on behalf of the client.
	Prefix IDP Cookies	(If you selected **Disable Client Redirects**) Determine if you want to prevent the appliance from adding a BCSI-SWR- prefix to IDP cookies; this prefix prevents IDP cookies and OCS cookies from having the same name. Clear the **Prefix IDP Cookies** option if you do not want to add the prefix to cookies.
Specify SSL device profile for when the redirect/POST URL uses HTTPS	**SSL Device Profile**	(Optional) Select an SSL device profile to use to communicate with the IDP when the redirect/POST URL uses HTTPS.

See Also:

- "Configure a SAML Realm " on page 286

Configure SAML Authorization

You can authorize users through one or more SAML realms and specify the criteria to use when searching for users.

To configure SAML authorization:

1. In the Admin Console, select **Configuration > Authentication > Realms and Domains**.
2. In **Realms**:
   - Select **Add Realm > SAML** to add a new realm.
   - Select an existing realm to edit it.

   The console opens the Add/Edit SAML Realm page.
3. In the Authorization section, the Authorization Realm option is set to **Self** by default. To authorize with the current realm, leave this option selected. When set to **Self**, the authorization username is set automatically to **Use FQDN**.

> **Note:** If you use LDAP for authorization and **Use FQDN** is selected, ensure that the SAML fullname attribute (see "Configure SAML Realm Details" on page 289) contains the user's distinguished name. Later, you must also configure the IDP to send the distinguished name in assertions.

To authorize with another realm, select **Selected**. Enter or edit the following as needed:

Task	Setting or field	Description
Select the realm with which to authorize.	**Authorization Realm Name**	Select a configured authorization realm.
Select one of the options for the authorization username.	**Custom username**	Enter the username in the **Username** field.
	Use FQDN	Use the fully-qualified domain name (FQDN).
	Determine by search	Determine the username by LDAP search. Specify the following. ■ **LDAP Search Realm Name**: Enter the name of the LDAP search realm. ■ **Search Filter**: Specify the LDAP search filter. ■ **User Attribute** or **FQDN**: Specify either the LDAP attribute name or the FQDN as the username attribute for search results. ■ **Add User to Ignore**: Add, edit, or remove usernames from the list of users to ignore when determining authorization.

See Also:

- "Configure a SAML Realm " on page 286

Policy Conditions

The appliance uses existing policy conditions to make authorization decisions for the user.

- `group=`

  This condition maps to the values of the SAML group attribute setting specified in the realm.

- `attribute.<name>=`

  This condition maps to the values of the Attribute name setting specified in the realm.

- `saml.attribute.<saml_attribute_name>=`

  This condition compares strings with the value of the SAML assertion attribute obtained from the user's entry.

  In the VPM, you can use the **SAML Attribute** and **Control Request/Response Header** objects in conjunction to forward SAML assertion attributes via custom headers to back-end or front-end servers. Refer to "Configure SAML Attributes" on page 287 for instructions.

Considerations

Refer to the following considerations when using a SAML realm.

Preventing Dropped Connections When Policy is Set to Deny

Users might experience dropped connections if the appliance uses the default policy of Deny. When the policy is set to Deny, the appliance intercepts and denies requests to the IDP.

To prevent dropped connections due to this limitation, install the following policy to allow requests to the IDP:

```
<Proxy>
 authenticate(<realm-name>)

<Proxy>
   allow group=saml_users

<Proxy>
   allow url.host=<hostname>
```

In the policy example above, <*realm-name*> is the name of the SAML realm and <*hostname*> is the hostname of the IDP.

Backing Up Configuration

You may need to back up the ProxySG configuration and save the backup file (called an archive) on a remote system, which you can restore in the unlikely event of system failure or replacement. For more information on configuration backups, refer to "Backing Up the Configuration" in the *SGOS Administration Guide*.

Save Keyrings Before Backing up the Configuration

Backing up the configuration does not automatically save keyring data; thus, you must save all keyrings and certificates before creating a configuration archive.

To save keyrings before backup:

1. In the command line interface (CLI), enter configuration mode and issue the following command:

   `#(config ssl)`**`show ssl keypair`** `<keyring_name>`

2. Copy and paste the output from the command into a text editor. You will copy and paste this text into the CLI after you restore the appliance.

Import Saved Keyrings After Restoring the Configuration

After restoring the appliance, but before applying the archived configuration, issue the following CLI command:

```
#(config ssl)inline keyring show <keyring_name> <eof marker>
 <copy and paste text here>
 <eof marker>
```

You should now be able to apply the archived configuration without having to create and import a new keyring.

Re-Import the Appliance Certificate to the Trust List (AD FS)

In order to establish an HTTPS connection to the appliance, AD FS must trust the default certificate set in the Admin Console; however, after you back up and restore a configuration archive, you must re-import the certificate to AD FS. For instructions, refer to "Import the Appliance Certificate to AD FS's Trust List" in the *SGOS Administration Guide*.

Authenticate Outbound SSH Client Connections

You must add SSH ciphers, Hash-based Message Authentication Code algorithms (HMACs), and known hosts before the appliance can make an outbound connection over SSH. As a best practice, make sure that the remote SSH server and the appliance share at least one cipher and one HMAC in common.

All ciphers, HMACs, and known hosts for outbound SSH connections stored on the appliance are available for selection and review in the Admin Console (**Configuration > Authentication > SSH Outbound Connections**).

See Also:

- "Manage Host Keys for Outbound SSH Connections" below
- "Manage SSH Client Keys for Outbound Connections" on page 298
- "Manage SSH Ciphers for Outbound Connections" on page 299
- "Manage SSH HMACs for Outbound Connections" on page 301

Manage Host Keys for Outbound SSH Connections

The ProxySG appliance allows you to store the public keys of known hosts, which are used to uniquely identify a server for authentication. When you add a host key to the appliance, the host key data is saved in OpenSSH key file format.

Each configured known host has the following information in the **Known Hosts** section (**Configuration > Authentication > SSH Outbound Connections**). When you look up existing hosts on the appliance, the system searches on the criteria specified in some of these fields.

Host Entry Field	Description
Host pattern	Comma-separated list of canonical host name or IP address patterns. This field accepts wildcard characters when you are adding a known host manually.
Unique ID	A number that the system assigns upon host creation, used in the CLI for deleting specific hosts. ID numbers are not reused when individual hosts are deleted; thus, an ID number can tell you when the host was added to the system, relative to other existing entries. For more information, refer to `#(config ssh-client)` **`known-hosts`** in the *Command Line Interface Reference*.
Key type	The host key type (algorithm), such as ssh-rsa.
Marker	(Optional) One of the following: ■ @cert-authority: The key includes a certification authority (CA) key and is trusted by the CA to validate signed certificates. ■ @revoked: The key is revoked by a CA and must never be accepted for authentication or certification. If no marker is specified, the value is (None).
Comment	(Optional) Enter your own details about the SSH host key.
Key	The entire host key in OpenSSH format.

Add a Known Host

If the appliance has not yet connected to the host, or if an existing host has changed, you can add a new host. Use one of the following methods:

- Obtain the host key from a remote host. See "Fetch host key" on the facing page.
- Copy and paste entire host entries, for example, from an OpenSSH known_hosts file. See "Paste known hosts entry" on the facing page.

> **Note:** Symantec recommends that you use the copy-and-paste method if you want to add known hosts that are not in use yet; however, you should exercise caution because the appliance does not validate your entry using this method.

- Specify individual field values for the known host. See "Manually specify a known entry" on page 298.

> **Note:** If you attempt to add a known host entry that already exists on the appliance, you receive a “Host key already stored” message. If you are copying and pasting a host or specifying it manually, correct the entry before trying to add it again.

Fetch host key

Obtain the host key from a remote host:

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections**.
2. In the Known Hosts section, select **Add Host**. The console displays a Create Host Key dialog.
3. Select **Fetch**.
4. In the **Host** field, enter either the host name or the IP address.
5. In the **Port** field, enter the port number or accept the default (22).
6. Select **Fetch Key From Server**.The Key field shows the connection status.
7. Verify the information. The **Comment** field is optional but helpful for identifying hosts.
8. Select **Create** to confirm the host details. The host is added. To view the entry, locate it in the list (scroll down if necessary). Select and expand the entry to see full details.

Paste known hosts entry

Copy and paste an entire host entry:

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections**.
2. In the Known Hosts section, select **Add Host**. The console displays a Create Host Key dialog.
3. Select **Paste**.

   Add one host key entry or add multiple entries one at a time. Copy and paste a single host key entry into the field,.

   A pasted host key entry must be in the following format:

   marker comma_separated_list_of_host_pattern(s) key_type key #comment

   If the entry does not comply with this format, fields might contain incorrect values or be empty.

   > **Note:** The marker is optional, but if specified it must be @revoked or @certauthority. The comment is optional but helpful for identifying hosts.

   The following is an example of an entry with all values specified:

   ```
   @cert-authority 10.169.2.228,symantec.* ssh-rsa AB3Nz...= #my_comment
   ```

> **Note:** In this example, the host key is abbreviated. When you perform this step on the appliance, paste the host key in its entirety.

4. Select **Create** to confirm the host details. The host is added. To view the entry, locate it in the list (scroll down if necessary). Select and expand the entry to see full details.

Manually specify a known entry

Specify individual field values for the known host:

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections**.
2. In the Known Hosts section, select **Add Host**. The console displays a Create Host Key dialog.
3. Select **Manual**.
4. Specify known host settings; see the table at the top of this page for descriptions.
5. Select **Create** to confirm the host details. The host is added. To view the entry, locate it in the list (scroll down if necessary). Select and expand the entry to see full details.

Delete a Known Host

You can delete a selected host entry, for example, if the appliance no longer needs to connect to the known host or an entry has changed.

To delete a known host, select the check box for each host you want to delete. Then, select **Delete** under Actions.

Manage SSH Client Keys for Outbound Connections

To manage SSH client keys for outbound client connections, in the Admin Console, select **Configuration > Authentication > SSH Outbound Connection** and go to the **Client Keys** section.

> **Note:** For best security, key pairs are SSHv2 only; make sure that port 22 is available. You do not need to create a new host key unless you want to change the existing configuration.

Creating a Client Key Pair

To create an SSHv2 client key pair:

1. Select **Generate**. A Generate Host Key dialog appears.
2. In the dialog, select the key type and strength:

Client Key Type	Strength
RSA	Select one of: ■ 2048 (default) ■ 3071 ■ 4096
ECDSA	Select one of: ■ nistp256 ■ nistp384 ■ nistp521 (default)
Ed25519	N/A

3. Select **Generate**. The client key is created and saved immediately.

You can create a client key pair in the CLI using `#(config ssh-client client-keys)` commands. Refer to the *Command Line Interface Reference* for details.

Deleting a Client Key Pair

To delete a client key pair, select the delete option in the respective section. On the dialog that appears, select **Delete** to confirm the deletion. The key pair is deleted immediately.

Manage SSH Ciphers for Outbound Connections

To manage SSH ciphers for outbound connections, select **Configuration > Authentication > SSH Outbound Connections** and go to the **Ciphers** section.

The Ciphers section shows current list of ciphers, including any ciphers you added explicitly and excluding any you removed explicitly.

To display all available ciphers that the SGOS version supports, select **Select Ciphers**. The Select Ciphers dialog shows all ciphers; however, note that:

- Fewer ciphers are available when the appliance is in FIPS mode.
- A marked checkbox indicates that the cipher is currently selected.

After an upgrade or downgrade, the selected ciphers may change. If you modify the selections, the changes persist after system upgrades, downgrades, and reboots; however, selections will not be identical to the list prior to upgrade/downgrade if the system must consider deprecated ciphers. To understand the behavior after upgrade/downgrade:

- Ciphers that were previously added explicitly are selected if they are supported after upgrade/downgrade. They are not selected if they are deprecated in the current version.
- Ciphers that were previously removed explicitly are not selected even if they are supported in the current version.
- Ciphers that were neither added nor removed explicitly are selected if supported in the current version and not selected if deprecated.
- If you upgrade to a release that supports only ciphers that you previously removed, resulting in no selected ciphers, the appliance warns you that the list is empty and event-logs the occurrence.

For example, if you upgrade to a version of SGOS in which an added cipher is deprecated, the cipher is not selected. Downgrading to the previous SGOS version selects the cipher again.

Add Ciphers

The appliance selects a number of ciphers by default. You can add more ciphers from a list of available ciphers, and also specify the order in which the appliance should use ciphers for SSH connections.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **Ciphers** section.
2. Select **Select Ciphers**. The Select Ciphers dialog shows all available ciphers. If necessary, toggle the **Show selected only** option.
3. Add ciphers by selecting the checkboxes beside them.
4. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any ciphers are added.

Remove Ciphers

You can remove ciphers from the selections list. Removing a cipher means it will not be available for SSH connections unless you add it again.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **Ciphers** section.
2. Do one of the following:
   - Remove ciphers by selecting **Remove** beside them, and then save your pending changes to the list.
   - Select **Select Ciphers**. The Select Ciphers dialog shows selected ciphers. Remove ciphers by clearing the checkmarks beside them. Use the **Select** option to save changes.

Tip: The event log indicates when any ciphers are removed.

Set the Preferred Order of Ciphers

When the appliance sends its list of cipher suites for SSH connections, it uses the order specified on the Ciphers list. You can change the preferred order of ciphers using the arrows to the right of the list.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **Ciphers** section.
2. To promote a cipher on the list, select **Move Up** as many times as required.
3. To demote a cipher on the list, select **Move Down** as many times as required.
4. Save your pending changes to the list.

Restore the Default List

You can restore the default list of ciphers that the SGOS version supports. Using this option resets both the default cipher selections and the default preferred order.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **Ciphers** section.
2. Select **Revert to Default**.

Manage SSH HMACs for Outbound Connections

To manage SSH HMACs for outbound connections, select **Configuration > Authentication > SSH Outbound Connections** and go to the **HMACs** section.

The HMACs section shows current list of HMACs, including any HMACs you added explicitly and excluding any you removed explicitly.

To display all available HMACs that the SGOS version supports, select **Select HMACs**. The Select HMACs dialog shows all HMACs; however, note that:

- Fewer HMACs are available when the appliance is in FIPS mode.
- A marked checkbox indicates that the HMAC is currently selected.

After an upgrade or downgrade, the selected HMACs may change. If you modify the selections, the changes persist after system upgrades, downgrades, and reboots; however, selections will not be identical to the list prior to upgrade/downgrade if the system must consider deprecated HMACs. To understand the behavior after upgrade/downgrade:

- HMACs that were previously added explicitly are selected if they are supported after upgrade/downgrade. They are not selected if they are deprecated in the current version.
- HMACs that were previously removed explicitly are not selected even if they are supported in the current version.
- HMACs that were neither added nor removed explicitly are selected if supported in the current version and not selected if deprecated.
- If you upgrade to a release that supports only HMACs that you previously removed, resulting in no selected HMACs, the appliance warns you that the list is empty and event-logs the occurrence.

For example, if you upgrade to a version of SGOS in which an added HMAC is deprecated, the HMAC is not selected. Downgrading to the previous SGOS version selects the HMAC again.

Add HMACs

The appliance selects a number of HMACs by default. You can add more HMACs from a list of available HMACs, and also specify the order in which the appliance should use HMACs for SSH connections.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **HMACs** section.
2. Select **Select HMACs**. The Select HMACs dialog shows all available HMACs. If necessary, toggle the **Show selected only** option.
3. Add HMACs by selecting the checkboxes beside them.
4. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any HMACs are added.

Remove HMACs

You can remove HMACs from the selections list. Removing an HMAC means it will not be available for SSH connections unless you add it again.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **HMACs** section.
2. Do one of the following:
   - Remove HMACs by selecting **Remove** beside them, and then save your pending changes to the list.
   - Select **Select HMACs**. The Select HMACs dialog shows selected HMACs. Remove HMACs by clearing the checkmarks beside them. Use the **Select** option to save changes.

Tip: The event log indicates when any HMACs are removed.

Set the Preferred Order of HMACs

When the appliance sends its list of HMACs for SSH connections, it uses the order specified on the HMACs list. You can change the preferred order of HMACs using the arrows to the right of the list.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **HMACs** section.
2. To promote an HMACon the list, select **Move Up** as many times as required.
3. To demote an HMACon the list, select **Move Down** as many times as required.
4. Save your pending changes to the list.

Restore the Default List

You can restore the default list of HMACs that the SGOS version supports. Using this option resets both the default HMAC selections and the default preferred order.

1. In the Admin Console, select **Configuration > Authentication > SSH Outbound Connections** and go to the **HMACs** section.
2. Select **Revert to Default**.

Manage the SSH Console

By default, the appliance uses Secure Shell (SSH) and password authentication so administrators can access the CLI securely. SSH is a protocol for secure remote logon over an insecure network.

Authentication for administrative users connecting to the appliance via SSH can be controlled with the following authentication types:

- The predefined admin account
- Local authentication realm
- IWA authentication realm
- LDAP authentication realm
- Localized RSA key

When managing the SSH console, you can:

- Generate or re-generate SSH host keys
- Specify the SSHv2 algorithm
- Create or remove client keys and Director keys
- Specify a welcome message for clients accessing the appliance using SSHv2

See Also:

- "Manage the Host Key Pairs" below

Manage Host Key Pairs for the SSH Console

You can create SSHv2 RSA, ECDSA, or Ed25519 keys to secure the SSH console. See "Manage the Host Key Pairs" below.

Create a Welcome Banner for SSH

A welcome banner provides notice to users of computer networks, computers, and other systems and resources. Users are required to accept the terms in the banner prior to authentication. The banner is often presented to users before a login process, and it requires users to acknowledge and agree to the message before they can log in or access resources on the network.

To create a welcome banner for the SSH console, enter the text in the **Welcome Banner** field.

Refer to the *Notice and Consent Banner Configuration Guide* for more information, including samples of CPL that you can use:

https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/notice-and-consent-banner-config.html

Manage Ciphers and HMACs

The Admin Console (select **Configuration > Authentication > SSH Inbound Connections** and go to **Ciphers**) shows the current list of SSH ciphers for inbound connections, including any ciphers you added explicitly and excluding any you removed explicitly. The **HMACs** section shows the current list of Hash-based Message Authentication Code algorithms (HMACs), including any HMACs you added explicitly and excluding any you removed explicitly.

To make add, remove, or make other changes to the selected ciphers or HMACs, see "Manage SSH Ciphers for Inbound Connections" on page 307 and "Manage SSH HMACs for Inbound Connections" on page 309.

Manage the Host Key Pairs

To manage the SSH host, select **Configuration > Authentication > SSH Inbound Connections**, and go to the **Host Keys** section.

> **Note:** For best security, key pairs are SSHv2 only; make sure that port 22 is available. You do not need to create a new host key unless you want to change the existing configuration.

Create a Host Key Pair

To create an SSHv2 host key pair:

1. Select **Generate**. A Generate Host Key dialog appears.
2. In the dialog, select the key type and strength:

Host Key Type	Strength
RSA	Select one of: ■ 2048 (default) ■ 3071 ■ 4096
ECDSA	Select one of: ■ nistp256 ■ nistp384 ■ nistp521 (default)
Ed25519	N/A

3. Select **Generate**. The host key is created and saved immediately.

> **Note:** If you receive an error message when attempting to log in to the system after regenerating the host key pair, locate the ssh known hosts file and delete the previous entry for the ProxySG appliance.

You can create a host key pair in the CLI using the `#(config ssh-console) create host-keypair` command. Refer to the *Command Line Interface Reference* for details on this command.

Delete a Host Key Pair

To delete a host key pair, select the delete option in the respective section. On the dialog that appears, select **Delete** to confirm the deletion. The key pair is deleted immediately.

Disabling all key pairs might lock you out of the ProxySG CLI (for example you receive a message such as “could not agree on a host key algorithm”). If this occurs, re-create an SSH key pair through a serial connection or through the Management Console.

Manage SSH Client Keys for Inbound Connections

You can import multiple RSA client keys on the appliance to provide public key authentication, an alternative to using password authentication. An RSA client key can only be created by an SSH client and then imported onto the appliance. Many SSH clients are commercially available for UNIX and Windows.

To import the client keys, copy them from your SSH client and import them to the appliance via the CLI. Refer to the *Command Line Interface Reference*.

> **Note:** The appliance cannot create client keys. You must use your SSH client to create a key.

After you create an RSA client key following the instructions of your SSH client, you can import the key onto the appliance using either the Admin Console or the CLI.

About the OpenSSH.pub Format

Symantec supports the OpenSSH.pub format. Keys created in other formats will not work.

An OpenSSH.pub public key is similar to the following:

```
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAwFI78MKyvL8DrFgcVxpNRHMFKJrBMeBn
2PKcv5oAJ2qz+uZ7hiv7Zn43A6hXwY+DekhtNLOk3HCWmgsrDBE/NOOEnDpLQjBC6t/
T3cSQKZjh3NmBbpE4U49rPduiiufvWkuoEiHUb5ylzRGdXRSNJHxxmg5LiGEiKaoELJfs
DMc= user@machine
```

The OpenSSH.pub format appends a space and a user ID to the end of the client key.

The user ID used for each key must be unique.

> **Note:** 4096 bits is the maximum supported key size.
> The `ssh-rsa` prefix is required
> Trailing newline characters must be removed from the key before it is imported.

Viewing RSA Client Keys

You can view imported RSA client keys in the Admin Console. Client keys provide more secure authentication compared to user name/password authentication.

To view RSA client keys:

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections**.
2. The Client Keys section displays the following information:

- **Username**: The client key is associated with the specified user.
- **Key ID**: The ID is included in the imported RSA key.
- **Fingerprint**: The fingerprint is a unique ID of the imported key.

Manage SSH Ciphers for Inbound Connections

To manage SSH ciphers for inbound connections, select **Configuration > Authentication > SSH Inbound Connections** and go to the **Ciphers** section.

The Ciphers section shows current list of ciphers, including any ciphers you added explicitly and excluding any you removed explicitly.

To display all available ciphers that the SGOS version supports, select **Select Ciphers**. The Select Ciphers dialog shows all ciphers; however, note that:

- Fewer ciphers are available when the appliance is in FIPS mode.
- A marked checkbox indicates that the cipher is currently selected.

After an upgrade or downgrade, the selected ciphers may change. If you modify the selections, the changes persist after system upgrades, downgrades, and reboots; however, selections will not be identical to the list prior to upgrade/downgrade if the system must consider deprecated ciphers. To understand the behavior after upgrade/downgrade:

- Ciphers that were previously added explicitly are selected if they are supported after upgrade/downgrade. They are not selected if they are deprecated in the current version.
- Ciphers that were previously removed explicitly are not selected even if they are supported in the current version.
- Ciphers that were neither added nor removed explicitly are selected if supported in the current version and not selected if deprecated.
- If you upgrade to a release that supports only ciphers that you previously removed, resulting in no selected ciphers, the appliance warns you that the list is empty and event-logs the occurrence.

For example, if you upgrade to a version of SGOS in which an added cipher is deprecated, the cipher is not selected. Downgrading to the previous SGOS version selects the cipher again.

Add Ciphers

The appliance selects a number of ciphers by default. You can add more ciphers from a list of available ciphers, and also specify the order in which the appliance should use ciphers for SSH connections.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **Ciphers** section.
2. Select **Select Ciphers**. The Select Ciphers dialog shows all available ciphers. If necessary, toggle the **Show selected only** option.

3. Add ciphers by selecting the checkboxes beside them.
4. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any ciphers are added.

Remove Ciphers

You can remove ciphers from the selections list. Removing a cipher means it will not be available for SSH connections unless you add it again.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **Ciphers** section.
2. Do one of the following:
   - Remove ciphers by selecting **Remove** beside them, and then save your pending changes to the list.
   - Select **Select Ciphers**. The Select Ciphers dialog shows selected ciphers. Remove ciphers by clearing the checkmarks beside them. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any ciphers are removed.

Set the Preferred Order of Ciphers

When the appliance sends its list of cipher suites for SSH connections, it uses the order specified on the Ciphers list. You can change the preferred order of ciphers using the arrows to the right of the list.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **Ciphers** section.
2. To promote a cipher on the list, select **Move Up** as many times as required.
3. To demote a cipher on the list, select **Move Down** as many times as required.
4. Save your pending changes to the list.

Restore the Default List

You can restore the default list of ciphers that the SGOS version supports. Using this option resets both the default cipher selections and the default preferred order.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **Ciphers** section.

2. Select **Revert to Default**.

> **Note:** If SSH connections fail and you receive an error stating that no ciphers are available, refer to https://knowledge.broadcom.com/external/article/169231.

Manage SSH HMACs for Inbound Connections

To manage SSH HMACs for inbound connections, select **Configuration > Authentication > SSH Inbound Connections** and go to the **HMACs** section.

The HMACs section shows current list of HMACs, including any HMACs you added explicitly and excluding any you removed explicitly.

To display all available HMACs that the SGOS version supports, select **Select HMACs**. The Select HMACs dialog shows all HMACs; however, note that:

- Fewer HMACs are available when the appliance is in FIPS mode.
- A marked checkbox indicates that the HMAC is currently selected.

After an upgrade or downgrade, the selected HMACs may change. If you modify the selections, the changes persist after system upgrades, downgrades, and reboots; however, selections will not be identical to the list prior to upgrade/downgrade if the system must consider deprecated HMACs. To understand the behavior after upgrade/downgrade:

- HMACs that were previously added explicitly are selected if they are supported after upgrade/downgrade. They are not selected if they are deprecated in the current version.
- HMACs that were previously removed explicitly are not selected even if they are supported in the current version.
- HMACs that were neither added nor removed explicitly are selected if supported in the current version and not selected if deprecated.
- If you upgrade to a release that supports only HMACs that you previously removed, resulting in no selected HMACs, the appliance warns you that the list is empty and event-logs the occurrence.

For example, if you upgrade to a version of SGOS in which an added HMAC is deprecated, the HMAC is not selected. Downgrading to the previous SGOS version selects the HMAC again.

Add HMACs

The appliance selects a number of HMACs by default. You can add more HMACs from a list of available HMACs, and also specify the order in which the appliance should use HMACs for SSH connections.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **HMACs** section.

2. Select **Select HMACs**. The Select HMACs dialog shows all available HMACs. If necessary, toggle the **Show selected only** option.
3. Add HMACs by selecting the checkboxes beside them.
4. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any HMACs are added.

Remove HMACs

You can remove HMACs from the selections list. Removing an HMAC means it will not be available for SSH connections unless you add it again.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **HMACs** section.
2. Do one of the following:
   - Remove HMACs by selecting **Remove** beside them, and then save your pending changes to the list.
   - Select **Select HMACs**. The Select HMACs dialog shows selected HMACs. Remove HMACs by clearing the checkmarks beside them. Use the **Select** option to save changes.

> **Tip:** The event log indicates when any HMACs are removed.

Set the Preferred Order of HMACs

When the appliance sends its list of HMACs for SSH connections, it uses the order specified on the HMACs list. You can change the preferred order of HMACs using the arrows to the right of the list.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **HMACs** section.
2. To promote an HMACon the list, select **Move Up** as many times as required.
3. To demote an HMACon the list, select **Move Down** as many times as required.
4. Save your pending changes to the list.

Restore the Default List

You can restore the default list of HMACs that the SGOS version supports. Using this option resets both the default HMAC selections and the default preferred order.

1. In the Admin Console, select **Configuration > Authentication > SSH Inbound Connections** and go to the **HMACs** section.
2. Select **Revert to Default**.

> **Note:** If SSH connections fail and you receive an error stating that no HMACs are available, refer to https://knowledge.broadcom.com/external/article/169231.

Monitor the Appliance

Refer to the following topics for ProxySG appliance statistics:

- "Advanced Statistics" on the facing page

Advanced Statistics

In the Admin Console, select **Dashboard > Advanced URLs** to display a list of advanced statistics information:

- **Groups**: URL groups are listed in alphabetical order. Select a group to display the URLs.
- **All URLs**: All URLs are listed in alphabetical order.

Caution: Proceed with caution. Symantec recommends that you refer to Symantec Support for assistance with these URLs. Support might direct you to these links to provide additional information during troubleshooting.

View System Reports

You can review session statistics as XML data for some system operations:

- "View Active Sessions (Per-Session Connection Statistics)" on the facing page

> **Tip:** For detailed information about statistics, or to see graphical reports, refer to the *SGOS Administration Guide*.

View Active Sessions (Per-Session Connection Statistics)

View active sessions for the following statistics:

- **Active sessions memory statistics**: View statistics for memory used by active sessions reports.
- **Bypassed connections**: View data for all unintercepted TCP traffic. When the appliance is first installed in an in-path deployment, all services are bypassed by default. By analyzing the connection data in the Bypassed Connections page, you can review the types of traffic flowing through the appliance to identify traffic flows that would benefit from optimization. The Bypassed Connections page is also useful for identifying new types of traffic flowing through the appliance.
- **Proxied connections**: View an immediate picture of the sessions and the protocol types, services, bytes, savings, and other statistics. These statistics are derived from WAN optimization and object caching and are associated with client traffic.
- **ADN connections**: View inbound ADN connection data. These connections are not persistent. When a connection completes, the statistics for the connection are no longer displayed.
- **Active sessions**: View all active sessions.

> **Tip:** For detailed information about statistics, or to see graphical reports, refer to the *SGOS Administration Guide*.

Administer the Appliance

Refer to the following topics to learn about ProxySG administration using the Admin Console.

- "Configure General ProxySG Settings" on the facing page
- "Configure Data Services" on page 323
- "Upgrade and Manage Systems " on page 371

> **Note:** These upgrade options are not available in the Admin Console launched from Management Center. If you're using Management Center, upgrade the ProxySG version using jobs. Refer to the Management Center documentation for details.

> **Note:** Before upgrading the appliance, refer to the *SGOS Upgrade/Downgrade Guide* to determine your upgrade path: https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/upgrade-downgrade-guide.html

- "About Licensing" on page 374
- "Configure ProxySG Access Logs" on page 383
- "Malicious Content Scanning Services " on page 418
- "Verify Service Health and Status" on page 446
- "Service Information" on page 467

For more information, refer to the *SGOS Administration Guide* at MySymatec.

Configure General ProxySG Settings

Refer to the following topics to learn about configuring general ProxySG appliance settings:

- "Identification and General Information" on the next page
- "Configure System Time" on page 319
  - "Select a Time Zone" on page 321
- "Required Ports, Protocols, and Services" on page 322

Identification and General Information

You can assign any name to the ProxySG appliance. A descriptive name helps identify the system, for example, when managing multiple appliances from Symantec Management Center.

To set the appliance name:

1. In the Admin Console, select **Administration > General > Identification**.
2. In the Device Information section, in the **Appliance name** field, enter a unique name for the appliance.

   The section displays other basic information about the appliance.

Review Third-Party Software Attributions

To review the third-party software attributions used in the ProxySG Admin Console, select **Administration > General > Identification**

In the Open Source Attributions section, select the **3rd Party Copyright Attributions** link. The list of attributions opens in a new browser window.

Review Copyright Information

To review full product copyright information, see "Legal Notice" on page 477.

Configure System Time

The ProxySG appliance uses NTP and the UTC to keep the system time accurate. To manage objects, the ProxySG appliance must know the current UTC, which is the international time standard and is based on a 24-hour clock. The ProxySG appliance accesses the NTP servers to obtain accurate UTC time and synchronizes its time clock. You can also record time stamps in the current local time or use the browser's time.

Set Time Source

By default, the ProxySG appliance connects to an NTP server in the order they are listed under **NTP Servers** and acquires the UTC. If the appliance cannot access any of the listed NTP servers, you must manually set the UTC.

To set the time source:

1. In the Admin Console, select **Administration > General > Time**.
2. In **Time Source**, do one of the following:
   - Select **Acquire from NTP Server** (default setting) to use NTP. See "Manage NTP Servers " on the next page for information.
   - Select **Use Manual Setup** to disable NTP.

Set Local Time

To record time stamps in current local time, you must set the local time based on your time zone. You can set local time using the time zone database, or by using browser time.

The ProxySG appliance ships with a limited list of time zones. If a specific time zone is missing from the included list, you can update the list at your discretion (see "Update the Time Zone Database" on page 321). The time zone database might need to be updated if the Daylight Savings rules change in your area.

To set local time :

1. In the Admin Console, select **Administration > General > Time**.
2. In **Current Time**:
   - To set local time based on browser time, select **Set to Browser Time**. After the local time is set to browser time, the option is unavailable (grayed out) unless you make a change to the time source.
   - To set local time using a time zone database:
     a. Beside Local Timezone, select **Edit**.
     b. On the dropdown, select **Set Timezone**. The Select Time Zone page opens.
     c. Select a time zone. See "Select a Time Zone" on page 321.

Manage NTP Servers

NTP is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver or modem. There are more than 230 primary time servers, synchronized by radio, satellite and modem.

The ProxySG appliance ships with a list of NTP servers available on the Internet, and attempts to connect to them in the specified order in the NTP Servers section. You can add and remove entries on the list.

To manage NTP servers:

1. In the Admin Console, select **Administration > General > Time**.
2. In **NTP Servers**, select **Add**.
3. In the new Name field, enter a new NTP server using one of the following methods:
    - Domain name: Enter a domain name of an NTP server that resolves to an IPv4 or IPv6 address.
    - IP address: Enter an IPv4 or IPv6 address of an NTP server.

To remove an NTP server, delete it in the Actions column.

Update NTP

To force an NTP update:

1. In the Admin Console, select **Administration > General > Time**.
2. In **NTP Servers**, select **Query NTP**. A dialog notifies you that the update is started and will continue in the background.

Select a Time Zone

The **Select Time Zone** page displays all time zones, organized by geographical region, in the currently installed time zone database.

To add more time zones to the list or update the appliance's time zone database, see "Update the Time Zone Database" below.

To select the time zone that represents your local time:

1. In the Admin Console, select **Administration > General > Time**.
2. In Current Time, under Local Timezone, select **Edit**.
3. On the dropdown, select **Set Timezone**. The Select Time Zone page opens.
4. (If needed) Select **Go to Selection** to locate the currently selected time zone.
5. Locate the country/region for the local time zone, expand it, and select the time zone. If needed, use **Expand All** and **Collapse All** to expand and collapse the list of time zones, respectively.

   After you select a time zone, event logs record the local time instead of GMT.

Update the Time Zone Database

The ProxySG appliance ships with a limited list of time zones. If a specific time zone is missing from the included list, you can update the list at your discretion.

You might also want to update the time zone database if the Daylight Savings rules change in your area. You can download an updated database from the default Symantec URL, or specify a different URL for your database downloads.

To update the time zone database:

1. In the Admin Console, select **Administration > General > Time**.
2. In **Current Time**, beside Local Timezone, select **Edit**.
3. On the dropdown, select **Update Timezone Database**. The Load new time zone database dialog opens.
4. Enter a URL from which the database will be downloaded or accept the default URL (http://download.bluecoat.com/release/timezones.tar).
5. Select **Download**.

Required Ports, Protocols, and Services

Depending on your ProxySG appliance configuration, you must open certain ports and protocols on your firewalls for the appliance to function as intended, or to allow connectivity to various components and data centers. For full details, refer to the following knowledge base article:

https://www.symantec.com/docs/INFO5294

Configure Data Services

Refer to the following topics to learn about setting up content filtering providers on the ProxySG appliance:

- "Content Filtering Overview" on the next page
  - "About Symantec WebFilter and the WebPulse Service" on page 329
- "Supported Content Filter Providers" on page 333
  - "Configure Blue Coat Content Filtering" on page 339
  - "Configure a Local Database" on page 336
  - "Configure Categories for YouTube" on page 342
- "View and Verify Threat Protection Services" on page 343
  - "Considerations Before Configuring WebPulse" on page 344
  - "Configure WebPulse" on page 348
  - "Analyze the Threat Risk of a URL" on page 350
  - "Configure Threat Risk Levels" on page 352
- "Use Intelligence Services to Classify Applications" on page 353
- "About Geolocation " on page 357
  - "Using Geolocation " on page 358
  - "Test Outbound Connections Based on Geographic Location" on page 359
  - "Determine Locations of IP Addresses for Incoming Connections" on page 362
  - "Troubleshoot Geolocation Errors" on page 364
  - "Remove Geolocation Settings" on page 366
- "About Web Isolation" on page 367
  - "Configure Web Isolation" on page 367

Content Filtering Overview

Content Filtering allows you to categorize and analyze Web content. With policy controls, content filtering can support your organization's Web access rules by managing or restricting access to Web content and blocking downloads from suspicious and unrated Web sites, thereby helping protect your network from undesirable or malicious Web content.

Content filtering is a method for screening access to web content. It allows you to control access to web sites based on their perceived content. On the ProxySG appliance, using a content filtering database in conjunction with policy allows you to manage employee access to web content and to restrict access to unsuitable content. Restricting access or blocking web content helps reduce the risk of malware infections caused by visiting questionable sites.

Content filtering categories comprehensively classify the vast and constantly growing number of URLs that are found on the web into a relatively small number of groups or categories. These categories then allow you to control access to web content through policy.

A content filtering database:

- Has a pre-defined set of categories provided by the content filtering vendor. Individual content filter providers such as Symantec Intelligence Services define the content-filtering categories and their meanings. Depending on the vendor, a URL is listed under one or more categories. Each URL can support a maximum of 16 categories.
- Does not block any web site or category. The role of the database is to offer additional information to the proxy server and to the administrator about a client request. After you configure your content filter provider and download the database, you can map the URLs to the list of categories. You can then reference these categories in policy and limit, allow, or block requests. Client access to a web request depends on the rules and policies that you implement in accordance with company standards.

Some policies that you might create, for example, are as follows:

- Block or unblock specific sites, categories, or specific file types such as executables.
- Apply different filtering policy for each site or group within your organization, by IP address or subnet. If you wish to use password authentication to grant or deny access to the requested content, you must have configured authentication realms and groups on the ProxySG appliance. For information about configuring authentication, see " About Appliance Security" on page 232.
- Allow schedule-based filtering to groups within your organization.

A valid vendor subscription or license is required to download a content filter database. For example, Symantec WebFilter is licensed whereas some supported third-party vendors require a subscription.

For information on integrating your local appliance content filtering policy with Symantec Cloud Service policy, refer to Universal Policy: Auto Policy Synchronization WebGuide.

Web filtering process flow

The following diagram illustrates the process flow when web content filtering is employed in the network. This diagram does not include the dynamic categorization process; for details, see "About Dynamic Categorization " on page 329.

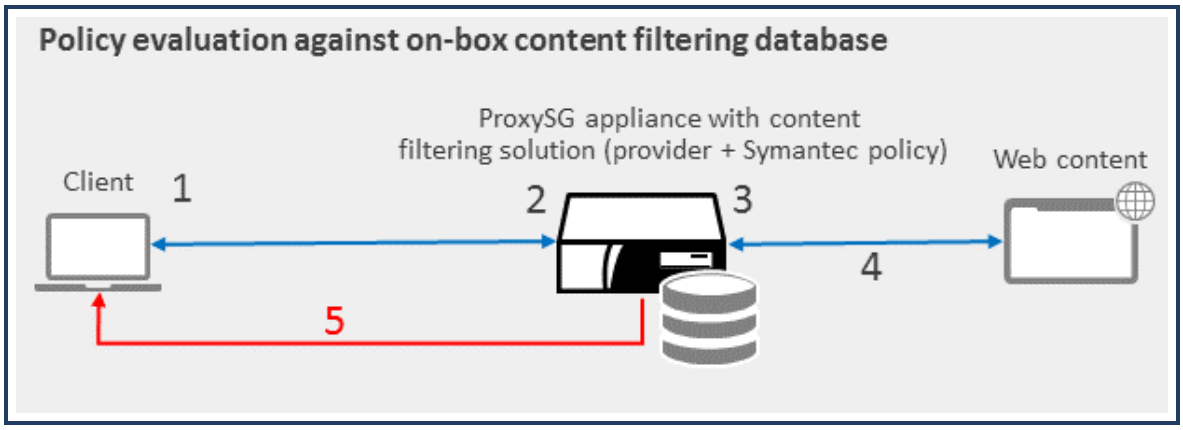

1. The client requests a Web page.
2. The ProxySG appliance checks the requested URL against the on-box content filtering database to determine the categorization.
3. After the URL is categorized, the policy engine determines if the URL is allowable or not.
4. The URL is allowed and the request continues to its destination
5. The policy denies the request and returns a message concerning corporate Web compliance.

Enable a Content Filter

Make sure that you have a valid license or subscription (if needed) for the provider before enabling a provider; otherwise, health errors could occur.

The ProxySG appliance supports several content filter providers. To enable a provider, you can select the **Enable** switch beside the provider name (if available); however, Broadcom recommends that you configure and verify provider settings first:

- Local database: Create and upload your custom content filtering database to the ProxySG appliance. This database must be in a text file format.

  (SGOS 6.7.4 and later) You can configure up to seven additional local databases using the CLI. Refer to the `# (config local` *`database_name`*`)` commands in the *Command Line Interface Reference*.

- Third-party content filtering vendor database: Internet Watch Foundation (IWF), Proventia, and Optenet are supported.

  Only IWF can be enabled in the Admin Console. For more information, refer to https://www.iwf.org.uk/.

  To configure Proventia and Optenet, refer to the *Command Line Interface Reference*.

> **Note:** You cannot use multiple third-party content filtering vendor databases simultaneously.

- Intelligence Services: A framework for the delivery of data feeds to Symantec platforms. Multiple data feeds are entitled by subscription to an Intelligence Services solution bundle. These data feeds are delivered and made available to the ProxySG appliance through the Intelligence Services framework. You can obtain a license for one or more bundles, and also enable or disable data feeds in your solution bundle as your requirements change. See "Configure Blue Coat Content Filtering" on page 339.
- (If available) Symantec WebFilter: Provides both an on-box content filtering database and the WebPulse service, a cloud-based threat-protection feature. See "Configure Blue Coat Content Filtering" on page 339 to configure settings before enabling the service for the first time.
- YouTube: The appliance obtains video categories from the YouTube Data API v3.0. After you enable YouTube categories, you can reference these categories in policy to control YouTube traffic. Before you can enable YouTube, you require a server key from Google. See "Set the YouTube Server Key " on page 342 for information to configure settings before enabling the service for the first time.

> **Note:** This feature is provided on an "as-is" basis. Symantec has no control of, and is not responsible for, information and content provided (or not) by YouTube. Customer is required to apply and use its own API key in order to activate this feature, and therefore obligated to comply with all terms of use regarding the foregoing (for example, see https://developers.google.com/youtube/terms), including quotas, restrictions and limits on use that may be imposed by YouTube. Symantec shall not be liable for any change, discontinuance, availability or functionality of the features described herein.

Add Category Review Link to Exception Messages

(Symantec WebFilter/Intelligence Services only) Add a link to the default content filter exception page when a user is denied a request for a web page.

Exception pages are customized web pages (or messages) sent to users under specific conditions defined by a company and its security policies. An exception page is served, for example, when a category is blocked by company policy. The ProxySG appliance offers multiple built-in exception pages that can be modified to meet your enterprise needs. For content filtering, the appliance includes the `content_filter_denied` and `content_filter_unavailable` built-in exception pages.

The `content_filter_denied` exception page includes the following information:

- an exception page message that includes the content filtering category affecting the exception.
- a category review URL, where content categorizations can be reviewed and/or disputed.

Typically, the exception page informs the user why a URL request is denied. When you add the link to exception pages, the user can select it to request a review of the category assigned to the blocked URL. To customize the exception page, refer to the *Web Visual Policy Manager Reference* or *Legacy Visual Policy Reference*.

To add the category review link to exception messages:

1. In the Admin Console, select **Administration > Data Services > Content Filtering**.
2. In the Options section, select **Enable category review message in exceptions**.

The `content_filter_unavailable` exception page includes a message that states the reason for the denial and provides a probable cause. For example, the request was denied because an external content filtering service was not available owing to transient network problems, or a configuration error.

Set Memory Allocation

The default memory allocation setting (Normal) is ideal for most deployments; increasing the setting is not recommended in most situations because content filtering databases can be very large and require significant resources to process.

You might be able to adjust the amount of memory allocated to the database in the following cases:

- If you are not using ADN and have a high transaction rate for content filtering, you can increase the memory allocation setting to High. This helps content filtering run more efficiently.
- If you are using both ADN and content filtering but the transaction rate for content filtering is not very high, you can reduce the memory allocation setting to Low. This makes more resources available for ADN, allowing it to support a larger number of concurrent connections.

To set memory allocation:

1. In the Admin Console, select **Administration > Data Services > Content Filtering**.
2. Under Options, select the memory allocation level.

Test the Category, Application, and Operation for a URL

If you're using Intelligence Services or WebFilter for content filtering, you have the additional ability to deny or allow access to certain web applications and/or operations; this is done via policy (either VPM or CPL). To determine the application or operation name associated with a URL so that you can create policy to block or allow it, you can do a URL test.

In addition to URL category filtering, you can filter content by Web application and/or specific operations or actions done within those applications. For example, you can create policy to:

- Allow users to access all social networking sites, except for Facebook. Conversely, block access to all social networking sites except for LinkedIn.
- Allow users to post comments and chat in Facebook, but block uploading of pictures and videos.

- Prevent the uploading of videos to YouTube, but allow all other YouTube operations such as viewing videos others have posted. Conversely, preventing uploading but block access to some videos according to the video's category.
- Allow users to access their personal email accounts, but prevent them from sending email attachments.

This feature allows administrators to block actions in accordance with company policy to avoid data loss accidents, prevent security threats, or increase employee productivity.

To test the category, application, and operation associated with a URL:

1. In the Admin Console, select **Administration > Data Services > Content Filtering**.
2. In the Test Category Match section, enter the URL and select **Test URL**. The browser opens a new page with the category, application, and operation assigned to the URL.

   > **Note:** The maximum number of categories for any single URL is 16. If more than 16 categories are specified, the ProxySG appliance arbitrarily matches against 16 out of the total number specified.

   The URL facebook.com/video/upload_giver.php is categorized as follows:

   - Social Networking; Audio/Video Clips
   - Facebook
   - Upload Videos

   The test results in this example indicate that the URL has two categories (Social Networking and Audio/Video Clips), is the Facebook application, and is the Upload Videos operation.

   Note that not all URLs have applications and operations associated with them. For URLs that WebFilter has not assigned an application or operation, the test results indicate none.

   If you have a license, you can also test applications and operations using Application Classification. Refer to the *SGOS Administration Guide* for details. See "Use Intelligence Services to Classify Applications" on page 353.

About Symantec WebFilter and the WebPulse Service

Symantec WebFilter, in conjunction with the WebPulse service, offers a comprehensive URL-filtering solution. Symantec WebFilter provides an on-box content filtering database and WebPulse provides an off-box dynamic categorization service for real-time categorization of URLs that are not categorized in the on-box database. WebPulse dynamic categorization includes both traditional content evaluation, for categories such as pornography, as well as realtime malware and phishing threat detection capabilities. WebPulse services are offered to all customers using WebFilter.

WebPulse is a cloud service that allows inputs from multiple enterprise gateways and clients and creates a computing grid. Based on the analysis of this large volume of requests, the computing grid continuously updates the master WebFilter database, and the ProxySG appliance expediently updates its on-box copy of the Symantec WebFilter database. About 95% of the Web requests made by a typical enterprise user (for the English language) are present in the onbox Symantec WebFilter database, thereby minimizing bandwidth usage and maintaining quick response times.

By default, the WebPulse service is enabled and configured to dynamically categorize unrated and newWeb content for immediate enforcement of policy. Typically, the response time from the dynamic categorization service is about 500 milliseconds and is subject to the response/performance of the site in question. However, if this service is causing significant delays to your enterprise web communications, you can run it in Background mode. If dynamic categorization is disabled, proactive threat detection, content and reputation ratings are also disabled. The appliance contacts the WebPulse service using mutual endpoint authentication over TLS. For best security, the connection is always encrypted. If you opt to use a non-secure connection, all data is sent over the connection as plain text. For information, see "Configure WebPulse" on page 348.

About Dynamic Categorization

The dynamic categorization service analyzes and categorizes new or previously unknown URLs, which are not in the on-box Symantec WebFilter database. Dynamic categorization can be processed in two modes—immediately or in the background.

By default, dynamic categorization is set to be performed immediately, which is in real time. When a user requests a URL that has not already been categorized by the Symantec WebFilter database (for example, a new web site), the WebPulse dynamic categorization service queries the targetWeb site and retrieves the page's content. WebPulse analyzes the page's content and context in search of malicious content. If malicious content is found, an appropriate category (for example, Spyware/Malware sources or Phishing) is returned.

If no malicious content is found, WebPulse's dynamic real time rating service determines the language of the page, a category for the page, and a confidence factor that the category is correct. If the confidence factor is high, the calculated category is returned. In situations where the dynamic categorization service cannot categorize a URL with enough confidence to dynamically return a category with a high confidence level, the category rating request for the particular page is labeled none. All URLs received by WebPulse that are not categorized in the Symantec WebFilter database are logged and forwarded to Symantec's centralized processing center, where they are prioritized for rating by a series of automated URL analysis tools and/or human analysis. These ratings are then used to update the master Symantec WebFilter database, and the automatic database update feature then refreshes the local Symantec WebFilter database on the ProxySG appliance.

When dynamic categorization is performed in Background mode, the ProxySG appliance continues to service the URL request without waiting for a response from the WebPulse dynamic categorization service. The system category *pending* is assigned to the request, indicating that the policy was evaluated with potentially incomplete category information. When WebPulse

returns a category rating, the rating is stored in a dynamic categorization cache so that the next time the URL is accessed, WebPulse will not be required to determine its category.

> **Note:** The dynamic service is consulted only when the installed Symantec WebFilter database does not contain authoritative category information for a requested URL. If the category returned by the WebPulse service is blocked by policy, the offending material does not re-enter the network.

The URL is first looked up in the local Symantec WebFilter database. The expected results are shown in the following table.

Found in Local WebFilter Database	Found in Rating Cache	Process Mode	Result / Description
Yes	Any	Any	The corresponding list of categories in the database is returned.
No	Yes	Any	The corresponding list of categories in the ratings cache is returned.
No	No	Dynamic Categorization disabled	**None**. No categories are available for the URL.
No	No	Dynamic Categorization in Background Mode	**Pending**. The ProxySG appliance continues to service the URL request without waiting for a response from WebPulse. If a response is received, it is added to the rating cache, so future requests for that same URL will have the appropriate list of categories returned immediately. Reference to the site is recorded for future categorization in the WebFilter database by automated background URL analysis or human analysis. If a response is not received in a timely manner, or the request results cannot be categorized, nothing is added to the rating cache. **Note:** It is possible that multiple requests for the same content can result in a Pending status if WebPulse has not completed processing the first request before subsequent requests for the same URL are received by the ProxySG appliance.

Found in Local WebFilter Database	Found in Rating Cache	Process Mode	Result / Description
No	No	Dynamic Categorization in Real-time Mode and categories returned by WebPulse	A request to categorize the URL is sent to WebPulse and the ProxySG appliance waits for a response. The response is added to the rating cache and also used as the list of categories for the current request.
No	No	Dynamic Categorization in Real-time Mode and categories not returned by WebPulse	**None**. Categories might not be returned because: ■ The ProxySG appliance did not get a response from the WebPulse service. ■ The WebPulse service was unable to retrieve the requested URL in a timely manner. ■ The WebPulse service cannot categorize the request with high confidence. References to all URLs requested in WebPulse are recorded for future categorization in WebFilter by automated background analysis or human analysis. **Note:** Timeout is currently set to three seconds. Average response time for WebPulse to retrieve the content and perform real-time analysis is under 500 milliseconds.
Any	Any	Any	**Unlicensed**. A problem exists with the WebFilter license.
Any	Any	Any	**Unavailable**. A problem (other than licensing) exists with the local WebFilter database or accessing the WebPulse service.

About the Dynamic Categorization Process

Dynamic analysis of content is performed through theWebPulse cloud service and not locally on the ProxySG appliance. There is a small amount of bandwidth used for the round-trip request and response, and a slight amount of time waiting for the service to provide results. As the service is only consulted for URLs that cannot be locally categorized using the Symantec WebFilter database and WebPulse results are cached on the appliance, the user experience is generally not affected.

To avoid per-request latency, you might want to run dynamic categorization in background mode. For modifying the default, see "Configure WebPulse" on page 348.

The following diagram illustrates Symantec WebFilter's content filtering flow when dynamic categorization is enabled.

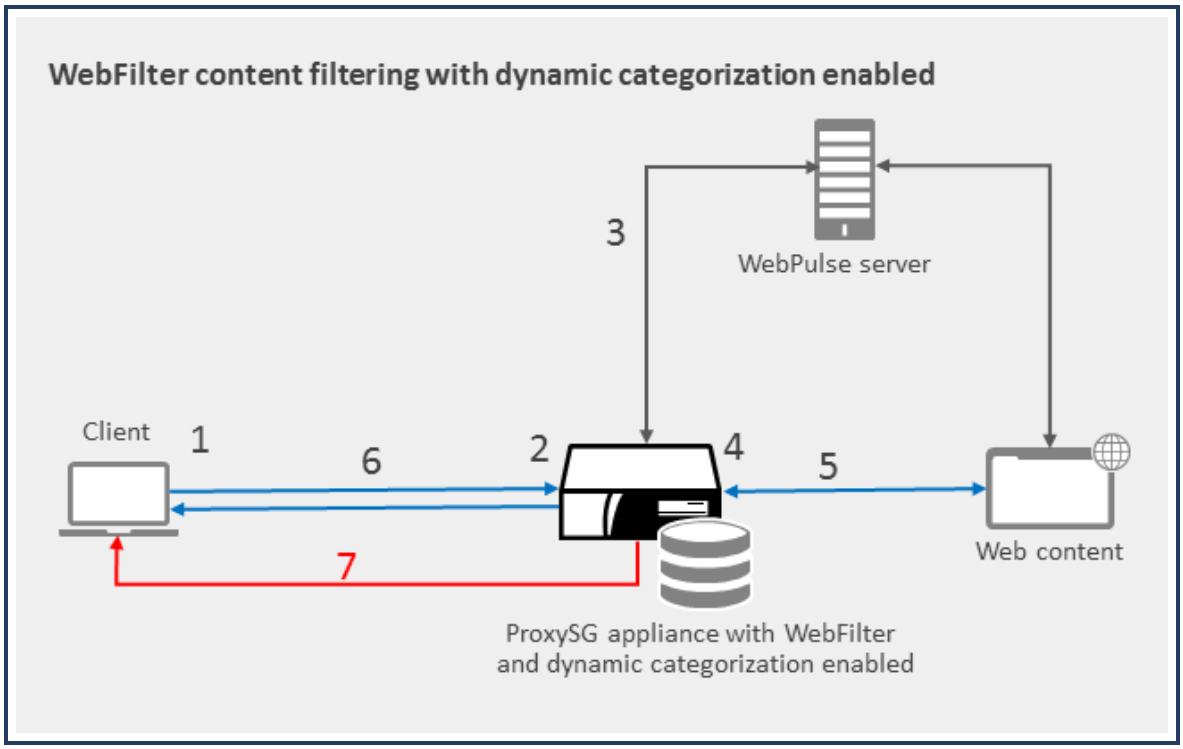

Clear the WebPulse Cache

Use the following CLI command:

```
# (config bluecoat) service clear-cache
```

Dynamic Categorization States

Dynamic categorization has three states:

- Enabled: The service attempts to categorize unrated web sites. This is the default state.
- Disabled: If the service is disabled, the ProxySG appliance does not contact the WebPulse service, regardless of any policy that might be installed.
- Suspended: Categorization from the database continues, but the service is no longer employed. This occurs when the installed database is over 30 days old due to the expiration ofWebFilter download credentials or network problems. After credentials are renewed or network problems are resolved, the service returns to Enabled.

Supported Content Filter Providers

Before you begin setting up content filtering, ensure that you have a valid subscription from a content filter provider of your choice. Only YouTube and the local database do not require a subscription or license. Content filtering configuration on the appliance consists of:

- Enabling the content filter provider
- Downloading the content filter database
- Applying policy

The following content filter providers are supported:

- " Local Database " below
- "Intelligence Services and WebFilter" on the next page
- "Categories for YouTube" on the next page
- "Third-Party Content Providers" on page 335

See Also:

- "Downloading the Content Filter Database" on page 335

Local Database

A local database is a text file that consists of `define category` CPL definitions. Defining categories in a local database instead of a policy file is useful because:

- A local database is more efficient than policy if you have a large number of URLs
- A local database separates administration of categories from policy. This separation is useful for three reasons:
    - It allows different individuals or groups to be responsible for administering the local database and policy.
    - It keeps the policy file from getting cluttered.
    - It allows the local database to share categories across multiple boxes that have different policy.

However, some restrictions apply to a local database that do not apply to policy definitions:

- No more than 200 separate categories are allowed.
- Category names must be 32 characters or less.
- A given URL pattern can appear in no more than four category definitions.
- The local database produces only the most specific URL match and returns a single category.

The same policy syntax will produce a different match. If more that one category is provided, policy processing may match more than one category and hence will return more than one category. See "Local Database Matching Example" on page 338 for more information.

You can use any combination of the local database, policy files, or the VPM to manage your category definitions. You can also use both a local database and a third party vendor for your content filtering needs.

(SGOS 6.7.4.107 and later) The appliance supports up to seven additional local databases, which you configure in the CLI. Refer to the `#(config local database_name)` commands in the *Command Line Interface Reference* for details.

To configure a local database, see "Configure a Local Database" on page 336.

> **Tip:** Symantec recommends putting your local database on the same server as any policy files you are using.

Intelligence Services and WebFilter

Symantec provides two on-box content filtering data sources:

- Intelligence Services: Subscription service that delivers existing and new data services such as Threat Risk Levels. To configure Intelligence Services, see "Intelligence Services" on page 339.
- WebFilter: Content filtering database that protects data and users from network attacks. To configure WebFilter, see "Blue Coat WebFilter" on page 339.

> **Tip:** The WebFilter data source will be removed in a future release of SGOS 7.x.

You can't use both data sources at once; Intelligence Services is recommended for its advanced capabilities.

See Also:

- "Configure Threat Risk Levels" on page 352

Categories for YouTube

Before enabling categories for YouTube, you require an API server key. For instructions, see "Configure Categories for YouTube" on page 342.

Note: This feature is provided on an "as-is" basis. Symantec has no control of, and is not responsible for, information and content provided (or not) by YouTube. Customer is required to apply and use its own API key in order to activate this feature, and therefore obligated to comply with all terms of use regarding the foregoing (for example, see https://developers.google.com/youtube/terms), including quotas, restrictions and limits on use that may be imposed by YouTube. Symantec shall not be liable for any change, discontinuance, availability or functionality of the features described herein.

Third-Party Content Providers

Currently, the ProxySG appliance supports the Proventia and Optenet third-party providers. These content filtering services are configured through the CLI only. Refer to the *Command Line Interface Reference* for details.

Downloading the Content Filter Database

Note: To download the content filter database, use the Management Console. Refer to the *SGOS Administration Guide* for details.

The dynamic nature of the Internet makes it impossible to categorize web content in a static database. With the constant flow of new URLs, URLs of lesser-known sites, and updated Web content, maintaining a current database presents a challenge. To counter this challenge, the ProxySG appliance supports frequent content filter database downloads.

For more information about the Symantec WebFilter database, see "About Symantec WebFilter and the WebPulse Service" on page 329.

About Database Updates

Symantec enables all customers with a valid content filtering license to schedule automatic downloads of content filter databases. By default, automatic updates are enabled; The ProxySG appliance checks for updates once in every five minutes and downloads an incremental update when available.

After selecting your provider(s) of choice, you must enter the license information and download the database(s) on the ProxySG appliance. You can download a database on demand or schedule a periodic download using the automatic download feature.

Typically, a complete database download occurs when you enable the provider and add the license key for the first time. Thereafter, the ProxySG appliance periodically checks the download server for updates to the installed database. If the database is current, no download is performed.

When an update is available, it is automatically downloaded and applied. An update provides the most current categorization of URLs and contains only the changes between the current installed version and the latest published version of the database,

and hence is much smaller than a full copy of the database. In the unlikely event that this conditional download fails, the ProxySG appliance downloads the latest published version of the complete database.

> **Note:** By default, the ProxySG appliance checks for database updates once in every five minutes. While you can schedule the time interval for an automatic database update, the frequency of checks is not configurable.

Configure a Local Database

The local database is a text file that must be located on a web server that is accessible by the ProxySG appliance on which you want it configured. You can't upload the local database from a local file. The local database file allows `define category`definition blocks only. Refer to the *Content Policy Language Reference* for details on using this policy gesture.

To configure a local database:

1. Create a text file in the following format:

   ```
   define category category_name
     url1
     url2
     urln
   end

   define category category_name
     url1
     url2
     urln
   end
   ```

   Each category can have an unlimited number of URLs.

2. Upload the text file to a web server that the ProxySG appliance can access.
3. In the Admin Console, select **Administration > Data Services > Content Filtering**.
4. In the Providers section, select **Local Database (default)** to configure the default local database.

   (In SGOS 6.7.4. and later) To configure additional local databases, select **Add Local Database**. You can configure up to seven additional local databases. The browser displays the Add/Edit Local Database page.

5. In the Settings section, enter or edit the following settings as needed:

Setting or field	Description
Enable Local Database	Enables or disables the local database. **Note:** When you enable the service for the first time and apply changes, a database download is initiated.
Lookup Mode	(Optional) For a web request, the lookup mode determines the databases that the ProxySG appliance searches for a category match. To perform a lookup, the database must be enabled. The lookup sequence executed is policy, local database, IWF, Blue Coat content filter, and third-party database. ■ **Always**: (Default) Database is always consulted for category information. If a URL is categorized under more than one category in different databases, policy is checked against each category listed. ■ **Uncategorized**: Database lookup is skipped if the URL match is found in policy, Blue Coat content filter, or IWF database.
Show Local Database Source	Display the contents of the local database file.

6. In the Download section, enter or edit the following settings as needed:

Setting or field	Description
URL	The database download location.
Username and **Password**	If the database is located on a server that requires a password for access, you must configure the appliance to use that password when accessing the database:
Automatically check for updates	By default, the ProxySG appliance checks for database updates once in every five minutes. To schedule the time interval for an automatic database update, select **between** and enter the start and end times of the interval. You can't specify the frequency of checks.
Download Now	If needed, select this to download the database immediately. The Download Status field shows the download progress.
Cancel Download	(Available while a download is in progress) To stop any download of the content filtering database that is currently in progress (including a download initiated from the CLI), click Cancel Download. The console displays a “Canceling download” dialog. When the download is canceled, the dialog message changes to “Download Canceled”.

Setting or field	Description
Refresh Status	(Available while a download is in progress) When the database is downloaded, the download log includes detailed information on the database. If you have just configured content filtering and are downloading the database for the first time, the ProxySG appliance downloads the latest published version of the complete database. Subsequent database updates occur incrementally.

Note: Incremental updates are not available for local databases.

Local Database Matching Example

The local database produces only the most specific URL match and returns a single category. Consider the following examples.

Local Database

Consider the following syntax.

```
define category no_detect_protocol
  mail.google.com
end

define category google
  google.com
end
```

Local database result:

```
https://<proxy>:8082/ContentFilter/TestUrl/mail.google.com/
  Local: no_detect_protocol
  Blue Coat: Mail
Gmail
none
```

Policy

This example uses the same syntax as the local database example.

```
<proxy>
  ALLOW

define category no_detect_protocol
  mail.google.com
end
```

```
define category google
  google.com
end
```

Policy Result:

```
https://<proxy>:8082/ContentFilter/TestUrl/mail.google.com/
  Policy: no_detect_protocol; google
  Blue Coat: Mail, Search Engine
Gmail
none
```

As shown, policy returns both categories; whereas, the local database returns only the URL match.

Configure Blue Coat Content Filtering

You can use one of two Blue Coat content filtering services for content filtering and Application Classification: Blue Coat WebFilter or Intelligence Services.

To configure a Blue Coat content filter, you must select the appropriate "Configure the Blue Coat Data Source" on the next page.

Blue Coat WebFilter

Blue Coat WebFilter (BCWF) is an on-box content filtering database that protects data and users from network attacks. All Symantec WebFilter subscribers are a part of the WebPulse cloud service, which continuously updates the on-box database. Symantec WebFilter and WebPulse provide dynamic real-time rating, which can instantly categorize Web sites when a user attempts to access them.

BCWF requires a valid license, a downloaded database, and a constant connection to Symantec servers to maintain license validity and download regular database updates. If you enable the feature but do not have a valid license, you will be unable to download the BCWF database.

Intelligence Services

- (Applicable to SGOS 7.1.x and earlier and some upgrades to SGOS 7.2.x and later) If it isn't already set, make sure that Intelligence Services is the data source (see "Configure the Blue Coat Data Source" on the next page. Then, the appliance can perform categorization using the Intelligence Services database.

  In addition, you can download a new version of the database when required, and also check the feature's health status.

- Verify that your subscribed bundles are listed in the Admin Console. See "Verify Subscribed Intelligence Services Bundles" on page 378.

Intelligence Services requires a valid license for the bundles that include the data feeds you want to use, a downloaded database, and a constant connection to Symantec servers to maintain license validity and download regular database updates. If you enable the feature but do not have a valid license, you will be unable to download the Intelligence Services database.

Configure the Blue Coat Data Source

> **Note:** This procedure is applicable only if you have the option to select either WebFilter or Intelligence Services as a content filtering data source.

To configure the data source:

1. In the Admin Console, select **Administration > Data Services> Content Filtering**.
2. In the Providers section, select **Blue Coat**. The browser displays the Configure Blue Coat Content Filtering page.
3. In the Settings section, enter or edit the following settings as needed:

Setting or field	Description
Enable Blue Coat Filtering	Enables or disables the service. **Note:** When you enable the service for the first time and apply changes, a database download is initiated.
WebPulse Status	Select **Edit** to configure WebPulse. See step 3 in "Configure WebPulse" on page 348.
Lookup Mode	(Optional) For a web request, the lookup mode determines the databases that the ProxySG appliance searches for a category match. To perform a lookup, the database must be enabled. The lookup sequence executed is policy, local database, IWF, Blue Coat content filter, and third-party database. ■ **Always**: (Default) Database is always consulted for category information. If a URL is categorized under more than one category in different databases, policy is checked against each category listed. ■ **Uncategorized**: Database lookup is skipped if the URL match is found in policy, a local database, or IWF database.
Data Source	Select **WebFilter** or **Intelligence Services**. **Tip:** If a Symantec WebFilter username and password are configured on the appliance, but you save a configuration archive (or Symantec Management Center backs up the configuration) while the data source is set to Intelligence Services, the archive does not save the WebFilter username/ password. To archive the WebFilter username and password, switch the data source back to WebFilter and save a separate configuration file.

4. In the Download section, enter or edit the following settings as needed:

Setting or field	Description
URL	The default database download location. If you have been instructed to use a different URL, enter it here.
Username and **Password**	Enter valid subscription credentials to download the database. If the database has previously been downloaded on a local web server that requires authentication, you must configure the ProxySG appliance to use credentials that allow access to the Web server, which hosts the content filter database.
Automatically check for updates	By default, the ProxySG appliance checks for database updates once in every five minutes. To schedule the time interval for an automatic database update, select **between** and enter the start and end times of the interval. You can't specify the frequency of checks.
Download Now	If needed, select this to download the database immediately. The Download Status field shows the download progress. **Note:** When you enable the service for the first time, a database download is initiated.
Cancel Download	(Available while a download is in progress) To stop any download of the content filtering database that is currently in progress (including a download initiated from the CLI), click Cancel Download. The console displays a “Canceling download” dialog. When the download is canceled, the dialog message changes to “Download Canceled”.
Refresh Status	(Available while a download is in progress) When the database is downloaded, the download log includes detailed information on the database. If you have just configured content filtering and are downloading the database for the first time, the ProxySG appliance downloads the latest published version of the complete database. Subsequent database updates occur incrementally.

See "Downloading the Content Filter Database" on page 335 for information on database downloads.

5. (If needed) Enable the service. In the Settings section, select **Enable Blue Coat Filtering**.

License Validity and Database Expiration

A valid subscription is required to update the database. Each time a database download is triggered manually or using automatic download, the validity of the database is reset.

When your license expires, you can no longer download the latest version. The expiry of a database license does not have an immediate effect on performing category lookups for the on-box categories. You can continue to use the on-box database until the expiry of the database.

When the database expires, the category unlicensed is assigned to all URLs and no lookups occur on the database.

See Also:

- "About Symantec WebFilter and the WebPulse Service" on page 329

Configure Categories for YouTube

The appliance recognizes three types of YouTube URLs. The effectiveness of policy and coaching pages differs amongst the different URL types. Refer to the following table.

URL Type	Example(s) of User Action	Deny policy works	Coaching page works
YouTube homepage	User plays a video at www.youtube.com within a desktop web browser. The URL starts with: http://www.youtube.com/watch?v=	Yes	Yes
YouTube mobile website	User plays a video at www.youtube.com within a mobile web browser (the URL redirects to http://m.youtube.com/index). The URL starts with: http://m.youtube.com/watch?v=	Yes	Yes
Embedded in an <iframe>	User plays a video that has been embedded in a blog post. The URL could start with: http://www.youtube.com/embed/	Yes. The deny page is confined within the <iframe>.	Yes. The coaching page is confined within the <iframe>.
Video transport	User plays a YouTube playlist within a desktop web browser. The URL could start with: *.c.youtube.com/videoplayback?	Yes. The user may see an error message for the blocked video.	No

For information on implementing coaching pages, refer to the **Notify User** action in the *Web Visual Policy Manager Reference*.

The list of categories is static. In the Visual Policy Manager, you can view the categories in the category list (**Configuration > Edit Categories**) and in the **Request URL Category** object, but you can't add, rename, edit, or remove them.

Set the YouTube Server Key

Before you can enable YouTube as a provider, you must obtain an API server key and set it on the ProxySG appliance. For instructions, refer to the following article:

http://www.symantec.com/docs/TECH241321

After you set the server key, select **YouTube** in the Admin Console at **Configuration > Content Filtering > General**.

View YouTube Categories in the Access Log

If the feature is enabled and categories are selected, and if you use an ELFF format for access logs, you can use the existing category access-log field to report on specified categories; however, the categories will be logged without the provider name. In addition, some categories share names with Symantec categories, such as Entertainment.

To distinguish between Symantec-defined categories and YouTube categories in the access log, specify the ELFF field `cs-categories-qualified`. This field provides a list of all content categories of the request URL, qualified by the provider. For example, traffic matching YouTube's Entertainment category would be logged as `Entertainment@YouTube`.

For information on access log formats, see "About Access Log Formats" on page 402.

View and Verify Threat Protection Services

The appliance supports the following threat protection solutions:

- WebPulse: An off-box dynamic categorization service for real-time categorization of URLs that are not rated or categorized in the on-box database. See "About WebPulse" below.
- Threat Risk Levels: A service that uses WebPulse to perform real-time lookups to analyze the threat risk level of URLs. See "About Threat Risk Levels" below.

In the Admin Console, the **Administration > Data Services > Threat Protection** page provides an overview of threat protection services. From here, you can:

- Enable or disable a service by switching the **Enable** toggle.
- Verify successful database download in the **Last Download** field.
- Determine if any downloads failed in the **Failed Attempts** field; then, troubleshoot connection or licensing issues if needed.
- Configure service settings; see below for more information.

About WebPulse

WebPulse is a cloud service that provides the off-box component of Symantec's complete content filtering solution. The WebPulse cloud service blocks malware hosts, ratesWeb content and protects ProxySG appliance web gateways. For more information, see "About Symantec WebFilter and the WebPulse Service" on page 329.

See Also:

- "Configure WebPulse" on page 348

About Threat Risk Levels

The Threat Risk Levels service analyzes a requested URL's potential risk and summarizes it in the form of a numeric value (see "About WebPulse" above). You can reference these values in policy to protect your network and your users from potentially malicious web content. WebPulse is used for real-time Threat Risk Level lookups.

See Also:

- "Configure Threat Risk Levels" on page 352

Considerations Before Configuring WebPulse

The WebPulse protocol regulates the communication between the dynamic categorization client on the ProxySG appliance and the WebPulse cloud service.

Before configuring WebPulse services (see "Configure WebPulse" on page 348), answer the following questions:

- Do you use proxy chaining or SOCKS gateways?

  If you use a forwarding host for forwarding dynamic categorization requests through upstream proxies or use SOCKS gateways, see "About Proxy Chaining Support for WebPulse Services" below for more information on the forwarding options in WebPulse.

- Would you like to configure private networks to identify traffic relating to your internal networks while using WebPulse for content rating accuracy?

  If you specify your private networks, information pertaining to the configured internal network is removed by the ProxySG appliance, prior to sending a dynamic categorization request across to theWebPulse service. For understanding the interaction between private networks and dynamic categorization, see "About Private Information Sent to WebPulse" on the facing page.

- Would you like to provide malware feedback notification to the WebPulse community?

  This option is applicable only if you have a ProxyAV configured on the ProxySG appliance, and malware scanning and WebFilter are enabled. When the appliance is integrated with the ProxyAV, the appliance monitors the results of the ProxyAV scan and notifies the WebPulse service when a new virus or malware is found. This feedback helps update the malware and content ratings and protects the entire community of WebPulse users. For information on adding a ProxyAV and enabling malware scanning, see "Malicious Content Scanning Services " on page 418.

About Proxy Chaining Support for WebPulse Services

Proxy chaining is a method for routing client requests through a chain of ProxySG appliances until the requested information is either found in cache or is serviced by the OCS. The ProxySG allows you to forward dynamic categorization requests through upstream proxies and SOCKS gateways.

Forwarding Hosts and Dynamic Categorization

To forward dynamic categorization requests through an upstream HTTP proxy, configure a forwarding host that is defined as a proxy and specify the HTTP port for the connection. You can then select that forwarding host in the WebPulse configuration.

If forwarding is configured, you cannot enable secure dynamic categorization; if secure dynamic categorization is enabled, you cannot select a forwarding host.

SOCKS Gateways

If you use proxy chaining for load balancing or for forwarding the dynamic categorization request through an upstream SOCKS gateway, you must configure the SOCKS gateway before configuring the WebPulse service.

> **Note:** Before configuring the SOCKS gateway target for WebPulse, verify that the SOCKS gateway is operating correctly.

When both SOCKS and forwarding are configured, the ProxySG connects to the SOCKS gateway first, then to the forwarding host, and then to theWebPulse service.

About Private Information Sent to WebPulse

A private network is an internal network that uses private subnets and domains, for example, your intranet. For information on configuring private subnets or private domains, refer to the *SGOS Administration Guide*.

By default, dynamic categorization is enabled on the ProxySG appliance. When a requested URL is not in the dynamic categorization ratings cache or categorized in the WebFilter database, the request is sent to the WebPulse cloud service for dynamic categorization. By configuring private subnets and private domains within your network, you can use dynamic categorization to ensure accuracy of content filter ratings, while preserving the security of sensitive information relating to your private networks.

Before a request is sent for content rating to the WebPulse cloud service, the following conditions are verified on the appliance:

- Is WebPulse service and dynamic categorization enabled?
- Is dynamic categorization permitted by policy?
- Is the host specified in the private domain or private subnet list?

  Any request that is determined to be part of your configured private network is not sent to WebPulse.

The ProxySG appliance might send information from HTTP and HTTPS requests to the WebPulse service if they are not directed to hosts that are part of the configured private network.

> **Note:** Private network domain names and IP subnets can be user-defined.

Customer information sent to the WebPulse service is controlled by user-defined policy, although you can still use the default policy and configuration settings provided by the ProxySG appliance. Overriding the default settings with your organization's policy definitions results in more control of the type of information that is sent to the WebPulse service.

When WebPulse service is enabled, the default configuration settings send the fixed customer data and the following information:

- Customer License Key (Example: QA852-KL3RA)
- Scheme (Example: HTTP, HTTPS)
- Method (Examples: GET, POST)
- URL Host
- URL Port
- URL Path
- URL query string
- Referer header
- User-Agent header

However, this additional information can be controlled by policy and/or configuration settings:

- If the service send-request-info setting is set to disable, by default, only the customer license key, URL scheme, method, host, port, and path are sent to the WebPulse service; URL query string, and the Referer and User-Agent headers are not sent.
- If the service send-request-info setting is set to enable, by default, referrer and user agent information and so on can be sent to the server. Symantec gathers this customer information from back-end logs and analyzes the data to help improve its threat protection. In its analysis, Symantec does not consider the source of the data; that is, customer information is anonymous.

> **Note:** Be aware that personal information might be included in the URL query string. If this information is sent to WebPulse, Symantec might use it when accessing content from the web site to categorize it.

You can further control whether to include the URL path and query string, and individually control whether the Referer or User-Agent headers are sent for specific requests. Restrictions are accomplished through the use of policies that can be defined from the Admin Console or CLI.

The following table lists the type of information that is sent to the WebPulse service based on default settings for all SGOS versions supporting WebPulse.

Information Sent to the WebPulse Service	service send request- info disable	service send request- info enable
Customer License Key (Example: QA852-KL3RA)	Yes	Yes
Scheme (Examples: HTTP, HTTPS)	Yes	Yes
Method (Examples: GET, POST)	Yes	Yes
URL Host/Port	Yes	Yes

Information Sent to the WebPulse Service	service send request- info disable	service send request- info enable
URL Path[1]	Yes	Yes[2]
URL Query String	No	Yes[2]
Referer Header	No	Yes[2]
User-Agent Header	No	Yes[2]
Content-Type	No	Yes[2]
Content-Length	No	Yes[2]

1. Path = URL minus any query string.
2. Can be controlled using CPL.

Configure WebPulse

To configure the WebPulse service:

1. In the Admin Console, select **Administration > Data Services > Threat Protection**.
2. In Threat Protection, select **WebPulse**.

   The WebPulse page displays service settings.
3. Enter or edit the following settings as needed:

Task	Setting or field	Description
Check or change the status of the service, the Blue Coat content filter, and download	**Enable WebPulse Service**	Switch the toggle to enable or disable the service. If you disable the WebPulse service, dynamic categorization and malware feedback are also disabled. References to perform dynamic categorization in policy are also disregarded. For more information, see "About Symantec WebFilter and the WebPulse Service" on page 329.
	Blue Coat Content Filtering	Verify that Blue Coat is enabled as your content filter provider. If disabled, switch the toggle to enable the setting. When this option is disabled, the WebPulse option is unavailable.
	Last Download Status	Confirm that the last download succeeded. The field displays statistics about the last download, whether or not it succeeded. If you receive a download error, check your network configuration and make sure that the appliance can connect to the Internet or URL to the database.
Change WebPulse protocol settings	**Forwarding Target**	(Optional) Select a forwarding host.
	SOCKS gateway target	(Optional) Select a SOCKS gateway target.

Task	Setting or field	Description
Configure dynamic analysis	Select one of: ■ **Real-time** ■ **Background** ■ **Disable** If Symantec WebFilter license has expired and dynamic categorization is enabled, the service enters a suspended state. For more information, see "Dynamic Categorization States" on page 332.	Specifies a categorization mode or disables dynamic categorization. If enabled, verify that Blue Coat content filtering is enabled. ■ **Real-time**: (Default) If the category of the request is not already known, the URL request will wait for the WebPulse service to respond with the categorization before proceeding. The advantage of real-time mode categorization is that Symantec policy has access to the results, allowing policy decisions to be made immediately after receiving all available information. ■ **Background**: When dynamic categorization is triggered, the URL request continues to be serviced without waiting for a response from the WebPulse service. The system category "pending" is assigned to the request, indicating that the policy was evaluated with potentially incomplete category information. The result of the categorization response is entered into a categorization cache. This cache ensures that any subsequent requests for the same or similar URLs can be categorized quickly, without needing to query the WebPulse cloud service again. ■ **Disable**: Disable dynamic categorization. The appliance does not contact the WebPulse service when a category match for a URL is not found in the on-box database.
Enable or disable sending malware feedback	**Send potential malware sources to WebPulse**	Switch the toggle to enable or disable sending malware feedback to WebPulse. If the appliance is integrated with the ProxyAV or Content Analysis for ICAP scanning while Blue Coat content filtering and WebPulse are enabled, and a malware threat is detected, the appliance issues a malware notification to WebPulse to update the WebFilter database. When this option is disabled, the ProxySG appliance does not notify WebPulse about malware URLs that Content Analysis or WebFilter detects; however, you can use policy to override the default malware feedback settings.

Analyze the Threat Risk of a URL

The Threat Risk Levels service analyzes a requested URL's potential risk and summarizes it in the form of a numeric value (see "About the Threat Risk Levels" below). You can reference these values in policy to protect your network and your users from potentially malicious web content. Threat Risk Levels are calculated based on numerous factors that measure current site behavior, site history, and potential of future malicious activities. This service is part of Intelligence Services. For information on Intelligence Services, refer to the "Content Filtering Overview" on page 324.

How the Threat Risk Service Works with WebPulse

To have the Threat Risk Levels feature return both risk levels and category information for requests, the ProxySG appliance must have valid Intelligence Services and Threat Risk Levels licenses. Although it is not required, Symantec recommends that you also enable the WebPulse categorization service on the appliance. Enabling WebPulse ensures that a risk level is returned for every user request, provided there is no connection problem. If the risk level for a request does not exist in the on-box database, the appliance queries WebPulse for information and forwards the request to WebPulse for risk level analysis and categorization. The appliance then caches the information that it retrieves from WebPulse.

See Also:

- "About Symantec WebFilter and the WebPulse Service" on page 329

About the Threat Risk Levels

The Threat Risk Levels service assigns threat risk levels to URLs according to specific criteria:

Level	Description
Low (Levels 1-2)	The URL has an established history of normal behavior and has no future predictors of threats; however, this level should be evaluated by other layers of defense (such as Content Analysis and Malware Analysis).
Medium-Low (Levels 3-4)	The URL has an established history of normal behavior, but is less established than URLs in the Low group. This level should be evaluated by other layers of defense (such as Content Analysis and Malware Analysis).
Medium (Levels 5-6)	The URL is unproven; there is not an established history of normal behavior. This level should be evaluated by other layers of defense (such as Content Analysis and Malware Analysis) and considered for more restrictive policy.
Medium-High (Levels 7-9)	The URL is suspicious; there is an elevated risk. Symantec recommends blocking at this level.
High (Level 10)	The URL is confirmed to be malicious. Symantec recommends blocking at this level.

The service returns a system-defined value in the rare instance when a threat risk level does not apply:

CPL Use for `url.threat_risk.level=`	Use in Threat Risk Level objects	Description
none	Threat Risk Level not available	The URL does not have a threat risk level assigned to it and the request is not forwarded to WebPulse (or it has been forwarded and there is still no data).
pending	Threat Risk Level data pending background WebPulse analysis	No risk level is assigned to the URL on the appliance, but the appliance performed a WebPulse request in the background. Subsequent requests for the URL could match a result from the WebPulse cache.
unavailable	Threat Risk Level database or service not accessible	A non-licensing problem exists with the WebFilter database or with accessing the WebPulse service.
unlicensed	Threat Risk Level feature not licensed	The Threat Risk Levels license is not valid.

See Also:

- "Configure Threat Risk Levels" on the next page

Configure Threat Risk Levels

Before using Threat Risk Levels, make sure you have the following:

- A constant connection to Symantec servers to maintain license validity and download periodic database updates.
- A valid license for the solution bundles that include Threat Risk Levels. Refer to your Symantec point of contact for more information.

If you enable the feature but do not have a valid license:

- You cannot download the database.
- Any existing threat risk policy will not work.

After verifying that you have a license, you can make sure that you have the Threat Risk Levels service. See "Verify Subscribed Intelligence Services Bundles" on page 378.

Enable Threat Risk Levels

Before you can use Threat Risk Levels or related policy, you must enable the service on the appliance.

1. In the Admin Console, select **Administration > Data Services > Threat Protection**.
2. In Threat Protection, select **Threat Risk Levels**.

   The Threat Risk Levels page displays service settings.
3. In Threat Risk Levels, switch the Enable Threat Risk Levels toggle.
4. Make sure that WebPulse and dynamic categorization are both enabled. If the console displays a message "WebPulse Status: Enabled, Dynamic Categorization Status: Enabled", you don't have to make any more changes here.

   Otherwise, select **Edit WebPulse Settings** to enable WebPulse and Dynamic Categorization. For instructions, see "Configure WebPulse" on page 348.

Check Threat Risk Levels Database and License Status

When you enable the Threat Risk Levels service, the appliance starts to download the database in the background. The Admin Console displays the download in progress; wait for the download to complete before attempting to use the feature. The License and Download Status field displays statistics when download is complete.

1. In the Admin Console, select **Administration > Data Services > Threat Protection**.
2. In Threat Protection, select **Threat Risk Levels**.

3. In License and Download Status, the Status field displays statistics about the previous successful and unsuccessful downloads. If download was successful, the field shows the database version amongst other information. If the last download was unsuccessful, the field contains an error.
4. (If needed) Select one of the following
   - **Download Now** to force a database download.
   - **Refresh Status** to view the status of the download in progress.
   - **Cancel Download** to cancel a download in progress (including a download initiated from the CLI).

Look up the Threat Risk Level of a URL

Determine the threat risk of a specific web page.

1. In the Admin Console, select **Administration > Data Services > Threat Protection**.
2. In Threat Protection, select **Threat Risk Levels**.
3. In **Check Risk Level Match**, enter a valid URL in the URL to Test field and select **Test URL**.

   The console displays the threat risk level associated with the URL, as described in "Analyze the Threat Risk of a URL" on page 350.

Use Intelligence Services to Classify Applications

The Application Classification service uses the Intelligence Services database.

Before you can use the feature:

- Make sure that Intelligence Services is set as the data source. See "Configure the Blue Coat Data Source" on page 340.
- Ensure that your license is valid.

After enabling the feature, you can review:

- Applications and operations for a URL
- Application groups for a URL

> **Note:** Some application groups and content filter categories have similar names, such as “Social Media”, and you can use both groups and categories in policy; however, they are used for different functions and each should be controlled using appropriate VPM objects

> or policy gestures.

- Attributes associated with a web application

In addition, you can download new database versions when required, and also check the feature's health status.

Enable Application Classification

To enable application classification:

1. In the Admin Console, select **Administration > Data Services > Application Classification**.

   The browser displays the Application Classification page.

2. In the Application Classification section, select **Enable Application Classification**.

   Before you can use Intelligence Services to classify applications or related policy, you must enable the feature on the appliance.

   You can use either WebFilter or the Intelligence Services categorization data feed for content filtering. When Blue Coat content filtering is enabled (**Administration > Data Services > Content Filtering**), the appliance uses whatever data source is selected as the non-third-party on-box content filtering database for Application Classification.

> **Note:** When you enable the service for the first time and apply changes, a database download is initiated.

> **Tip:** To disable this service, make sure that Application Attributes is disabled. See "Review Application Attributes" on page 356.

Look up the Application, Group(s), and Operation for a URL

Use the Application Classification feature to review the applications and operations for a web page. Application groups allow you to apply policy actions to groups of similar applications.

To look up web applications, groups, and operations for a URL:

1. In the Admin Console, select **Administration > Data Services > Application Classification**.

   The browser displays the Application Classification page.

2. In the Application Classification section:

- To look up the application, group(s), and operation for a URL, enter a URL and select **Check URL**. The console displays the lookup results. Refer to "Classification Lookup Results" below to determine what the messages mean.

  You can also use WebFilter to review the applications and operations (but not application groups) for a URL. See "Test the Category, Application, and Operation for a URL" on page 327.

- (Optional) To view the list of supported applications, groups, or operations in the downloaded database, select **All Applications**, **All Groups**, or **All Operations** respectively. The list appears in a new browser window. If the page says "unavailable", the database has not been downloaded.

Classification Lookup Results

Message Text	Meaning
Application: <*application_name*>	The URL is associated with the specified application. To obtain more detailed information about the application, see "Review Application Attributes" on the next page.
Application: none	The URL is not associated with any application.
Group: <*group_name*>	The URL is associated with the specified application group(s).
Group: none	The URL is not associated with any defined application group.
Operation: <*operation_name*>	The URL is associated with the specified operation.
Operation: none	The URL is not associated with any operation.

Download the Application Classification Database

When you enable the Application Classification service, the appliance starts to download the database in the background. The Admin Console displays the download in progress; wait for the download to complete before attempting to use the feature. The console displays statistics when download is complete.

To download the Application Classification database:

1. In the Admin Console, select **Administration > Application Classification**.
2. In the Download Application Classification section, select **Download Now**. The console shows a download is in progress.

To cancel any download of the Application Classification database that is currently in progress (including a download initiated from the CLI), select **Cancel Download**. The License and Download Status indicates that download was canceled by the user.

Review Application Attributes

You can find out more information about a web application by looking at the attributes and their values. Attributes can provide insight into a web application and its governance, risk management, and compliance.

To use the Application Attributes feature, you must have a valid Application Classification subscription, included in an Intelligence Services bundle, that includes attributes.

Enable Application Attributes

Before enabling Application Attributes, verify that Application Classification is enabled.

To enable application attributes:

1. In the Admin Console, select **Administration > Data Services > Application Classification**.

   The browser displays the Application Classification page.

2. In the Application Attributes section, select **Enable Application Attributes**.

> **Tip:** You must disable Application Attributes before attempting to disable Application Classification.

Download the Application Attributes Database

> **Note:** Application Classification and Application Attributes share the same database.

To download the Application Attributes database:

1. In the Admin Console, select **Administration > Application Classification**.
2. In the Download Application Attributes section, select **Download Now**. The console shows a download is in progress.

To cancel any download of the Application Attributes database that is currently in progress (including a download initiated from the CLI), select **Cancel Download**. The License and Download Status indicates that download was canceled by the user.

To display current download statistics in License and Download Status, select **Refresh Status**. The License and Download Status indicates that download was canceled by the user.

See Also:

- "Configure Blue Coat Content Filtering" on page 339

About Geolocation

To comply with local regulations, assist with traffic analysis, or reduce the risk of fraud and other security issues, you may need to know the origin of traffic in your network, or restrict outbound connections to specific countries. SGOS

supports geolocation in both reverse proxy and forward proxy modes:

- Geolocation in forward proxy mode: Restrict by country the outbound connections that your users make, based on the geographic region of URLs.
- Client geolocation in reverse proxy mode: Identify the source of traffic through the ProxySG appliance based on IP address (and when applicable, the effective client IP address; refer to the Visual Policy Manager Reference for information on effective client IP address). Both IPv4 and IPv6 are supported.

Before using geolocation, you must download a database that maps IPv4 and IPv6 addresses to the countries with which they are associated. It also provides the supported names and ISO codes for countries, such as “United States [US]”. You can use country names or codes in policy to perform actions such as denying traffic to or from specific countries.

Prerequisites for Using Geolocation

Before you can use geolocation:

- Verify that you have a valid license for the feature. In the Management Console, select **Maintenance > Licensing > View** and look for license details in the Intelligence Service Bundles section.

  If you do not have a valid license, the appliance is unable to download the database and the Management Console might display Health Monitoring errors. The access logs might also display error messages about the subscription. Review "Troubleshoot Geolocation Errors" on page 364 for more information.

- Enable the geolocation service. See "Enable Geolocation" on the next page.
- Download and verify the geolocation database. See "Download Database Updates" on the next page.

Using Geolocation

Before you can use geolocation features, you must enable the service on the appliance. When the database is downloaded successfully, you can use inbound and outbound geolocation features and write policy using the country name or country code as defined in the geolocation database. You will also be able to see the supported country names and codes:

- In the Management Console (**Administration > Data Services > Geolocation**).
- In output for the `#show geolocation countries` and `#(config geolocation) view countries` CLI commands.
- When you add geolocation objects in the Visual Policy Manager (VPM).

Enable Geolocation

Enable geolocation:

1. In the Admin Console, select **Administration > Data Services > Geolocation**.
2. In the Geolocation section, select **Enable Geolocation**.

When you enable the geolocation service, the appliance starts to download the database in the background. The Admin Console displays the download in progress; wait for the download to complete before attempting to use the feature. The **License and download status** field on the Geolocation page displays statistics when download is complete.

What if the Initial Download is Not Successful?

If you receive a download error and the appliance health status is Critical shortly after you attempt to download the database, the download might have failed. To confirm if this is the case, select Statistics > Health Monitoring > Subscription in the Management Console and look for a "failed on initial download" error for Geolocation Communication Status.

> **Note:** The Critical error appears if the initial download attempt fails. After the database downloads successfully, the service periodically checks for a newer version of the database. If several update checks fail to connect, a Warning error appears in Health Monitoring until the failure is corrected.

Download Database Updates

To download database updates:

1. In the Admin Console, select **Administration > Data Services > Geolocation**.
2. In the Download Database section, select **Download Now**.

Verify the Database Download

To verify that the database was downloaded, in the Geolocation Lookup section, select Show list of countries in Geolocation database to display the countries in a new browser tab.

In addition, the **License and download status** field in the Download Database section shows statistics about the previous successful and unsuccessful downloads. If the last download was unsuccessful, the field contains an error.

If you receive a download error, check your network configuration and make sure that the appliance can connect to the Internet or URL to the database.

Cancel a Database Download in Progress

To stop any download of the Geolocation database that is currently in progress (including a download initiated from the CLI), select **Cancel Download** in the Download Database section. The Admin Console displays a “Canceling download” dialog. When the download is canceled, the dialog message changes to “Download Canceled”.

Use Geolocation Lookup

You can use geolocation lookup to test inbound and outbound connections. For example:

- In forward proxy, you can write policy that allows or denies outbound connections to specific geographic locations associated with a given URL. See "Using Geolocation " on the previous page .
- In reverse proxy, you can identify the country that is associated with the origin of a client connection and use the country code or name in policy. See "Determine Locations of IP Addresses for Incoming Connections" on page 362.

Test Outbound Connections Based on Geographic Location

If you intend to write policy that allows or denies outbound connections to specific geographic locations, you can test what policy decisions would be made for a given URL based on location criteria.

To test a URL, select at least one country and specify whether to allow or deny connections to the selected region(s). The appliance determines the IPv4 and IPv6 address (or addresses) to which the URL resolves and returns the associated country for each address. The test results show whether each IP address is allowed or denied based on your criteria.

> **Note:** For better security, the allow/deny determination is made during DNS resolution, before a connection is attempted.

To test URLs, ensure that:

- The geolocation service is enabled; see "Using Geolocation " on the previous page.
- The geolocation database is downloaded to the appliance; see "Using Geolocation " on the previous page.
- Primary and alternate DNS servers are configured; see "Configure DNS" on page 112 .

Test outbound connections based on geolocation:

1. Specify an IP address or URL to test.
    a. In the Admin Console, select **Administration > Data Services > Geolocation**.
    b. In the Geolocation Lookup section, enter an IP address of URL. URLs must be in the form of a fully-qualified domain name (FQDN).
2. Select **Evaluate Geolocation Result Against A Test Policy Rule**. The Geolocation Condition and Policy Action options appear.
3. Select the countries to allow or deny in policy. In the Geolocation Condition dialog that opens, select countries. Geolocation Condition displays the selected locations.
4. Specify whether to allow or deny connections to the regions. For Policy Actions, select one of the following:
    a. **Deny selected geolocation**: Denies connections to the selected locations.
    b. **Allow selected geolocation**: Allows connections to the selected locations.

    If the DNS lookup returns countries that are among the ones you selected in step 2, the selected verdict (Deny/Allow) is applied. If the lookup returns countries that are not in the selected list, the opposite verdict is applied.
5. Select **Lookup Geolocation**. The Geolocation Results section shows:
    - IP Address: The resolved IP addresses for the URL.
    - Geolocation: The geographical location of each IP address and its two-letter country code.

    > **Note:** IP address mappings to locations change over time, and periodic geolocation database updates reflect these changes. A changed IP address mapping can cause the following behavior:
    > - DNS lookup results are different from a previous lookup
    > - geolocation policy is no longer working as expected
    >
    > Use current lookup results to update policy as appropriate.

    - Policy Action: Whether policy would deny or allow the connection to the IP address.

Example: Test Policy Based on Geolocation

Given geolocation policy that allows connections only to IP addresses in the United States, you want to determine if any connections to web pages under torproject.org would be allowed.

1. In the Admin Console, select **Administration > Data Services > Geolocation**.
2. In the IP Address/URL field, type “torproject.org”.

3. Select **Evaluate Geolocation Result Against A Test Policy Rule**. The Geolocation Condition and Policy Action options appear.
4. **Select** Germany in the Geolocation Condition dialog.
5. For Policy Actions, select **Allow selected geolocation**.
6. Select **Lookup Geolocation**. The Geolocation Results section shows the resolved IP addresses for torproject.org, the geographical location of each IP address, and whether policy would deny or allow the connection.

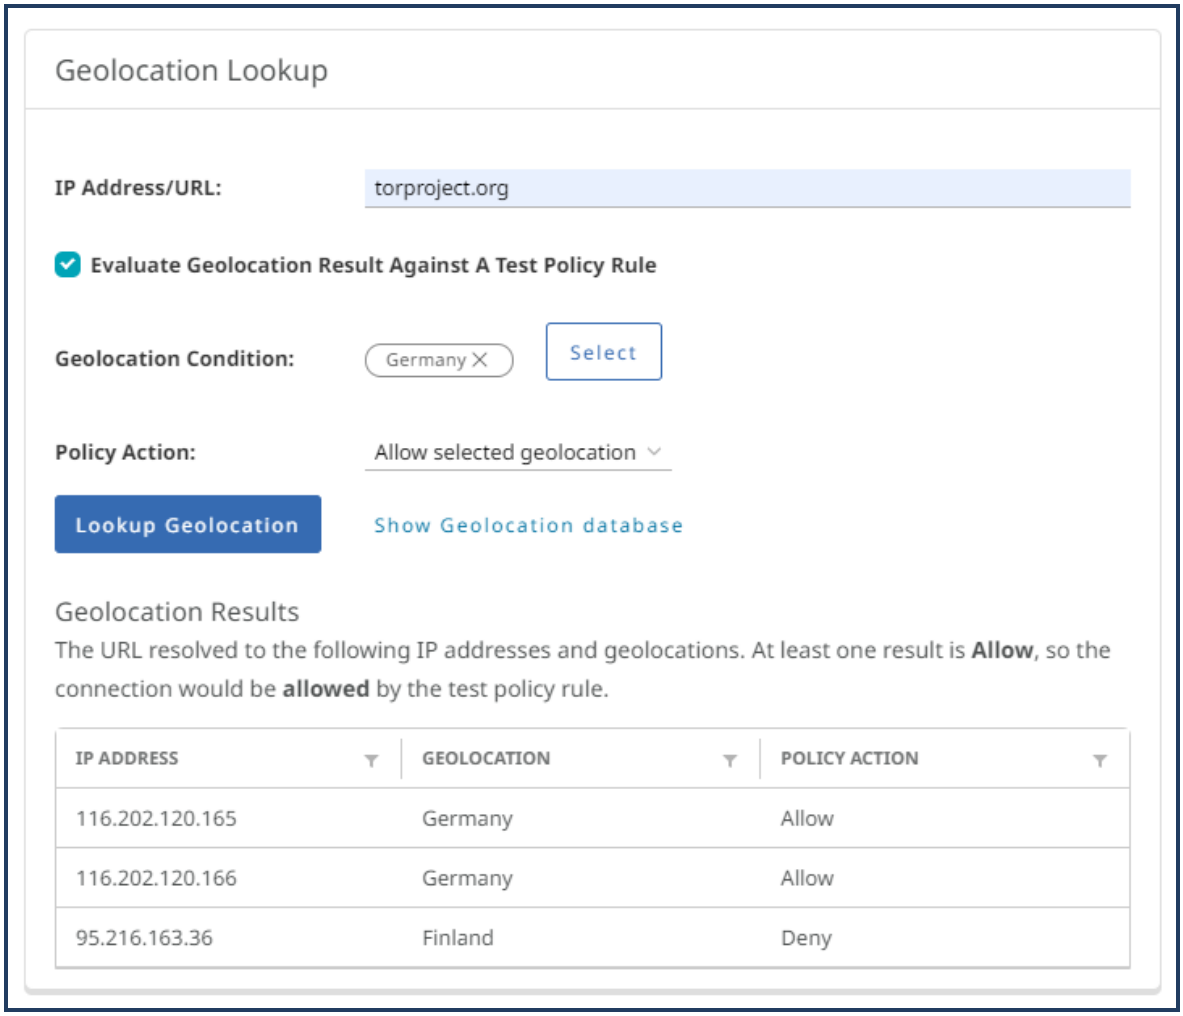

Note: If you perform the lookup in this example and receive different location results, the discrepancy might be due to IP address mappings changing over time.

If any region's verdict is Allowed, the connection will be allowed. The first server in the list that is in a permitted region and which responds to the connection request will be used for the connection.

For example, in the example above, IP addresses resolve to Germany and Finland. The connection permission for each IP address in the Germany is Allow, and all others are Deny. The IP address 116.202.120.165 in Germany would be used for connections to torproject.org, unless subsequent policy rules result in a Deny.

Forward Proxy Use Cases: Write Geolocation Policy

Refer to the following examples of writing geolocation policy for outgoing connections. For detailed usage information on policy gestures, refer to the Content Policy Language Reference.

Use Case 1

You require policy that scans files from IP addresses in Russia. You can use the following CPL:

```
; scan connections to Russia using the specified ICAP service
<cache>
supplier.country=RU response.icap_service(AV_scan)
```

Use the `supplier.country=` condition for policy decisions based on the response payload. In this case, policy evaluation occurs after a connection is made.

In the VPM, use the **Resolved Country Destination** object for the `supplier.country=` condition. Refer to the legacy or web *Visual Policy Manager Reference* for information.

Use Case 2

You require policy that denies connections based on both geolocation and type of requested URL. You can use the following CPL:

```
; 1. allow all countries, except the following override
; 2. deny connections to China when a batch file is requested
<proxy>
supplier.allowed_countries(all)
<proxy>
url.extension=BAT supplier.allowed_countries[CN] (deny)
```

Use the `supplier.allowed_countries()` action when policy evaluation must occur before a connection is made.

In the VPM, use the **Set Geolocation Restriction** Action object for the `supplier.allowed_countries()` property. Refer to the legacy or web *Visual Policy Manager Reference* for information.

Determine Locations of IP Addresses for Incoming Connections

After you download a geolocation database, you can identify the country that is associated with a given IP address. You can then use the country code or country name in policy.

Identify a country by IP address:

1. In the Admin Console, select **Administration > Data Services > Geolocation**.
2. In the Geolocation Lookup section, in the **IP Address/URL** field, enter a valid URL or IP address.
3. Select **Lookup Geolocation**. The country associated with the IP address, as defined in the geolocation database, displays next to the IP Address field.

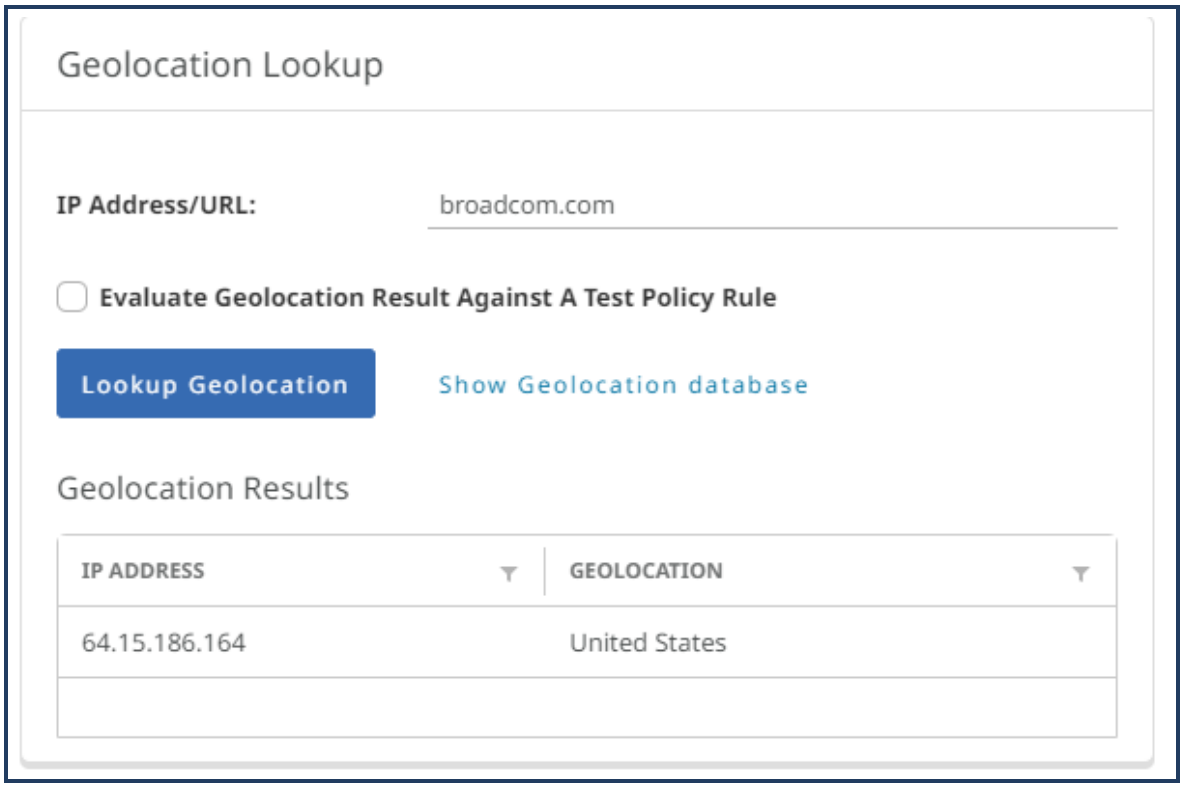

Note: IP address mappings to locations change over time, and periodic geolocation database updates reflect these changes. A changed IP address mapping can cause the following behavior:
- geolocation lookup results are different from a previous lookup
- geolocation policy is no longer working as expected

Use current lookup results to update policy as appropriate.

If the IP address you enter is not valid, an error appears. See "Troubleshoot Geolocation Errors" on the next page for more information.

To write policy about a specific country, use the country code specified in square brackets beside the country name in the country list. For example, the country code for the United States is “US”. To view the list of countries in the geolocation database, select **Show Geolocation database**. The list opens in a separate browser window.

Reverse Proxy Use Cases: Write Geolocation Policy

Refer to the following examples of writing geolocation policy for incoming connections.

The `client.address.country=<"country_name">` condition returns the country from which traffic originates, based on the client IP address. For detailed usage information on policy gestures, refer to the *Content Policy Language Reference*.

Use Case 1

You require policy to allow client connections only from North America. You can use the following CPL:

```
; only accept client connections from North America
<proxy>
allow client.address.country=(US, CA)
deny("Restricted location: $( x-cs-client-ip-country)")
```

Use Case 2

You require policy to allow client connections only from North America, including proxied traffic as specified. You can use the following CPL:

```
; only accept traffic from North America with support for proxied traffic
; with client address in X-Forwarded-For
<proxy>
client.effective_address("$(request.header.X-Forwarded-For)")

<proxy>
client.effective_address.country=(US, CA) ok
deny("Restricted location: $( x-cs-client-effective-ip-country)")
```

In the VPM, the **Client Geolocation** object is available in the Source column in policy layers. Refer to the *Visual Policy Manager Reference* for information.

Troubleshoot Geolocation Errors

You might encounter the following geolocation messages. Refer to the specified troubleshooting steps.

Note: This is not a valid IP

Cause: You have entered an invalid IP address in the IP field when performing geolocation lookups.

Fix: Correct the IP address and look it up again.

Note: The geolocation database is currently unavailable

Cause: No geolocation database is installed.

Fix: Download the database. See "Using Geolocation " on page 358.

This device does not have a valid geolocation license

Cause: The appliance does not have a valid geolocation license.

Fix: Ensure that the following are true:

- Each ProxySG appliance in your deployment has its own license.
- The licensing status for your appliance is in good health.

Warning: The geolocation database is not installed

This error appears in the browser window that opens when you select the Show Geolocation database link in the Geolocation Lookup section. It can also appear in the CLI.

Cause: No geolocation database is installed.

Fix: Download the database. See "Using Geolocation " on page 358.

Warning: Obsolete country name

This message appears when you try to install CPL policy. An example of the message is, “Warning: Obsolete country name: 'Russian Federation' is now 'Russia'”.

Cause: CPL includes an outdated country name.

Fix: Use the suggested country name. You can also refer to the geolocation database for current country names and codes. In the Admin Console, select **Administration > Data Services > Geolocation**. In the Geolocation Lookup section, select the **Show Geolocation database** link to display the list.

Access Log Errors

The following access log errors may appear in the access log if there is a problem with your subscription.

- The Geolocation subscription file is out of date
- Failed trying to get the subscription settings from the Geolocation subscription file
- Failed trying to download the Geolocation subscription file
- Failed trying to extract and activate the Geolocation payload file

Remove Geolocation Settings

To remove geolocation settings, purge the geolocation database and remove policy.

Remove the Database

Remove the database:

1. Disable geolocation.
   a. In the Admin Console, select **Administration > Data Services > Geolocation**.
   b. In the General section, disable the **Enable Geolocation** switch.
   c. Select **Save** to apply changes.
2. Log in to the CLI and enter the following command:

   ```
   #(config geolocation)purge
   ```

   The CLI returns to the `#(config geolocation)` node.

You can issue the view countries command to verify that the geolocation database has been removed. The output should show no list of countries and warn that the database is not installed:

```
#(config geolocation)purge
#(config geolocation)view countries
Countries defined by system:
Invalid
None
Unavailable
Unlicensed
Additional locations:
Countries defined by geolocation database:
Warning: The geolocation database is not installed
```

Remove Policy

Remove geolocation policy:

1. In the Management Console, select **Configuration > Policy > Visual Policy Manager**.
2. Launch the web or legacy VPM.
3. Remove any geolocation objects or rules. Refer to the *Visual Policy Manager Reference* for instructions.
4. Select **Install Policy**.

About Web Isolation

The most common element of cyber-security solutions is detection. Anti-virus/malware, network sandboxes, next-generation firewalls, and web application firewalls all depend on detection. In some cases, such as sandboxes, notification of malicious content arrives to the administrator after the user has received the content because of the time required to ascertain the verdict.

The Symantec Web Isolation solution addresses this security weakness. Instead of relying on malware detection, Web Isolation protects users from cyber attacks by isolating malware and preventing it from reaching users' browsers. A common use case is to protect employees who browse uncategorized and potentially malicious sites.

Web Isolation is a client-less solution that enables users to safely browse the internet on any device using any browser. The zero footprint negates the need for software installation on the clients.

Web Isolation:

- Provides a safe visual stream of the original web site in the user's browser.
    - End users browse the site as if the site was running directly in the browser.
- Preserves original browsing experience with full usability and control.
    - Only safely rendered information arrives at the user’s browser.
- Executes web sessions away from endpoints.

Requirements

To use the Web Isolation feature, you require:

- SGOS 7.3.x and later
- configured isolation policy

As an alternative to the Symantec Web Isolation service, you can use a third-party isolation service; however, you can't enable both services at the same time. See "Configure Web Isolation" below for configuration instructions for both options.

Configure Web Isolation

The ProxySG appliance can integrate with cloud-based Symantec Web Isolation service or a custom isolation service; you can't enable both services at the same time. When isolation policy is installed and enabled, the appliance forwards matching traffic to the configured isolation service.

For more information about this feature, refer to https://knowledge.broadcom.com/external/article/201609.

> **Tip:** Web Isolation requires SGOS 7.3.x and later and correctly configured isolation policy.

Configure web isolation:

1. (If using a custom isolation service) Obtain the hostname/IP address and port to use for communication with the isolation service.
2. In the Admin Console, select **Administration > Data Services > Isolation**.

3. In the Isolation section, specify settings for the service:

Field/Setting	Description
Destination Type	Select **SYMC Isolation Service** to forward matching traffic to the cloud-based Symantec Web Isolation service. Select **Custom Destination** to forward matching traffic to a configured custom isolation service.
Hostname	(Custom Destination only) Specify a hostname or IPv4/IPv6 address for the custom isolation service.
Port	(Custom Destination only) Specify the port for the custom isolation server. By default, the port is 8080.
Enable Secure Connection	(Custom Destination only) Select to secure connections to the isolation server.
SSL Device Profile	(If **Enable Secure Connection** is selected) Specify an existing SSL device profile to use for securing the HTTP/S requests to the isolation service over a secure tunnel. This device profile specifies the client certificate and the CA to use for validating the certificate from the isolation service SSL tunnel terminator. **Note:** This device profile is not used to validate the isolated traffic. After the SSL tunnel to the isolation service is terminated, the HTTP/S requests are forwarded to the isolation service and the ProxySG appliance will use the **bluecoat-isolation** or specified **Server CCL** to validate the isolation proxy certificate. See **Server CCL** for more information.
Include Headers	Send originating client IP address, authenticated user, or authenticated group details in the HTTP headers of traffic forwarded to the isolation service: ■ **User**: Authenticated user (`x-authenticated-user`). ■ **Group**: Group of the authenticated user (`x-authenticated-groups`). ■ **X-Forwarded-For**: Originating client IP address (`X-Forwarded-For`).

Field/Setting	Description
SSL Issuer Keyring	Specify a keyring for SSL interception of Shared Isolation Domains. Otherwise, the appliance uses the **default** keyring. **Note:** To isolate HTTPS requests, enable SSL interception and write policy that includes SSL intercept rules for the traffic you want to isolate. SSL intercept policy should use this keyring for the isolated traffic.
Server CCL	If not using the cloud isolation service, specify a CCL to validate the isolation server certificate used for isolated requests. For instructions on creating a CCL and adding CA certificates to the list, refer to "Manage CA Certificate Lists" on page 218.

4. (If needed) Select **Reset to default** to set all options to their default selections.

Manage System Images

Refer to the following topics to manage the ProxySG appliance system images:

- "Upgrade and Manage Systems " on the facing page

Upgrade and Manage Systems

The Admin Console displays up to five available systems.

To view system information, in the Admin Console, select **Administration > Upgrade**. Look in **Systems** for the following information:

Field or Setting	Description
System Version	Shows SGOS version, build number, and other system information. Expand a system to view more details, including boot status.
Signed	Indicates whether the system image is signed.
Locked	Indicates if the system is locked. Locked systems can't be removed or replaced until unlocked.
Replace	When there are five systems on the appliance, indicates that the system is to be replaced when a new system is loaded.
Default	Indicates that the system will be booted on the next hardware restart.
Running	The system that is currently running. It can't be deleted or replaced while running.

Refer to the following sections to manage systems.

Set the Default Boot System

This setting allows you to select the system to be booted on the next hardware restart. If a system starts successfully, it is set as the default boot system. If a system fails to boot, the next most recent system that booted successfully becomes the default boot system.

1. In the Admin Console, select **Administration > Upgrade**.
2. In **Systems**, select the preferred system version in the **Default** column.

Upgrade the System Version

When you install a new system image, it becomes the default system. Perform a hardware reboot to run the new version.

1. In the Admin Console, select **Administration > Upgrade**.
2. Download a new version of SGOS from MySymantec:
    a. Select **Available Systems**. The browser opens the download page.
    b. Log in to MySymantec and download the upgrade version.
    c. If needed, move the system image to a location that the appliance can access.

3. Upload the system image:
   a. Select **New Upload**. The Upload and install the system image dialog opens.
   b. To upload the image from the local workstation, select **Local > the system image from local file**.

      To upload the image from a remote location, select **URL**. In the fields that appear, enter the URL of the system image (include the entire path and filename) and specify which system to replace.

      Optionally, switch the **Enforce installation of signed images** toggle to specify that only signed system images can be installed.

      Select **Upload** to upload the system image.
4. In **Systems**, select **Reboot and Load Default** to perform a hardware reboot. A System Restart dialog appears as follows.

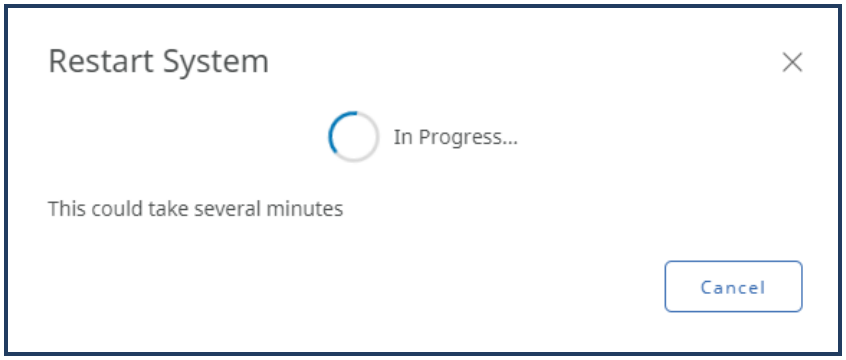

   The system indicates when it is restarted successfully:

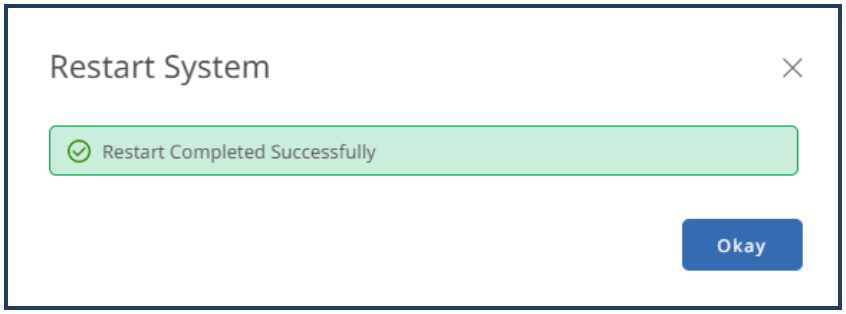

   On the Restart System dialog, select **Okay**.
5. To check the version after the reboot, select **Administration > Upgrade** and look for the system specified as **Default**.

Lock or Unlock a System

Any system can be locked, except a system that has been selected for replacement. If all systems, or all systems except the current system, are locked, the appliance cannot load a new system. If a system is locked, it cannot be replaced or deleted.

1. In the Admin Console, select **Administration > Upgrade**.
2. In **Systems**:
   - Select the system(s) to lock in the **Lock** column.
   - Clear the system(s) to unlock in the **Lock** column.

Replace a System

You can specify the system to be replaced when a new system is downloaded. If no system is specified, the oldest unlocked system is replaced by default. You cannot specify a locked system for replacement.

1. In the Admin Console, select **Administration > Upgrade**.
2. In **Systems**, select the system to replace in the **Replace** column.

Delete a System

You can delete any of the system versions except the current running system. A locked system must be unlocked before it can be deleted. If the system you want to delete is the default boot system, you need to select a new default boot system before the system can be deleted. You can't delete a system version through the Admin Console; you must use the CLI.

To delete a system, type the following commands:

```
# (config) installed-systems
# (config installed-systems) delete system_number
```

where *system_number* is the system you want to delete.

Manage Licensing

Refer to the following topics to manage ProxySG appliance licenses:

- "About Licensing" on the next page
  - "Manage ProxySG Licenses" on page 377
  - "Install the Base License" on page 380

About Licensing

Each ProxySG appliance requires a license to function. The license is associated with an individual serial number and determines what software features are available and the number of concurrent users that are supported. When you configure a new hardware appliance, the initial configuration wizard automatically installs a trial license that allows you to use all software features with support for an unlimited number of concurrent users for 60 days. (Trial periods are not applicable to virtual appliances.)

> **Note:** This information does not apply to the SWG VA. For licensing and upgrade information specific to the SWG VA, refer to the appropriate virtual appliance documentation.

Locate the ProxySG Serial Number

Each ProxySG serial number is the appliance identifier used to assign a license key file. The appliance contains an EEPROM with the serial number encoded, and it recognizes the serial number upon system boot-up.

To locate the serial number, in the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**. Look in the View section for the serial number.

License Types

The following license types are available:

- Trial: The 60-day license that ships with new physical appliances. All licensable components for the trial edition are active and available to use. In addition, the base SGOS user limit is unlimited. When a full license is installed, any user limits imposed by that license are enforced, even if the trial period is still valid.
- Demo: A temporary license that can be requested from Symantec to extend the evaluation period.
- Permanent: A license for hardware platforms that permanently unlocks the software features you have purchased. When a permanent license is installed, any user limits imposed by that license are enforced, even if the trial period is still valid.
- Subscription-based: A license that is valid for a set period of time. After you have installed the license, the ProxySG appliance will have full functionality, and you will have access to software upgrades and product support for the subscription period.

When a full license (permanent or subscription-based) or demo license is installed during the trial period, components previously available in the trial period, but not part of that license, remain available and active for the remainder of the trial period. However, if the license edition is different than the trial edition you selected, only functionality available in the edition specified in the license remains available for trial. If you do not want the trial components to be available after you install a full license, you can disable them.

Licensing Terms

Within sixty (60) days of the date from which the user powers up the ProxySG appliance (“Activation Period”), the administrator must complete the licensing requirements as instructed by the appliance to continue to use all of the features. Prior to the expiration of the Activation Period, the SGOS software will deliver notices to install the license each time the administrator logs in to manage the product. Failure to install the license prior to the expiration of the Activation Period may result in some features becoming inoperable until the administrator has completed licensing.

License Expiration

When the base license expires, the Admin Console banner displays a "No license" message:

Selecting **Show Details** brings up information such as "A valid license file is required."

When the appliance doesn't have a valid license, it stops processing requests. In addition, for services set to Intercept:

- In a transparent deployment, if the default policy is set to Allow, the appliance acts as if all services are set to Bypass, passing traffic through without examining it. If default policy is set to Deny, traffic to these services is denied with an exception. For details, see "Exceptions Due to Base License Expiration" below.
- In an explicit deployment, regardless of the default policy setting, traffic to these services is denied with an exception. For details, see "Exceptions Due to Base License Expiration" below.

Exceptions Due to Base License Expiration

In some cases, exceptions occur when the base license expires. Exceptions include:

- HTTP (Web browsers): An HTML page is displayed stating the license has expired.
- SSL: An exception page appears when an HTTPS connection is attempted, but only if the appliance is deployed explicitly or in the case of transparent proxy deployments, SSL interception is configured.
- FTP clients: If the FTP client supports it, a message is displayed stating the license has expired.
- Streaming media clients: If the Windows Media Player, RealPlayer, or QuickTime player version supports it, a message is displayed stating the license has expired.
- You can still perform configuration tasks through the CLI, SSH console, serial console, or Telnet connection. Although a component is disabled, feature configurations are not altered. Also, policy restrictions remain independent of component availability.

To correct license issues, refer to "Install the Base License" on page 380 or your Symantec point of contact.

See Also:

- "Manage ProxySG Licenses" on the facing page
- "Install the Base License" on page 380

Manage ProxySG Licenses

Manage the appliance base license and licensed components.

> **Note:** The license is associated with an individual serial number and determines what software features are available and the number of concurrent users that are supported. See "Locate the ProxySG Serial Number" on page 374 for details.

View License Status and Components

To view the license status in the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**. The View section displays base license information such as serial number, license date, and maximum number of concurrent users. The Licensed Components section other license components, indicates whether or not they are valid, and shows their expiration dates.

Review Licensed Components

Each component that can be licensed is listed under Licensed Components, along with its validity and its expiration date.

- To view the latest information after applying a license or update, select **Refresh**.
- Expand the name of a component to display details, as follows:

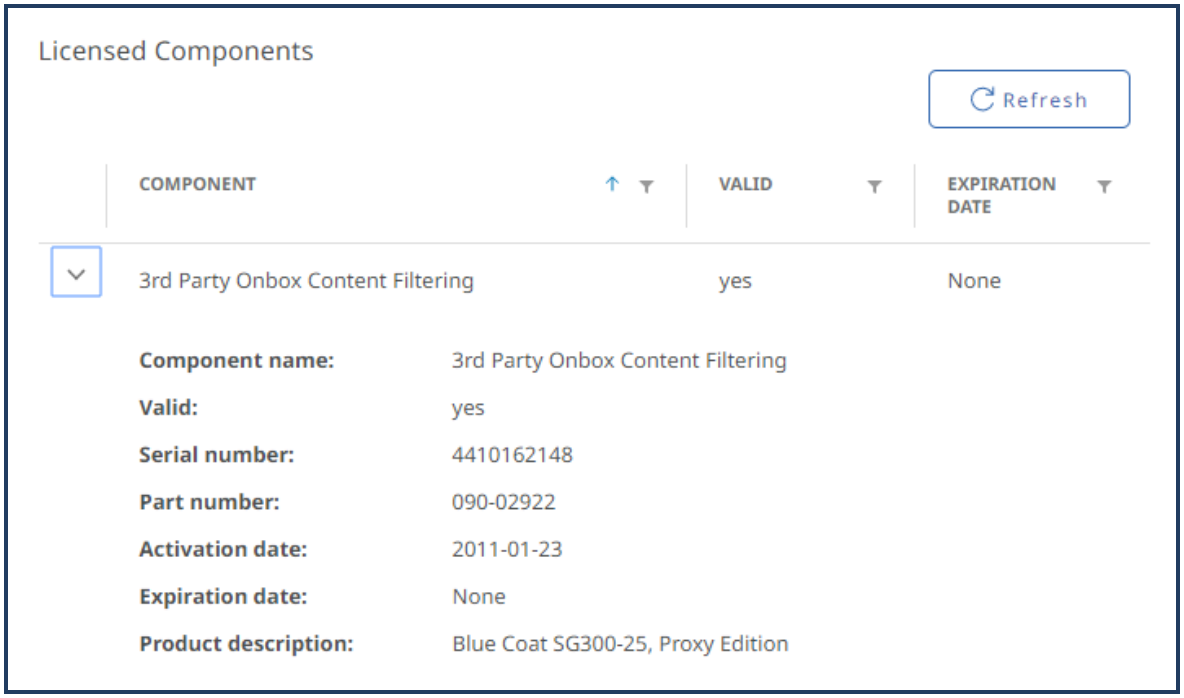

Install the ProxySG License

There are multiple ways to install the base ProxySG license. Use the method that is appropriate:

If the following is true:	Use this method:
The appliance and software are not registered yet.	"Register and License the Appliance and Software" on page 380
The appliance is a new system and is already registered.	"Retrieve the License on a Registered System" on page 381
The ProxySG serial number is not associated with a software license (you registered the hardware separately). OR The appliance is unable to access the licensing portal.	"Install the License Manually on a Registered System" on page 381

Verify Subscribed Intelligence Services Bundles

(If you have an Intelligence Services license) After you verify requirements, make sure your subscribed bundles appear in the Admin Console as expected. Verify that the bundles are correct and valid.

1. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
2. In Subscription, look for the names of your subscribed bundles and their expiration dates.

   Expand the name of a subscription to displays the feeds that it includes, as follows:

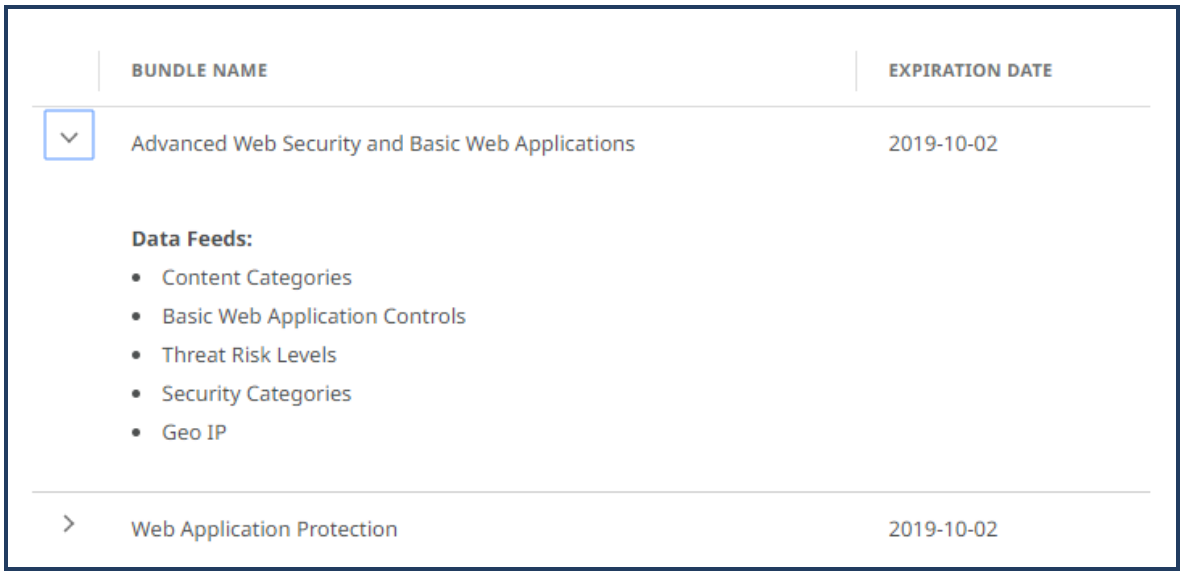

See Also:

- "Install the Base License" on the next page

Install the Base License

Before you can register and license your appliance, you must have the following:

- The serial number of your appliance. See "Locate the ProxySG Serial Number" on page 374.
- A MySymantec account. See "About the ProxySG Admin Console " on page 4.

There are multiple ways to install the base ProxySG license. Use the method that is appropriate:

If this is your appliance/license status	Complete this task
Appliance is not registered or licensed	"Register and License the Appliance and Software" below
New appliance is registered but not licensed	"Retrieve the License on a Registered System" on the facing page
Previously-registered older appliance is not licensed	"Install the License Manually on a Registered System" on the facing page
Registered appliance can't access the licensing portal/internet	

Notes

- After the initial license installation, you might decide to use another feature that requires a license. The license must be updated to support the new feature.
- A message is written to the event log when you install a license through the appliance.

Register and License the Appliance and Software

If you have not manually registered the appliance, you can automatically register the appliance and install the software license in one step.

> **Note:** Locate the email from Symantec that contains the activation code(s) for your software. You require these activation codes, as well as your appliance serial number, to complete the licensing process. Refer to Knowledge Base article 145885: https://knowledge.broadcom.com/external/article?articleId=145885

To register and license the appliance and software:

1. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
2. In the Install section, select **Register/Manage**. The licensing portal opens in a browser window and prompts you for your MySymantec login information.
3. Enter your login credentials and log in. The portal prompts you to enter your activation code.

4. Enter the activation code and follow the prompts to complete the process.
5. When prompted to accept the license agreement, read and accept the terms. The software license is now associated with the appliance.
6. (If necessary) Repeat the previous steps for your other activation codes.
7. Restart the appliance.

Retrieve the License on a Registered System

If the ProxySG appliance is a new system and the appliance has been registered, retrieve the associated license.

To retrieve the software license:

1. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
2. In the Install section, select **Retrieve**.
3. In the Retrieve License Key dialog, enter your MySymantec account login information and select **Request license**.
4. (Optional) Select **Show results** to verify a successful retrieval. If any errors occur, check the ability for the appliance to connect to Internet.
5. Restart the appliance.

Install the License Manually on a Registered System

Install the license manually if:

- The ProxySG serial number is not associated with a software license (you have registered the hardware separately)
- The appliance is unable to access the licensing portal (NPLP).

Install the license manually:

1. Follow the steps in "Register and License the Appliance and Software" on the previous page to log in to the NPLP. Select **License Download**. The portal prompts you for your appliance serial number.
2. Follow the prompts to enter your serial number and download the license file.
3. Save the license file to a location that your appliance can access.
4. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
5. In the Install section, under License Key Manual Installation, select an installation method:

Method	Instructions
Text editor	1. Select **Install from text editor**. 2. In the Load Edit and Install from text editor dialog, enter the contents manually or copy and paste the contents from another file. 3. Select **Install**.
Local file	1. Select **... > Local File**. 2. In the Load Edit and Install from local file dialog, browse to the directory of the file on the local system. 3. Select **Install**.
Remote URL	1. Select **... > Remote URL**. 2. In the Load Edit and Install from remote URL dialog, enter the fully-qualified URL, including the filename, where the file is located. 3. Select **Download**.

Configure License Updates

You can update the license manually (such as when changes are made to add-on components) or configure the appliance to automatically update the license, where the appliance contacts the Symantec licensing server 30 days before the license is to expire.

With the auto-update feature:

- If a new license has been purchased and authorized, the license is automatically downloaded.
- If a new license is not available, the appliance continues to contact the server daily for a new license until the current license expires.

Outside the license expiration window, the appliance makes this connection once every 30 days to check for new license authorizations.

To manually update the license:

1. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
2. In the Install section, switch the **Use Auto-Update** toggle.

To configure license auto-update:

1. In the Admin Console, select **Administration > Licensing > Licensed Components and Subscriptions**.
2. In the Install section, select **Update**.

Configure ProxySG Access Logs

Access logging allows you to trackWeb usage for the entire network or specific information on user or department usage patterns. These logs and reports can be made available in real-time or on a scheduled basis. The following topics provide procedures for enabling access logging and configuring upload schedules.

- "General Logging Configuration" on the next page
- "Configure an Access Log" on page 389
- "Configure an Upload Client" on page 393
- "About Access Log Formats" on page 402
  - "Default Access Log Formats" on page 410
  - "Configure an Access Log Format" on page 415

Fields Available for Creating Access Log Formats

Refer to the *ProxySG Log Fields and CPL Substitutions Reference*:
https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/proxysg/7-2/proxysg-log-fields-and-substitutions.html

General Logging Configuration

The ProxySG appliance can create access logs for the traffic flowing through the system; in fact, each protocol can create an access log record at the end of each transaction for that protocol (such as for each HTTP request).

These log records can be directed to one or more log facilities, which associates the logs with their configured log formats, upload schedules, and other customizable components. In addition, access logs can be encrypted and digitally signed before uploading.

Data stored in log facilities can be automatically uploaded to a remote location for analysis and archive purposes. Logs are uploaded using supported protocols. After logs are uploaded, reporting tools such as Symantec Reporter can be used to analyze the log files. For information on using Symantec Reporter, refer to the *Symantec Reporter WebGuide*.

Enable Access Logging

You can enable or disable access logging globally. When access logging is disabled, logging is turned off for all log objects, even if logging policy exists or logging configurations are set. When access logging is enabled, connection information is sent to the default log facility for the service (for example, HTTP traffic is logged to the main file). Logging begins immediately for all configured protocols.

> **Tip:** By default, protocols are associated with specific logs to allow you to start access logging as soon as you enable the feature. If you want to change the defaults, make the changes before enabling access logging. See the following sections.

To enable access logging:

1. Select **Administration > Access Logging > General**.
2. In **Default Logging Policy**, switch the **Enable** toggle.

Set Logging Defaults

To set logging defaults, select **Administration > Access Logging > General**. In Default Logging Policy, do any of the following as needed:

- Change a protocol's default log by selecting it under Default Log.
- Enable or disable log configuration for a protocol by switching the **Enable** toggle.

Specify Global Logging Settings

To determine which access logs contribute to high disk usage, to look for a trend in log sizes, or for other troubleshooting scenarios, you can set the ProxySG appliance to take certain actions when the combined size of all access logs reaches specified global limits:

- Stop all access logging.
- Delete the oldest entries (overflow) from any of the log facilities regardless of which log caused the total size to reach the global limit.
- Attempt an early upload of the log that caused the total size to reach the global limit.

You can also upload all configured access logs immediately.

After monitoring which logs the appliance uploads over a period of time, you can modify the amount of space allocated to individual logs for early upload.

To specify global logging settings:

1. Select **Administration > Access Logging > General**.
2. In **Global Settings**, specify the following as needed:

Section	Setting or field	Description
Global log file limit	**Limit total combined log sizes to**	Maximum total size of all log files, i.e., the sum of the sizes of all the individual logs.
	If limit reached	Specify what the appliance should do when the total log size reaches the maximum: ■ **Stop logging**: Stop all access logging and attempt an immediate upload. ■ **Delete oldest log entries**: Delete the oldest entries (overflow) from any of the log facilities regardless of which log caused the total size to reach the global limit. **Note:** To ensure that the appliance uploads older log entries before they are deleted, make sure that you have correctly specified an upload server, an early upload threshold for each individual log, and a global early upload threshold. Inability to connect to the upload server (for example, due to incorrect settings or network issues) will result in data loss.
Global Early Upload	**Start early upload when sum of all log sizes reaches**	Specify the sum total of all log sizes that triggers an early upload. The appliance attempts to upload the log that caused the total log size to reach the global limit. Individual logs each have their own early upload setting (see "Configure an Access Log" on page 389), which remain in effect even if you specify a global value. This means that, provided that upload is configured and working, the global early upload threshold could trigger a log upload before the file size reaches the threshold defined in the specific log facility.
Global Upload	**Upload All**	Upload all logs immediately.
	Cancel All	Prevent further attempts to upload all log facilities and cancel all uploads that are in progress.

Configure an Access Log

The appliance ships with several pre-defined access logs. Before creating a new one, check the existing ones to see if they fit your needs. If you want to use a custom log format with the new log facility, you must create the log format before associating it with a log (see "Configure an Access Log Format" on page 415).

To add a new access log or edit an existing access log, see "Configure an Access Log" on page 389. During configuration, you can opt to digitally sign access logs and set up the upload schedule.

Digitally Sign Access Logs

You can digitally sign access logs to certify that a particular ProxySG appliance wrote and uploaded this log file. Signing is supported for both content types– text and gzip–and for both upload types–continuous and periodic. Each log file has a signature file associated with it that contains the certificate and the digital signature for verifying the log file. The signature file has the same name as the access log file but with a .sig extension; that is, filename.log.sig, if the access log is a text file, or filename.log.gzip.sig, if the access log is a gzip file.

> **Note:** Signing is disabled by default.

You can digitally sign your access log files with or without encryption. If the log is both signed and encrypted, the signing operation is done first, meaning that the signature is calculated on the unencrypted version of the file. You must decrypt the log file before verifying the file. Attempting to verify an encrypted file fails.

When you create a signing keyring (which must be done before you enable digital signing), keep in mind the following:

- The keyring must include an external certificate. (An external certificate is one for which the ProxySG appliance does not have the private key.)
- The certificate purpose must be set for smime signing. If the certificate purpose is set to anything else, you cannot use the certificate for signing.
- Add the %c parameter in the filenames format string to identify the keyring used for signing. If encryption is enabled along with signing, the %c parameter expands to keyringName_Certname.

> **Note:** The signing feature is not available for custom clients.

Verify a Digital Signature

If the file whose digital signature you want to verify is also encrypted, you must decrypt the file prior to verifying the signature. You can use a program such as OpenSSL to verify the signature. For example, use the following command in OpenSSL:

```
openssl smime -CAfile cacrt -verify -in filename.sig -content filename.log -inform DER -out logFile
```

where:

- cacrt: CA certificate used to issue the certificate in the signature file.
- filename.sig: File containing the digital signature of the log file.
- filename.log: Log file generated after decryption. If the access log is a gzip file, it contains a .gz extension.
- logFile: Filename that is generated after signature verification

See Also:

- "Configure an Access Log" on page 389

Configure an Access Log Upload Schedule

The upload schedule defines the frequency of the access logging upload to a remote server, the time between connection attempts, the time between keep-alive packets, the time at which the access log is uploaded, and the protocol that is used. When configuring an upload schedule, you can specify either periodic uploading or continuous uploading. Both periodic and continuous uploading can send log information from an ProxySG appliance farm to a single log analysis tool. This allows you to treat multiple appliances as a single entity and to review combined information from a single log file or series of related log files.

With periodic uploading, the SGOS software transmits log entries on a scheduled basis (for example, once daily or at specified intervals) as entries are batched, saved to disk, and uploaded to a remote server.

> **Note:** When you configure a log for continuous uploading, it continues to upload until you stop it. To stop continuous uploading, switch to periodic uploading temporarily. This is sometimes required for gzip or encrypted files, which must stop uploading before you can view them.

With continuous uploading, the ProxySG continuously streams new access log entries from the device memory to a remote server. Here, streaming refers to the real-time transmission of access log information. SGOS transmits access log entries using the specified client, such as FTP client. A keep-alive data packet is sent to keep the data connection open.

Continuous uploading allows you to view the latest logging information almost immediately, send log information to a log analysis tool for real-time processing and reporting, maintain the ProxySG performance by sending log information to a remote server (avoiding disk writes), and save device disk space by saving log information on the remote server.

> **Note:** If the remote server is unavailable to receive continuous upload log entries, SGOS saves the log information on the device disk. When the remote server is available again, the appliance resumes continuous uploading.
>
> If you do not need to analyze the upload entries in real time, use periodic uploading because it is more reliable than continuous uploading. If there is a problem configuring continuous uploading to Microsoft Internet Information Server (IIS), use periodic uploading instead.

See Also:

- "Configure an Access Log" on the facing page

Configure an Access Log Format

To add a new access log format or edit an existing access log format, see "Configure an Access Log Format" on page 415.

See Also:

- "Configure an Access Log" on the facing page

Encrypt the Access Log

To encrypt access log files, you must first place an external certificate on the ProxySG appliance. The device derives a session key from the public key in the external certificate and uses it to encrypt the log. When an access log is encrypted, two access log files are produced: an ENC file (extension .enc), which is the encrypted access log file, and a DER file (extension .der), which contains the appliance session key and other information. You need four things to decrypt an encrypted access log:

- The ENC file
- The DER file
- The external (public key) certificate
- The corresponding private key

The encryption feature is not available for custom clients.

Troubleshoot Access Logging

Problem: The ProxySG appliance is uploading logs more frequently than expected.

Description: If access logging is enabled, logs can accrue on the ProxySG appliance's hard drive even if the upload client is not configured for specific protocols (often the case if you configured streaming or P2P). Eventually the size of these combined logs, triggers the **Global Early Upload** threshold in Global Settings (**Administration > Access Logging > General**). The appliance attempts to upload all configured logs more often than expected. For example, a main log that is configured for upload every 24 hours starts to upload small portions of the main log every 10 minutes.

Solution: To prevent the access logs that do not have an upload client configured from triggering the Start an Early upload threshold, edit the default logs for each protocol that you do not need uploaded. Set them to <None> in **Upload Client** on the Edit Log page (select **Administration > Access Logging > General** and select a log to edit).

Configure an Access Log

To configure an access log:

1. Select **Administration > Access Logging > General**.
2. In **Logs**, do the following under Logs:
   - To add a log: Click **Add New Log**.
   - To edit an existing log, select the name of a log in the list.
3. On the Add New Log or Edit log page, specify or change the log properties as follows.

Section	Setting or field	Description
Log Setting	**Log Name**	Name of the log. If editing an existing log, this is a read-only field. The name can include specifiers from "Access Log Filename Formats" on page 408. For example, if you name the file: ■ AccLog, the name is AccLog ■ AccLog%C%m%d%H%M%S, the name is *AccLog appliance_name month day hour min sec* ■ C%m%d, the name is *appliance_name month day* ■ Y%m%d%C, the name is *year month day appliance_name*
	Log Format	Supported log format. For a list of formats, see "Configure an Access Log Format" on page 415. **Note:** If you change the log format of a log, remember that ELFF formats require an ELFF header in the log (the list of fields being logged are mentioned in the header) and that non-ELFF formats do not require this header. The format of data written to the log changes as soon as the format change is applied; for best practices, do a log upload before the format change and immediately after (to minimize the number of log lines in a file with mixed log formats). Upload the log facility before you switch the format.
	Description	(Optional) Description of log, used for identification purposes only.

Section	Setting or field	Description
Log File Limits	**The maximum size of each remote file**	The maximum size for each remote log file (the file on the upload server) defaults to 0, meaning that all data is sent to the same log file. If you set a maximum size, a new log file opens when the file reaches that size. This setting is valid for both periodic and continuous uploads.
	Start early upload when sum of all log sizes reaches	Log size that triggers an early upload. The appliance attempts to upload the log when the limit is reached. This early upload setting remains in effect even if you specify a global value (see "Specify Global Logging Settings" on page 384). This means that, provided that upload is configured and working, the global early upload threshold could trigger a log upload before the file size reaches the threshold defined in the specific log facility.
Upload Client	**Upload Client**	Supported upload client. See "Configure an Upload Client" on page 393. Select **None** to disable log uploads.
	Open Settings	Displays client settings for the current log.
	Test Upload	Test uploading the current access log. Before you begin, make sure you have configured the upload client completely. 1. Select **Test Upload**. The appliance displays a "Format Syntax correct" or "Format Syntax incorrect" message. 2. Check the event log for upload results. For the duration of the test, configure the event log to use the verbose event level. This logs more complete log information. After you test uploading, you can check the event log for the test upload event and determine whether any errors occurred.

Section	Setting or field	Description
Transmission Parameters	**Encryption Certificate**	The appliance uses the selected external certificate to encrypt the uploaded log. The external certificate must be imported to the appliance before you can use it here. Not available for custom clients. (Optional setting for digitally signing access logs.)
	Signing Keyring	The access log is digitally signed using the selected keyring. The signing keyring, with a certificate set to smime, must already exist on the appliance. A certificate set to any other purpose cannot be used for digital signatures. Not available for custom clients. (Optional setting for digitally signing access logs.)
	Save the log file as	Specify whether or not to compress the uploaded access log: ■ **gzip file**: (Default) Compress the log. ■ **text file**: Don't compress the log. If you select gzip file for the Kafka Client, you can configure the appliance to add a MessageSet header to compressed log entries. Doing so allows for better performance when the compressed data is sent to the broker. For details, see "Ensure Logs are Processed Correctly" on page 398.
	Send partial buffer after	Applies only if **text file** is selected above. Specify the maximum time between text log packets, meaning that it forces a text upload after the specified length of time even if the internal log buffer is not full. If the buffer fills up before the time specified in this setting, the text uploads right away, and is not affected by this maximum setting. This setting is valid only for continuous uploading.
	Bandwidth Class	(Optional) Select a bandwidth class. The default setting is none, which means that bandwidth management is disabled for this log facility by default. Before you can manage the bandwidth for this log facility, you must first create a bandwidth-management class. It is the log facility that is bandwidth-managed–the upload client type does not affect this setting. Less bandwidth slows down the upload, while more could flood the network.

Section	Setting or field	Description
Upload Schedule See "Configure an Access Log Upload Schedule" on page 387 for details.	**Upload the access log**	Upload type: ■ **continuously**: Stream access log entries to a remote server. ■ **periodically**: Transmit on a scheduled basis.
	Wait between connect attempts	Number of seconds between connection attempts.
	Time between keep-alive log packets	Applies only if **continuously** is selected for upload type above. Number of seconds between keep-alive packets. Keep-alives maintain the connection to FTP and HTTP servers during low periods of system usage. When no logging information is being uploaded, the SGOS software sends a keep-alive packet to the remote server at the interval you specify, from 1 to 65535 seconds. If you set this to 0 (zero), you effectively disable the connection during low usage periods. The next time that access log information needs to be uploaded, the ProxySG automatically reestablishes the connection.
	Daily or **In Intervals**	Specify when logs are uploaded or rotated: ■ If **Daily**, specify the time of day to upload logs (for periodic) or rotate logs (for continuous) in the **Upload at** field. ■ If **In Intervals**, specify the amount of time, in the **Hours** and **Minutes** fields, between uploads.
	Upload Now	Applies only if **periodically** is selected for upload type above. You can upload the access logs now or you can cancel any access-log upload currently in progress (if you are doing periodic uploads). You can rotate the access logs now (if you are doing continuous uploads). These actions do not affect the next scheduled upload time.
	Cancel Upload	Applies only if **periodically** is selected for upload type above. Allows you to stop repeated uploadattempts if theWeb server becomes unreachable while an upload is in progress. Clicking this sets log uploading back to idle if the log is waiting to retry the upload. If the log file is in the process of uploading, it takes time for it to take effect.
	Rotate Now	Applies only if **continuously** is selected for upload type above. Log rotation helps prevent logs from growing excessively large. Especially with a busy site, logs can grow quickly and become too big for easy analysis. With log rotation, SGOS periodically creates a new log file, and archives the older one without disturbing the current log file.

Configure an Upload Client

The ProxySG appliance supports the following upload clients:

- FTP client, the default
- HTTP client
- Custom client

> **Note:** You must have a socket server to use the Custom client.

- Symantec Reporter client
- Kafka client

> **Note:** You can configure the appliance to add a Kafka MessageSet header to compressed log entries. Making any change to an access log's upload client configuration that reverses the MessageSet header state (that is, the header's presence or absence in the log files) can cause future log uploads to fail. To prevent these issues from occurring, see "Ensure Logs are Processed Correctly" on page 398.

- SCP client

Symantec also supports secure FTP (FTPS), secure HTTP (HTTPS), secure Custom client, and secure Kafka client. The Custom client is based on plain sockets.

The general options you enter in the Upload Client tab affect all clients. Specific options that affect individual clients are discussed in the FTP client, HTTP client, or Custom client, or the access-log ftp-client, https-client, or custom-client CLI commands.

Only one client can be used at any one time. All four can be configured, but only the selected client is used.

The appliance provides access logging with two types of uploads to a remote server:

- Continuous uploading, where the device continuously streams new access log entries from the device memory to a remote server.
- Scheduled (periodic) uploading, where the device transmits log entries on a scheduled basis. See "General Logging Configuration" on page 384.

The appliance allows you to upload either compressed access logs or plain-text access logs. The device uses the gzip format to compress access logs. Gzip compressed files allow more log entries to be stored in the device. Advantages of using file compression include:

- Reduces the time and resources used to produce a log file because fewer disk writes are required for each megabyte of log-entry text.
- Uses less bandwidth when the device sends access logs to an upload server.
- Requires less disk space.

Compressed log files have the extension .log.gz. Text log files have the extension .log. For greater security, you can configure the SGOS software to encrypt and sign the access log.

Symantec supports several upload clients for the appliance. Multiple upload clients can be configured per log facility, but only one can be enabled and used per upload.

> **Note:** For information on the Reporter Client, refer to the *Symantec Reporter WebGuide*.

To configure an upload client:

1. In the Admin Console, select **Administration > Access Logging > General**.
2. In **Logs**, select a configured log.
3. On the Edit Log page, in **Upload Client**, select an upload client:
    - "Custom Client Settings" below
    - "FTP Client Settings" on the facing page
    - "HTTP Client Settings" on page 396
    - "Kafka Client Settings" on page 397
    - "SCP Client Settings" on page 400

Custom Client Settings

Specify the following as needed:

Area	Settings
Primary Server:	■ **Host**: Specify the host IP address (in IPv4 format) of the upload destination. ■ **Port**: The default port is 69. Enter a different port if needed.
Alternate Server	■ **Host**: Specify the host IP address (in IPv4 format) of the alternate upload destination. ■ **Use port from primary server**: Selected by default; clear the selection and specify a different port if needed.
Other Settings	**Use secure connections (SSL)**: Select if you are using secure connections. If this setting is enabled, the IP address must match the IP address in the certificate presented by the server.

FTP Client Settings

Specify the following as needed:

Area	Settings
Primary Server	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) or domain name that resolves to an IPv4/IPv6 address of the upload destination. ■ **Port**: The default port is 21. Enter a different port if needed. ■ **Upload Path**: Directory path where the access log is uploaded on the server. ■ **Username**: User name that is known on the host. ■ **Change password**: Enable to change the password on the FTP server. Enter and confirm the password.
Alternate Server:	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) of the alternate upload destination. ■ **Use port, upload path and username from primary server**: Selected by default; clear the selection and specify a different port, path, and user name if needed.
Other Settings:	■ **Filename**: Log filename format, which supports the following specifiers and text: ◦ `%f` for the log name ◦ `%c` for the name of the external certificate used for encryption, if applicable. If you use more than one external certificate to encrypt logs, include this specifier to keep track of which external certificate was used to encrypt the uploaded log file. ◦ `%l` for the fourth parameter of the ProxySG appliance IP address ◦ `%m` and `%d` are date specifiers (month and day, respectively) ◦ `%H%` , `%M`, and `%S`, are time specifiers (hour, minute, and second, respectively) If you configure logs for continuous upload and do not have time specifiers in the Filename field, each access log file produced overwrites the previous file. ◦ `.gzip.log` or `.log` for the log file extension For additional specifiers that this log format might support, see "Access Log Filename Formats" on page 408. ■ **Timestamp**: The time stamp format used in access log entries and filename. Select **UTC** to reflect the UTC standard. Select **Local time** to reflect the local time zone. ■ **Use secure connections (SSL)**: Select if you are using FTPS. ■ **Use passive FTP**: Enabled by default, the ProxySG appliance connects to the FTP server. Otherwise, the FTP server uses the PORT command to connect to the appliance.

HTTP Client Settings

Access log uploads done through an HTTP/HTTPS client use the HTTP PUT method. The destination HTTP server (where the access logs are being uploaded) must support this method. Microsoft's IIS allows the server to be directly configured for write (PUT/DELETE) access. Other servers, such as Apache, require installing a new module for the PUT method for access log client uploads.

You can create either an HTTP or an HTTPS upload client through the HTTP client dialog. (Create an HTTPS client by selecting **Use secure connections (SSL)**.)

> **Note:** To create an HTTPS client, you must also import the appropriate CA Certificate.

Specify the following as needed:

Area	Settings
Primary Server	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) or domain name that resolves to an IPv4/IPv6 address of the upload destination. ■ **Port**: The default port is 80. Enter a different port if needed. For HTTPS, use port 443. ■ **Upload Path**: Directory path where the access log is uploaded on the server. ■ **Username**: User name that is known on the host. ■ **Change password**: Enable to change the password on the HTTP/S server. Enter and confirm the password.
Alternate Server:	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) of the alternate upload destination. ■ **Use port, upload path and username from primary server**: Selected by default; clear the selection and specify a different port, path, and user name if needed.

Area	Settings
Other Settings:	■ **Filename**: Log filename format, which supports the following specifiers and text: ◦ `%f` for the log name ◦ `%c` for the name of the external certificate used for encryption, if applicable. If you use more than one external certificate to encrypt logs, include this specifier to keep track of which external certificate was used to encrypt the uploaded log file. ◦ `%l` for the fourth parameter of the ProxySG appliance IP address ◦ `%m` and `%d` are date specifiers (month and day, respectively) ◦ `%H%` , `%M`, and `%S`, are time specifiers (hour, minute, and second, respectively) If you configure logs for continuous upload and do not have time specifiers in the Filename field, each access log file produced overwrites the previous file. ◦ `.gzip.log` or `.log` for the log file extension For additional specifiers that this log format might support, see "About Access Log Formats" on page 402 ■ **Timestamp**: The time stamp format used in access log entries and filename. Select **UTC** to reflect the UTC standard. Select **Local time** to reflect the local time zone. ■ **Use secure connections (SSL)**: Select if you are creating an HTTPS client. You must also create a key pair, import or create a certificate, and, if necessary, associate the key pair and certificate (called a keyring), with the SSL device profile. If this setting is enabled, the IP address must match the IP address in the certificate presented by the server.

Kafka Client Settings

You can use the Apache Kafka access log upload client to upload logs from the appliance to a Kafka broker cluster. A Kafka consumer client can consume the log data and process them to make them suitable for import into a log analysis tool.

To use Apache terminology, the appliance is the Kafka producer client, the Symantec Web Security Service or your own client is the Kafka consumer client, and the cluster of servers running Kafka broker software is the Kafka broker cluster. The appliance publishes log entries to consumer-specific feeds called topics. For more information on Kafka concepts and terminology, refer to Apache documentation: http://kafka.apache.org/documentation.html

> **Note:** As of this writing, the Kafka broker does not natively support SSL encryption. To have a secure communication channel from the ProxySG appliance, you must provide an authentication adapter between the appliance and the Kafka broker cluster.

The appliance can add a MessageSet header to gzipped log files before they are uploaded to the Kafka broker. This allows Kafka to process compressed data correctly, rather than as a large uncompressed data file. For details, see "Ensure Logs are Processed Correctly" on the next page.

You require the following to use Kafka as the upload client:

- The appliance must be able to access the Kafka broker cluster.
- A Kafka consumer client must exist to consume and process log data. If the appliance is running in FIPS mode, the key for mutual authentication must be FIPS-compliant.

For more information on FIPs mode, refer to the *ProxySG FIPS Mode WebGuide*.

Specify the following as needed:

Area	Settings
Primary Server	▪ **Host**: Specify the host IP address (in IPv4 or IPv6 format) or domain name that resolves to an IPv4/IPv6 address of the upload destination. ▪ **Port**: The default port is 9092. ▪ **SSL Device Profile**: Select an existing SSL device profile to use for mutual authentication.
Alternate Server:	▪ **Host**: Specify the host IP address (in IPv4 or IPv6 format) of the alternate upload destination. ▪ **Use port and SSL device profile from primary server**: Selected by default; clear the selection and specify a different port and SSL device profile if needed.
Other Settings:	▪ **Use secure connections (SSL)**: Select an existing SSL device prrofile if you want to secure the connection between the appliance and the broker. ▪ **Topic**: Topic name for this specific log facility. You can use any string-format combination, but the default is `SG_%f`, where `%f` is automatically replaced with the name of the log facility. Other supported format characters are: ◦ `%c` for the client name ◦ `%i` for the IP address For additional specifiers that this log format might support, see "About Access Log Formats" on page 402

Ensure Logs are Processed Correctly

If you selected gzip file for Kafka, you can configure the appliance to add a MessageSet header to the compressed log files so that the Kafka broker processes the data as gzip-compressed data. Use the following command to enable/disable the setting (by default, the setting is disabled):

```
# (config log log_name) kafka-client [no] message-set-codec
```

Refer to the *Command Line Interface Reference* for details on this command.

When all of the following conditions are true, the MessageSet header is added to log files:

- Kafka is the upload client for the log
- gzip is selected as the log file type
- the Kafka codec is enabled in the CLI

When one or more of the following conditions are true, the MessageSet header is not added to log files:

- Kafka is not the upload client for the log
- Kafka is the upload client for the log but gzip is not selected as the file type
- Kafka is the upload client but the codec setting is disabled

Making any change to an access log's upload client configuration that reverses the previous MessageSet header state (that is, the header's presence or absence in the log files) can cause future log uploads to fail. For example, if you change a log's upload client from FTP to Kafka with gzip and codec enabled, the header did not exist in previous log files but it is added to the files after the upload client change. If you then change the Kafka file type from gzip to text, the header is no longer added to the log files. In both of these scenarios, you must take additional steps to ensure that logs are processed correctly.

To prevent problems from arising when you make any changes that affect the MessageSet header state, clear the write buffer and remove any stored logs from the appliance. Symantec recommends that you perform these steps in the specified order: manually upload the logs, make your desired log changes, and then delete the access logs from the appliance.

Example: Changing the codec for Kafka client

1. Upload the access logs. In the Admin Console, select **Administration > Access Logging > General** and edit a log to configure the upload schedule.

   Example: Your dns log uses the Kafka client and the gzip file format. In **Upload Client**, select the dns log and click **Test Upload**.

2. Change the codec setting as appropriate. In the CLI, issue one of the following commands:

   ```
   # (config log log_name) kafka-client message-set-codec
   ```

   The header is added to the start of compressed log entries.

   ```
   # (config log log_name) kafka-client no message-set-codec
   ```

   Headers are not added to the start of any log entries.

   Example: You are enabling the header for the first time. Issue the following command for your dns log, named dns:

   ```
   # (config log dns) kafka-client message-set-codec
   ```

3. Delete the access logs. In the CLI, issue the following command for each log using the Kafka client:

   ```
   # (config log log_name) commands delete-logs
   ```

   Example: Issue the following command for the log named dns:

```
# (config log dns) commands delete-logs
```

The dns logs stored on the appliance are deleted. Subsequent uploads of the dns logs, whether system-initiated or user-initiated, will have the MessageSet header unless you do any of the following:

- select a different upload client
- change the log file format to text
- disable the codec setting

Whenever you intend to make any of these modifications, follow the previous procedure to ensure that logs continue to be processed correctly.

SCP Client Settings

The SCP protocol allows you to transfer data securely over an SSH connection. If your site has an SSH server, you can configure an SCP client for secure access log uploads. Using the SCP client is useful if your organization does not permit other secure methods such as HTTPS or FTPS.

Before setting up the SCP client, you must configure the hosts, client keys, HMACs, and ciphers for outbound SSH connections.

> **Note:** Continuous upload is not supported over SCP client.

Specify the following as needed:

Area	Settings
Primary Server	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) or domain name that resolves to an IPv4/IPv6 address of the upload destination. ■ **Port**: The default port is 22. Enter a different port if needed. ■ **Upload Path**: Directory path where the access log is uploaded on the server. ■ **Username**: User name that is known on the host. ■ **Change password**: Enable to change the password on the SCP server. Enter and confirm the password.
Alternate Server:	■ **Host**: Specify the host IP address (in IPv4 or IPv6 format) of the alternate upload destination. ■ **Use port, upload path and username from primary server**: Selected by default; clear the selection and specify a different port, path, and user name if needed.

Area	Settings
Other Settings:	■ **Filename**: Log filename format, which supports the following specifiers and text: ◦ `%f` for the log name ◦ `%c` for the name of the external certificate used for encryption, if applicable. If you use more than one external certificate to encrypt logs, include this specifier to keep track of which external certificate was used to encrypt the uploaded log file. ◦ `%l` for the fourth parameter of the ProxySG appliance IP address ◦ `%m` and `%d` are date specifiers (month and day, respectively) ◦ `%H%` , `%M`, and `%S`, are time specifiers (hour, minute, and second, respectively) If you configure logs for continuous upload and do not have time specifiers in the Filename field, each access log file produced overwrites the previous file. ◦ `.gzip.log` or `.log` for the log file extension For additional specifiers that this log format might support, see "About Access Log Formats" on the next page. ■ **Timestamp**: The time stamp format used in access log entries and filename. Select **UTC** to reflect the UTC standard. Select **Local time** to reflect the local time zone.

About Access Log Formats

The appliance supports the following access log formats:

- "Custom or W3C ELFF Format" below
- "SQUID-Compatible Format " on page 407
- "NCSA Common Access Log Format" on page 407

The ProxySG appliance can create access logs with any one of six formats. Four of the six are reserved formats and cannot be configured; however, you can create additional logs using custom or ELFF format strings. When using an ELFF or custom format, a blank field is represented by a dash character.

When using the SQUID or NCSA log format, a blank field is represented according to the standard of the format.

See "Access Log Filename Formats" on page 408 for supported specifiers for the access log filename.

Custom or W3C ELFF Format

The W3C ELFF is a subset of the Symantec format. ELFF is specified as a series of space delimited fields. Each field is described using a text string. The types of fields are described in the following table.

Field Type	Description
Identifier	A type unrelated to a specific party, such as date and time.
prefix-identifier	Describes information related to a party or a transfer, such as c-ip (client's IP) or sc-bytes (how many bytes were sent from the server to the client)
prefix (header)	Describes a header data field. The valid prefixes are: ■ c = Client ■ s = Server ■ r = Remote ■ sr = Server to Remote ■ cs = Client to Server ■ sc = Server to Client ■ rs = Remote to Server

ELFF formats are created by selecting a corresponding custom log format using the table below. Unlike the Symantec custom format, ELFF does not support character strings and requires a space between fields.

Selecting the ELFF format does the following:

- Puts one or more W3C headers into the log file. Each header contains the following lines:

```
#Software: SGOS x.x.x
#Version: 1.0
#Date: 2002-06-06 12:12:34
#Fields: date time cs-ip…
```

- Changes all spaces within fields to + or %20. The ELFF standard requires that spaces only be present between fields. ELFF formats are described in the following table.

Symantec Custom Format	ELFF	Description
space character	N/A	Multiple consecutive spaces are compressed to a single space
%	-	Expansion field
%%	-	'%' character
%a	c-ip	IP address of the client
%b	sc-bytes	Number of bytes sent from appliance to client
%c	rs(Content-Type)	Response header: Content-Type
%d	s-suppliername	Hostname of the upstream host (not available for a cache hit)
%e	time-taken	Time taken (in milliseconds) to process the request
%f	sc-filtercategory	Content filtering category of the request URL
%g	timestamp	Unix type timestamp
%h	c-dns	Hostname of the client (uses the client's IP address to avoid reverse DNS)
%i	cs-uri	The log URL
%j	-	Not used.
%k	-	Not used.
%l	x-bluecoatspecial-empty	Resolves to an empty string
%m	cs-method	Request method used from client to appliance
%n	-	Not used.
%o	-	Not used.
%p	r-port	Port from the outbound server URL
%q	-	Not used.
%r	cs-requestline	First line of the client's request

Symantec Custom Format	ELFF	Description
%s	sc-status	Protocol status code from appliance to client
%t	gmttime	GMT date and time of the user request in format: [DD/MM/YYYY:hh:mm:ss GMT]
%u	cs-user	Qualified username for NTLM. Relative username for other protocols
%v	cs-host	Hostname from the client's request URL. If URL rewrite policies are used, this field's value is derived from the 'log' URL
%w	s-action	Type of action the ProxySG appliance took to process this request (see "Action Field Values" on the facing page).
%x	date	GMT Date in YYYY-MM-DD format
%y	time	GMT time in HH:MM:SS format
%z	s-icap-status	ICAP response status
%A	cs(User-Agent)	Request header: User-Agent
%B	cs-bytes	Number of bytes sent from client to appliance
%C	cs(Cookie)	Request header: Cookie
%D	s-supplier-ip	IP address used to contact the upstream host (not available for a cache hit)
%E	-	Not used.
%F	-	Not used.
%G	-	Not used.
%H	s-hierarchy	How and where the object was retrieved in the cache hierarchy.
%I	s-ip	IP address of the appliance on which the client established its connection
%J	-	Not used.
%K	-	Not used.
%L	localtime	Local date and time of the user request in format: [DD/MMM/YYYY:hh:mm:ss +nnnn]
%M	-	Not used.
%N	scomputername	Configured name of the appliance
%O	-	Not used.
%P	s-port	Port of the appliance on which the client established its connection
%Q	cs-uri-query	Query from the log URL
%R	cs(Referer)	Request header: Referer
%S	s-sitename	Service type used to process the transaction
%T	duration	Time taken (in seconds) to process the request

Symantec Custom Format	ELFF	Description
%U	cs-uri-path	Path from the log URL. Does not include query.
%V	cs-version	Protocol and version from the client's request, e.g. HTTP/1.1
%W	sc-filterresult	Content filtering result: Denied, Proxied or Observed
%X	cs(XForwarded-For)	Request header: X-Forwarded-For
%Y	-	Not used.
%Z	s-icap-info	ICAP response information

Example Access Log Formats

```
Squid log format: %g %e %a %w/%s %b %m %i %u %H/%d %c
NCSA common log format: %h %l %u %t "%r" %s %b
NCSA extended log format: %h %l %u %L "%r" %s %b "%R" "%A"
Microsoft IIS format: %a, -, %x, %y, %S, %N, %I, %e, %b, %B, %s, 0, %m,
%U, -
```

The Symantec custom format allows any combination of characters and format fields. Multiple spaces are compressed to a single space in the actual access log. You can also enter a string, such as `My default is %d`. The appliance goes through such strings and finds the relevant information. In this case, that information is %d.

Action Field Values

The following table describes the possible values for the `s-action` field.

Value	Description
ACCELERATED	(SOCKS only) The request was handed to the appropriate protocol agent for handling.
ALLOWED	An FTP method (other than the data transfer method) is successful.
DENIED	Policy denies a method. A DENIED s-action value is returned for CIFS, Endpoint Mapper, MAPI, FTP, P2P, Shell Proxy, SOCKS proxy, streaming proxies, and SSL proxy when policy denies a request. When the same kind of denial happens in the HTTP proxy, TCP_DENIED is reported.
FAILED	An error or failure occurred.
LICENSE_ EXPIRED	(SOCKS only) The request could not be handled because the associated license has expired.
TCP_ ACCELERATED	For CONNECT tunnels that are handed off to the following proxies: HTTP, SSL, Endpoint mapper, and P2P for BitTorrent/EDonkey/Gnutella.
TCP_AUTH_HIT	The requested object requires upstream authentication, and was served from the cache.
TCP_AUTH_ HIT_RST	The requested object requires upstream authentication, but the client connection was reset before the complete response was delivered.

Value	Description
TCP_AUTH_MISS	The requested object requires upstream authentication, and was not served from the cache. This is part of CAD (Cached Authenticated Data).
TCP_AUTH_MISS_RST	The requested object requires upstream authentication, and was not served from the cache; the client connection was reset before the complete response was delivered.
TCP_AUTH_FORM	Forms-based authentication is being used and a form challenging the user for credentials is served in place of the requested content. **Note:** Upon submission of the form, another access log entry is generated to indicate the status of the initial request.
TCP_AUTH_REDIRECT	The client was redirected to another URL for authentication.
TCP_BYPASSED	A TCP-Tunnel connection was bypassed because an upstream ADN concentrator was not discovered; this can occur only when the bypass-if-no-concentrator feature is enabled and all conditions for activating the feature are met.
TCP_CLIENT_REFRESH	The client forces a revalidation with the origin server with a Pragma: no-cache. If the server returns 304 Not Modified, this appears in the Statistics:Efficiency file as In Cache, verified Fresh.
TCP_CLIENT_REFRESH_RST	The client forces a revalidation with the origin server, but the client connection was reset before the complete response was delivered. TCP_DENIED Access to the requested object was denied by a filter.
TCP_ERR_MISS	An error occurred while retrieving the object from the origin server.
TCP_HIT	A valid copy of the requested object was in the cache.
TCP_HIT_RST	A valid copy of the requested object was in the cache, but the client connection was reset before the complete response was delivered.
TCP_LOOP	The current connection is dropped because the upstream connection would result in a looped connection.
TCP_MEM_HIT	The requested object was, in its entirety, in RAM.
TCP_MISS	The requested object was not in the cache.
TCP_MISS_RST	The requested object was not in the cache; the client connection was reset before the complete response was delivered.
TCP_NC_MISS	The object returned from the origin server was noncacheable.
TCP_NC_MISS_RST	The object returned from the origin server was noncacheable; the client connection was reset before the complete response was delivered.
TCP_PARTIAL_MISS	The object is in the cache, but retrieval from the origin server is in progress.
TCP_PARTIAL_MISS_RST	The object is in the cache, but retrieval from the origin server is in progress; the client connection was reset before the complete response was delivered.

Value	Description
TCP_POLICY_REDIRECT	The client was redirected to another URL due to policy.
TCP_REFRESH_HIT	A GIMS request to the server was forced and the response was 304 Not Modified; this appears in the Statistics:Efficiency file as In Cache, verified Fresh.
TCP_REFRESH_HIT_RST	A GIMS request to the server was forced and the response was 304 Not Modified; the client connection was reset before the complete response was delivered.
TCP_REFRESH_MISS	A GIMS request to the server was forced and new content was returned.
TCP_REFRESH_MISS_RST	A GIMS request to the server was forced and new content was returned, but the client connection was reset before the complete response was delivered.
TCP_RESCAN_HIT	The requested object was found in the cache but was rescanned because the virus-scanner-tag-id in the object was different from the current scanner tag.
TCP_RESCAN_HIT_RST	The requested object was rescanned (see TCP_RESCAN_HIT) but the client connection was reset before the complete response was delivered.
TCP_SPLASHED	The user was redirected to a splash page.
TCP_SWAPFAIL	The object was believed to be in the cache, but could not be accessed.
TCP_TUNNELED	The CONNECT method was used to tunnel this request (generally proxied HTTPS).
TCP_WEBSOCKET	The request was a WebSocket upgrade request. You can determine if the traffic was plain WebSocket or secure WebSocket by looking at the scheme (HTTP or HTTPS).
TUNNELED	Successful data transfer operation. TCP_ Refers to requests on the HTTP port.

SQUID-Compatible Format

The SQUID-compatible format contains one line for each request. For SQUID-1.1, the format is: time elapsed remotehost code/status bytes method URL rfc931 peerstatus/peerhost type For SQUID-2, the columns stay the same, though the content within might change a little.

NCSA Common Access Log Format

The common log format contains one line for each request. The format of each log entry is shown below:

```
remotehost rfc931 authuser [date] "request" status bytes
```

Each field is described in the following table.

Field	Description
`authuser`	The username of the authenticated user.
`bytes`	The content length of the document transferred.
`[date]`	Date and time of the request.
`remotehost`	DNS hostname or IP address of remote server.
`"request"`	The request line exactly as it came from the client.
`rfc931`	The remote log name of the user. This field is always –.
`status`	The HTTP status code returned to the client.

Access Log Filename Formats

The following table details the specifiers for the access log upload filenames.

Specifier	Description
%%	Percent sign.
%a	Abbreviated weekday name.
%A	Full weekday name.
%b	Abbreviated month name.
%B	Full month name.
%c	The certificate name used for encrypting the log file (expands to nothing in non-encrypted case).
%C	The appliance name.
%d	Day of month as decimal number (01 – 31).
%f	The log name.
%H	Hour in 24-hour format (00 – 23).
%i	First IP address of the appliance, displayed in x_x_x_x format, with leading zeros removed.
%I	Hour in 12-hour format (01 – 12).
%j	Day of year as decimal number (001 – 366).
%l	The fourth (last) octet in the appliance IP address For example, for the IP address 10.11.12.13, %l would be 13.
%m	Month as decimal number (01 – 12).
%M	Minute as decimal number (00 – 59).
%p	Current locale's A.M./P.M. indicator for 12-hour clock.
%S	Second as decimal number (00 – 59).

Specifier	Description
%U	Week of year as decimal number, with Sunday as first day of week (00 – 53).
%v	Milliseconds; usually used in conjunction with %H%M%S to get more accuracy in the log filename.
%w	Weekday as decimal number (0 – 6; Sunday is 0).
%W	Week of year as decimal number, with Monday as first day of week (00 – 53).
%y	Year without century, as decimal number (00 – 99).
%Y	Year with century, as decimal number.
%z, %Z	Timezone name or abbreviation; no characters if time zone is unknown.

Default Access Log Formats

The following describes the default access log formats and how to create customized access log formats. For a description of each value in the log, see "About Access Log Formats" on page 402.

bcreportercifs_v1

Designed to for Proxy deployments that use ADN to transfer data with CIFS, and send that access information to Symantec Reporter. This format includes the following fields:

```
date time c-ip c-port r-ip r-port s-action s-ip cs-auth-group csusername x-client-connection-bytes x-
server-connection-bytes xserver-adn-connection-bytes
x-cifs-method x-cifs-client-readoperations x-cifs-client-write-operations x-cifs-client-
otheroperations x-cifs-server-operations x-cifs-error-code
x-cifs-server x-cifs-share x-cifs-path x-cifs-orig-path x-cifs-client-bytes-read x-cifs-server-bytes-
read x-cifs-bytes-written x-cifs-uid x-cifs-tid
x-cifs-fid x-cifs-file-size x-cifs-file-type
```

This is a reserved format and cannot be edited.

bcreportermain_v1

Designed to send HTTP access information to Symantec Reporter. This format includes the following fields:

```
date time time-taken c-ip cs-username cs-auth-group s-supplier-name ssupplier-ip s-supplier-country s-
supplier-failures x-exception-id scfilter-result
cs-categories cs(Referer) sc-status s-action cs-rs(Content-Type) cs-uri-scheme cs-host cs-uri-port cs-
uri-path cs-uriquery cs-uri-extension cs(User-Agent)
s-ip sc-bytes cs-bytes x-virusid x-bluecoat-application-name x-bluecoat-application-operation x-
bluecoat-application-groups cs-threat-risk
x-bluecoat-transaction-uuid x-icap-reqmod-header(X-ICAP-Metadata) x-icap-respmod-header(X-
ICAPMetadata)
```

This is a reserved format and cannot be edited.

bcreporterssl_v1

Designed to send HTTPS access information to Symantec Reporter. This format includes the following fields:

```
date time time-taken c-ip cs-username cs-auth-group s-supplier-name ssupplier-ip s-supplier-country s-
supplier-failures x-exception-id scfilter-result
cs-categories sc-status s-action cs-method rs(Content-Type) cs-uri-scheme cs-host cs-uri-port cs-uri-
extension cs(User-Agent) s-ip sc-bytes cs-bytes x-virus-id
x-rs-certificate-observederrors x-cs-ocsp-error x-rs-ocsp-error x-rs-connection-negotiatedcipher-
strength x-rs-certificate-hostname
x-rs-certificate-hostnamecategory cs-threat-risk x-rs-certificate-hostname-threat-risk
```

This is a reserved format and cannot be edited.

bcreporterstreaming_v1

Designed to send streaming media access information to Symantec Reporter. This format includes the following fields:

```
date time time-taken c-ip sc-status s-action sc-bytes rs-bytes csmethodcs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query csusername cs-auth-group
cs(Referer) cs(User-Agent) c-starttime filelength filesize avgbandwidth x-rs-streaming-content x-streamingrtmp-app-name x-streaming-rtmp-stream-name
x-streaming-rtmp-swf-url xstreaming-rtmp-page-url s-ip s-dns s-session-id x-cache-info
```

This is a reserved format and cannot be edited.

bcreporterwarp_v1

Designed to log reverse proxy traffic data.With this data, reverse proxy administrators can run reports for inspecting traffic flow between the appliance and backend web application servers. The bcreporterwarp_v1 access log format is recommended for WAF deployments because it includes detailed information for requests that WAF blocks (or monitors). To populate the Web Application Protection and Geolocation access log fields, you must have valid subscriptions for the respective service and enable the feature on the appliance. This format includes the following fields:

```
date time time-taken c-ip cs-username cs-auth-group x-bluecoattransaction-uuid x-exception-id cs(Referer) sc-status s-action csmethod rs(Content-Type)
cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-uri-extension cs(User-Agent) s-ip sc-bytes cs-bytes xvirus-id x-cs-client-ip-country
x-user-x509-serial-number x-user-x509-subject rs-bytes x-cs-client-effective-ip x-cs-client-effective-ipcountry cs(X-Forwarded-For) rs-service-latency
r-ip x-bluecoatapplication- name x-bluecoat-waf-attack-family x-risk-score x-bluecoatwaf- block-details x-bluecoat-waf-monitor-details
x-bluecoat-requestdetails-header x-bluecoat-request-details-body x-bluecoat-waf-scaninfo
```

> **Note:** `x-bluecoat-request-details-header` includes all header content up to the body of an HTTP request, not just common fields like host and referrer.

This is a reserved format and cannot be edited.

bcsecurityanalytics_v1

Designed to send appropriate log entries to the Security Analytics Platform. This format includes the following fields:

```
date time c-ip c-port s-ip s-port s-source-ip s-source-port r-ip rport s-supplier-ip s-supplier-port r-supplier-ip r-supplier-port rsupplier-dns x-virus-id
cs(Referer) x-cs(Referer)-uri-category xbluecoat-application-name x-bluecoat-application-operation x-bluecoatapplication-groups
```

This is a reserved format and cannot be edited.

For more information on the Security Analytics Platform, refer to documentation: https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/security-analytics/8-1.html

cifs

An ELFF format that includes the following fields:

```
date time c-ip c-port r-ip r-port s-action s-ip cs-auth-group csusername x-client-connection-bytes x-server-connection-bytes x-serveradn-connection-bytes
x-cifs-method x-cifs-client-read-operations xcifs-client-write-operations x-cifs-client-other-operations x-cifsserver-operations x-cifs-error-code x-cifs-server
x-cifs-share x-cifspath x-cifs-orig-path x-cifs-client-bytes-read x-cifs-server-bytesread x-cifs-bytes-written x-cifs-uid x-cifs-tid x-cifs-fid x-cifsfile-size
x-cifs-file-type
```

collaboration

Designed to log WebEx actions. It is an ELFF format that includes the following fields:

```
date time c-ip r-dns duration x-collaboration-method s-action xcollaboration-user-id x-collaboration-meeting-id x-webex-site
```

dns

An ELFF format that includes the following fields.

```
date time time-taken c-ip x-dns-cs-transport x-dns-cs-opcode x-dns-csqtype x-dns-cs-qclass x-dns-cs-dns x-dns-cs-address x-dns-rs-rcode xdns- rs-a-records
x-dns-rs-cname-records x-dns-rs-ptr-records s-ip
```

main

An ELFF format that includes the following fields:

```
date time time-taken c-ip sc-status s-action sc-bytes cs-bytes csmethod cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query csusername cs-auth-group
s-supplier-name rs(Content-Type) cs(Referer) cs(User-Agent) sc-filter-result cs-categories x-virus-id s-ip
```

mapi

An ELFF format that includes the following fields:

```
date time c-ip c-port r-ip r-port x-mapi-user x-mapi-method cs-bytes sr-bytes rs-bytes sc-bytes x-mapi-cs-rpc-count x-mapi-sr-rpc-count xmapi-rs-rpc-count
x-mapi-sc-rpc-count s-action cs-username cs-authgroup s-ip
```

mapi-http

Associated with the MAPI-HTTP protocol (Office 365 traffic) by default. It is an ELFF format that includes the following fields:

```
date time c-ip c-port r-ip r-port x-mail-message-id x-mail-user xmail- operation x-mail-from x-mail-to
x-mail-cc s-action x-mailattachments
x-mail-attachments-removed s-icap-info cs-icap-status rsicap-status x-virus-id x-virus-details x-icap-
error-code x-icap-errordetail
```

ncsa

Reserved format that cannot be edited. The NCSA/Common format contains the following strings:

```
remotehost rfc931 authuser [date] "request" status bytes
```

The ELFF/custom access log format strings that represent the previous strings are:

```
$(c-ip) - $(cs-username) $(localtime) $(cs-request-line) $(sc-status) $(sc-bytes)
```

p2p

An ELFF format that includes the following fields:

```
date time c-ip c-dns cs-username cs-auth-group cs-protocol x-p2pclient-type x-p2p-client-info x-p2p-
client-bytes x-p2p-peer-bytes duration s-action
```

squid

Reserved format that cannot be edited. You can create a new SQUID log format using custom strings. The default SQUID format is SQUID-1.1 and SQUID-2 compatible.

SQUID uses several definitions for its field formats:

```
SQUID-1:time elapsed remotehost code/status/peerstatus bytes method URL
```

```
SQUID-1.1: time elapsed remotehost code/status bytes method URL rfc931 peerstatus/peerhost type
```

SQUID-2 has the same fields as SQUID-1.1, although some of the field values have changed.

ssl

An ELFF format that includes the following fields:

```
date time time-taken c-ip s-action x-rs-certificate-validate-status xrs-certificate-observed-errors x-
cs-ocsp-error x-rs-ocsp-error cs-host s-supplier-name
x-rs-connection-negotiated-ssl-version x-rsconnection-negotiated-cipher x-rs-connection-negotiated-
cipher-size xrs-certificate-hostname
x-rs-certificate-hostname-category x-csconnection-negotiated-ssl-version x-cs-connection-negotiated-
cipher xcs-connection-negotiated-cipher-size
x-cs-certificate-subject s-ip ssitename
```

streaming

An ELFF format that includes the following fields:

```
c-ip date time c-dns cs-uri-scheme cs-host cs-uri-port cs-uri-path csuri-query c-starttime x-duration
c-rate c-status c-playerid cplayerversion c-playerlanguage
cs(User-Agent) cs(Referer) c-hostexe chostexever c-os c-osversion c-cpu filelength filesize
avgbandwidth protocol transport audiocodec videocodec channelURL
sc-bytes c-bytes s-pkts-sent c-pkts-received c-pkts-lost-client c-pkts-lost-net c-pktslost- cont-net
c-resendreqs c-pkts-recovered-ECC c-pkts-recoveredresent
c-buffercount c-totalbuffertime c-quality s-ip s-dns stotalclients s-cpu-util x-cache-user s-session-
id x-cache-info xclient-address s-action
```

Configure an Access Log Format

The appliance ships with several default log formats, which might be sufficient for your needs. If these formats do not meet your needs, you can create a custom or ELFF format and specify the string and other qualifiers used.

To create a custom log format based on a reserved log format, see "Configure a Custom or ELFF Log Format" below. See "Default Access Log Formats" on page 410 to determine which log formats are reserved. Reserved log formats can't be edited or modified.

To view the default and custom log formats that exist on the system, use the following procedure.

Configure a Custom or ELFF Log Format

First, decide what protocols and log formats to use, and determine the logging policy and the upload schedule. Then, perform the following:

- Associate a log format with the log facility.
- Associate a log facility with a protocol and/or create policies for protocol association and to manage the access logs and generate entries in them (if you do both, policy takes precedence).
- Determine the upload parameters for the log facility. For more information, see "Default Access Log Formats" on page 410.

To create a format based on a reserved format, see "Create a Format Based on a Reserved Format" on the next page.

To configure a custom or ELFF format:

1. In the Admin Console, select **Administration > Access Logging > General**.
2. In **Log Formats**, select **Add New Format** to create a new format or select an existing format to edit it.
3. On the Create/Edit Format page, specify the following as needed:

> **Note:** If the page says View Format instead of Create/Edit Format, the format is reserved and can't be edited. See "Create a Format Based on a Reserved Format" on the next page for details.

Setting or field	Description
Format Name	

Setting or field	Description
Format type: ■ **Custom format string** ■ **W3C Extended Log File Format (ELFF) string**	Specifies either a custom log format or an ELFF format. **Note:** ELFF strings can't start with spaces. The access log ignores any ELFF or custom format fields it doesn't understand. In a downgrade, the format still contains all the fields used in the upgraded version, but only the valid fields for the downgraded version display any information.
Log fields	Valid fields to include in the format, separated by spaces.
Test Format	Validates the specified format-string syntax. Test the format every time you add or remove log fields. When testing is complete, the console displays a message indicating whether the syntax is valid, including deprecated fields.
Multiple-valued header policy	Specifies what happens with HTTP headers that have multiple instances: ■ Log all headers ■ Log last header ■ Log all header

Create a Format Based on a Reserved Format

There might be instances where the reserved log format is insufficient for your purposes and requires either a log format extension or reduction. Although the reserved log formats cannot be directly manipulated, you can create new custom log formats based on these reserved log formats.

To create a format based on a reserved format:

1. In the Admin Console, select **Administration > Access Logging > General**.
2. In **Log Formats**, select an existing reserved format.
3. On the View Format page, select the access log fields and copy them to the clipboard. (The log fields appear to be unavailable, but you will be able to select them.)
4. Create a new format using steps in "Configure a Custom or ELFF Log Format" on the previous page, and paste the access log fields to the new format. Make additional edits as needed.

Configure Malicious Content Scanning Services

Refer to the following topics to configure malicious content scanning services:

- "Malicious Content Scanning Services " on the next page
  - "Content Processing Modes " on page 421
  - "Improve User Experience" on page 424
- "ICAP Configuration Overview" on page 430
  - "Configure an ICAP Service" on page 431
- "Configure an ICAP Service" on page 431
- "About ICAP Service Groups" on page 437
  - "Configure an ICAP Service Group" on page 440

"Configure ICAP Feedback and Patience Pages" on page 442

Malicious Content Scanning Services

You can configure the ProxySG appliance to interact with ICAP servers to provide content scanning. The appliance supports ICAP connections with external Content Analysis, ProxyAV, Symantec DLP, and other third party ICAP services. For ease of use, Symantec recommends that you use the Policy Services subscription for content scanning policy.

About Content Scanning

ICAP is an open standard protocol that allows content engines to send HTTP based content to an ICAP server for performing value added services.

An ICAP server can filter, modify, or adaptWeb content to meet the needs of your enterprise. The appliance, when integrated with a supported ICAP server such as Content Analysis and ProxyAV, provides content scanning, filtering, and repair service for Internet-based malicious code, in addition to reducing bandwidth usage and latency.

To eliminate threats to the network, the ProxySG appliance forwards a web request and/or the response to the ICAP server. The ICAP server filters and adapts the requested content, based on your needs, then returns the content to the ProxySG appliance. The scanned and adapted content is then served to the user who requested the content, and stored on the ProxySG appliance object store. For frequently accessed web content, this integrated solution provides defense against malware, along with the benefits of limiting bandwidth usage and latency in the network.

Plain ICAP and Secure ICAP

The transaction between the ProxySG appliance and the ICAP server can be executed using plain ICAP, secure ICAP or both. Plain ICAP is useful for scanning non-confidential data (HTTP).

Secure ICAP is SSL encrypted ICAP and requires an SSL license; both the ProxySG appliance and the ICAP server must support secure ICAP. While Secure ICAP can be used for both HTTP and HTTPS traffic, plain ICAP is faster than secure ICAP because it does not have to deal with any encryption overhead. Therefore, Symantec recommends that you only use secure ICAP when scanning confidential data.

> **Note:** The appliance does not support secure ICAP for the internal Content Analysis service.

- "ICAP Configuration Overview" on page 430

Request Modification and Response Modification

An ICAP server processes web content that is directed to it during proxy policy evaluation. The content that the ICAP server receives can be processed in two modes: request modification and response modification.

See Also:

- "Content Processing Modes " on page 421

Cache and Serve the Object

After an object has been scanned and is determined to be cacheable, the ProxySG appliance caches it and serves it for a subsequent request. When the appliance detects that the cached content has changed on the origin server, it fetches a fresh version, then forwards it to the ICAP server for scanning. If the ProxySG appliance uses policies in the ICAP configuration, the policy applies to content fetches, distributions, refreshes, and pipelined requests. For more information on the <Cache> layer, refer to the *Content Policy Language Reference*.

ICAP v1.0 Features

This section describes the options for the ICAP v1.0 protocol that are provided on the ProxySG appliance.

Sense Settings

The Sense Settings feature allows the ProxySG appliance to query any identified ICAP server running v1.0, detect the parameters, and configure the ICAP service as appropriate. See "Configure an ICAP Service" on page 431. ISTags ISTags eliminates the need to designate artificial pattern version numbers, as was required in v0.95.

Every response from an ICAP v1.0 server must contain an ISTag value that indicates the current state of the ICAP server. For instance, when the pattern/ scanner version of a virus scanner on the ICAP server changes, the ISTag value changes. This change invalidates all content cached with the previous ISTag value and a subsequent request for any content in cache must be refetched from the origin content server and scanned by the ICAP server.

Backing out a virus pattern on the ICAP server can revert ISTags to previous values that are ignored by the ProxySG appliance. To force the appliance to recognize the old values, use the Sense Settings option. See "Configure an ICAP Service" on page 431.

Persistent Connections

New ICAP connections are created dynamically as ICAP requests are received (up to the defined maximum connection limit). The connection remains open to receive subsequent requests. If a connection error occurs, the connection closes to prevent more errors.

Determine Which Files to Scan

In determining which files to scan, this integrated solution uses the content scanning server's filtering in addition to proxy capabilities. The following table describes the supported content types and protocols.

ICAP Server supported content types	Proxy supported protocols	Unsupported content protocols
All or specified file types, based on the file extension, as configured on the server. Examples: .exe (executable programs), .bat (batch files), .doc and .rtf (document files), and .zip (archive files); or specific MIME types.	■ All HTTP objects (uploaded or downloaded) ■ All FTP over HTTP (webftp) objects (uploaded or downloaded) ■ All native FTP objects (uploaded or downloaded) The above is true for both transparent and explicit proxies.	■ Streaming content (for example, RTSP and MMS) ■ Live HTTP streams (for example, HTTP radio streams) ■ CIFS ■ MAPI ■ IM ■ TCP tunnel traffic
	HTTPS connections terminated at a ProxySG appliance	HTTPS connections tunneled through a ProxySG appliance

Whenever an object is requested or being refreshed and it was previously scanned, the appliance verifies whether the pattern file has been updated since it was last scanned. If it was, the object is scanned again, even if the content has not changed. If the content has changed, the object is rescanned.

With the appliance, you can define flexible, yet enterprise-specific content scanning policies, which are discussed in the following two sections.

Content Processing Modes

An ICAP server processes web content that is directed to it during proxy policy evaluation. The content that the ICAP server receives can be processed in two modes – request modification and response modification.

- Request modification (REQMOD): Allows modification of outbound client requests. These requests are sent from the ProxySG appliance to the ICAP server on their way to the origin content server. This is represented in the VPM as **Perform Request Analysis**.
- Response modification (RESPMOD): Allows modification of inbound client requests. These requests are sent from the ProxySG appliance to the ICAP server after the requested content is retrieved from the OCS.
- REQMOD or RESPMOD is an attribute that is specified in the ICAP service, which is configured between the ProxySG appliance and the ICAP server. This is represented in the VPM as **Perform Response Analysis**.

About Response Modification

The ProxySG appliance sends the first part (a preview) of the object to the ICAP server that supports response modification. The object preview includes the HTTP request and response headers, and the first few bytes of the object. After checking those bytes, the ICAP server either continues with the transaction (that is, asks the ProxySG appliance to send the remainder of the object for scanning) or sends a notification to the appliance that the object is clean and opts out of the transaction.

The response modification mode enables scanning of HTTP responses, remote system file retrieval or FTP RETR responses, FTP over HTTP, and SSL-intercepted response data.

Returning the Object to the ProxySG Appliance

For response modification, the returned object can be the original unchanged object, a repaired version of the original object minus a virus, or an error message indicating that the object contained a virus. Each of these responses is configured on the ICAP server, independent of the appliance and the ICAP protocol. If the appliance receives the error message, it forwards the error message to the client and does not save the infected file.

The following diagram illustrates the response modification process flow.

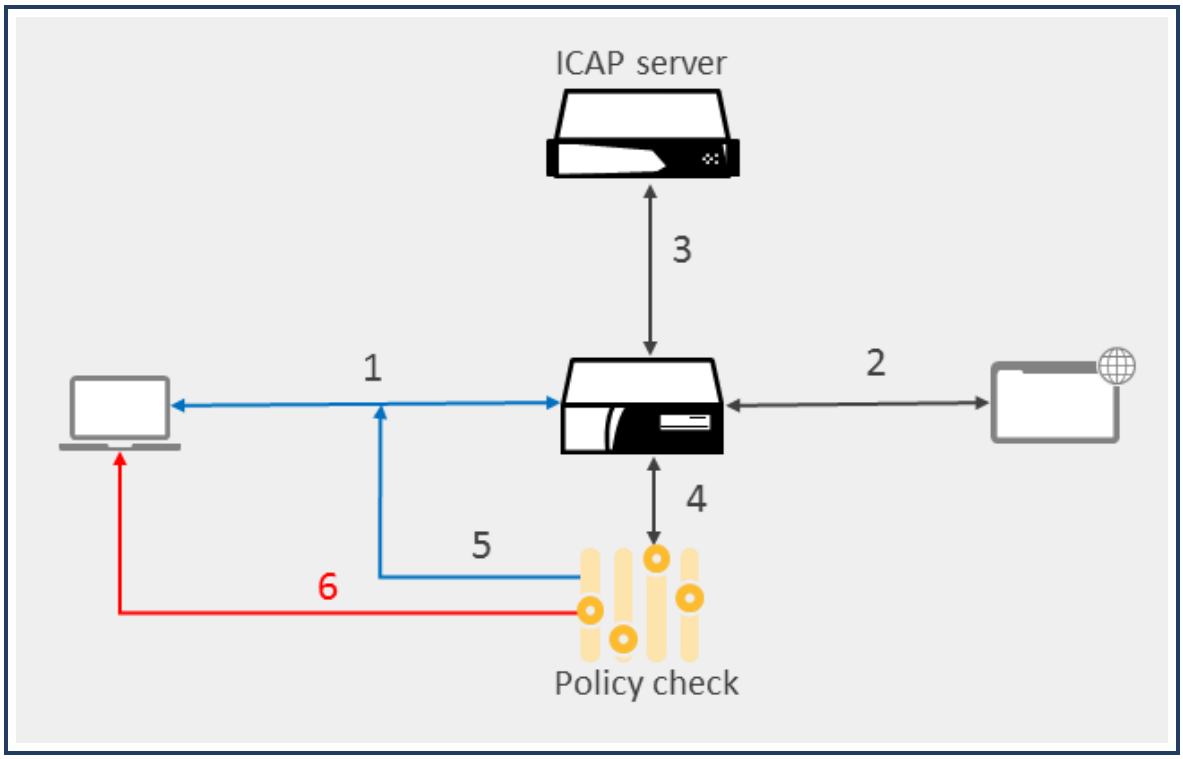

Process flow:

1. (Blue arrow) Client requests a web page.
2. The ProxySG appliance retrieves the page from an OCS.
3. The appliance forwards the web page to the ICAP server for content scanning. The contents (typically modified if exceptions occur) is returned to the appliance. Clean content is stored in the cache, from which future requests are served.
4. Policy check evaluates how exceptions and other errors are processed.
5. (Blue arrow) The content is clean and allowable. The request is allowed to continue to its destination for full retrieval.
6. (Red arrow) A virus or other exception occurred and the client is notified with an HTTP message.

About Request Modification

Request modification means the ICAP server scans contents that a client is attempting to send outside the network. This prevents unaware users from forwarding corrupted files or webmail attachments. Request modification is also a method of content filtering and request transformation, which is used to protect network identification. Based on the results of the scan, the server might return an HTTP response to the client (for example, sports not allowed); or the client request might be modified, such as stripping a referer header, before continuing to the origin content server. Request modification mode enables scanning of HTTP GET requests, PUT requests and POST requests, FTP upload requests and outgoing Webmail.

Note: Some ICAP servers do not support virus scanning for request modification, but support only content filtering.

The following diagram illustrates the request modification process flow.

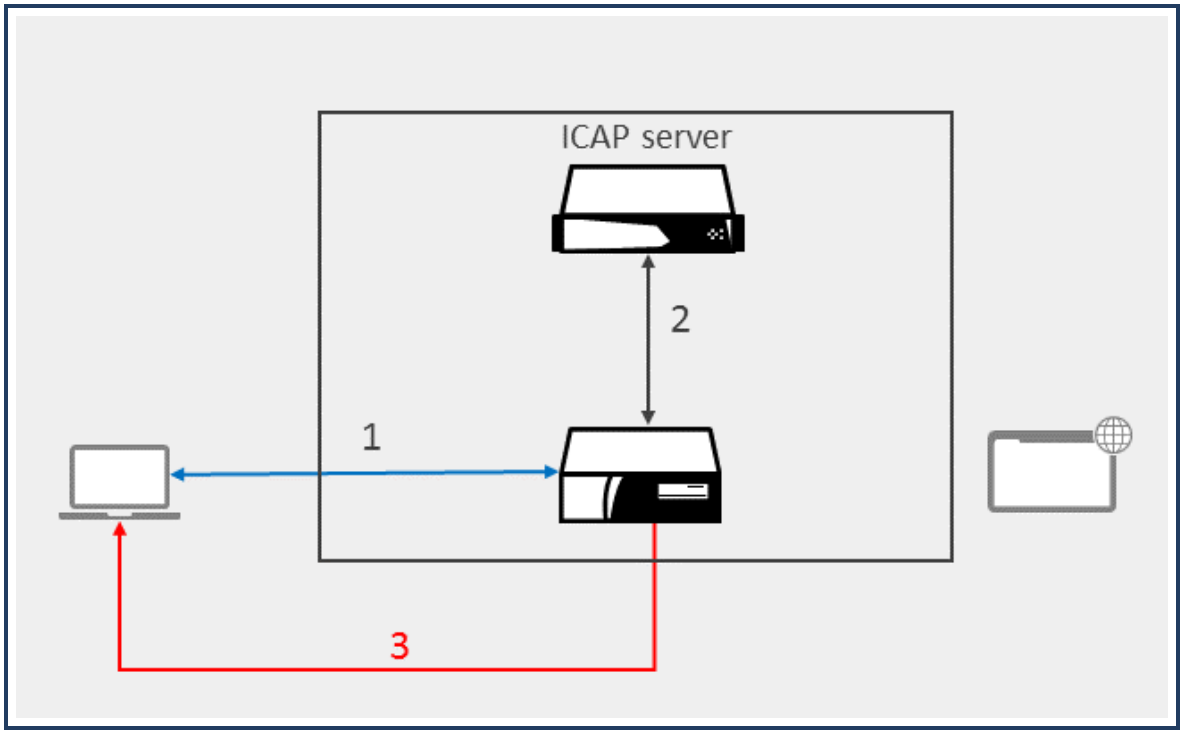

Process flow:

1. (Blue arrow) Client requests a web page.
2. The ProxySG appliance forwards content to the ICAP server for scanning services (which can vary depending on ICAP server configuration). The ICAP server identifies a virus in an attachment and returns the results to the ProxySG appliance.
3. (Red arrow) Client receives an exception message that a infected attachment was identified and stripped.

Improve User Experience

Object scanning adds another operation to the user process of requesting and receiving web content. Therefore, the user might experience extremely slight noticeable delays during web browsing as ICAP servers scan content. The ProxySG appliance allows you to mitigate slower browse times and educate your users about what is occurring on their systems.

About Patience Pages

Patience pages are HTML pages displayed to the user if an ICAP content scan exceeds the specified duration (seconds). You can configure the content of these pages to include a custom message and a help link. Patience pages refresh every five seconds and disappear when object scanning is complete.

Notes

- Patience pages are not compatible with infinite stream connections—or live content streamed over HTTP—such as a cam or video feed. ICAP scanning cannot begin until the object download completes. Because this never occurs with this type of content, the ProxySG appliance continues downloading until the maximum ICAP file size limit is breached. At that point, the ProxySG appliance either returns an error or attempts to serve the content to the client (depending on fail open/closed policy). However, even when configured to fail open and serve the content, the delay added to downloading this large amount of data is often enough to cause the user give up before reaching that point.
- Patience pages are limited to web browsers.

See Also:

- "Configure ICAP Feedback and Patience Pages" on page 442

About Data Trickling

Patience pages provide a solution to appease users during relatively short delays in object scans. However, scanning relatively large objects, scanning objects over a smaller bandwidth pipe, or high loads on servers might disrupt the user experience because connection time-outs occur. To prevent such time-outs, you can allow data trickling to occur. Depending on the trickling mode you enable, the ProxySG appliance either trickles—or allows at a very slow rate—bytes to the client at the beginning of the scan or near the very end.

The appliance begins serving server content without waiting for the ICAP scan result. However, to maintain security, the full object is not delivered until the results of the content scan are complete (and the object is determined to not be infected).

> **Note:** This feature is supported for the HTTP proxy only; FTP connections are not supported.

Trickling Data From the Start

In trickle from start mode, the ProxySG appliance buffers a small amount of the beginning of the response body. As the ICAP server continues to scan the response, the appliance allows one byte per second to the client.

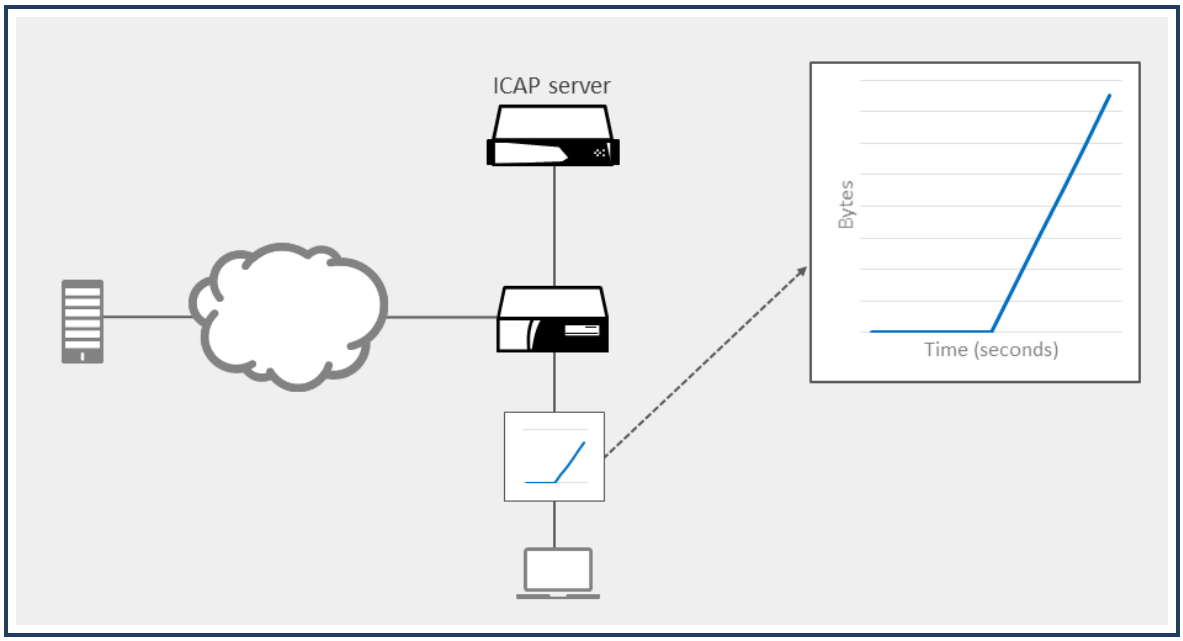

After the ICAP server completes its scan:

- If the object is deemed to be clean (no response modification is required), the ProxySG appliance sends the rest of the object bytes to the client at the best speed allowed by the connection.
- If the object is deemed to be malicious, the ProxySG appliance terminates the connection and the remainder of the response object bytes—which in this case are the majority of the bytes—are not sent to the client.

Deployment Notes

- This method is the more secure option because the client receives only a small amount of data pending the outcome of the virus scan.
- One drawback is that users might become impatient, especially if they notice the browser display of bytes received. They might assume the connection is poor or the server is busy, close the client, and restart a connection.

Trickling Data at the End

In trickle at end mode, the ProxySG appliance sends the response to the client at the best speed allowed by the connection, except for the last 16 KB of data. As the ICAP server performs the content scan, the appliance allows one byte per second to the client.

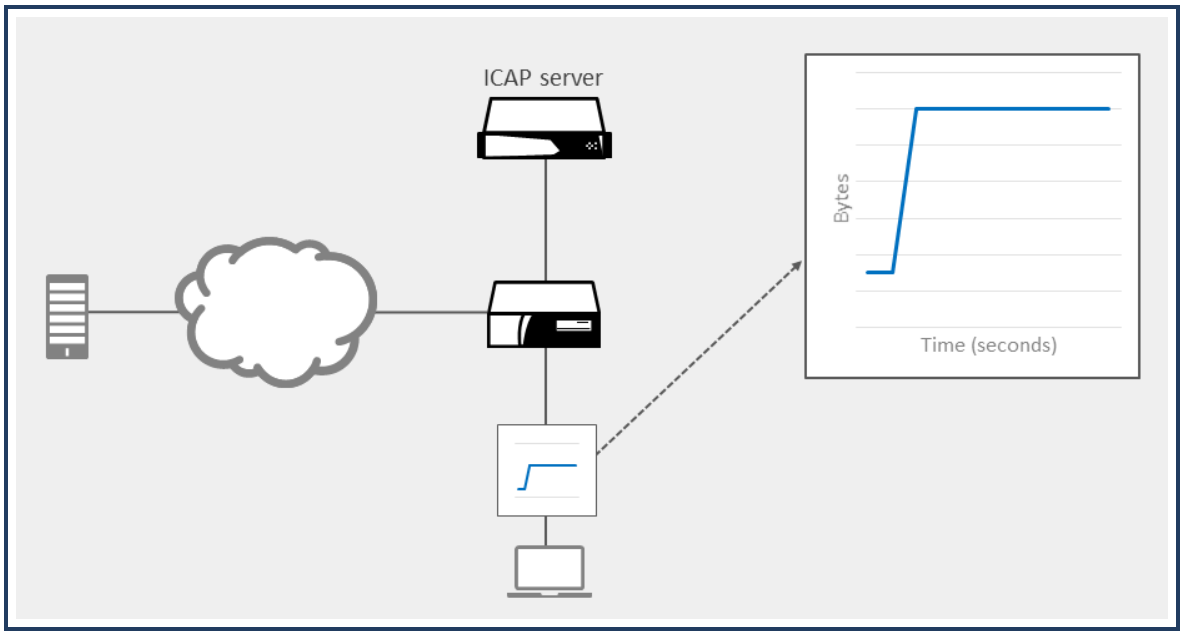

After the ICAP server completes its scan, the behavior is the same as described in "Trickling Data From the Start " on page 424.

Deployment Notes

- Symantec recommends this method for media content, such as flash objects.
- This method is more user-friendly than trickle at start. This is because users tend to be more patient when they notice that 99% of the object is downloaded versus 1%, and are less likely to perform a connection restart. However, network administrators might perceive this method as the less secure method, as a majority of the object is delivered before the results of the ICAP scan.

Deciding between Data Trickling and Patience Pages

ProxySG appliance configuration options plus policy allow you to provide different ICAP feedback actions depending upon the type of traffic detected:

- Symantec defines interactive as the request involving a web browser. web browsers support data trickling and patience pages.
- Non-interactive traffic originates from non-browser applications, such as automatic software download or update clients. Such clients are not compatible with patience pages; therefore, data trickling or no feedback are the only supported options.

Based on whether the requirements of your enterprise places a higher value either on security or availability, the ProxySG appliance allows you to specify the appropriate policy. However, you must also consider the user agents involved when determining the appropriate feedback method. For example, streaming clients cannot deliver patience pages, but they are susceptible to connection timeouts. Therefore, trickling is the suggested method. The following diagram provides basic guidelines for deciding which feedback method to implement.

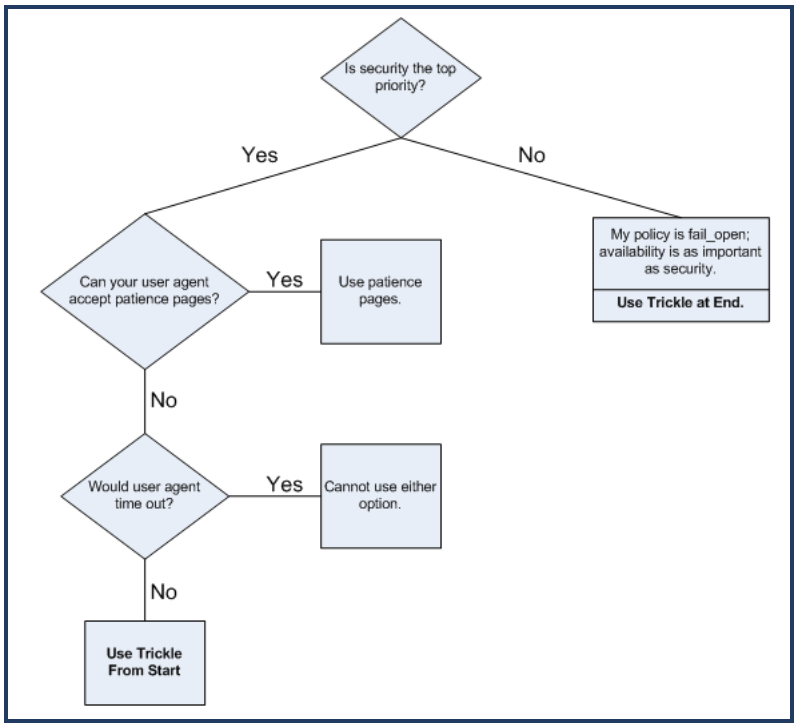

Recommendations for Proxy Chaining Deployments

Proxy chaining deployments are common in enterprises, especially in core/ branch office scenarios. Data trickling is achievable, but behavior is dependent upon how the ProxySG appliances are configured. The following are common deployment scenarios.

- The downstream appliance is performing ICAP scanning, and the upstream appliance is not: Data trickling and patience pages are not affected in this scenario.
- The upstream appliance is performing ICAP scanning, and the downstream appliance is not: The only issue with this deployment is that user agent-specific policy cannot be applied at the core ProxySG appliance because the branch appliance consolidates multiple client requests in one out-going request to the upstream appliance. If data trickling is employed at the upstream appliance and if ICAP scanning detects a virus, the upstream appliance resets the client connection. This also deletes the corrupted object from the downstream appliance cache.
- Both ProxySG appliances (upstream and downstream) are scanning: Behavior is mostly determined by the configuration of the upstream ProxySG appliance.
    - If the upstream appliance is configured to deliver patience pages, then the downstream appliance also attempts to serve patience pages, including to non-graphical user agents. Therefore, this method is not recommended.
    - If the upstream appliance employs data trickle from start, the downstream appliance is not able to send any bytes to the client for a long period of time. If a patience page is not configured on the downstream appliance, users might experience connection time-outs.

- If the upstream appliance employs trickle at end, the downstream appliance allows for all options of patience page and data trickling.

Avoid Network Outages due to Infinite Streaming Issues

Infinite streams are connections such as web cams or flash media–traffic over an HTTP connection–that conceivably have no end. Characteristics of infinite streams may include no content length, slow data rate and long response time. Because the object cannot be fully downloaded, the ICAP content scan cannot start; however, the connection between the ProxySG appliance and the Content Analysis or ProxyAV appliance remains, which wastes finite connection resources.

The deferred scanning feature (enabled by default) solves the infinite streaming issue by detecting ICAP requests that are unnecessarily holding up ICAP connections and defers those requests until the full object has been received.

How Deferred Scanning Works

Deferred scanning detects the possibility of infinite streams by the fact that the number of ICAP resources in use has reached a certain threshold. It then defers the scanning of those streams by deferring the oldest, outstanding ICAP requests first. For every new ICAP request, the ProxySG appliance does the following:

- If the total number of outstanding ICAP actions for the current server has reached the defer threshold, the ProxySG appliance defers the oldest ICAP connection that has not yet received a full object.
- The defer threshold is specified by the administrator as a percentage. For example, if the defer threshold is set to 70 percent and the maximum connections are set to 100, then up to 70 connections are allowed before the ProxySG appliance begins to defer connection which have not finished downloading a complete object.

  > **Note:** See "Configure an ICAP Service" on page 431 for information about setting the defer scanning threshold value in the Admin Console.

- When an ICAP connection is deferred, the connection to the ICAP server is closed. The application response continues to be received and when the download is complete the ICAP request is restarted. The new ICAP request may still be queued if there are no available ICAP connections. After a request is deferred, ICAP waits to receive the full object before restarting the request. If there is a queue when a deferred action has received a complete object, that action is queued behind other deferred actions that have finished. However it will be queued before other new requests.

Deferred Scanning and Setting the Feedback Options

Depending on how you configure the ICAP feedback option (patience page or data trickling) and the size of the object, deferred scanning might cause a delay in ICAP response because the entire response must be sent to the ICAP server at once. The feedback option allows you to specify the type of feedback you want to receive during an ICAP scan. For information about setting feedback options, see "Configure ICAP Feedback and Patience Pages" on page 442.

If a patience page is configured, the browser continues to receive a patience page until the object is fully received and the outstanding ICAP actions have completed.

If the data trickle options are configured, the object continues to trickle during deferred scanning. However, because of the trickle buffer requirement, there might be a delay, with or without deferred scanning, before the ProxySG appliance starts sending a response.

About ICAP Server Failover

When creating an ICAP action, you can specify a list of ICAP servers or groups to use, in order of preference. If the first server or group in the list does not pass the health checks, the ProxySG appliance moves down the list until it finds a server or group that is healthy and uses that to perform the scanning. The primary server resumes ICAP processing when the next health check is successful; the standby server or server group does not retain the primary responsibility.

Notes

- Failover is configured as part of the ICAP policy definition.
- You cannot configure failover policy until ICAP services are configured on the ProxySG appliance.
- To avoid errors, ICAP service names cannot be named fail_open or fail_closed (the CLI commands prevent these names from being created).

ICAP Configuration Overview

Configuring ICAP on the appliance consists of the following tasks:

Task	Description
Install and configure the ICAP server.	Follow the manufacturer instructions for installing the ICAP server, including any configuration necessary to work with the ProxySG appliance. Based on your network environment, you might use the ProxySG appliance with multiple ICAP servers or multiple scanning services on the same server. Configure options as needed, including the exception message displayed to end users when a requested object is modified or blocked.
Determine whether to scan data using plain ICAP or secure ICAP.	Scan data using the plain ICAP method, secure ICAP method, or both. ■ Plain ICAP should be used only for non-confidential data. In particular, if plain ICAP is used for intercepted HTTPS traffic, then data intended to be cryptographically secured would be transmitted in plain text on the local network. ■ Secure ICAP send data through a secure data channel. This method protects the integrity of messages that are sent between the ProxySG appliance and the ICAP server while it allows users to authenticate ICAP servers by enabling certificate verification.
(Optional; secure ICAP only) Select the default SSL profile or create an SSL device profile on the ProxySG appliance.	An SSL device profile is required to authorize the ICAP server, if you use secure ICAP. **Note:** If you create an SSL device profile, instead of using a built-in device profile, ensure that the ICAP server certificate is installed as trusted under External certificates. Otherwise, when the Verify Peer option is enabled, the ProxySG appliance fails to verify the ICAP server as a trusted server.
Create and configure new or existing ICAP services on the ProxySG appliance. For information on ICAP content processing modes, see "Content Processing Modes " on page 421.	Create an ICAP service that specifies the ICAP server IP address, supported connection method, content processing mode and select deferred scanning, if desired. See "Configure an ICAP Service" on the facing page.
Specify the feedback method.	Select patience pages or data trickling for feedback method. See "Configure ICAP Feedback and Patience Pages" on page 442.
Add ICAP rules to policy.	Depending on your network needs, add ICAP rules to policy and install the policy file on the ProxySG appliance.

Configure an ICAP Service

An ICAP service is a collection of attributes that defines the communication between the ProxySG appliance and the external ICAP server. It includes the server IP address or hostname, ICAP scanning method, and a host of other options including the supported number of connections.

You must create an ICAP service for each ICAP server or scanning service. For example, if you are using the appliance with multiple ICAP servers or multiple scanning services (RESPMOD or REQMOD) on the same server, add an ICAP service for each server and RESPMOD or REQMOD service.

Similar ICAP scanning services can then be grouped together to create a service group that helps distribute and balance the load of scanning requests. Further, each ICAP service or service group can be accessed through VPM or CPL to configure policy for better administrative control.

The following instructions describe how to create an ICAP service for any supported third-party ICAP server.

To configure an ICAP service:

1. In the Admin Console, select **Administration > ICAP > ICAP Services**.
2. In ICAP Services, select **Add > ICAP Service** to add a new service, or select an existing ICAP service to edit.
3. On the Add/Edit Service page, enter or edit the following settings as required:

Task	Setting or field	Description
Configure ICAP communication settings	**Service Name**	Unique name for the ICAP service.

Task	Setting or field	Description
	Service URL	ICAP server URL, which includes the URL schema, ICAP server hostname, or IP address. For example, the ProxyAV and Content Analysis appliances support this type of URL: icap://10.x.x.x/avscan If your ICAP server is something other than ProxyAV or Content Analysis, check with your ICAP vendor for the appropriate URL format.
	ICAP Version	The default ICAP version is 1.0 and can't be changed.
	Service Type	Identify the ICAP service type as being **Threat Protection**, **DLP**, or **Other**. Use this service if you are migrating appliance policy to the Symantec Web Security Service; however, policy rules that reference an Other service are not enforceable in the Web Security Service, with the exception of rules that reference an Other service that has the name ProxyAV. Services with the name ProxyAV are Threat Protection policy and can be migrated.
	Maximum number of connections	Maximum possible connections at any given time that can occur between the ProxySG appliance and the ICAP server. The range is a number from 5 to 4096. The default is 25. The number of recommended connections depends on the capabilities of the ICAP server. Refer to the vendor's product information.

Task	Setting or field	Description
	Connection timeout	Number of seconds the ProxySG appliance waits for replies from the ICAP server. This timeout is the duration for which the TCP connection between the ProxySG appliance and the ICAP server is maintained. It helps verify the responsiveness of the ICAP server and prevents users from experiencing unnecessary delays. The default timeout is 70 seconds, and you can enter a value in the range of 5 to 4096. **Note:** The connection timeout value does not measure how much of the scanning process is complete, it is a mechanism for ensuring that the communication between the appliances is alive and healthy. The details of the interaction between the ProxySG appliance and the ICAP server can only be viewed through a packet capture. If the ICAP server does not respond within the configured timeout value, by default, the user will not receive the requested content.
	Defer scanning at threshold	Switch the toggle to set the threshold at which the ProxySG appliance defers the oldest ICAP connection that has not yet received a full object. In the **Threshold** field, enter a range between 0 and 100 percent. The default is 80 percent. By default, the deferred scanning threshold is enabled when an ICAP service is created.
	Notify administrator when virus detected (Email msg)	Switch the toggle to notify the administrator when a virus is found.
	"Virus found" page	Page to show the client when a virus is detected: ■ **Use vendor's "virus found" page from the ICAP server**: Display the default vendor error exception page. ■ **Use ProxySG's exception page**: Display the appliance's exception page.

Task	Setting or field	Description
Configure ports At least one of these options must be selected.	**Support plain ICAP connections**	Service supports plain ICAP connections to use plain ICAP. Use plain ICAP when you are scanning plain data (HTTP). In this case, if the HTTPS proxy is enabled on the ProxySG appliance, the data is decrypted first on the ProxySG appliance and then sent to the ICAP server. In the **Plain ICAP port** field, enter a port number. The default port is 1344.
	Support secure ICAP connections	Service supports secure ICAP connections to use secure ICAP. Use secure ICAP when you are scanning sensitive or confidential data (HTTPS). In the **Secure ICAP port** field, enter a port number. The default port is 11344. For secure ICAP, a valid SSL profile is required. Select one in the **SSL Device Profile** field. This associates an SSL device profile with the secure ICAP service. **Note:** If you do not select an SSL device profile you cannot use secure ICAP connections. The SSL device profile can be customized for your environment.

Task	Setting or field	Description
Configure server settings	**Sense Settings**	Automatically configure the ICAP service using the ICAP server parameters. **Note:** This function assumes a deployment with one ProxySG appliance sending data to one ICAP server. If your deployment has multiple proxies using one ICAP server, you might have to modify the **Maximum number of connections** value accordingly.
	Notification method	Select the ICAP method: **Response** modification or **Request** modification. An ICAP server might have separate URLs for response modification and request modification services.
	Enable Preview	Enable only in one of the following situations: ■ (Only for RESPMOD service) If you are using file scanning policies based on file extensions on the ProxyAV appliance, enter 0 in the **Preview size** field, and select enabled. With a 0 bytes preview size, only response headers are sent to the ICAP server; more object data is only sent if requested by the ICAP server. ■ If you have enabled the Kaspersky Apparent Data Types feature on the ProxyAV appliance, enter a value (512 is recommended) in the **Preview size** field, and select enabled. The ICAP server reads the object up to the specified byte total. The ICAP server either continues with the transaction (that is, receives the remainder of the object for scanning) or opts out of the transaction.
	Send	(Optional) Allow additional information to be forwarded to the ICAP server. Select one or more of the following: ■ **Client Address** ■ **Server Address** ■ **Authenticated User** ■ **Authenticated Groups**
Perform health check	**Perform health check**	Perform an immediate health check on this service.

See Also:

- "Configure an ICAP Service Group" on page 440

About ICAP Service Groups

You can create and manage external ICAP service groups. In high-traffic network environments, a service group accelerates response time by a performing a higher volume of scanning.

A ProxySG ICAP service is a named entity that identifies the ICAP server, the ICAP method, and the supported number of connections. A service group is a named set of ICAP services. You will need to create service groups when you are using multiple ICAP servers to process a large volume of scanning requests. The following example shows a service group of three ICAP servers:

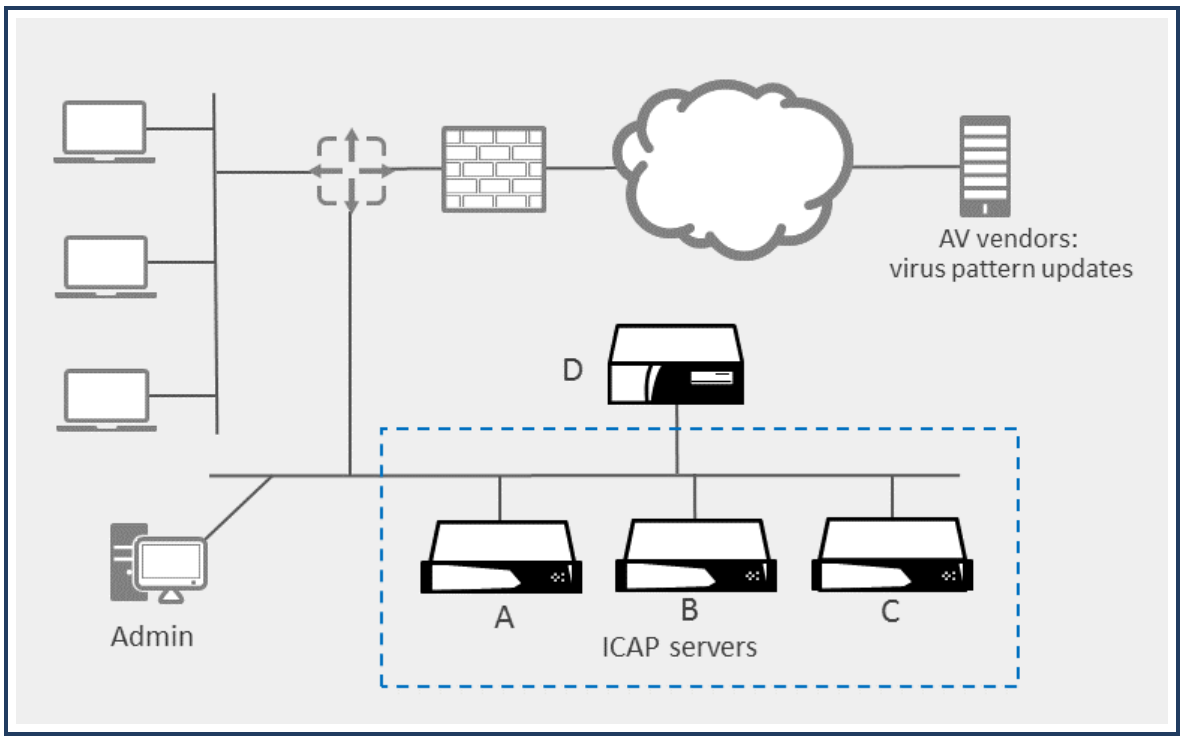

In the previous example:

a. ICAP server has a maximum of 10 connections and a specified weight of 1.

b. ICAP server has a maximum of 10 connections and a specified weight of 1.

c. ICAP server has a maximum of 25 connections and a specified weight of 3.

d. ProxySG appliance has a Service Group that contains ICAP servers A, B, and C.

To help distribute and balance the load of scanning requests when the ProxySG is forwarding requests to multiple services within a service group, the ProxySG uses an intelligent load balancing algorithm. When deciding which service in the service group to send a scanning request, this algorithm takes into consideration the following factors:

- Number of requests that are in a waiting state on each service (a request is in this state when it has been sent to the service but the response hasn't been received)
- Number of unused connections available on each service (calculated by subtracting the number of active transactions from the connection maximum on the server)
- The user-assigned weight given to each server (see below)

Weighting

Weighting determines what proportion of the load one server bears relative to the others when transactions are waiting to be scanned. (The waiting transactions are typically large file downloads.) If all servers have either the default weight (1) or the same weight, each share an equal proportion of the load when transactions are waiting. If one server has weight 25 and all other servers have weight 50, the 25-weight server processes half as much as any other server.

Before configuring weights, consider the capacity of each server. The processing capacity of the server hardware in relationship to other servers (for example, the number and performance of CPUs or the number of network interface cards) could affect assigned weight of a ICAP server.

Having appropriate weights assigned to your services is critical when all servers in a service group have waiting transactions. As servers reach their capacity, proper weighting is important because requests are queued according to weight.

One technique for determining weight assignments is to start out by setting equal weights to each service in a group; then, after several thousand requests, make note of how many requests were handled by each service. For example, support there are two services in a group: Service A handled 1212 requests, Service B handled 2323. These numbers imply that the second service is twice as powerful as the first. So, the weights would be 1 for Service A and 2 for Service B.

Setting the weight value to 0 (zero) disables weighted load balancing for the ICAP service. Therefore, if one ICAP server of a two-server group has a weight value of 1 and the second a weight value of 0, should the first server go down, a communication error results because the second server cannot process the request.

Load Balancing Examples

When load balancing between services, how does the ProxySG appliance decide which ICAP service to send a scanning request to? For each service, it calculates an index by dividing the number of waiting transactions by the server weight (think of this as wait/weight). The ICAP service with the lowest index value handles the new ICAP action, assuming that the service has an available connection to use. If it does not, it sends the request to the service with the next lowest index value that has a free connection.

> **Note:** If there are no transactions waiting, load balancing using the assigned weights does not take effect.

Load will be distributed among services proportionally according to their configured weights until the maximum connection limit is reached on all services.

Example 1

Service A and B are in the same service group.

- Service A can handle up to 50 connections, is assigned a weight of 1, has 17 active transactions, with 5 transactions in the waiting state. The index is calculated by dividing the wait by the weight: 5/1 = 5.
- Service B can handle up to 100 connections, is assigned a weight of 2, has 17 active connections, with 15 waiting transactions. The index is 15/2 = 7.5.

The appliance assigns the next ICAP action to Service A because it has a lower index.

Example 2

Service C and D are in the same service group.

- Service C can handle up to 5 connections, is assigned a weight of 1, has 5 active transactions, with 1 transaction in the waiting state. The index is 1/1=1.
- Service D can handle up to 10 connections, is assigned a weight of 1, has 7 active transactions, with 5 waiting transactions. The index is 5/1=5.

Although Service C has a lower index than Service D, it does not have any available connections; therefore, the appliance assigns the next ICAP action to Service D which has several free connections.

Configure an ICAP Service Group

You can create a service group and add the relevant ICAP services to the group. When instructed by configured ICAP policies, the appliance sends ICAP response modification requests to ICAP servers in the service group. The load carried by each service in the group is determined by the weight values.

To configure an ICAP service group:

1. In the Admin Console, select **Administration > ICAP > ICAP Services**.
2. In ICAP Services:
   - To add a new group, select **Add > ICAP Service Group**.
   - To edit an existing group, select it in the list. Group names are in the format "Group: *name*".

   > **Note:** You can expand a group to show the services that are members of it.

3. (If editing a service group) In **ICAP Services**, the **Group Members** section shows all ICAP services that belong to the group. Edit the group as needed:
   - In the **Weight** field of each ICAP service, assign a weight value. The valid range is 0-255.

     See "Weighting" on page 438 for information on using weight to configure load balancing.
   - Select **Remove** to remove a service from the group.
4. (When editing a group or adding a new one) On the Add/Edit Service page, enter or edit the following settings as required:

Task	Setting or field	Description
Set group properties	**Group Name**	Unique name for the group. On the **ICAP Services** page, the group name is displayed as "Group: *name*", where *name* is the group name specified here.
Configure group membership	**Select Services**	Selecting this option opens the **Select ICAP Services** dialog. The dialog displays all configured services on the appliance. - To remove a service from the group, clear its checkbox. - To add a service to the group, select its checkbox.

Display Content Analysis and Group Information

After configuring a service group, you can display aggregate service group (and other Content Analysis) information.

To display information about all Content Analysis services and groups, use the following CLI commands:

```
# (config) content-analysis
# (config content-analysis) view
```

Individual service information is displayed first, followed by service group information, such as:

```
; Content Analysis
ICAP-Version: 1.0
URL: icap://10.9.59.100/
Plain-ICAP-enabled: yes
Plain-ICAP-port: 1344
Secure-ICAP-enabled: no
Secure-ICAP-port: none
Ssl-device-profile: none
Max-conn: 25
Timeout(secs): 70
Defer-threshold: 80%
Notification: virus-detected
Use ICAP Vendor's virus page: disabled
Event-log: connection-failure
Methods: RESPMOD
Preview-size: 0
Send: nothing
ISTag:
Last-ISTag-change: never
```

Configure ICAP Feedback and Patience Pages

You can specify what type of feedback is provided to users during an ICAP scan and customize patience pages, as described in "Improve User Experience" on page 424.

Configure ICAP Feedback for Interactive and Non-Interactive Traffic

To configure ICAP feedback:

1. In the Admin Console, select **Administration > ICAP > ICAP Services**.
2. In ICAP Feedback, enter or edit the following settings as required. These options apply to both interactive and non-interactive traffic.

Setting or field	Description
Provide ICAP feedback during ICAP scanning for interactive traffic and **Provide ICAP feedback during ICAP scanning for non-interactive traffic**	If this option isn't selected, users are not notified that delays in receiving content are due to ICAP scanning. When this option is selected, the following fields become available: ■ **Provide feedback after**: Amount of time to wait before notifying client. Range is between 5 and 65535 seconds. If you make the value too large, users might become impatient and manually close the client, believing the connection has stopped responding. Default is 5 seconds. ■ **Feedback type**: One of the following feedback methods: ◦ **Return patience page**: (Available for interactive feedback only) The client displays a web page to the user providing a description of the delay (ICAP scanning). This page is customizable, as described in "Customize Patience Pages" on the next page. **Note:** When the deferred scanning option is enabled and a patience page is configured, the browser continues to receive a patience page until the object is fully received and the outstanding ICAP actions have completed. ◦ **Trickle object data from start**: The client receives 1 byte per second, which should prevent connection time-outs while the ICAP server performs the scan. If the response from the ICAP server is clean, the client receives the rest of the object data at the best connection speed possible. If the scan detects malicious content, the connection is dropped. This is the more secure method ◦ **Trickle object data at end**: The client receives most (99%) of the object data, but the final bytes are sent at the rate of one per second while the ICAP scanner performs the scan. If the response from the ICAP server is clean, the client receives the rest of the object data at the best connection speed possible. If the scan detects malicious content, the connection is dropped. This is the least secure method, as most of the data has already been delivered to the client. However, this method provides the best user experience because there most of the object is already delivered. **Note:** When deferred scanning is enabled and the data trickle options are configured, the object continues to trickle during deferred scanning. However, due to the trickle buffer requirement, there may be a delay before the ProxySG appliance begins sending a response.

> **Note:** These configurations are global. You can define further feedback policy that applies to specific user and conditional subsets. In the VPM, the object is located in the Web Access Layer: **Return ICAP Feedback**.

Customize Patience Pages

This section describes how to customize text displayed during ICAP scanning. Patience pages are displayed if the appropriate option is selected, as described in "Improve User Experience" on page 424.

To customize HTTP/FTP patience pages:

1. In the Admin Console, select **Administration > ICAP > ICAP Services**.
2. In **ICAP Patience Page**, you can edit the FTP and HTTP patience pages. Select an element and edit it as follows:

Patience page element	Example
FTP Patience Page	For content over FTP, the patience text displayed to FTP clients during an ICAP scan can be modified.
HTTP Patience Page - Header	Contains HTML tags that define what displays in the dialog title bar.
HTTP Patience Page - Summary	HTML and text that informs users that a content scan is occurring.
HTTP Patience Page - Details	Uses data to indicate scanning progress. The information includes the URL currently being scanned, the number of bytes processed, and the elapsed time of the scan.
HTTP Patience Page - Help	Displays instructions for users should they experience a problem with the patience page.

Interactivity Notes

- When ICAP scanning is enabled and a patience page is triggered, a unique URL is dynamically generated and sent to the browser to access the patience page. This unique URL might contain a modified version of the original URL. This is expected behavior.
- Patience pages and exceptions can only be triggered by left-clicking a link. If a user right-clicks a link and attempts to save it, it is not possible to display patience pages. If this action causes a problem, the user might see browser-specific errors (for example, an Internet site not found error); however, ICAP policy is still in effect.
- A patience page is not displayed if a client object request results in an HTTP 302 response and the ProxySG appliance

pipelines the object in the Location header. After the appliance receives the client request for the object, the client enters a waiting state because a server-side retrieval of the object is already in progress. The wait status of the client request prevents the patience page from displaying. To prevent the appliance from pipelining these requests (which decreases performance) and to retain the ability to provide a patience page, configure HTTP as follows:

```
# (config) http no pipeline client redirects
```

- The status bar update does not work if it is disabled or if the JavaScript does not have sufficient rights to update it.
- Looping: Certain conditions cause browsers to re-spawn patience pages. For example, a site states it will begin a download in 10 seconds, initiates a pop-up download window, and returns to the root window. If the download window allows pop-ups, the patience page displays in a separate window. The automatic return to the root window initiates the download sequence again, spawning another patience page. If unnoticed, this loop could cause a system hang. The same behavior occurs if the user clicks the back button to return to the root window. For known and used download sites, you can create policy that redirects the page so that it doesn't return to the root window after a download starts.

Verify Service Health and Status

Refer to the following to learn about appliance health checks:

- "Overview of Health Checks" on the facing page
  - "Background DNS Resolution Tests" on page 452
  - "Default Health Check Settings" on page 453
  - "Override Health Check Default Settings" on page 455
  - "Configure Health Check Notifications" on page 458
- "Composite Health Checks" on page 462
- "Verify Health Check Settings" on page 465

Overview of Health Checks

> **Note:** Not all health check tests, types, and configuration are available in the Admin Console. Refer to the Management Console for full access.

The ProxySG appliance performs health checks to test for network connectivity and to determine the responsiveness of external resources. Examples of external resources include: DNS servers, forwarding hosts, SOCKS gateways, authentication servers, and ICAP services (for example, anti-virus scanning services).

System-defined health checks are based on:

- Forwarding configuration
- SOCKS gateways configuration
- DNS server configuration
- ICAP service configuration
- Authentication realm configuration
- Whether WebPulse is enabled

You also can create user-defined health checks, including a composite health check that combines the results of multiple other health check tests. For information on health check types, see "About Health Check Components" on the next page.

Health checks fall into three broad categories:

- Determining if the IP address can be reached. Health check types that fall into this category are:
  - Forwarding hosts
  - SOCKS gateways
  - User-defined host health checks
- Determining if a service is responsive. Health check types that fall into this category are:
  - Authentication servers
  - DNS server
  - WebPulse service
  - ICAP services
- Determining if a group is healthy. Group tests are compilations of individual health checks, and the health of the group is determined by the status of the group members. Health check types that fall into this category are:

- Forwarding groups
- SOCKS gateway groups
- ICAP service groups
- User-defined composite health checks

Information provided by health checks allows you to accomplish the following:

- Detect potential network issues before they become critical. For example, if the health check for an individual host fails, the appliance sends an alert (using e-mail, SNMP, or by writing to an event log) to the designated recipients, if configured.
- Track response times and report failures. For example, if the DNS server performance suffers a reduction, the users experience response time delays. The DNS health check records the average response time (in milliseconds) and allows you to interpret the reason for the performance reduction. Should the DNS server become unavailable, the failed health check triggers an alert.

Furthermore, the appliance uses health check information to accomplish the following:

- When combined with failover configurations, health checks redirect traffic when a device or service failure occurs. For example, a health check detects an unhealthy server and a forwarding rule redirects traffic to a healthy server or proxy.
- Monitor the impact of health check states on the overall health of the appliance.

About Health Check Components

Health checks have two components:

- Health check type, which is the kind of device or service the specific health check tests. The following types are supported:
  - Forwarding host and forwarding group
  - SOCKS gateway and SOCKS gateway group
  - DNS servers
  - External Authentication servers
  - ICAP service and ICAP service group
  - Dynamic Real-Time Rating Service
  - User-defined host and composite health checks

  See "Health Check Types" on the facing page for more information.

- Health check test, which is the method of determining network connectivity, target responsiveness, and basic functionality:

- Health checks (external targets):
  - Authentication
  - ICMP
  - DNS
  - TCP
  - SSL
  - HTTP
  - HTTPS
  - ICAP
  - WebPulse
- Health checks (group targets):
  - Groups
  - Composite

> **Note:** Some health checks (such as forwarding hosts and SOCKS gateways) can be configured to report the result of a composite health check instead of their own test.

Some health check types only have one matching test, while others have a selection. For more information about health check types and tests, see "Health Check Tests" on the next page.

Health Check Types

Most health checks are automatically created and deleted when the underlying entity being checked is created or deleted. When a forwarding host is created, for example, a health check for that host is created. Later, if the forwarding host is deleted, the health check for it is deleted as well. User interaction is not required, except to change or customize the health check behavior if necessary. However, if a health-check is referenced in policy, you cannot delete the corresponding host or the health check itself until the reference in policy is deleted. In addition to the automatically generated health checks generated, run, and deleted, Symantec also supports two types of user-defined health checks. These health checks are manually created, configured, and deleted.

- Composite health checks: A method to take the results from a set of health checks (automatically generated or user-defined health checks) and combine the results.
- Host health checks: A method to test a server, using a selection of ICMP, TCP, SSL, HTTP, and HTTPS tests.

> **Note:** Although a host health check tests an upstream server, it can also be used to test whether a proxy is working correctly. To test HTTP/HTTPS proxy behavior, for example, you can set up a host beyond the proxy, and then use forwarding rules so the health check passes through the proxy to the host, allowing the proxy to be tested.

User-defined health checks allow you to test for attributes that the appliance does not test automatically. For example, for a forwarding host, you could perform three user-defined tests – an HTTP test, an HTTPS test, and a TCP test of other ports. Then, you can set up a composite health check that combines the results of these user-defined tests to represent the health of the forwarding host. The appliance reports the status of the (user-defined) composite health check as the forwarding host's health, instead of the default forwarding host health check. All health check types are given standardized names, based on the name of the target. For example:

- Forwarding hosts and groups have a prefix of `fwd`
- DNS servers have a prefix of `dns`
- SOCKS gateways and gateway groups have a prefix of `socks`
- Authentication realms have a prefix of `auth`
- Content Analysis services have prefixes of `icap`
- User-defined or composite health checks have a prefix of `user`

Health Check Tests

Based on the health check type, the appliance periodically tests the health status, and thus the availability of the host. You can configure the time interval between tests. If the health check test is successful, the appliance considers the host available. The health check tests are described in the table below.

Health check test	Description	Used with health check type
Response Times	The minimum, maximum, and average response times are tracked, with their values being cleared whenever the health check changes state.	All
ICMP Test (Layer 3)	The basic connection between the appliance and the origin server is confirmed. The server must recognize ICMP echoing, and any intervening networking equipment must support ICMP. The appliance sends a ping (three ICMP echo requests) to the host. ICMP tests do not support policy for SOCKS gateways or forwarding.	Forwarding hosts, SOCKS gateways, or user-defined hosts

Health check test	Description	Used with health check type
TCP Socket Connection Test (Layer 4)	A TCP test establishes that a TCP layer connection can be established to a port on the host. Then the connection is dropped. TCP tests for a SOCKS gateway do not support policy for SOCKS gateways or forwarding. TCP tests for a forwarding host or a userdefined health check support SOCKS gateways policy but not forwarding policy.	Forwarding hosts, SOCKS gateways, or user-defined hosts
SSL Test	A connection is made to a target and the full SSL handshake is conducted. Then, much like the TCP test, the connection is dropped. For a forwarding host, a terminating HTTPS port must be defined or the test fails. SSL tests for a forwarding host or a userdefined health check support SOCKS gateways policy. The SSL tests do not support forwarding policy. An SSL test executes the SSL layer in policy and obeys any settings that apply to server-side certificates, overriding any settings obtained from a forwarding host.	Forwarding hosts or user-defined hosts
HTTP/HTTPS Tests for Servers and Proxies	HTTP/HTTPS tests execute differently depending on whether the upstream target is a server or a proxy. For a forwarding host, the server or a proxy is defined as part of the forwarding host configuration. For a userdefined health check, the target is always assumed to be a server. For a server: The HTTP test sends an HTTP GET request containing only the URL path to an HTTP port. ■ The HTTPS test sends an HTTPS GET request containing only the URL path over an SSL connection to a terminating HTTPS port. If an appropriate port is not available on the target, the test fails. For a proxy: ■ The HTTP test sends an HTTP GET request containing the full URL to an HTTP port. ■ Since a server is required to terminate HTTPS, the HTTPS test sends an HTTP CONNECT request to the HTTP port. If an appropriate HTTP port is not available on the proxy, either test fails. An HTTP/HTTPS test requires a full URL for configuration. The HTTP/HTTPS tests for a forwarding host support SOCKS gateway policy but not forwarding policy. The HTTP/HTTPS tests for a user-defined health check support SOCKS gateway and forwarding policy. An HTTPS test executes the SSL layer in policy and obeys any settings that apply to server-side certificates, overriding any settings obtained from a forwarding host.	Forwarding hosts or user-defined hosts
HTTP/HTTPS Authentication	For HTTP/HTTPS tests, you can test authentication using a configured username and password. The passwords are stored securely in the registry.	Forwarding hosts or user-defined hosts
HTTP/HTTPS Allowed Responses	For an HTTP or HTTPS test, this is the set of HTTP response codes that indicate success. The default is to accept only a 200 response as successful. You can specify the sets of response codes to be considered successful.	Forwarding hosts or user-defined hosts

Health check test	Description	Used with health check type
Content Analysis Tests	The tests for Content Analysis are specialized tests devised for each particular kind of Content Analysis service. The health check system conducts by sending requests to the configured services, which reports back a health check result.	ICAP, WebPulse services
Group	Individual tests that are combined for any of the four different available groups (forwarding, SOCKS gateways, and ICAP services). If any of the members is healthy, then the group as a whole is considered healthy. These settings are configurable. By default, group health tests are used for two purposes: ■ Monitoring and notification ■ Policy	Forwarding groups, SOCKS gateways groups, and ICAP external service groups
DNS Server	The DNS server maps the hostname, default is www.bluecoat.com, to an IP address. The health check is successful if the hostname can be resolved to an IP address by the DNS server.	DNS
Authentication	Authentication health checks assess the realm's health using data maintained by the realm during active use. Authentication health checks do not probe the authentication server with an authentication request.	Authentication

Background DNS Resolution Tests

To keep the list of IP addresses associated with a hostname current, the appliance tests DNS resolution for all resolvable hostnames used in the health check system, including forwarding hosts, WebPulse, and SOCKS gateways. The DNS system is checked whenever the time-to-live (TTL) value of the DNS entry expires.

> **Note:** If a hostname consists of a dotted IP address, no DNS resolution is performed.

When a host is resolved by DNS to multiple IP addresses, health checks keep those addresses current through background updates. You can specify the timing for the updates. After the test or tests are conducted for each IP address, the results are combined. If the result for any of the resolved IP addresses is healthy, then the host is considered healthy because a healthy connection to that target can be made.

To specify the intervals for background DNS testing:

1. In the Admin Console, select **Administration > Health Checks > Health Checks**.
2. In the Background DNS section, specify the following settings as needed:

Setting or field	Description
Minimum time to live for DNS results	Cannot be zero (0). Test results are valid for this length of time. Retests can occur any time after this value.
Maximum time to live for DNS results	(Optional) How long the DNS test results remain valid before a retest is required.
Interval to wait after DNS resolution error	If the background DNS test discovers errors, this value specifies how long to wait before retesting. If a specific error repeatedly displays in the event logs, further network troubleshooting is required.

Default Health Check Settings

You can edit health check intervals, severity, thresholds, and notifications for automatically generated health checks in two ways:

- Setting the global defaults. These settings affect all health checks, unless overridden by explicit settings. See "Override Health Check Default Settings" on page 455 for the steps to override settings per health check.
- Setting explicit values on each health check.

The default health check values are:

- Ten seconds for healthy and sick intervals (an interval is the period between the completion of one health check and the start of the next health check).
- One for healthy and sick thresholds. A healthy threshold is the number of successful health checks before an entry is considered healthy; a sick threshold is the number of unsuccessful health checks before an entry is considered sick.
- Warning for the severity notification, which governs the effect that a health check has on the overall health status of the appliance.
- Disabled for logging health check status using e-mails, event logs, or SNMP traps.

> **Note:** All health checks are initially configured to use global defaults. The only exception is the DRTR service, which has the healthy interval set to 10800 seconds (3 hours), and the failure trigger set to 1.

Change Default Health Check Settings

To change default health check settings:

1. In the Admin Console, select **Administration > Health Checks > Health Checks**.
2. In **Default Settings**, specify the following settings as needed:

Setting or field	Description
Healthy Interval	Interval, in seconds, between health checks to a service that was determined to be healthy. The default is 10. The healthy interval can be between 1 second and 31536000 seconds (about one year).
Healthy Threshold	Threshold for the number of successful health checks before an entry is considered healthy. The threshold can be between 1 and 65535. The default is 1.
Sick Interval	Interval, in seconds, between health checks to a service that was determined to be unhealthy or out of service. The default is 10. The sick interval can be between 1 second and 31536000 seconds (about 1 year).
Sick Threshold	Threshold for the number of failed health checks before an entry is considered unhealthy. The threshold can be between 1 and 65535. The default is 1.
Failure trigger threshold	(Disabled by default) If enabled, the threshold for the number of failed connections to the server before a health check is triggered. Valid values can be between 1 and 2147483647. The failures are reported back to the health check as a result of either a connection failure or a response error. The number of these external failures is cleared every time a health check is completed. If the number of failures meets or exceeds the threshold and the health check is idle, the health of the device or service is immediately checked.
Response time threshold	(Disabled by default) Specify the maximum response time threshold, in milliseconds. The threshold time can be between 1 and 65535.

Explicit health settings override the global defaults. To override the default settings for a health check, see "Override Health Check Default Settings" on the facing page.

To execute a health check immediately, see "Execute an Instant Health Check" on page 465.

Override Health Check Default Settings

You can enable or disable health checks and configure them to report as healthy or unhealthy during the time they are disabled.

Setting a health check as disabled but reporting healthy allows the appliance to use the device or service without performing health checks on it. If, for example, you have configured a forwarding host on the appliance, a health check for the forwarding host is automatically created. If you then configure the health check as disabled reporting healthy, the appliance considers the forwarding host as healthy without performing periodic health checks on it.

If the case of a group health check that is disabled but reporting healthy, all members of the group are treated as healthy regardless of the status of the members' individual health check result.

> **Note:** Individual health checks for members of a group remain active; they can be used apart from the group.

Setting a health check as disabled but reporting sick is useful to remove an upstream device for servicing, testing, or replacement. This setting takes the device offline after it completes processing pre-existing traffic. Then the device can be safely disconnected from the network without altering any other configuration. You cannot enable or disable all health checks at once.

See Also:

- "Override Default Notification Settings" on page 459

User-Defined Health Checks

You can manually create and manage ICMP, TCP, HTTP, HTTPS, or SSL health check tests for any upstream TCP/IP device. You can use these user-defined health check types to send notifications of health check state changes.

Under most circumstances, you do not need to create user-defined health checks because the automatically generated health checks meet most needs. However, to check for things that Symantec does not test for automatically—for example, the health of the Internet or of the router, you might create user-defined heath checks.

If, for example, you want to control Web traffic based on the apparent health of the Internet, you can create a user-defined health check to target known Internet sites. As long as a certain number of the sites are healthy, you can consider the Internet as healthy.

Further, you can use policy to configure forwarding rules on the appliance. Subsequently, if the user-defined health check determining internet accessibility transitions to unhealthy, all requests directed to the appliance will be forwarded to the alternate appliance until the primary appliance transitions to healthy again.

The appliance supports two types of user-defined health checks:

- Host: This health check type is for any upstream TCP/IP device. You can't create user-defined health checks for external service tests, such as authentication servers, ICAP, and the WebPulse service. See "Configure Health Check Settings" below.
- Composite: This health check type combines the results of other existing health checks. It can include other composite health checks, health checks for user-defined hosts, and any automatically generated health checks. See "Composite Health Checks" on page 462.

For information about configuring parameter and notification settings for automatically generated health check types, see "Default Health Check Settings" on page 453.

Configure Health Check Settings

You can configure settings for user-defined health checks and system-defined health checks. Some settings are available depending on the health check type. See "Health Check Tests" on page 450 for more information.

To configure health check settings:

1. In the Admin Console, select **Administration > Health Checks > Health Checks**.
2. In the **Health Checks List**:
   - Select **Add > Add Health Checks** to add a new health check.
   - Select an existing health check to edit it.
3. On the Add/Edit Health Check page, in **User Defined Host Health Check**, enter or edit the following as needed:

Health check type	Setting or field	Description
All Available when creating a user-defined health check	**Service Name**	Unique name for the health check.

Health check type	Setting or field	Description
N/A Available when creating a user-defined health check	**Type of Test**	Select the type of test: ■ HTTP ■ HTTPS ■ ICMP ■ SSL ■ TCP See "Health Check Types" on page 449 for more information on these types.
All	**Enable Health Check**	Select the enabled or disabled behavior: ■ **Enabled**: Enable the health check and report changes in the health state. ■ **Disabled: Healthy**: Disable the health check and always report the service as healthy. ■ **Disabled: Unhealthy**: Disable the health check and always report the service as unhealthy.
HTTP/HTTPS	**URL**	URL of the target.
	Enable basic authentication	(For Basic user authentication) Select and enter the username and password of the target.
	Enable basic proxy authentication	(For Basic proxy authentication when intermediate proxies exist between appliance and the target) Select and enter the username and password of the target.
	Allowed Response Codes	List of HTTP/HTTPS response codes that are considered successes, separated by semi-colons. If one of them is received by the health check, the health check considers the HTTP(S) test to have been successful. **Note:** The 200 response code is added by default.
ICMP	**Host**	Hostname of the target. The hostname can be an IPv4 or IPv6 host or address.
SSL	**Host**	Hostname of the target. The hostname can be an IPv4 or IPv6 host or address.
	Port	Port to use for the target.
TCP	**Host**	Hostname of the target. The hostname can be an IPv4 or IPv6 host or address.
	Port	Port to use for the target.

4. Specify whether to keep the default settings or to override them. The **Settings** section shows the current interval and threshold settings for this health check. By default, new health check settings are all set to **Default**. For descriptions of the default settings, see "Default Health Check Settings" on page 453

    - To override a default setting:

        a. Select **Default** and change it to **Override**.

        b. (Applicable to the**Failure trigger threshold**and **Response time threshold** options, which are disabled by default) select the threshold toggle.

        c. Enter a new value in the field.

    - To revert an overridden setting to the default, select **Override**and change it to **Default**. (If applicable, the **Default** option shows the default value; for example, **Default (10)** indicates that the default setting is 10 seconds.)

Override Default Notification Settings

See "Override Default Notification Settings" on the facing page.

See Also:

- "Overview of Health Checks" on page 447
- "Override Health Check Default Settings" on page 455
- "Composite Health Checks" on page 462

Configure Health Check Notifications

The appliance allows you to configure notifications that alert you to changes in health status and to emerging issues. By default, notifications for health check events and status are disabled.

You can set up health check notifications globally and per individual health check. Explicit health settings override the global defaults.

To configure health check notifications globally:

1. Select **Administration > Health Checks > Health Checks**.
2. In **Default Notifications**, select the **Severity** level for the health check:
    - **No effect**: The health check has no impact on the overall health of the device.
    - **Warning**: If the health check fails, the device needs to be monitored and the health check status displays as Warning. This is the default setting.
    - **Critical**: If the health check fails, the device is in critical condition

3. Select the options to enable notifications:

Notification type	Setting	Description
Email and SNMP	**Notify on a transition to healthy**	Send a notification when the health check changes from any state to a healthy state.
Email and SNMP	**Notify on a transition to sick**	Send a notification when the health check changes from any state to a sick state.
Event log	**Log a transition to healthy as**	Log the event when the health check changes from any state to a healthy state with the selected severity level (Informational or Severe).
	Log a transition to sick as	Log the event when the health check changes from any state to a sick state with the selected severity level (Informational or Severe).
Email, event log, and SNMP	**Notify/log for each IP address**	Notify/log any change in health state for each IP address within the health check. For example, if you create a health check “user.google” with the URL “google.com”, the appliance checks the health for each IP address that the Google URL/domain resolves to. Note that these checks are internal and not visible in the console. If any one of these internal health checks is healthy, the overall “user.google” health check is healthy.

Tip: For email notifications, recipients are specified in the Management Console (select **Maintenance > Event Logging > Mail**).
For event logging, messages can be logged as either informational or severe.

Override Default Notification Settings

All notification settings are set to defaults in a new health check:

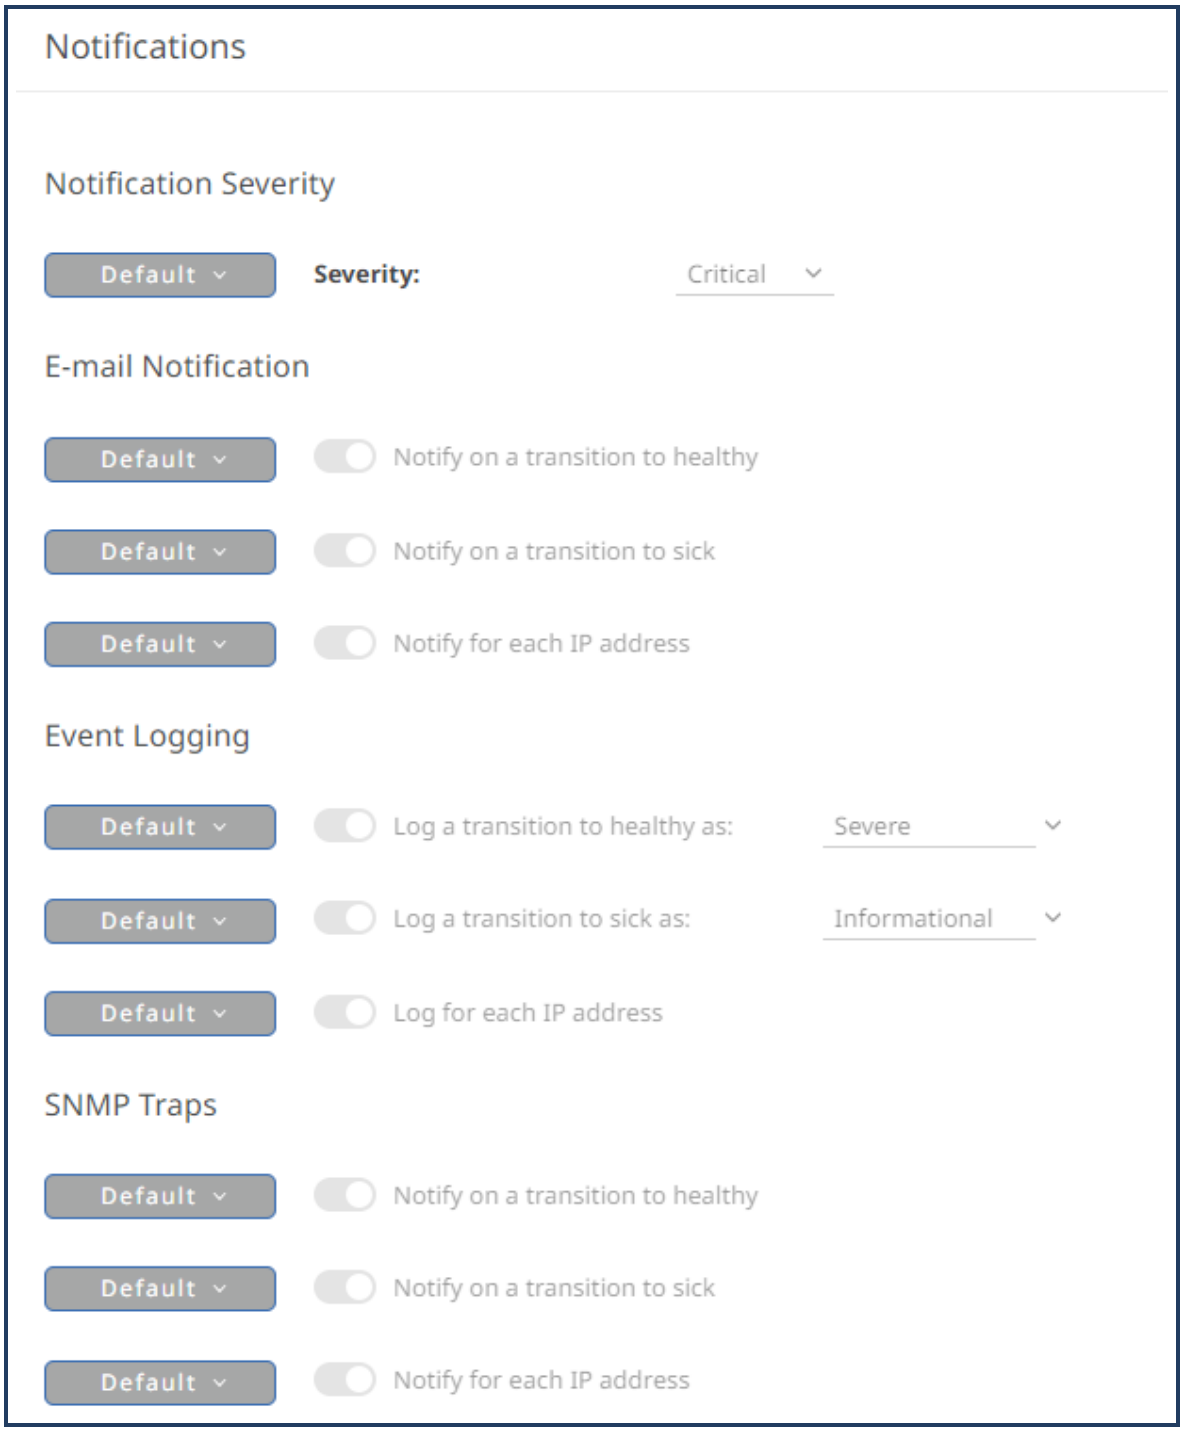

For each notification setting, use the toggle to select **Override**; the **Default (*value*)** setting indicates the default behavior. For example, by default, notification is emailed upon transition to healthy:

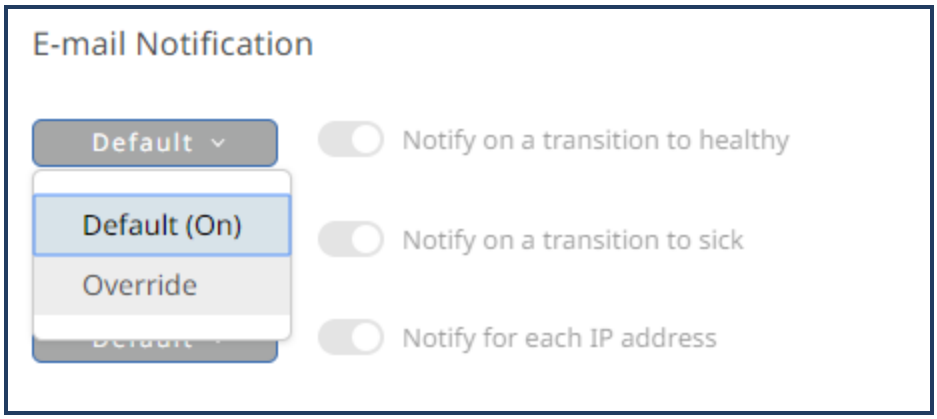

When a setting is overridden, you can specify the behavior. For example, override the global email notification setting and do not notify on a transition to healthy:

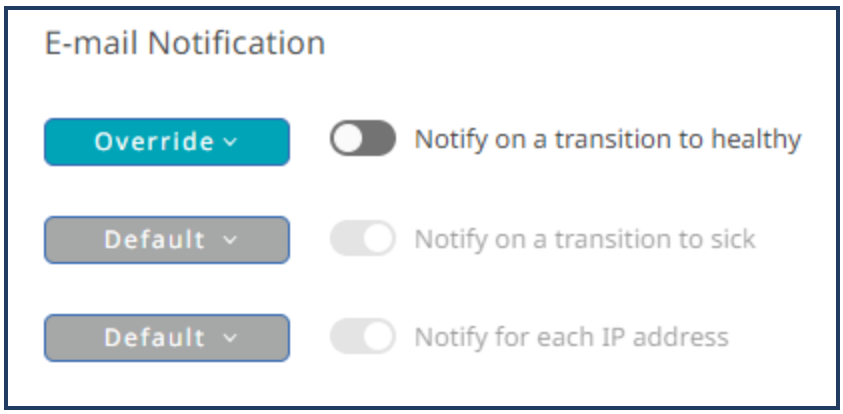

To override notification settings for the health check:

1. Select **Administration > Health Checks > Health Checks**.
2. In **Health Checks**:
   - Select **Add > Add Health Checks** to add a new user-defined health check.
   - Select **Add > Add Composite Health Checks** to add a new composite health check.
   - Select an existing health check or composite health check to edit it.
3. On the Add/Edit Health Check page, in **Notifications**:
   - To override a default setting, select **Override**. The Default option shows the default behavior, such as "Default (On)".
   - To revert a notification setting that was overridden, select **Default**.

For descriptions of the notification settings, see steps 2 and 3 in "Configure Health Check Notifications" on page 458.

Composite Health Checks

You can create a composite health check to combine the results of multiple health checks. A composite health check can contain any number of individual health checks. Further, forwarding host and SOCKS gateway health checks can be configured to use the result of a composite health check.

By default, to report healthy, all members of a composite health check must be healthy. However, you can configure the number of members that must be healthy for the composite result to report healthy.

Composite health checks with no members always appear unhealthy.

> **Tip:** User-defined composite health checks are not the same as automatically generated group tests. Group tests are automatically generated and can't be deleted. Some editing is permitted, but you cannot add or remove members of the group through the health checks module. You must modify the forwarding or SOCKS gateways groups to update the automatically generated group tests. For a group test, the default is for the group to be healthy if any member is healthy. For a composite test, the default is for the group to be healthy if all members are healthy. (The default is configurable.)

Configure Composite Health Check Settings

To configure a composite health check:

1. In the Admin Console, select **Administration > Health Checks > Health Checks**.
2. In **Health Checks**:
   - Select **Add > Add Composite Health Checks** to add a new composite health check.
   - Select an existing composite health check to edit it.

3. On the Add/Edit Composite Health Check page, in **Properties**, enter or edit the following as needed:

Setting or field	Description
Name	(Available when creating a composite health check) Unique name for the health check. The name is automatically prepended with "user."
Enable Health Check	Select the enabled or disabled behavior: ■ **Enabled**: Enable the health check and report changes in the health state. ■ **Disabled: Healthy**: Disable the health check and always report the service as healthy. ■ **Disabled: Unhealthy**: Disable the health check and always report the service as unhealthy.
This health check is healthy when	Select whether any, all, or specific members must be healthy for the composite health check to be healthy.

To add or remove services, see " Modify Composite Health Check Membership" below.

Modify Composite Health Check Membership

To modify composite health check membership:

1. In the Admin Console, select **Administration > Health Checks > Health Checks**.
2. In **Health Checks**:
   - Select **Add > Add Composite Health Checks** to add a new composite health check.
   - Select an existing composite health check to edit it.
3. On the Add/Edit Composite Health Check page, in **Selected Health Checks**:
   - To add members, select **Select Health Checks**. In the Select Health Check Members dialog, select services to add to the composite health check and confirm your action. The services appear in the Selected Members list.
   - To remove members, select **Remove** for the services to remove.

     Alternatively, select **Select Health Checks**. In the Select Health Check Members dialog, clear services to remove them from the composite health check and confirm your action. The services are removed from the Selected Members list.

Override Default Notification Settings

See "Override Default Notification Settings" on page 459.

See Also:

- "Overview of Health Checks" on page 447
- "Override Health Check Default Settings" on page 455

Verify Health Check Settings

You can verify settings of the health checks configured on the appliance. In the Admin Console, select **Administration > Health Checks > Health Checks**. The Health Checks List shows all configured health checks.

Setting or field	Description
Enable	The enabled or disabled state of the health check.
Name	Unique name of the health check.
State	State of the service. See "Enabled Health Check States" below.
Type	Health check type: ■ **Composite**: User-defined composite health check. ■ **User defined host**: User-defined health check. ■ (If the services are defined on the appliance) System-defined health checks including Authentication, DNS Server, ICAP. See "Health Check Types" on page 449 for information.
Test	Health check test. See "Health Check Tests" on page 450 for information.
Settings	Whether notification settings for intervals and thresholds are the system default or overridden. See "Default Health Check Settings" on page 453.
Notifications	Whether notification methods are the system default or overridden. See "Configure Health Check Notifications" on page 458.

Execute an Instant Health Check

Although the appliance automatically executes health checks, you can perform an instant health check. In the Admin Console, select **Administration > Health Checks > Health Checks**. In the **Health Checks**, select **Perform health check** for the health check. Composite health checks with a large number of members can take longer to complete.

Enabled Health Check States

If a health check is enabled, the Admin Console displays the health check state (select **Administration > Health Checks > Health Checks**). Refer to the following:

State	Description	Health
OK	Target device or service is healthy, or all selected members in the composite health check are healthy.	Healthy
Unknown	Health has not yet been tested successfully.	Healthy
OK for some IPs	Some selected members in the composite health check are healthy. Not applicable to health checks that require all members be healthy for a healthy state.	Healthy
Check failed	Device or service can't be used, or all selected members in the composite health check are unhealthy.	Unhealthy

See Also:

- "Override Health Check Default Settings" on page 455

Service Information

- "Service Information Summaries" on the next page
- "Send Service Information to Support" on page 469
- "About Snapshots" on page 471
  - "Configure Snapshot Jobs" on page 472
- "About Packet Capture" on page 473

Service Information Summaries

In the Admin Console, you can use snapshots and packet captures to provide diagnostic information.

- "View and Download Snapshots" below: SysInfo provides a snapshot of statistics and events on the appliance.
- "View, Run, and Download Packet Captures" below: Use this onboard packet capture utility to capture packets of Ethernet frames going in or out of an appliance.

View and Download Snapshots

To configure snapshot jobs, see "About Snapshots" on page 471.

To view and download snapshots:

1. In the Admin Console, select **Administration > Service Information > Summary**.
2. In **Snapshots Link**, select **View all snapshots**.

   The browser opens a new window with a list of the 100 most recent snapshots.
3. To view a snapshot, select the **View** link for a snapshot in the list.

   To download a snapshot, select the **Download** link for a snapshot in the list.

View, Run, and Download Packet Captures

To configure packet captures, see "About Packet Capture" on page 473.

To view, run, and download packet captures:

1. In the Admin Console, select **Administration > Service Information > Summary**.
2. In **Packet Capture Links**, select the appropriate link for a task.

Link	Description
Show packet capture statistics	Shows current packet capture data.
Start packet capture or **Stop packet capture**	Starts the packet capture or stops one in progress.
Download packet capture file	Downloads packet capture data.

Send Service Information to Support

The service information options allow you to send service information to Symantec using either the Admin Console or the CLI. You can select the information to send, send the information, view the status of current transactions, and cancel current transactions. You can also send service information automatically in case of a crash.

You can send service information to as needed or automatically.

Send Service Information Immediately

You can send service information to Symantec using either the Admin Console or the CLI. For information about sending service information automatically, see "Send Service Information Automatically" on the next page.

You must specify a case number before you can send service information. See Symantec Support for details on opening a service request ticket: https://support.symantec.com/en_US/contact-support.html

The following list details information that you can send:

- Packet capture
- Event log
- Memory core
- Policy trace file
- SysInfo
- Access logs
- Snapshots
- Contexts

To send service information immediately:

1. In the Admin Console, select **Administration > Service Information > Send Information**.
2. In **Send Selected Info Now**, expand **Select show and send**.
3. Enter a valid case number in the **Auto-Send Request Number** field.
4. In the list of service information types, select the information to send.

   Options for items that you do not have on your system are unavailable and can't be selected.
5. (Optional) If you select Access Logs, Snapshots, or Contexts, you must also specify the files to send.

In the dialog that appears, select specific files or select all files.

6. Select **Send**.

Send Service Information Automatically

Enabling automatic service information allows you to enable the transfer of relevant service information automatically whenever a crash occurs. This saves you from initiating the transfer, and increases the amount of service information that Symantec can use to solve the problem. The core image, system configuration, and event log are system-use statistics that are sent for analysis. If a packet capture exists, it is also sent. The auto-send feature requires that a valid Service Request is entered. If you do not have an open case you must first contact Symantec Technical Support.

> **Caution:** A core image and packet capture can contain sensitive information– for example, parts of an HTTP request or response. The transfer to Symantec is encrypted, and therefore secure; however, if you do not want potentially sensitive information to be sent to Symantec automatically, do not enable the automatic service information feature.

To send service information automatically:

1. In the Admin Console, select **Administration > Service Information > Send Information**.
2. In **Automatic Operation**, enter a valid case number in the **Auto-Send Request Number** field.
3. Switch the **Automatically send service information** on crash toggle.

Manage the Bandwidth for Service Information

You can control the allocation of available bandwidth for sending service information. Some service information items are large, and you might want to limit the bandwidth used by the transfer. Changing to a new bandwidth management class does not affect service information transfers already in progress. However, changing the details of the bandwidth management class used for service information, such as changing the minimum or maximum bandwidth settings, affects transfers already in progress if that class was selected prior to initiating the transfer.

> **Note:** Before you can manage the bandwidth for the automatic service information feature, you must first create an appropriate bandwidth management class.

To manage bandwidth for service information:

1. In the Admin Console, select **Administration > Service Information > Send Information**.
2. In **Bandwidth Classes**, switch the **Use a Bandwidth Class to send service information** toggle.
3. In the **Bandwidth Class** list, select a bandwidth class.

About Snapshots

The snapshot subsystem periodically pulls a specified console URL and stores it in a repository, offering valuable resources for Symantec customer support in diagnosing problems. By default, two System Information (SysInfo) snapshots are defined:

- `sysinfo`: Takes a snapshot of the system information URL once every 24 hours. This snapshot job keeps the last 100 snapshots.
- `sysinfo-stats`: Takes an hourly snapshot of the system information statistics (sysinfo-stats). This snapshot job keeps the last 168 snapshots. Determining which console URL to poll, the time period between snapshots, and how many snapshots to keep are all configurable options for each snapshot job.

You can also add snapshot jobs for any advanced statistic on the appliance. See "Advanced Statistics" on page 313 for details.

Snapshot jobs configured on the appliance are available in the Admin Console at **Administration > Service Information > Snapshots**. The Snapshots Jobs section shows the following:

Setting or field	Description
Enable	(Not available for system-defined jobs, which are always enabled) Shows whether the snapshot job is enabled. When enabled, the job runs at the set **Interval**.
Name	Name of the snapshot job.
Target	Source SysInfo or advanced statistic URL of the snapshot.
Interval	Intervening time between each job, expressed in minutes.
Take	Total number of snapshots to take for the job. If there is no limit, takes are **Infinite**.
Keep	Maximum number of snapshots to store for the job.
Next Snapshot	The time that the next job is configured to start, expressed in UTC.
Edit	Edit the current job. See "Configure Snapshot Jobs" on the next page .
Download All Reports	Download all snapshots taken for the current job. The system downloads the snapshots as a gzip file named snapshot_*name*.txt.gz.
Clear All Reports	Delete all saved snapshots for the current job. On the Clear Snapshot Reports dialog, select **Clear** to confirm deletion.
Delete	(Not available for system-defined jobs) Delete the current job.

> **Tip:** When troubleshooting issues with Symantec Technical Support, you might be requested to configure snapshot jobs with non-default settings.

To view all snapshots information, select the **Show All Snapshots** link. The browser opens a new window with a list of the 100 most recent snapshots.

See Also:

- "Configure Snapshot Jobs" below

Configure Snapshot Jobs

To add or edit snapshot jobs:

1. In the Admin Console, select **Administration > Service Information > Snapshots**.
2. In the Snapshot Jobs section:
   - To create a new job, select **Add Job**.
   - To edit an existing job, select it or select **Edit** under Actions.

   The Add/Edit Snapshot Job page displays job settings.
3. Enter or edit the following settings as needed:

Setting or field	Description
Enable job executions	(Not available for system-defined jobs, which are always enabled) Enable the job to run. When enabled, the job runs at the set **Interval**.
Name	Specify a unique name for the snapshot job.
Target	**Select** the source SysInfo or advanced statistic URL of the snapshot. The Select Target dialog displays all available system targets.
Interval	Specify how often to execute the job, in minutes. Default is 1440.
Total number to take	Specify the total number of snapshots to take for the job: ■ Default is **Infinite**; snapshots are taken unless job execution is disabled. ■ Enter a number.
Maximum number to keep	Specify the maximum number of snapshots to store for the job. Default is 30.

About Packet Capture

A packet capture (PCAP) captures all received packets that are either directly addressed to the appliance through an interface's MAC address or through an interface's broadcast address. The utility also captures transmitted packets that are sent from the appliance. The collected data can then be transferred to the desktop or to Broadcom for analysis.

Running PCAP increases the amount of processor usage performed in TCP/IP. To analyze captured packet data, you must have a tool that reads Packet Sniffer Pro 1.1 files (such as Wireshark or Packet Sniffer Pro 3.0).

You can capture packets of Ethernet frames going into or leaving an appliance. Packet capturing allows filtering on various attributes of the frame to limit the amount of data collected. Any packet filters must be defined before a capture is initiated, and the current packet filter can only be modified if no capture is in progress.

Configure PCAP

You can configure packet capturing and start a capture. To control which packets are captured, specify a filter expression.

> **Note:** To analyze captured packet data, you must have a tool that reads Packet Sniffer Pro 1.1 files (for example, Ethereal or Packet Sniffer Pro 3.0).

In the Admin Console, select **Administration > Service Information > Packet Capture**. The Capture section displays options for configuring and running captures.

Setting or field	Description
Direction	Capture incoming, outgoing, or all packets.
Interface	Capture packets through a specified interface or all interfaces.
Filter	Define or change the PCAP filter expression. Type an expression in the field, or select **Expressions** to select from a list of templates. In the list of Common PCAP Filter Expressions dialog that opens, select an expression to insert into the **Filter** field. To remove the filter, select **Clear**. For more information, see "Common PCAP Filter Expressions" on the next page.
Start Capture	Specify settings for the packet capture; see "Packet Capture Settings" on page 475. Then, select **Start Capture**. When you start the capture, the console displays a "Capture in Progress" message and options to stop the capture.
Download	Download the last PCAP.

Setting or field	Description
Stop **Stop and Download**	(Available when a packet capture is in progress) Stop the capture without downloading, or stop the capture and download the file. The name of a downloaded packet capture file has the format: `bluecoat_date_filter-expression.cap` where *date* is the UTC date and time of the packet capture and *filter-expression* is the filter expression(s) used. **Note:** Because the filter expression can contain characters that are not supported by a file system, a translation can occur. The following characters are not translated: - Alphanumeric characters (a-z, A-Z, 0-9) - Periods (.) The following characters are translated: - Space (replaced by an underscore) - All other characters (including the underscore and dash) are replaced by a dash followed by the ASCII equivalent; for example, a dash is translated to -2D and an ampersand (&) to -26. After the file is downloaded, you can upload it to a remote server using FTP. See "Upload Packet Capture Data Via FTP" on page 476.
Refresh Status	Update the details in the **Status** field.
Status	Shows the status of the last or current capture.

Common PCAP Filter Expressions

When running a PCAP, you can filter on various attributes of the frame to limit the amount of data collected.

Specify filters in Berkeley Packet Filter (BPF) format , which is also used by the tcpdump program. A few simple examples are provided below. If filters with greater complexity are required, refer to sources on the Internet and in books that describe the BPF syntax.

Note: Some qualifiers must be escaped with a backslash because their identifiers are also keywords within the filter expression parser:
- `ip proto protocol`
where `protocol` is a number or name (icmp, udp, tcp)
- `ether proto protocol`
where `protocol` can be a number or name (ip, arp, rarp)

PCAP Filter Expressions

See more filter expressions in the Admin Console by selecting **Administration > Service Information > Packet Capture**. In the Capture section, select **Expressions**.

Filter Expression	Packets Captured
`ip host 10.25.36.47`	Captures packets from a specific host with IP address 10.25.36.47.
`not ip host 10.25.36.47`	Captures packets from all IP addresses except 10.25.36.47.
`not port 8082`	Captures packets from all ports except 8082.
`ip host 10.25.36.47 and ip host 10.25.36.48`	Captures packets sent between two IP addresses: 10.25.36.47 and 10.25.36.48. Packets sent from one of these addresses to other IP addresses are not filtered.
`ether host 00:e0:81:01:f8:fc`	Captures packets to or from MAC address 00:e0:81:01:f8:fc:.
`port 80`	Captures packets to or from port 80.

Using Filter Expressions in the CLI

To specify filters using CLI, issue the following command:

```
# pcap filter expr parameters
```

> **Tip:** Define parameters within double quotation marks to avoid confusion with special characters. For example, a space is interpreted by the CLI as an additional parameter, but the CLI accepts only one parameter for the filter expression. Enclosing the entire filter expression in quotations allows multiple spaces in the filter expression.

To remove a filter, use the command:

```
# pcap filter
```

Packet Capture Settings

On the Start Packet Capture dialog, specify the following settings:

Setting	Description
Buffer size	Capture **All** matching packets or **Selected** packets. If capturing selected packets, specify one of the following: ■ first *n* matching packets. If the number of packets reaches this limit, packet capturing stops automatically. *n* is a value between 1 and 1000000. ■ last *n* matching packets. Any packet received after the memory limit is reached results in the discarding of the oldest saved packet prior to saving the new packet. The saved packets in memory are written to disk when the capture is stopped. *n* is a value between 1 and 1000000. ■ first *n* matching kbytes. Enter the number of kilobytes to capture. If the buffer reaches this limit, packet capturing stops automatically. The value is limited to 3% of the available system memory at startup (not to exceed 4GB). If a value is not specified, the default packet capture file size is 100MB. ■ last *n* matching kbytes. Any packet received after the memory limit is reached results in the discarding of the oldest saved packet prior to saving the new packet. The saved packets in memory are written to disk when the capture is stopped. The value is limited to 3% of the available system memory at startup (not to exceed 4GB). If a value is not specified, the default packet capture file size is 100MB.
Byte limit per packet	(Optional) Truncate the number of bytes saved in each frame. Select **Save First** and enter a number between 1 and 65535 in the **bytes of each packet** field. Packet capture collects no more than the specified bytes of packets from each frame when writing to disk.
Bytes to keep in core image	(Optional) Number of kilobytes of packets kept in a core image. Select **Include** and enter a value in the **K Bytes in core image** field. You can capture packets and include them along with a core image. This is extremely useful if a certain pattern of packets causes the unit to restart unexpectedly. The core image size The value is limited to 3% of the available system memory at startup (not to exceed 4GB). By default, no packets are kept in the core image.

Upload Packet Capture Data Via FTP

Use the following command to transfer packet capture data from the appliance to an FTP site. After uploading is complete, you can analyze the packet capture data.

```
# pcap transferftp://url/path/filename.cap username password
```

If the FTP server requires a username and password, specify them.

Legal Notice

Broadcom, the pulse logo, Connecting everything, and Symantec are among the trademarks of Broadcom. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Copyright © 2020 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Wednesday, November 11, 2020