Forcepoint

Web Security On-prem

v8.5.x

Content Gateway Manager Online Help

© 2025 Forcepoint Forcepoint and the FORCEPOINT logo are trademarks of Forcepoint. All other trademarks used in this document are the property of their respective owners.

Published 27 June 2025

Every effort has been made to ensure the accuracy of this document. However, Forcepoint makes no warranties with respect to this documentation and disclaims any implied warranties of merchantability and fitness for a particular purpose. Forcepoint shall not be liable for any error or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual or the examples herein. The information in this documentation is subject to change without notice.

Contents

1 Forcepoint Content Gateway Overview	7
Introduction	7
Content Gateway deployment options	8
Content Gateway components	10
Content Gateway processes	11
Content Gateway administration tools	
Proxy traffic analysis features	
Technical Support	
2 Getting Started with Content Gateway	15
Introduction	
Accessing the Content Gateway manager	
Content Gateway online Help	
Entering your subscription key	
Providing system information	
Verifying that the proxy is processing Internet requests	
Using the command-line interface	
Starting and stopping Content Gateway on the command line	
3 Explicit Proxy	
Manual browser configuration.	28
Lising a PAC file	28
Using WPAD	30
Configuring FTP clients in an explicit proxy environment	
A Transparent Provy and APM	25
4 Indisparent Floxy and ARM	
The Content Cateway APM	
Transport interception strategies	
Interception bypage	
Connection load chedding	
Connection load shedding	
Reducing DNS lookups	
5 Additional Proxy Configuration	63
Introduction	63
Content Gateway IP spoofing	63
Content Gateway support for IPv6	
6 Clusters	
Introduction	71
Management clustering	
Changing clustering configuration	72
Adding nodes to a cluster	73
Deleting nodes from a cluster	
Virtual IP failover	
7 Hierarchical Caching	79
Introduction	

HTTP cache hierarchies	
8 Configuring the RAM Cache	83
Introduction	
Configuring cache object size limit	
Changing the size of the RAM cache	
9 DNS Proxy Caching	
Introduction	
Configuring DNS proxy caching	
10 Saving and Restoring Configurations	89
Introduction	89
Taking configuration snapshots	89
Restoring configuration snapshots	90
Deleting configuration snapshots	
11 Monitoring Traffic	
Viewing statistics in the Content Gateway manager	
Viewing statistics from the command line	
working with alarms	
Using Performance graphs	
Creating SSL certificate authorities reports	
Creating an SSL incidents report	
12 Working With Web DLP	101
Introduction	
How the Forcepoint Web Security DLP Module works	
Deploying Content Gateway to work with Forcepoint DLP	
Registering Content Gateway with Forcepoint DLP	
Stopping and starting Forcepoint DLP processes	
Configuring the ICAP client	106
13 Working With Encrypted Data	
Introduction	
Enabling SSL support	
Initial SSL configuration tasks	
Certificates	
Internal Root CA	
Backing up your internal Root CA	
Managing certificates	
Decryption and Encryption	
Validating certificates	
Bypassing verification	
Keeping revocation information up to date	
Directing SSL traffic to Content Gateway via explicit proxy	
Managing HTTPS website access	
Client certificates	
Customizing SSL connection failure messages	
Custom certificate key	
SSL decryption port mirroring (appliance deployments)	
14 Content Gateway Security	140
ra content outeway occurry	

Introduction	
Controlling client access to the proxy	150
Controlling access to the Content Gateway manager	151
FIPS 140-2 Mode	154
Content Gateway filtering rules	155
Configuring SOCKS firewall integration	
Setting SOCKS server bypass	163
Using the Split DNS option	
Content Gateway user authentication	164
15 Working With Log Files	
Introduction	
Event log files	218
Managing event log files	219
Event log file formats	
Rolling event log files	227
Splitting event log files	230
Collating event log files	232
Viewing logging statistics	
Viewing log files	
Example event log file entries	
A Statistics	
My Proxy	
Protocols	
Security	
Subsystems	
Networking	
Performance	
SSL	
B Commands and Variables	
Content Gateway commands	
Content Gateway variables	
O Confirmation Ontions	075
C Configuration Options	
My Ploxy	
Flotocols	
Security	
Subsystems	
Networking	324
SSI	
002	
D Event Logging Formats	
Custom logging fields	355
Logging format cross-reference	
E Content Gateway Configuration Files	
Specifying URL regular expressions (url_regex)	
auth_domains.config	
auth_rules.config	
bypass.config	

filter.config	372
ip_allow.config	
ipnat.conf	
log_hosts.config	
logs_xml.config	
mgmt_allow.config	
parent.config	
records.config	
remap.config	
socks.config	
socks_server.config	
splitdns.config	
wccp.config	
F Content Gateway Error Messages	
Error messages in log files	
Content Gateway alarm messages	
Content Gateway HTML messages sent to clients	
Content Gateway standard HTTP response messages	

Chapter 1 Forcepoint Content Gateway Overview

Contents

- Introduction on page 7
- Content Gateway deployment options on page 8
- Content Gateway components on page 10
- Content Gateway processes on page 11
- Content Gateway administration tools on page 12
- Proxy traffic analysis features on page 12
- Technical Support on page 13

Introduction

Content Gateway is the web proxy component of Forcepoint[™] Web Security.

Content Gateway performs advanced content analysis precisely when it is needed—as the content flows through the proxy. The results of analysis are used by Forcepoint Web Security to protect you from malicious content and apply your Acceptable Use Policy (AUP). This on-demand analysis protects users and networks at the same time that it makes rapidly changing websites safe for your organization and users. Advanced analysis may be applied to HTTP, HTTPS, and FTP channels.

The precise application of advanced analysis is configured by the administrator for each Forcepoint Web Security deployment.

Content Gateway can also be deployed as the web proxy component of Forcepoint DLP (absent Web Security). A core version of Content Gateway is included in Forcepoint DLP Network licenses. Known as Forcepoint DLP Web Content Gateway, this core version is managed through Content Gateway and Forcepoint DLP managers, and allows Content Gateway to block traffic that matches the Forcepoint DLP web policies. Note that some features of Content Gateway are available only when Content Gateway is deployed with Forcepoint Web Security, and not in a standalone deployment with Forcepoint DLP Network. Those features are marked accordingly.

Content Gateway can be deployed as described in Content Gateway deployment options.

Content Gateway can also be configured to:

- Ensure that clients are authenticated before they access content. Content Gateway supports Integrated Windows Authentication, legacy NTLM (NTLMSSP), LDAP, and RADIUS. See, Content Gateway user authentication.
- Control client access to the proxy. See, Controlling client access to the proxy.
- Use different DNS servers, depending on whether the proxy needs to resolve host names located inside or outside a firewall. This enables you to keep your internal network configuration secure while providing transparent access to external sites on the Internet. See, Using the Split DNS option.

- Use the co-located Data policy engine or the ICAP interface to enable sites using Forcepoint DLP to examine outbound material such as web postings, and block or allow based on company policy. See Working With Web DLP.
- Control access to the Content Gateway manager using:
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access
 - User accounts that define which users can access the manager and which activities they can perform (for example, view statistics only or view statistics and configure Content Gateway).
- Integrate into your firewall and control traffic through a SOCKS server. See Content Gateway Security.

Related concepts

Content Gateway user authentication on page 164

Related tasks

Controlling client access to the proxy on page 150 Using the Split DNS option on page 163

Related information

Content Gateway deployment options on page 8 Working With Web DLP on page 101 Content Gateway Security on page 149

Content Gateway deployment options

SSL inspection

When the HTTPS option is enabled, HTTPS traffic is decrypted, inspected, and re-encrypted as it travels to and from the client and origin server.

Content Gateway includes a complete set of certificate-handling capabilities. See Working With Encrypted Data.



Important

Even when HTTPS is **not** enabled, Content Gateway still performs a URL lookup for HTTPS requests and applies policy accordingly.

In explicit proxy mode, when HTTPS is disabled, Content Gateway performs URL filtering based on the hostname in the request. If the site is blocked, Content Gateway serves a block page. Note that some browsers do not support display of the block page. To disable this feature, configure clients to not send HTTPS requests to the proxy.

In transparent proxy mode, when HTTPS is disabled, if there is an SNI in the request, Content Gateway gets the hostname from the SNI and performs URL filtering based on the hostname. Otherwise, Content Gateway uses the Common Name in the certificate of the destination server. However, if the Common Name contains a wildcard (*), the lookup is performed on the destination IP address. If the site is blocked, the connection with the client is dropped; no block page is served. To disable this feature when used with WCCP, do not create a service group for HTTPS.

Related information

Working With Encrypted Data on page 111

Cache hierarchy

Content Gateway can participate in flexible cache hierarchies, where Internet requests not fulfilled in one cache can be routed to other regional caches, taking advantage of their contents and proximity. In a hierarchy of proxy servers, Content Gateway can act either as a parent or child, either to other Content Gateway servers or to other caching products. See *Hierarchical Caching*.

Related information Hierarchical Caching on page 79

Managed cluster

Content Gateway scales from a single node to multiple nodes, with a maximum recommended limit of 16. This forms a managed cluster that improves system capacity, performance, and reliability.

- A managed cluster detects the addition and removal of nodes.
- Cluster nodes automatically share configuration information, allowing members of the cluster to all be administered at the same time.

If the virtual IP failover option is enabled, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes of the cluster. Content Gateway can detect node failures (such as power supply or CPU failures) and reassign IP addresses of the failed node to the operational nodes. See *Virtual IP failover*, for details.

If Content Gateway is configured as a transparent proxy with WCCP, failover is handled by WCCP and virtual IP failover should not be used. See WCCP load distribution.

For complete information, see Clusters.

Related concepts Virtual IP failover on page 76 WCCP load distribution on page 41

Related information Clusters on page 71

DNS proxy cache

As a DNS proxy cache, Content Gateway can resolve DNS requests for clients. This offloads remote DNS servers and reduces response times for DNS lookups. See *DNS Proxy Caching*.

Related information DNS Proxy Caching on page 85

Content Gateway components

RAM cache

Content Gateway maintains a small RAM memory cache of extremely popular objects. This RAM cache serves the most popular objects quickly. You can configure the RAM cache size. See *Changing the size of the RAM cache*.

Related tasks

Changing the size of the RAM cache on page 84

Adaptive Redirection Module

The Adaptive Redirection Module (ARM) provides several essential functions. One is to send device notifications for cluster communication interface failover. Another is to inspect incoming packets before a routing decision is made and redirect the packets to Content Gateway for processing.

The ARM:

- Is always active.
- Uses iptables, policy routing, and transparent sockets which are configured during product installation.
 The installation program also creates redirection rules to intercept packets.
- Supports automatic bypass of sites that do not transit properly through a proxy.
- Prevents client request overloads. When there are more client connections than the specified limit, the ARM forwards incoming requests directly to the origin server. See Connection load shedding.

Related tasks

Connection load shedding on page 61

Host database

The host database stores the Domain Name Server (DNS) entries of origin servers to which the proxy connects. Among other information, the host database tracks:

DNS information (for fast conversion of host names to IP addresses)

- The HTTP version of each host (so advanced protocol features can be used with hosts running modern servers)
- Host reliability and availability information (to avoid waits for non-functional servers)

DNS resolver

For transparent proxy deployments, the proxy includes an asynchronous DNS resolver to streamline conversion of host names to IP addresses. Content Gateway implements the DNS resolver natively, directly issuing DNS command packets, rather than relying on resolver libraries. Many DNS queries can be issued in parallel and a fast DNS cache maintains popular bindings in memory, reducing DNS traffic.



Important

Should the Linux system DNS server configuration change (/etc/resolv.conf), you must restart Content Gateway.

Content Gateway processes

Content Gateway has 4 primary processes:

Process name	Description
content_gateway	Accepts connections, processes protocol requests, and serves documents from the origin server.
content_manager	Launches, monitors, and reconfigures the content_gateway process.
	The content_manager process is also responsible for the Content Gateway manager user interface, the proxy auto- configuration port, the statistics interface, cluster administration, and virtual IP failover.
	If the content_manager process detects a content_gateway process failure, it restarts the process and also maintains a connection queue of all incoming requests. Incoming connections that arrive in the several seconds before server restart are saved in the connection queue and processed in sequence. This connection queuing shields users from server restart downtime.

Process name	Description
content_cop	Monitors the health of content_gateway and content_manager .
	The content_cop process periodically (several times each minute) queries content_gateway and content_manager by issuing heartbeat requests to fetch synthetic Web pages. If no response is received within the timeout interval or if an incorrect response is received, content_cop restarts content_manager and content_gateway.
	The content_cop process uses the port defined by proxy.config.admin.heartbeat_port.
analytics_server	Manages the requests made and processes spawned for analytics.

Content Gateway administration tools

The primary Content Gateway configuration and administration tool is the web-based graphical user interface that is accessible through your browser. The Content Gateway manager offers password-protected, SSL-encrypted, single-point administration for an entire Content Gateway cluster. The Content Gateway manager provides graphs and statistical displays for monitoring Content Gateway performance and network traffic, and options for configuring and fine-tuning the proxy.

Sometimes it is convenient or necessary to use the Content Gateway command-line interface. You can execute individual commands or script a series of commands in a shell. This facility is not available when Content Gateway is hosted on a Forcepoint appliance. Instead, use the Content Gateway manager and see your Forcepoint appliance documentation.

Like the command line interface, it is sometimes convenient or necessary to make configuration changes in Content Gateway configuration files. They support administration through a file-editing and signal-handling interface. Any changes you make through the Content Gateway manager or command-line interface are automatically made to the configuration files.

See:

Related concepts Accessing the Content Gateway manager on page 16

Related tasks

Using the command-line interface on page 23

Proxy traffic analysis features

Content Gateway provides options for network traffic analysis and monitoring:

- Manager statistics and graphs show network traffic information. View graphs and statistics from the Content Gateway manager, or collect and process statistics using the command-line interface.
- A variety of *Performance* graphs show historical information about virtual memory usage, client connections, and so on. View Performance graphs in the Content Gateway manager.
- Manager alarms are presented in the Content Gateway manager. Content Gateway signals an alarm for any
 detected failure condition. You can configure Content Gateway to send email or page support personnel when
 an alarm occurs.

Content Gateway also sends select alarms to the Forcepoint Security Manager, where they are referred to as **alerts**. Summary alert messages are displayed on the **Web** > **Status** > **Dashboard** > **System** page. The full alert message is displayed on the **Status** > **Alerts** page. Web Security administrators can configure which

Content Gateway conditions cause alert messages to be sent, and which methods (email or SNMP) are used to send the alert.

Transaction logging lets you record information in a log file about every request the proxy receives and every error it detects. Use the logs to determine how many people use the proxy, how much information each person requested, and which pages are most popular. You can see why a transaction was in error. For example, you can see that Content Gateway was restarted or that cluster communication timed out.

Content Gateway supports several standard log file formats, such as Squid and Netscape, and its own custom format. You can analyze the standard format log files with off-the-shelf analysis packages. To help with log file analysis, separate log files so that they contain information specific to protocol or hosts.

For traffic analysis options, see Monitoring Traffic. For logging options, see Working With Log Files.

Related information

Monitoring Traffic on page 93 Working With Log Files on page 217

Technical Support

Technical information about Forcepoint products is available 24 hours a day at: https://support.forcepoint.com In the Support site you will find:

in the Support site you

- Tips
- Customer Forums
- Latest release information
- Searchable Knowledge Base
- Latest hotfixes and patches
- Show-Me tutorials and videos
- Product documents
- Answers to frequently asked questions
- In-depth technical papers
- Monthly Support Webinars
- Technical Alerts
- Most Popular Solutions

The Support site offers access to all technical resources, including opening a case through the Service Request portal.

Chapter 2 Getting Started with Content Gateway

Contents

- Introduction on page 15
- Accessing the Content Gateway manager on page 16
- Content Gateway online Help on page 20
- Entering your subscription key on page 20
- Providing system information on page 21
- Verifying that the proxy is processing Internet requests on page 22
- Using the command-line interface on page 23
- Starting and stopping Content Gateway on the command line on page 24

Introduction

After you have installed Content Gateway on a system or on all of the nodes in your cluster, the proxy is ready for use.

You can configure Content Gateway via its web-based user interface: the Content Gateway manager.

To get started, see:

Related concepts

Accessing the Content Gateway manager on page 16

Related tasks

Entering your subscription key on page 20 Verifying that the proxy is processing Internet requests on page 22 Using the command-line interface on page 23 Starting and stopping Content Gateway on the command line on page 24

Accessing the Content Gateway manager

Content Gateway has a browser-based management console: the Content Gateway manager.

- See the Certified Product Matrix for a list of browsers that the console supports. Use of other browsers and versions may result in unexpected behavior.
- Java and JavaScript must be enabled in your browser. See your browser documentation for instructions.

There are 2 ways to access the Content Gateway manager:

From the Forcepoint Security Manager, using single sign-on (SSO)



Note

When SSO (not available with Forcepoint DLP Web Content Gateway) is used, the browser must be configured to allow pop-ups on the Content Gateway IP address.

- When two-factor authentication is enabled, this is the only method that can be used. See Configuring Content Gateway for two-factor authentication.
- For SSO configuration instructions, see the Forcepoint Web Security Administrator Help.
- If you log on to Content Gateway manager using SSO, when you log off of Content Gateway manager your session is closed.
- By entering the IP address and port of the Content Gateway host system in your browser:
 - 1) In the browser address bar, enter:

https://<nodename>:<port>

Here, *<nodename>* is the IP address of Content Gateway and *<port>* is the port number assigned to the Content Gateway manager (8081, by default).

- 2) On the logon page, enter your administrator ID (admin, by default) and password.
 - The Content Gateway manager password is set during installation.
 - You can change the ID and password, as well as create and modify user accounts. See *Controlling* access to the Content Gateway manager.

When you on to Content Gateway manager directly, when you click **Log Off**, your session is not closed until you close all open browser windows.

On launch, the Content Gateway manager displays the **Monitor > My Proxy > Summary** page. This page provides information on the features of your subscription and details of your Content Gateway system.

- For information about the Monitor tab, see *Viewing statistics in the Content Gateway manager*.
- Click the Configure tab to display the available configuration options.
 - This document provides instructions for the many tasks that can be performed via the options on the Configure tab.
 - A list describing all of the options available on the Configure tab appears in *Configuration Options*.

Related concepts Configuring Content Gateway for two-factor authentication on page 17 Controlling access to the Content Gateway manager on page 151 Viewing statistics in the Content Gateway manager on page 94

Related information Configuration Options on page 275

Security certificate alerts

An SSL connection is used for secure, browser-based communication with the Content Gateway manager. This connection uses a security certificate issued by Forcepoint LLC. Because the supported browsers do not recognize Forcepoint LLC as a known Certificate Authority, a certificate error displays the first time you launch the Content Gateway manager from a new browser. To avoid seeing this error, install or permanently accept the certificate within the browser. See your browser documentation for details.



Note

If you are using Internet Explorer, the certificate error will still be present after you accept the certificate. Close and reopen your browser to remove the error message.

Windows 7 considerations

If you are using Windows 7, you may need to run the browser as administrator for it to allow ActiveX controls.

Steps

- 1) Right-click the browser application and select Run as administrator.
- 2) Log on to the Content Gateway manager and accept the security certificate as described above.

Configuring Content Gateway for two-factor authentication

Two-factor (certificate) authentication (not available with Forcepoint DLP Web Content Gateway):

- Is configured for and applies to the Forcepoint Security Manager only.
- Requires administrators to provide 2 forms of identification to log on.
- Can be made to apply to the Content Gateway manager by forcing administrators to log on to the Forcepoint Security Manager before accessing the Content Gateway manager.
- Requires single sign-on to be configured for administrators allowed access to the Content Gateway manager.

Requires that the password logon capability be disabled on Content Gateway (see below), preventing administrators not configured for single sign-on from accessing the Content Gateway manager. If Content Gateway is deployed on an appliance, password access is disabled using an appliance command. See your Forcepoint appliance documentation.

For more information about configuring two-factor authentication, see "Configuring Certificate Authentication" in Forcepoint Security Manager Help.

Disabling and enabling Content Gateway password logon

The Content Gateway manager password logon can be disabled to allow two-factor authentication only, or single sign-on access from the Forcepoint Security Manager.



Important

If Content Gateway is installed on an appliance, see your appliance documentation for details.

To disable password logon

Steps

- 1) Make sure members of the Super Administrators group in the Web module of the Forcepoint Security Manager have Content Gateway Direct Access (single sign- on) permissions.
- 2) If two-factor authentication will be used, set up two-factor authentication in the Security Manager.
- 3) Log on to the Content Gateway host system and acquire root privileges.
- Change directory to "/etc" and check to see if there is a "websense" subdirectory. If not, create one ("mkdir websense").
- 5) Change directory to "websense" (path is now "/etc/websense") and check to see if the file "passwordlogon.conf" exists.
- 6) If not, create it ("touch password-logon.conf".
- 7) Edit "password-logon.conf".
- Add the line, or modify the existing line to: password-logon=disabled
- Write and exit the file.
 The change takes effect immediately. There is no need to restart Content Gateway.

To re-enable password logon for all administrators

Steps

- 1) Log on to the Content Gateway host system and acquire root privileges.
- 2) Navigate to the /etc/websense directory.
- Edit password-logon.conf and change: password-logon=disabled

to:

password-logon=enabled

4) Write and exit the file.

The change takes effect immediately. There is no need to restart Content Gateway.

Accessing the Content Gateway manager if you forget the master administrator password



Note

The following procedure applies to Content Gateway software installations.

If Content Gateway is running on an appliance, see your Forcepoint appliance documentation.

During installation, you specify an administrator password. The installer automatically encrypts the password and stores the encrypted password in the **records.config** file. Each time you change passwords in the Content Gateway manager, Content Gateway updates the **records.config** file.

If you forget the administrator password and cannot access the Content Gateway manager, you can clear the current password in the **records.config** file (set the value of the configuration variable to NULL) and then enter a new password in the Content Gateway manager. You cannot set passwords in the **records.config** file because the password variables can contain only password encryptions or the value NULL.

Steps

- 1) Open the records.config file in /opt/WCG/config.
- 2) Set the variable proxy.config.admin.admin_password to NULL to leave the password blank.



Note

Ensure that there are no trailing spaces after the word NULL.

- 3) Save and close the file.
- 4) From the Content Gateway bin directory (/opt/WCG/bin), run the following command to apply the changes: ./content_line -x

5) Log on to the Content Gateway manager. When prompted for the user name and password, enter the administrator ID. For the password, enter: Gateway#123

An alarm will display telling you that you are using the default password and reminding you to reset it.

- 6) Navigate to the Configure > My Proxy > UI Setup > Login tab.
- 7) In the Administrator section, enter Gateway#123 in the Old Password field. Enter new password in the New Password field, and then repeat it in the New Password (Retype) field. Passwords must be 8 to 15 characters and include at least one:
 - Uppercase character
 - Lowercase character
 - Number
 - Special character Supported characters include:

!#%&'()*+,-./;<=>?@[]^_{|}~

The following special characters are not supported:

Space \$:`\"

8) Click Apply.

The next time you access the Content Gateway manager, you must use the new password.

Content Gateway online Help

Click on **Get Help!** on any page in the Content Gateway manager to get detailed information about using the product.



Important

Default Internet Explorer settings may block operation of the Help system. If a security alert appears, select **Allow Blocked Content** to display Help.

If your organization's security standards permit, you can permanently disable the warning message on the Advanced tab of the Tools > Internet Options interface. (Mark Allow active content to run in files on My Computer under Security options.)

To access a PDF version of online help, or to access Release Notes, installation and deployment information, FAQs, tips, and other technical information, go to the Documentation Library.

Entering your subscription key

When Content Gateway is deployed with Forcepoint Web Security, there is no need to enter a subscription key in the Content Gateway manager. The Forcepoint Web Security key is automatically shared with Content Gateway.

Note

A key is associated with a Policy Server instance. If you have multiple keys, make sure that Content Gateway is connected to the correct Policy Server instance on the **More Details** view of the **Monitor > My Proxy > Summary** page.

To change which Policy Server instance Content Gateway uses:

- (Appliance) See your Forcepoint appliance documentation.
- (Software) Edit the /opt/WCG/websense.ini file to set the value of PolicyServerIP. After making the change, stop and start Content Gateway processes: /opt/WCG/WCGAdmin stop

/opt/WCG/WCGAdmin start

When Content Gateway is deployed with only Forcepoint DLP, you will need to enter your subscription key manually.

Steps

- Go to the Configure > My Proxy > Subscription > Subscription Management page of the Content Gateway manager.
- 2) Enter your key in the field provided.
- 3) Click Apply.
- 4) Click Restart on the Configure > My Proxy > Basic > General page.

Providing system information

Content Gateway communicates with Filtering Service to apply policies to web requests. You can configure the timeout period used to determine whether Content Gateway can connect to Filtering Service, and define what Content Gateway does when communication is interrupted.

To do this:

Steps

- 1) Log on to the Content Gateway manager.
- Go to the Configure > My Proxy > Subscription > Scanning tab and note the Filtering Service IP address and port. This is information that you entered when you installed Content Gateway.
- Review the Communication Timeout setting. This is the time, in milliseconds, that Content Gateway waits on communication with Policy Server or Filtering Service before timing out and triggering the Action for Communication Errors setting.

The default timeout value is 5000 ms (5 seconds). If you change the value, you must restart Content Gateway.

- 4) In the Action for Communication Errors section, specify whether to permit or block traffic if a communication timeout condition occurs. When a timeout occurs, Content Gateway applies the setting and regularly polls the services to detect their renewed availability.
- 5) Use the Scanning Data Files Update section to configure how long to wait after downloading security analytic data files before they are put into use. Select a Delay time from the drop-down list. Keep in mind that the longer the delay, the higher the security risk. The Suspend updates option is not recommended for extended use. Selecting it will prompt an alarm as a reminder that downloads have been suspended. It is recommended that you not clear the alarm until Delay time has been reset.

When a delay time is in place, there may be up to 2 sets of data files present on the Content Gateway machine.

- The current set of data files that are being used by the analytics.
- The set of data files whose complete download is being delayed.

Once the delay period is met, the delayed database is moved to the current set of files and the delay period is applied to next download.

This feature is typically used for a backup system.

6) When you have finished making changes, click Apply.

Verifying that the proxy is processing Internet requests

After you have installed the proxy, verify that it is processing requests for web content.

Steps

- 1) Log on to the Content Gateway manager.
- 2) Go to the **Monitor** > **My Proxy** > **Summary** page to view subscription detail, scanning data file status, and node details, including the number of objects served, and other basic proxy service information.
- 3) Navigate to Monitor > Protocol > HTTP > General to display the General HTTP Statistics table.

4) Note the current Total Document Bytes statistic in the Client section of the table.

General Transaction FTP over HTTP

General HTTP Statistics	I				
Attribute	Current Value				
Client					
Total Document Bytes	1.8 GB				
Total Header Bytes	1.7 MB				
Total Connections	34,758				
Current Connections 0					
Transactions in Progress	0				
Server					
Total Document Bytes	1.7 GB				
Total Header Bytes	1.3 MB				
Total Connections	35,776				
Current Connections	0				
Transactions in Progress	0				

- 5) Set your browser to the proxy port.
- 6) Browse the Internet.
- Recheck the Total Document Bytes statistic.
 This value increases as the proxy processes HTTP requests.

Using the command-line interface

The command-line interface provides a quick way to view proxy statistics and configure Content Gateway if you do not have access to a browser or if you prefer to use a UNIX shell-like command interface.



Note

This facility is not available when Content Gateway is hosted on a Forcepoint appliance. Instead, use the Content Gateway manager and see your appliance documentation.

You can execute individual commands or script multiple commands in a shell. See Content Gateway commands.

Steps

1) Become root:

su

- Change to the Content Gateway bin directory (/opt/WCG/bin). Run Content Gateway commands from this directory.
 - Commands take the form: content_line -<command_argument>
 - To view a configuration setting, enter the following command:

```
content_line -r <var>
```

Here, <*var>* is the variable associated with the configuration option (for a list of the variables, refer to *Configuration variables*).

• To change the value of a configuration setting, enter the following command:

content_line -s <var> -v <value>

Here, <*var*> is the variable associated with the configuration option and <*value*> is the value you want to use.

3) For a list of content_line commands, enter:

content_line -h



Note

If the Content Gateway bin directory is not in your path, prepend the command with "./".

For example:

```
./content_line -h
```

Related concepts

Configuration variables on page 388

Related information Commands and Variables on page 267

Starting and stopping Content Gateway on the command line

To stop or start Content Gateway from the command line:

Steps

1) Become root:

su

2) Change to the Content Gateway installation directory (/opt/WCG).

- 3) Do one of the following:
 - To start the proxy:
 ./WCGAdmin start
 - To stop the proxy:
 - ./WCGAdmin stop
 - To restart the proxy:
 ./WCGAdmin restart
 - To see which Content Gateway services are running:
 - ./WCGAdmin status

The no_cop file

The presence of the /opt/WCG/config/internal/no_cop file acts as an administrative control that instructs the **content_cop** process to exit immediately without starting **content_manager** or performing any health checks. The **no_cop** file prevents the proxy from starting automatically when it has been stopped with the ./WCGAdmin stop command.

Without such a static control, Content Gateway would restart automatically upon system reboot. The **no_cop** control keeps Content Gateway off until it is restarted with the ./WCGAdmin start command.

When the **no_cop** file prevents Content Gateway from starting, the following message is recorded in the system log file:

content_cop[16056]: encountered "config/internal/no_cop" file...exiting

Chapter 3 Explicit Proxy

Contents

- Introduction on page 27
- Manual browser configuration on page 28
- Using a PAC file on page 28
- Using WPAD on page 30
- Configuring FTP clients in an explicit proxy environment on page 32

Introduction

If Internet requests are not transparently routed to Content Gateway via a Layer 4 switch or router (see *Transparent Proxy and ARM*), traffic must be **explicitly** routed to Content Gateway by configuring the client's Internet browser. (This is sometimes referred to as an *explicit proxy deployment*.)

Clients can configure their web browsers in 1 of 3 ways:

- By directly configuring their browsers to send requests directly to the proxy. See Manual browser configuration.
- By configuring their browsers to download proxy configuration instructions from a proxy auto-config (PAC) file.
 See Using a PAC file.
- By using Web Proxy Auto-Discovery Protocol (WPAD) to download proxy configuration instructions from a WPAD server (Microsoft Internet Explorer only). See Using WPAD.

In addition, if Content Gateway is configured to proxy FTP traffic, FTP client applications, such as FileZilla or WS_FTP, must be configured to explicitly send requests to the proxy. See *Configuring FTP clients in an explicit proxy environment*.

Related concepts Manual browser configuration on page 28 Using WPAD on page 30 Configuring FTP clients in an explicit proxy environment on page 32

Related tasks

Using a PAC file on page 28

Related information

Transparent Proxy and ARM on page 35

Manual browser configuration

To configure a browser to send requests to Content Gateway, clients must provide the following information for each protocol they want the proxy to serve to their browsers:

The proxy's hostname or IP address.



Important

If Integrated Windows Authentication is configured for user authentication, the Fully Qualified Domain Name must be used. Specifying the IP address will result in authentication failure. See *Integrated Windows Authentication*.

• The proxy server port. The Content Gateway default proxy server port is 8080.



Important

Do not set up the IP address of the Content Gateway proxy to be a virtual IP address.

Although the Content Gateway manager does not prohibit the entry of a virtual IP address, the proxy does not function properly if a VIP is used.

In addition, clients can specify not to use the proxy for certain sites. Requests to the listed sites go directly to the origin server.

For Microsoft Internet Explorer, proxy configuration settings are in **Tools > Internet Options > Connections > LAN Settings**. By default, Microsoft Internet Explorer sets all protocols to the same proxy server. To configure each protocol separately, click **Advanced** in the **LAN Settings** section. See the browser documentation for complete proxy configuration instructions.

For Mozilla Firefox, proxy configuration settings are in **Tools > Options > Advanced > Network Setting > Connection Settings > Manual Proxy Configuration**. By default, you must configure each protocol separately. However, you can set all protocols to the same proxy server by selecting **Use this proxy server for all protocols**.

You do not have to set configuration options on the proxy to accept requests from manually configured browsers.

Related concepts

Integrated Windows Authentication on page 173

Using a PAC file

A PAC file is a JavaScript function definition that a browser calls to determine how requests are handled. Clients must specify in their browser settings the URL from which the PAC file is loaded.

You can store a PAC file on the proxy and provide the URL for this file to your clients.



Note

The PAC file can reside on any server in your network. Small networks may store the file on the proxy itself, but large, enterprise-class networks should use a separate server for storing the PAC file.

If the HTTPS protocol option is enabled, see *Enabling SSL support* for information about a PAC file to use with HTTPS traffic.

Steps

- 1) If you have an existing **proxy.pac** file, replace the **proxy.pac** file located in the Content Gateway **config** directory with your existing file.
- 2) Navigate to the Configure > Content Routing > Browser Auto-Config > PAC tab.
- In the Auto-Configuration Port field, specify the port that Content Gateway uses to serve the PAC file. The default port is 8083.
- 4) The PAC Settings area displays the proxy.pac file:
 - If you copied an existing PAC file into the Content Gateway config directory, the proxy.pac file contains your proxy configuration settings. Check the settings and make changes if necessary.
 - If you did not copy an existing PAC file into the Content Gateway config directory, the PAC Settings area is empty. Enter the script that provides the proxy server configuration settings. A sample script is provided in Sample PAC file. See, also, the article titled "PAC File Best Practices" in the Documentation Library.
- 5) Click Apply.
- 6) Go to the Configure > My Proxy > Basic > General tab and click Restart.
- 7) Inform your users to set their browsers to point to this PAC file.

For example, if the PAC file is located on the proxy server with the hostname **proxy1** and Content Gateway uses the default port 8083 to serve the file, users must specify the following URL in the proxy configuration settings:

http://proxy1.company.com:8083/proxy.pac

Next steps

The procedures for specifying the PAC file location vary among browsers.

Related concepts

Sample PAC file on page 30

Related tasks Enabling SSL support on page 113

Procedure to specify the PAC file location for Microsoft Internet Explorer

Steps

1) Go to Tools > Internet Options > Connections > LAN Settings.

- 2) Select Use automatic configuration script.
- In the Address field, enter: http://<proxy_host>:8083/proxy.pac
- 4) Click OK.

Procedure to specify the PAC file location for Mozilla Firefox

Steps

- 1) Go to Tools > Options > Advanced > Network > Connection > Settings.
- Select Automatic proxy configuration URL field, and enter http://<proxy_host>:8083/proxy.pac
- Click Reload, and then click OK. See the documentation for your browser for details.

Sample PAC file

The following sample PAC file instructs browsers to connect directly to all hosts without a fully qualified domain name and to all hosts in the local domain. All other requests go to the proxy server called **myproxy.company.com**.

```
function FindProxyForURL(url, host)
{
if (isPlainHostName(host) || dnsDomainIs(host,
".company.com"))
return "DIRECT";
else
return "PROXY myproxy.company.com:8080; DIRECT";
}
```

Using WPAD

WPAD allows Internet Explorer to automatically detect a server that can supply it with proxy server configuration settings. Clients do not have to configure their browsers to send requests to a proxy server: a single server provides the settings to all clients on the network.



Note

WPAD is incompatible with transparent proxy deployments.

When an Internet Explorer browser starts up, it searches for a WPAD server. It prepends the hostname WPAD to the current fully qualified domain name. For example, a client in **x.y.company.com** searches for a WPAD server at **wpad.x.y.company.com**. If unsuccessful, the browser removes the bottommost domain and tries again; for example, it tries **wpad.y.company.com**. The browser stops searching when it detects a WPAD server or reaches the third-level domain, **wpad.company.com**. The algorithm stops at the third level so that the browser does not search outside the current network.



Note

By default, Microsoft Internet Explorer are set to automatically detect WPAD servers. However, browser users can disable this setting.

To configure Content Gateway to be a WPAD server

Steps

- 1) If you have an existing **wpad.dat** file, replace the **wpad.dat** file located in the Content Gateway **config** directory with your existing file.
- Log on to the Content Gateway manager and go to Configure > Content Routing > Browser Auto-Config > WPAD to display the wpad.dat file.
- 3) The WPAD Settings area displays the wpad.dat file:
 - If you copied an existing wpad.dat file into the Content Gateway config directory, the file contains your proxy configuration settings. Check the settings and make changes if necessary.
 - If you did not copy an existing wpad.dat file into the Content Gateway config directory (/opt/WCG/ config), the WPAD Settings area is empty. Enter a script that will provide the proxy server configuration settings. A sample script is provided in Sample PAC file (a wpad.dat file can contain the same script as a proxy.pac file).
- 4) Click Apply.
- 5) Navigate to Configure > Networking > ARM.
- 6) In the Redirection Rules section, click Edit File to add a special remap rule to the ipnat.conf file.

- 7) Enter information in the fields provided, and then click Add:
 - a) Enter the Ethernet Interface that receives browser WPAD requests (for example hme0 or eth0).
 - b) From the Connection Type drop-down list, select tcp.
 - c) In the Destination IP field, enter the IP address of the Content Gateway server that will be resolved to the WPAD server name by the local name servers.
 - d) (*Optional*) In the Destination CIDR field, enter the CIDR mask value. If the Destination IP is in IPv4 format, enter 32. Enter 128 for an IPv6 Destination IP.
 - e) In the Destination Port field, enter 80.
 - f) In the Redirected Destination IP field enter the same IP address you entered in the Destination IP field.
 - g) In the Redirected Destination Port field, enter 8083.
 - h) (Optional) In the User Protocol field, select dns.
- 8) Click Add.
- 9) Use the arrow keys on the left to move the new rule to the first line in the file.
- 10) Click Apply, and then click Close.
- 11) Go to the Configure > My Proxy > Basic > General tab and click Restart.

Related concepts

Sample PAC file on page 30

Configuring FTP clients in an explicit proxy environment

When Content Gateway is configured to proxy FTP traffic (see *FTP*), also configure FTP client applications, such as FileZilla or WS_FTP, to send FTP requests to the proxy. After the initial configuration, users work with the FTP client application as if no proxy were present.

To connect to an FTP server, 4 pieces of information are usually needed. These pieces of information are mapped as follows:

From:	То:				
FTP server hostname	FTP proxy hostname				

From:	То:
FTP server port number	FTP proxy port number (default is 2121)
FTP server username	FTP_server_username@FTP_server_hostname For example: anon@ftp.abc.com
FTP server password	FTP server password

Some FTP client applications have a configuration page for specifying FTP proxy information. Update those settings to point to the Content Gateway FTP proxy. See your FTP client application documentation.

Settings	×
Select page: - Connection - FTP - Active mode - Passive mode - FTP Proxy - SFTP - Generic proxy - Transfers - File Types - File exists action - Interface - Themes - Date/time format - File lists - Language - File editing - File diting - File ype associations - Updates - Logging - Debug	FTP Proxy Type of FTP Proxy: None USER@HOST SITE OPEN Custom USER %u@%h PASS %p USER %u@%h PASS %p Format specifications: %h - Host %u - Username %p - Password %a - Account (Lines containing this will be omitted if not using Account logontype) %s - Proxy user %w - Proxy password Proxy host: Cghostname:2121 Proxy password: Proxy user: Proxy password: Note: This only works with plain, unencrypted FTP connections.
Cancel	

Here is an example configuration using a recent version of FileZilla.

In the FTP Proxy area:

1) Set FTP Proxy to Custom and enter the following definitions:

USER %u@%h

PASS %p

- 2) Set Proxy host to the Content Gateway FTP proxy hostname and port number.
- 3) Click OK.

The user then enters FTP connection information in the usual way, as if no proxy were present. For example:

Host: ftp.example.com

Username: anon

Password: 123abc

If the FTP client application is **not** configured, the user must enter FTP requests as shown below.

Host: Content Gateway proxy hostname

Username: anon@ftp.example.com

Password: 123abc

Port: 2121

FileZilla										
File Edit View 1	Transfer Server	Bookmark	s Help							
H - E			Īx Ī∕	≣ <u>Q́</u>	o 🦰					
Host: cghostname	Username	: anon@ft	p.example.c	Password:	•••••	Po	ort: 2121	Quickconnect	•	
										A
										-
Local site: C:\temp	p\				 Remote site: 					~
	sysprep	£								
Server/Local file		Direction	Remote file			Size	Priority	Status		
Queued files	Failed transfers	Success	ul transfers							
								Queues	empty	•• //.

Related reference

FTP on page 294

Chapter 4 Transparent Proxy and ARM

Contents

- Introduction on page 35
- The Content Gateway ARM on page 36
- Transparent interception strategies on page 37
- Interception bypass on page 58
- Connection load shedding on page 61
- Reducing DNS lookups on page 61

Introduction

The transparent proxy option enables Content Gateway to respond to client Internet requests without requiring users to reconfigure their browsers. It does this by redirecting the request flow to the proxy after the traffic has been intercepted, typically by a Layer 4 (L4) switch or router.

In a transparent proxy deployment:

- 1) The proxy intercepts client requests to origin servers via a switch or router. See *Transparent interception strategies*.
- The Adaptive Redirection Module (ARM) intercepts incoming packets and redirects them to the proxy. (The ARM is always enabled.)
- 3) The proxy receives and begins processing the intercepted client requests. The proxy retrieves the object from the origin server and serves it to the client.



Important

For transparent proxy configurations with multiple interfaces or gateways, Content Gateway must have proper routes to clients and the Internet in the operating system's routing table.

For HTTP, the proxy can identify problem clients and servers, and the ARM can disable interception for those clients and servers, passing their traffic directly to the origin server. You can also create ARM static bypass rules to exempt clients and servers from being redirected to the proxy. See *Interception bypass*.

Related concepts

Transparent interception strategies on page 37 Interception bypass on page 58

The Content Gateway ARM

The ARM inspects incoming packets before a routing decision is made and redirects the packets to Content Gateway for processing.

The ARM uses iptables, policy routing, and transparent sockets configured during product installation. The installation process also creates redirection rules to intercept packets. The ARM is always active.

To ensure that the proxy can serve HTTP, HTTPS, FTP, and DNS requests transparently, verify the redirection rules in the **ipnat.conf** file and edit them if necessary.

- If you are using WCCP for transparent interception, there must be a redirection rule for every port in every active service group.
- Rules for standard ports are included by default.

To review and edit the ARM redirection rules:

Steps

- 1) Log on to the Content Gateway manager and go to the Configure > Networking > ARM > General tab.
- 2) Verify the Redirection Rules (taken from the ipnat.conf file) and make any needed changes. To change a redirection rule:
 - a) Click Edit File to open the configuration file editor for the **ipnat.conf** file.
 - b) Select the rule you want to edit and modify the appropriate fields.
 - c) Click Set and then click Apply to apply your changes.
 - Click Close to exit the configuration file editor. All fields are described in ARM.
- 3) If you have made any changes, go to the Configure > My Proxy > Basic > General tab and click Restart.

Related reference ARM on page 397

Configuring a firewall with ARM

The ARM module uses a firewall. To facilitate traffic interception and redirection:

- IPTables rules are configured during Content Gateway installation and upgrade.
 - Forcepoint IPTables chains are inserted.
 - Forcepoint IPTables rules are inserted into existing chains.
 - Forcepoint chains and rules use "NC_" as a prefix for identification purposes.
- IPTables rules configured outside of the Content Gateway manager must:
 - Be inserted *after* most Forcepoint rules.
Customized rules should, however, be added *before* the NC_RESERVED_FWD_DEF rule. This Forcepoint rule was added to the forward chain so that traffic that is not specifically handled by the proxy is dropped and not forwarded. This rule should always be the last rule in the forward chain

- Never be added to Forcepoint chains.
- Never be modified on Appliance platforms.
- Forcepoint chains and rules should never be edited.
- If customized chains or rules impact the Forcepoint configuration, navigate to the Content Gateway bin directory (/opt/WCG/bin) and run the following command:

netcontrol.sh -r

This re-establishes the Forcepoint IPTables chains and rules.

Transparent interception strategies

Content Gateway supports the following transparent interception solutions:

- A Layer 4 switch. See Transparent interception with a Layer 4 switch.
- A router or switch that supports WCCP v2. Cisco IOS-based routers are the most common. See Transparent interception with WCCP v2 devices.
- Policy-based routing. See Transparent interception and multicast mode.
- Software routing. See Transparent interception with software-based routing.

Exactly how client requests reach the proxy depends on network topology. In a complex network, you must decide which clients are to be served transparently and make sure that network devices and the proxy are positioned to intercept their requests. Content Gateway, or routers or switches feeding Content Gateway, are often deployed at a major artery or aggregation pipe to the Internet.

To configure Content Gateway to serve only transparent traffic, see *Configuring Content Gateway to serve only transparent requests*.

Related concepts

Transparent interception with a Layer 4 switch on page 37 Transparent interception with WCCP v2 devices on page 38 Transparent interception with software-based routing on page 56 Configuring Content Gateway to serve only transparent requests on page 57

Related tasks

Transparent interception and multicast mode on page 55

Transparent interception with a Layer 4 switch

Layer 4 switches can redirect supported protocols to the proxy, while passing all other Internet traffic directly to its destination, as shown below for HTTP.



Layer 4 switches offer the following features, depending on the particular switch:

- A Layer 4 switch that can sense downed hosts on the network and redirect traffic adds reliability.
- If a single Layer 4 switch feeds several proxy servers, the switch handles load balancing among the Content Gateway nodes. Different switches might use different load-balancing methods, such as round-robin or hashing. If a node becomes unavailable, the switch redistributes the load. When the node returns to service, some switches return the node to its previous workload.

Note

It is recommended that you do **not** enable Content Gateway virtual IP failover when a switch is providing load balancing in a cluster configuration.

Transparent interception with WCCP v2 devices

Content Gateway supports transparent interception with WCCP v2-enabled routers and switches.

HTTP, HTTPS, FTP, and DNS protocols are supported. Default ARM redirection rules are included for HTTP, HTTPS, and FTP communicating on standard ports.



Important

The network clients, Content Gateway proxy servers, and destination web servers (default gateway) must reside on separate subnets.

WCCP v2 interception works as follows:

- 1) WCCP v2 devices send HTTP, HTTPS, FTP, and DNS traffic, per the configuration of the service group, to the proxy server or cluster of servers.
- 2) The ARM redirects traffic. For example, HTTP traffic on port 80 is redirected to Content Gateway port 8080.

- 3) The proxy processes the request as usual, sending the response back to the client.
- 4) The ARM redirects packets to the client. As a result, the user sees the response as if it had been sent directly from the origin server.



WCCP v2 setup outline

Steps

- 1) Install and configure your WCCP v2 devices. On each WCCP v2 device:
 - Configure the service groups.
 - Configure password security, if needed.
 - Configure multicast communication, if needed. See Configuring WCCP v2 routers.

- 2) Configure Content Gateway to work with your WCCP devices.
 - Define matching service groups.
 In addition to network interface, protocols, ports, authentication (if used), and multicast communication (if used), also configure:
 - The IP addresses of the WCCP v2 devices.
 - The Packet Forward Method and Packet Return Method.
 - If Content Gateway is deployed in a cluster, **assignment method** load distribution, if desired.
 - Create ARM redirect rules for non-standard ports.
 See Enabling WCCP v2 in Content Gateway and The Content Gateway ARM.
- 3) Validate the configuration with test traffic.

Related tasks Configuring WCCP v2 routers on page 42 Enabling WCCP v2 in Content Gateway on page 48 The Content Gateway ARM on page 36

WCCP v2 supported features

Content Gateway supports the following WCCP v2 features:

- Multiple routers in a proxy cluster
- Multiple ports per service group
- Multiple service groups per protocol. Sometimes it is necessary or convenient to have different service groups for different WCCP devices. For example, for Cisco ASA firewall, different service groups are required for each WCCP device in the network.
- Dynamic load distribution in a cluster through assignment method HASH or MASK. See WCCP load distribution.
- Packet Return Method and Packet Forward Method negotiation
- MD5 password security per service group
- Multicast mode

In a Content Gateway cluster, it is recommended that you **not** enable virtual IP failover in WCCP environments. WCCP v2 and the Content Gateway configuration handles node failures and restarts. (See *WCCP load distribution* and *Virtual IP failover*.) However, if a Content Gateway cluster uses WCCP exclusively, virtual IP failover can be used if no user authentication features are used. Note that the WCCP assignment method — not virtual IP failover — is the recommended method for managing load distribution. If a Content Gateway cluster receives requests both explicitly and transparently (the networks must be separate; this type of deployment is not recommended), virtual IP failover can be used on the explicit proxy network segment.

Related concepts WCCP load distribution on page 41 Virtual IP failover on page 76

ARM bypass and WCCP

If Content Gateway has an ARM bypass rule (discussed in *Interception bypass*), Content Gateway forwards particular client requests directly to the origin server, bypassing the proxy.

Bypassed requests are unchanged by the ARM.

With WCCP v2, you can exclude certain router interfaces from redirection. **Content Gateway ARM bypass rules work only if you exclude the router interface that Content Gateway is connected to from WCCP redirection**. You do this on the router by selecting the interface connected to Content Gateway and issuing the router

configuration command ip wccp redirect exclude in. This causes the router to exclude traffic inbound on the specified interface from all redirection rules.

Related concepts

Interception bypass on page 58

WCCP load distribution

The WCCP protocol provides the **assignment method** for dynamic symmetric and asymmetric load distribution in a cluster. WCCP detects node failures and performs redistribution based on the configuration communicated to it by Content Gateway.

- Load distribution is configured in Content Gateway Manager and is pushed to the WCCP devices.
- Load distribution is configured **per service group**. For each service group:
 - Participating cluster members must be registered to the service group. (The WCCP device makes no decisions about load balancing.)
 - The HASH or MASK assignment method is selected. HASH is typically used with the GRE forward/return method, and MASK with the L2 forward/return method.



Important

MASK was developed specifically for the Cisco Catalyst series switches, and is one of the key characteristics that enable WCCP interception to be performed completely in hardware on these platforms. It should be used only with devices for which there is documented support.

- One or more distribution attributes are selected. Typically the destination IP address is used.
- If load is to be distributed to different cluster members in different proportions, a weight value is set on each cluster member. These values determine the proportion of requests each will receive relative to other members of the cluster. This option is only useful if the Synchronize in the Cluster option is disabled. See Configuring service groups in the Content Gateway manager.

Asymmetric load distribution using the **weight** value is helpful when:

- There are multiple Content Gateway servers with different performance capabilities.
- The Internet traffic profile doesn't lend itself to even distribution due to preferences for specific origin servers (and therefore destination IP addresses).

Related tasks

Configuring service groups in the Content Gateway manager on page 49

How dynamic redistribution works

Dynamic redistribution is accomplished when the WCCP device detects that a cluster member is offline. It then automatically redistributes the load to the remaining cluster members based on the load distribution configuration. When a cluster member returns to service and is detected by the WCCP device, load distribution is, again, automatically adjusted based on the configuration.

For configuration steps, see Configuring service groups in the Content Gateway manager.

Related tasks

Configuring service groups in the Content Gateway manager on page 49

How the weight value supports asymmetric load distribution



Important

Weight is only useful if the **Synchronize in the Cluster** option is disabled. See *Configuring service* groups in the Content Gateway manager.

The weight value is unique to each service group and node. The weight value does not propagate around the cluster and must be set individually on every node in the cluster.

The value of weight, relative to the settings on other cluster members, determines the proportion of traffic that WCCP directs to the node.

By default, weight is set to 0, which results in equal distribution to all cluster members.

To achieve asymmetric distribution, weight is set relative to other members of the cluster. For example, assume a cluster of 3 nodes:

Node	Weight value	Load distribution
Node1	50	50%
Node2	25	25%
Node3	25	25%

If Node1 goes offline, Node2 and Node3 will get an equal amount of traffic. If Node3 goes offline, Node1 will get two thirds of the traffic and Node2 will get one third of the traffic.

Because the weight value is relative to the settings on other cluster nodes, the same distribution as above can be achieved with weight values of 10, 5, 5. (The valid range of weight is 0-255.)

If weight is changed from its default value of 0, it should be configured on all nodes in the cluster.

Related tasks

Configuring service groups in the Content Gateway manager on page 49

Configuring WCCP v2 routers

Consult the documentation for your WCCP v2 device, as well as the manufacturer's support site, for device configuration and performance information.

Most devices should be configured to take best advantage of hardware-based redirection.

With Cisco devices, the most recent version of IOS is usually the best.

To prepare WCCP v2 devices for use with the proxy:

Steps

- 1) Configure one or more service groups for the protocols you intend to use. A service group can handle one or multiple protocols. See *Configuring servicegroups on the WCCP device*.
- 2) Configure the router to enable WCCP processing for these service groups. See *Enabling WCCP processing for a service group*.
- 3) Optionally, enable router security. Router security must also be enabled for the service group in Content Gateway. See *Enabling WCCP v2 security on the router*.



Note

For instructions on configuring your specific router, please refer to the documentation provided by your hardware vendor. For Cisco routers, see http://www.cisco.com/cisco/web/psa/ default.html?mode=prod and search for your IOS and device version, for example, IOS 12.4.

4) When you are done configuring the router, you must also configure and enable WCCP in the Content Gateway manager. See *Enabling WCCP v2 in Content Gateway*.

Related tasks

Configuring service groups on the WCCP device on page 43 Enabling WCCP processing for a service group on page 44 Enabling WCCP v2 security on the router on page 48 Enabling WCCP v2 in Content Gateway on page 48

Configuring service groups on the WCCP device

WCCP uses service groups to specify the traffic that is redirected to Content Gateway (and other devices).

- A service group can intercept one or more protocols on one or more ports.
- Service groups are assigned a unique integer identifier (ID) from 0 to 255.
- Service groups IDs are user defined; they do not have a default port or traffic type.

The following table illustrates a set of service group definitions that are often found in networks. If you are configuring for IP spoofing, see the table in *Content Gateway IP spoofing* for common reverse service group IDs.

Service ID	Port	Traffic Type
0	80	HTTP
5	21	FTP
70	443	HTTPS (when HTTPS support is enabled)

Service groups must be configured on the router and in Content Gateway.

The best practice is to configure the routers first and Content Gateway second.

Follow the instructions in your router documentation for specifics, but in general:

Steps

- To see what has been configured on the router for WCCP, enter: show running-config | include wccp
- To enable WCCP v2, enter: ip wccp version 2
- 3) If you used another proxy cache with your router prior to Content Gateway, disable the service ID that was previously used. For example, if you have a Cisco router, disable the service ID web-cache as follows: no ip wccp web-cache
- Specify the service group IDs you will use with Content Gateway. For the specific commands to use, see your router documentation.

You must configure each service group supported by the router individually. You cannot configure a router globally.

Related concepts

Content Gateway IP spoofing on page 63

Enabling WCCP processing for a service group

For each WCCP v2 service group that you configure, you must enable WCCP processing.

WCCP v2 routers contain multiple network interfaces, including:

- one or more interfaces that receive inbound (ingress) client traffic
- one or more interfaces connected to Content Gateway
- an interface dedicated to outbound (egress) traffic that is aimed at the Internet



Following are some guidelines for enabling WCCP processing for a service group on a router. Consult the procedures in your router documentation for specifics.

Steps

1) Turn on the WCCP feature:

```
ip wccp <service group ID> password [0-7] <passwd>
```

 On the router or switch interface, enable redirection for incoming (ingress) packets or outgoing (egress) packets.



Note

Where your hardware and network topology support it, it is recommended that redirection be performed on the ingress interface (using the "redirect in" commands).

The following are examples. Be sure to substitute the service group IDs that you have established on your router(s).

First, select the interface to configure:

interface <type> <number>

Second, establish your redirection rules:

ip wccp <service group ID> redirect in

Examples for inbound redirection:

Run these commands for each protocol that you want to support, but only on the interfaces dedicated to **inbound (ingress)** traffic.

For example, to turn on redirection of HTTP destination port traffic, enter:

ip wccp 0 redirect in

To turn on redirection of HTTPS destination port traffic:

ip wccp 70 redirect in

To turn on redirection of FTP destination port traffic enter:

ip wccp 5 redirect in

To turn on redirection of HTTP source port traffic, which is required for IP spoofing, enter:

ip wccp 20 redirect in

Examples for outbound redirection:

Run these commands for each protocol that you want to support, but only on the interfaces dedicated to **outbound (egress)** traffic.

First, select the interface to configure:

interface <type> <number>

Second, establish your redirection rules:

ip wccp <service group ID> redirect out

For example, to turn on redirection for HTTP, enter:

ip wccp 0 redirect out

To turn on redirection for HTTPS:

ip wccp 70 redirect out

To turn on redirection for FTP enter:

ip wccp 5 redirect out

3) When ARM bypass occurs, or IP spoofing is enabled, the proxy sends traffic to the Internet with the original source IP address. The "redirect exclude in" command prevents the router from looping the traffic back to Content Gateway.



Important

When ARM dynamic or static bypass is enabled, or IP spoofing is enabled, and redirection is on the outbound (egress) interface, exclude redirection of Content Gateway outbound packets on the router interface that handles Content Gateway's egress traffic. See the illustration, below.

- a) Select the interface that handles Content Gateway egress traffic: interface <type> <number>
- b) Exclude Content Gateway outbound traffic on the interface from all redirection rules on the router: ip wccp redirect exclude in



Disabling WCCP processing for a service group

If you need to disable WCCP processing for any reason, issue this command to turn off the WCCP feature:

no ip wccp <service group ID> password [0-7] <passwd>

Enabling WCCP v2 security on the router

If you are running WCCP v2, you can enable security on the Content Gateway node so that the proxy and your routers can authenticate each other. You must individually enable security for each service group that the router supports. You cannot configure a router globally as you would Content Gateway.

You enable the security option and provide the authentication password in the Content Gateway manager.

The authentication password you specify must match the authentication password configured on the router for each service group being intercepted. The following procedure provides an example of how to set an authentication password for different service groups.

Steps

- 1) Telnet to the router and switch to **Enable** mode.
- 2) At the prompt, enter the following command to configure the router from the terminal: configure terminal
- 3) If you defined a password when you enabled WCCP on the router, skip to step 4. Otherwise, enter the following command for each service group that the router intercepts:

<hostname>(config)# ip wccp <service_group> password <pwd>

Here, *<hostname>* is the host name of the router you are configuring, *<service_group>* is the service group ID (for example, 0 for HTTP), and *<pwd>* is the password you want to use to authenticate Content Gateway. This password must match the password you specify in the Content Gateway configuration for this service group.

4) Exit and save the router configuration.

Enabling WCCP v2 in Content Gateway

After you have configured your WCCP v2 routers, these steps remain:

Steps

- 1) Enabling WCCP in the Content Gateway Manager:
 - a) Go to the Configure > My Proxy > Basic > General tab.
 - b) Under Features > Networking, locate WCCP and click On, then Apply. Do not restart Content Gateway.
- 2) Configuring service groups in the Content Gateway manager

3) Restarting Content Gateway.

Important

Before you restart Content Gateway, make sure that your configuration meets the following requirements:

- Cisco IOS devices are running a very recent version of IOS with all appropriate patches applied.
- WCCP routers are programmed with the correct service groups and other features.

Related tasks

Configuring service groups in the Content Gateway manager on page 49

Configuring service groups in the Content Gateway manager

Every WCCP service group that redirects traffic to a Content Gateway proxy must have a corresponding service group defined for it in the Content Gateway server or cluster.

To define service groups:

Steps

- 1) Go to the Configure > Networking > WCCP page.
- Review the existing entries in the Service Groups table and click Edit File to add, modify, delete, or reorder service groups.
 - Entries are stored in the wccp.config file.
 - Click Refresh to prompt the Content Gateway manager to reread the wccp.config file and update the table.
 - Detailed instructions can be found in *Configuring a service group (editing wccp.config)*.
- If Content Gateway is configured in a cluster, under Synchronize in the Cluster, Select Enabled (default) or Disabled. (The value of this option is always synchronized in the cluster.)
 - When this option is enabled, the WCCP configuration (stored in wccp.config) is synchronized in the cluster and configuration changes can be made on any node in the cluster.
 - When this option is disabled, the WCCP configuration is not synchronized in the cluster and changes to the WCCP configuration must be made individually on each node. A common use case for this is to control which service groups are enabled/disabled on each node, and/or to use proportional load distribution using weight.
 - If this option is disabled, and then later enabled, the configuration on the node on which the administrator enables the option is used to initially synchronize the cluster.



CAUTION

When **Synchronize in the Cluster** is **disabled**, you must visit each node in the cluster to examine and maintain your WCCP configuration. This can also make WCCP troubleshooting more difficult.

Related tasks

Configuring a service group (editing wccp.config) on page 50

Configuring a service group (editing wccp.config)

Open the file and define the service group

Steps

- 1) On the Configure > Networking > WCCP page, click Edit File to open wccp.config in the editor.
 - Defined service groups are summarized at the top of the page.
 - Click an entry in the list to view its complete details, modify, or reposition it.
 - When an entry is selected, the down and up arrows to the left of the list reposition the entry in the list.
 - Click X to delete a selected entry.
- 2) For each service group, enter the following information:
 - a) To enable a service group, set Service Group Status to Enabled. A service group can be defined but not active.
 - b) Specify a unique Service Group Name. The service group name is an aid to administration.
 - c) Specify a WCCP **Service Group ID** from 0-255. This ID must match a corresponding service group ID configured on the router. See *Configuring service groups on the WCCP device*.
 - d) Specify the network **Protocol** applicable to the service group (TCP or UDP).
 - e) Specify the Ports that this service group will use.
 Select Specify ports to enter up to 8 ports in a comma-separated list.

Select All ports to redirect traffic from all ports.



Important

Every port in the service group must have a corresponding ARM redirection rule to redirect the traffic to Content Gateway. See *The Content Gateway ARM*.

f) From the drop down list, select the **Network Interface** on the Content Gateway host system that this service group will use.



Note

When the FTP protocol is enabled, select eth0. Client machines must be reachable via the eth0 interface.

Related tasks

Configuring service groups on the WCCP device on page 43 The Content Gateway ARM on page 36

Configure mode negotiation

- The **Packet Forward Method** determines how traffic is transmitted from the WCCP router to the proxy.
- The Packet Return Method specifies the method used to return traffic back to the WCCP router.



Important

If you change the forward/return method configuration while there is an active connection with the WCCP device, in order to re-negotiated the method you must force the current connection to terminate. Typically, this means turning off the service group on the WCCP device for 60 seconds. See the documentation for your WCCP device.

If multiple proxies are installed in your environment, each with WCCP enabled, but configured with different **Packet Forward** and **Packet Return Methods**, traffic may not be processed. Some routers support only a single **Packet Forward Method** within a group and may forward packets to the other proxies using a method they do not support.

Typically the router supports only one method, and the forward and return methods match.

Steps

- If traffic is routed to the proxy by a Cisco ASA firewall, in the Special Device Profile drop down box select ASA Firewall. When this option is selected, GRE is automatically selected for both Packet Forward Method and Packet Return Method. These settings cannot be changed.
- If traffic is routed to the proxy by a router or switch, select the Packet Forward Method (L2 or GRE) and Packet Return Method that matches the capabilities and position of your router or switch.

If Content Gateway is configured with a Forward/Return method that the router does not support, the proxy negotiates the method supported by the router.

If L2 is selected, L2 is automatically selected as the return method (GRE is not an option).



Important

Selecting L2 requires that the router or switch be Layer 2-adjacent (in the same subnet) as Content Gateway.

If GRE is selected, for each router in the service group a unique Content Gateway tunnel endpoint IP address must be specified in the WCCP Routers section (see the "Provide router information" step, below).



Important

GRE cannot be used with WCCP multicast mode.

GRE return, as documented by Cisco (see this site), is fully functional in all deployments. GRE enhanced tunnel return, in which the proxy forwards traffic back to the router, is *only available on an appliance*. Contact Technical Support for information on how to enable the functionality.

Configure advanced settings

Steps

- Use Assignment Method to specify the parameters used to distribute intercepted traffic among multiple nodes in a cluster. For a description of the WCCP load distribution feature, see WCCP load distribution.
 HASH applies a hash operation to the selected distribution attributes.
 - With HASH, more than one distribution attribute can be selected.
 - The result of the hash operation determines the cluster member that receives the traffic.

MASK applies a mask operation to the selected distribution attribute.

- Only one distribution attribute can be selected, typically the destination IP address.
- The result of the mask operation determines the cluster member that receives the traffic.
- The following distribution attributes can be selected:
- Destination IP address
- Destination Port
- Source IP address
- Source Port

The MASK value is applied up to 6 significant bits (in a cluster, a total of 64 buckets are created). See your WCCP documentation for more information about assignment method HASH and MASK operations. Use the value recommended in the manufacturer's documentation for your device.

2) For proportional load distribution, specify a **weight** value from 0-255. The value determines the proportional distribution of load among servers in a cluster.

Weight is only useful when Synchronize in the Cluster is disabled.

- All cluster members have a value of 0 by default, which results in a balanced distribution of traffic.
- If weight is set to 1 or higher, the value guides proportional distribution among the nodes. For example, if there are 3 nodes in a cluster and Proxy1 has a weight of 20, Proxy2 has a weight of 10, and Proxy3 has a weight of 10, Proxy1 will get one half of the traffic, Proxy2 will get one-quarter of the traffic, and Proxy3 will get one-quarter of the traffic.



Important

When the value of **weight** is greater than 0 on any member of the cluster, any member of the cluster with a weight of 0 receives **no** traffic. If you plan to use weight, be sure to set a weight on every member of the cluster.

For more information about load distribution, see WCCP load distribution.

3) Specify a Reverse Service Group ID for IP spoofing.

When IP spoofing is enabled, you must define a reverse service group for each HTTP and HTTPS forward service group.



Note

Only HTTP and HTTPS are supported for IP Spoofing.

Using the specified ID, Content Gateway creates a reverse service group that is a mirror of the forward service group. For example, if the forward service group has assignment method based on destination IP address, the reverse service has an assignment method based on the source IP address.



Note

IP spoofing is not supported with service groups that use a hashing assignment method with both destination and source attributes. If IP spoofing is enabled on such a service group, an alarm is raised and IP spoofing is disabled.

Related concepts

WCCP load distribution on page 41

Provide router information



Note

It may take up to a minute for the router to report that a new proxy server has joined a service group.

Steps

- To use optional WCCP authentication, under Security, select Enabled and enter the same password used for service group authentication on the router. See Enabling WCCP v2 security on the router.
- To run in multicast mode, under Multicast, select Enabled and enter the multicast IP address. The multicast IP address must match the multicast IP address specified on the router. See *Transparent interception and multicast mode*.



Important

GRE packet Forward/Return method cannot be used with multicast mode.

 Under WCCP Routers, specify up to 10 Router IP Addresses. These routers must be configured with a corresponding service group.

If ASA_Firewall was selected as the **Service Device Profile**, enter both the router IP Address and the WCCP router ID, separated by /, in the **Router IP Address** column.

If GRE is selected for Packet Forward Method, also specify a unique Local GRE Tunnel Endpoint IP address for each router (not required for ASA firewall), and optionally, a GRE Tunnel Next Hop Router IP Address.

The Local GRE Tunnel Endpoint IP address is the Content Gateway tunnel endpoint for the associated Router IP Address.

The Local GRE Tunnel Endpoint IP Address:

- Must be unique and not assigned to any device
- Must be a routable IP address
- Should reside on the same subnet as the proxy. If it is not, you must define a route for it.
- Is not intended to be a client-facing proxy IP address
- Is bound to the physical interface specified for the service group (on Forcepoint appliances use the CLI command "show interface info" to view the logical name to physical interface bindings)

When GRE Packet Return Method is configured and Content Gateway does not have a route back to the WCCP router, specify a **GRE Tunnel Next Hop Router IP Address**. The IP address must be in IPv4 format.

You can use "ping" to test connectivity to the router.

- From Content Gateway, ping each router defined in the service group (in the Router IP Address field).
- If ping doesn't return a response, you need to define a GRE Tunnel Next Hop to that router. Intervening routers must have a route to the WCCP router, or a next hop.

E

Note

WCCP routers that have multiple interfaces assign the Router ID to the interface with the highest numeric value IP address. Content Gateway must be able to connect to the router ID to negotiate the method. To ensure connectivity and that the router ID doesn't change unexpectedly, it is a best practice to make the router loopback address the highest IP address. This also ensures that traffic and statistics reported on the **Monitor > Networking > WCCP** page are reported against a known router ID.

Related tasks

Enabling WCCP v2 security on the router on page 48 Transparent interception and multicast mode on page 55

Save your configuration changes

Steps

- 1) Click Add to add a new entry, or click Set to save changes to the selected entry.
- Click Apply and then Close to close the editor. Navigating away from the page before clicking Apply results in the loss of all changes.

 Restart the proxy to cause the changes to take effect. Navigate to the Configure > My Proxy > Basic > General tab and click Restart.



Note

To check that the router is sending traffic to the proxy, examine the statistics in the Content Gateway manager **Monitor** pane. For example, check that the **Objects Served** statistic in the **My Proxy > Summary** section increases.

Transparent interception and multicast mode

To configure Content Gateway to run in multicast mode, you must enable multicast mode and specify the multicast IP address in the Content Gateway manager.



Important

GRE packet Forward/Return method cannot be used with multicast mode.

In addition, you must set the multicast address on your routers for each service group being intercepted (HTTP, FTP, DNS, and SOCKS). The following procedure provides an example of how to set the multicast address for different service groups on a WCCP v2-enabled router.

Steps

- 1) Telnet to the router and switch to Enable mode.
- 2) At the prompt, enter the following command to configure the router from the terminal: configure terminal
- 3) At the prompt, enter the following command for each service group that the router intercepts: <hostname>(config)# ip wccp <service_group> group-address

<multicast_address>

Here, <hostname> is the hostname of the router you are configuring,

<service_group> is the service group ID (for example, 0 for HTTP), and

<multicast_address> is the IP multicast address.

- At the prompt, enter the following command to configure the network interface: interface <interface_name>
 Here, <interface_name> is the network interface on the router that is being intercepted and redirected.
- 5) At the prompt, enter the following command for each service group that the router intercepts: <hostname>(config-if)# ip wccp <service_group> group- listen
- 6) Exit and save the router configuration.

Transparent interception with policy-based routing

Instead of the WCCP protocol, you can use the policy routing capabilities of a router to send traffic to Content Gateway. WCCP or a Layer 4 switch are generally preferable to this configuration because policy-based routing has a performance impact on the router, and policy-based routing does not support load balancing or heartbeat messaging.

- All client Internet traffic is sent to a router that feeds Content Gateway.
- The router sends port 80 (HTTP) traffic to the proxy and sends the remaining traffic to the next hop router.
- The ARM redirects intercepted requests to Content Gateway.
- Web objects to be served transparently are redirected by the ARM on the return path to the client, so that the documents appear to have come from the origin server.

A Content Gateway cluster with virtual IP failover adds reliability; if one node fails, another node can take up its transparency requests. See *Virtual IP failover*.



Related concepts Virtual IP failover on page 76

Transparent interception with software-based routing

You can deploy Content Gateway without adding routers or switches by using routing software on the Content Gateway node. In this case, Content Gateway is a software router and directs all traffic through the proxy machine. This solution can be useful in low-traffic situations, where the performance cost of using the proxy machine as a router is not high.

On Linux systems, you can use the routed and gated daemons as a software-based routing solution.

- The **routed** daemon is a bundled part of all normal Linux distributions.
- The **gated** daemon is an extensible commercial software package from the Merit GateD Consortium.

When you use routing software with Content Gateway:

- All Internet traffic goes through Content Gateway from machines behind it in the network.
- The routing software routes all non-transparent requests to the Internet.
- The ARM redirects intercepted requests into proxy requests.
- Web objects to be served transparently are redirected by the ARM on the return path to the client, so that the objects appear to have come from the origin server.



Note

Although Content Gateway machines can function as routers, they are not expressly designed to be routers. For reliability, you can use a Content Gateway cluster with the virtual IP failover option. If one node fails, another cluster node takes over. (See *Virtual IP failover*.) The Content Gateway cluster failover mechanism is similar to the Hot Standby Router Protocol (HSRP).

Related concepts

Virtual IP failover on page 76

Configuring Content Gateway to serve only transparent requests

You can configure Content Gateway to serve only transparent requests and prevent explicit proxy requests from being served in the following ways:

- You can control client access to Content Gateway by specifying ranges of IP addresses that are allowed to connect to the proxy. If Content Gateway receives a request from an IP address not listed in a specified range, it discards the request. See Controlling client access to the proxy.
- If you do not know the ranges of client IP addresses allowed to access Content Gateway, you can add rules to the **ipnat.conf** file on the **Configure > Networking > ARM > General** tab in the Content Gateway manager so that only requests that have been redirected by your Layer 4 switch or WCCP router reach the proxy port. To make a transparent-only Content Gateway server, add rules in the **ipnat.conf** file before the normal redirect service rule to redirect explicit proxy traffic to a port on which no service is listening.

For example, if you want Content Gateway to ignore explicit HTTP requests, add rules above the normal HTTP redirect rule in the **ipnat.conf** file as shown below:

rdr hme0 <ipaddress> port 80 -> <ipaddress> port <port_number> tcp

rdr hme0 <ipaddress> port 8080 -> <ipaddress> port

<port_number> tcp

rdr hme0 0.0.0.0/0 port 80 -> <ipaddress> port 8080 tcp

Here, <ipaddress> is the IP address of your Content Gateway system and

<port_number> is a port number on which no service is listening.

Add equivalent rules to the **ipnat.conf** file for each protocol service port or separate network interface to be served. After you make changes to the **ipnat.conf** file, you must restart the proxy.

If your Content Gateway system has multiple network interfaces or if you configure the Content Gateway operating system to use virtual IP addresses, you can give Content Gateway 2 IP addresses. One address must be the real address that the proxy uses to communicate with origin servers and the other a private IP address (for example 10.0.0.1) for WCCP or switch redirection. After you configure the IP addresses, you must add the following variables to the end of the **records.config** file. Replace <private_ipaddress> with the private IP address used for WCCP or switch redirection and real_ipaddress> with the IP address the proxy uses to communicate with origin servers.

LOCAL proxy.local.incoming_ip_to_bind STRING

<private_ipaddress>

LOCAL proxy.local.outgoing_ip_to_bind STRING

<real_ipaddress>

Related tasks

Controlling client access to the proxy on page 150

Interception bypass

A small number of clients and servers do not work correctly with web proxies. Some reasons include:

- Client software irregularities (customized, non-commercial browsers)
- Server software irregularities
- Applications that send non-HTTP traffic over HTTP ports as a way of defeating security restrictions
- Server IP address authentication (the origin server limits access to a few client IP addresses, but the Content Gateway IP address is different, so it cannot get access)
 This is not in frequent use because many ISPs dynamically allocate client IP dial- up addresses, and more secure cryptographic protocols are now more often used.

Web proxies are very common in corporate and Internet use, so interoperability problems are rare. Nonetheless, Content Gateway contains an adaptive learning module that recognizes interoperability problems caused by transparent proxy processing and automatically bypasses the traffic around the proxy server without operator intervention.

Content Gateway follows 2 types of bypass rules:

- Dynamic (also called adaptive) bypass rules are generated dynamically if you configure Content Gateway to bypass when it detects non-HTTP traffic on port 80 or when it encounters certain HTTP errors. See Dynamic bypass rules.
- Static bypass rules must be manually configured in the bypass.config file. See Static bypass rules.



Note

Do not confuse ARM bypass rules with client access control lists. Bypass rules are created in response to interoperability problems. Client access control is simply restriction of the client IP addresses that can access the proxy, as described in *Controlling client access to the proxy*.

Related concepts

Dynamic bypass rules on page 59 Static bypass rules on page 60

Related tasks

Controlling client access to the proxy on page 150

Dynamic bypass rules

The proxy can be configured to watch for the following protocol interoperability errors and configure the ARM to bypass the proxy for the clients and servers causing the errors.

Error code	Description	
N/A	Non-HTTP traffic on port 80	
400	Bad Request	
401	Unauthorized	
403	Forbidden (authentication failed)	
405	Method Not Allowed	
406	Not Acceptable (access)	
408	Request Timeout	
500	Internal Server Error	

In this way, the small number of clients or servers that do not operate correctly through proxies are auto-detected and routed around the proxy so that they can continue to function.

For example:

- When Content Gateway is configured to bypass on authentication failure (403 Forbidden), if any request to an origin server returns a 403 error, Content Gateway generates a destination bypass rule for the origin server's IP address. All requests to that origin server are bypassed until you restart the proxy.
- If the ARM detects that a client is sending a non-HTTP request on port 80 to a particular origin server, Content Gateway generates a source/destination rule. All requests from that particular client to the origin server are bypassed; requests from other clients are not bypassed.

Procedure to enable dynamic bypass rules

Steps

- 1) In the Content Gateway manager, navigate to the **Configure > Networking > ARM > Dynamic Bypass** tab.
- 2) Under Dynamic Bypass, select Enabled.
- 3) Under Behavior, enable each dynamic bypass rule you want to use.
- 4) Click Apply.
- 5) Navigate to the **Configure > My Proxy > Basic > General** tab and click **Restart**.

Next steps

Bypass rules that are generated dynamically are purged after a Content Gateway restart. If you want to preserve dynamically generated rules, you can save a snapshot of the current set of bypass rules. See *Viewing the current set of bypass rules*.

To prevent Content Gateway from bypassing certain IP addresses dynamically, you can set dynamic deny bypass rules in the **bypass.config** file. Deny bypass rules can prevent the proxy from bypassing itself. For information about setting dynamic deny bypass rules, see *bypass.config*.

Content Gateway tallies bypassed requests for each type of dynamic bypass trigger (for example, requests bypassed in response to a 401 error). View these statistics on the **Monitor > Networking > ARM** page of the Content Gateway manager, under **HTTP Bypass Statistics**.

Related concepts

bypass.config on page 370

Related tasks

Viewing the current set of bypass rules on page 61

Static bypass rules

You can configure bypass rules to direct requests from certain clients or to particular origin servers around the proxy. Unlike dynamic bypass rules that are purged when you restart the proxy, these static bypass rules are saved in a configuration file.

You can configure 3 types of static bypass rules:

- Source bypass, in which Content Gateway bypasses a particular source IP address or range of IP addresses.
- Destination bypass, in which Content Gateway bypasses a particular destination IP address or range of IP addresses. For example, these could be origin servers that use IP authentication based on the client's real IP address.
- Source/destination pair bypass, in which Content Gateway bypasses requests that originate from the specified source to the specified destination. For example, you could route around specific client-server pairs that experience broken IP authentication or out of band HTTP traffic problems.
 Source/destination by pass rules might be preferable to destinction rules because they black a destinction.

Source/destination bypass rules might be preferable to destination rules because they block a destination server only for those particular users that experience problems.

To configure static bypass rules, edit the bypass.config file (See bypass.config).



Note

When Content Gateway bypass is enabled on the **Web** > **Settings** > **Scanning** > **Bypass Settings** page of the Forcepoint Security Manager, appropriate rules are added to bypass.config.

Related concepts bypass.config on page 370

Viewing the current set of bypass rules

The ARM has a supporting utility called **netcontrol** that allows you to view the current dynamic and static bypass rules.

To view all current dynamic and static bypass rules:

Steps

- Log on to a Content Gateway node and then change directory to the Content Gateway bin directory (/opt/ WCG/bin).
- 2) Enter the following command at the prompt and press Return:

```
./netcontrol.sh -B
```

All current static and dynamic bypass rules are displayed on screen. The rules are sorted by IP address. You can direct the output of **netcontrol** to a file and save it.

Connection load shedding

The load shedding feature prevents client request overloads. When there are more client connections than the specified limit, the ARM forwards incoming requests directly to the origin server. The default client connection limit is 1 million connections.

Steps

- In the Content Gateway manager, navigate to the Configure > Networking > Connection Management > Load Shedding page.
- In the Maximum Connections field, specify the maximum number of client connections allowed before the ARM starts forwarding requests directly to the origin server.
- 3) Click Apply.
- 4) Navigate to the Configure > My Proxy > Basic > General tab and click Restart.

Reducing DNS lookups

If you are running Content Gateway in transparent proxy mode, you can enable the **Always Query Destination** option to reduce the number of DNS lookups and improve response time. When enabled, the Always Query Destination option configures the proxy to always obtain the original destination IP address of incoming requests from the ARM. Content Gateway then uses that IP address to determine the origin server instead of doing a DNS lookup on the hostname of the request. Because the client already performed a DNS lookup, Content Gateway does not have to.

When Always Query Destination is enabled, the value defined for the variable proxy.config.arm.use_hostname_for_wisp_and_reporting determines whether IP address or hostname is captured for reporting purposes.



Important

It is recommended that you do not enable the Always Query Destination option if Content Gateway is running in both explicit and transparent proxy mode. In explicit proxy mode, the client does not perform a DNS lookup on the hostname of the origin server, so the proxy must perform a DNS lookup.

Also, the category lookup is performed based on the IP address, which is not always as accurate as a URL-based lookup.

In addition, do not enable the Always Query Destination option if you want domain names, rather than IP addresses, in Forcepoint Web Security transaction logs.

Procedure to enable Always Query Destination

Steps

- Navigate to the Content Gateway config directory (/opt/WCG/config) and open the records.config file in a text editor.
- Set the proxy.config.arm.always_query_dest variable to 1. This means that IP addresses are captured; domain names are not.

If you later need to disable Always Query Destination, change this setting to **0**. In this case, domain names are captured.

- 3) Save and close the file.
- 4) To apply the changes, run the following command from the Content Gateway bin directory: content_line -x

Chapter 5 Additional Proxy Configuration

Contents

- Introduction on page 63
- Content Gateway IP spoofing on page 63
- Content Gateway support for IPv6 on page 69

Introduction

Explicit and transparent proxy deployments can be used with:

IP spoofing

Ordinarily, when Content Gateway proxies requests for clients it communicates with origin servers using its own IP address in place of the client's IP address. This is the standard operation of forward proxies.

IP spoofing configures the proxy to use one of the following when communicating with the origin server:

- The IP address of the client (basic IP spoofing)
- A specified IP address (range-based IP spoofing)

For more information, see Content Gateway IP spoofing.

For configuration details, see Configuring IP spoofing.

IPv6 dual-stack networks
 See Content Gateway support for IPv6.

Related concepts

Content Gateway IP spoofing on page 63 Configuring IP spoofing on page 66 Content Gateway support for IPv6 on page 69

Content Gateway IP spoofing

IP spoofing is sometimes used to support upstream activities that require the client IP address or a specific IP address. It also results in origin servers seeing the client or specified IP address instead of the proxy IP address (although the proxy IP address can be a specified IP address; more below).

Content Gateway IP spoofing support has the following features and restrictions:

IP spoofing is supported for HTTP and HTTPS traffic only.

- When IP spoofing is enabled, it is applied to both HTTP and HTTPS. It cannot be configured for only one
 protocol.
- HTTPS traffic is spoofed whether SSL support is enabled or not.
- IP spoofing relies on the ARM.
- In transparent proxy deployments using WCCP and IP spoofing, with GRE or L2 mode negotiation, neither HASH nor MASK are supported on the source port or source port/source IP address.
- IP spoofing is not supported with edge devices such as a Cisco ASA or PIX firewall. When this is attempted, requests made by Content Gateway using the client IP address are looped back to Content Gateway.
- IP spoofing requires all IP addresses in the same routing path use the same format. That is, all IP addresses must be either IPv6 or IPv4. A combination of IPv6 and IPv4 addresses is not supported.



Warning

Deploying IP spoofing requires precise control of the routing paths on your network, overriding the normal routing process for traffic running on TCP port 80 and 443. When configured with either transparent or explicit proxy, return traffic must be routed back to the proxy.

For assistance, please contact your network equipment vendor or Technical Support.

With IP spoofing enabled, traditional debugging tools such as **traceroute** and **ping** have limited utility.



Important

For a discussion of how the proxy kernel routing table impacts transparent proxy deployment, see the Solution Center article titled, Web sites in the Static or Dynamic bypass list fail to connect.

Range-based IP spoofing

Range-based IP spoofing supports groupings of clients (IP addresses and IP address ranges) that are mapped to specified IP addresses.

Among other uses, range-based IP spoofing facilitates:

- The delivery of web-hosted services when the identification is by source IP address. For example, to receive a web-hosted service, an organization might be required to identify membership to the service via a known IP address.
- IP address-based authentication with an external service when a unique IP address represents a group of users.
- A way to configure traditional IP spoofing for some clients (source IP addresses that don't match any group are spoofed with their own IP address), range-based IP spoofing for some clients, and standard proxy IP address substitution for some clients. The latter is done by creating a group that specifies the proxy IP address.



Important

Range-based IP Spoofing is not supported on many older versions of Cisco IOS firmware. To avoid problems, update your Cisco device to the latest firmware.

IP Spoofing is supported for IPv6. However, range-based IP Spoofing is not supported for IPv6.

IP spoofing and the flow of traffic

When IP spoofing is used with WCCP, HTTP and HTTPS traffic flows as follows. The numbers in the diagram correspond to the actions described in the numbered list. (Note that policy-based routing can be implemented to achieve the same results.)



- 1) A client request arrives at a routed port or Switched Virtual Interface (SVI) looking for traffic with a destination port of HTTP (80) or HTTPS (443).
- The switch redirects the client request to Content Gateway. If needed, the proxy creates a connection to the origin server using the client IP address or specified IP address (range-based IP spoofing).
- 3) The request is sent to the origin server through the switch, NAT and/or firewall.
- 4) When the origin server response is returned, the IP packet has the substituted IP address as the destination (client or specified IP address).
- 5) The origin server response arrives at a routed port or Switched Virtual Interface (SVI) looking for traffic with a source port of HTTP (80) or HTTPS (443). See the note below.
- 6) The switch redirects the origin server response to the proxy, completing the proxy- to-origin server TCP connection.
- 7) A proxy response to the client is generated and returned to the client on the proxy-to-client TCP connection.



Note

When IP spoofing is enabled, the proxy advertises a reverse service group for each enabled WCCP service. The reverse service group must be applied along the return path of the proxy.

WCCP service group IDs are user defined and must be programmed on the WCCP devices and in Content Gateway (see *Configuring service groups on the WCCP device* and *Configuring service groups in the Content Gateway manager*).

The following definitions are suggested.

Service ID	Port	Traffic Type
0	destination port 80	HTTP
20	source port 80	HTTP
70	destination port 443	HTTPS (HTTPS support must be enabled)
90	source port 443	HTTPS

Policy-based routing (PBR) uses access control lists (ACL) to identify and redirect flows. In a PBR deployment, all of the configuration is done on the router and there is no corresponding Content Gateway configuration. PBR deployments have to redirect traffic returning from origin servers from port 80 and 443 to Content Gateway.

Related tasks

Configuring service groups on the WCCP device on page 43 Configuring service groups in the Content Gateway manager on page 49

Configuring IP spoofing

- To configure the proxy to use the IP address of the client for IP spoofing, see *Configure basic IP spoofing*.
- To configure the proxy to use a specified IP address for IP spoofing, see Configure range-based IP spoofing.
- For a list of features and restrictions for IP spoofing support, see Content Gateway IP spoofing.

Related concepts

Configure range-based IP spoofing on page 67 Content Gateway IP spoofing on page 63

Related tasks

Configure basic IP spoofing on page 66

Configure basic IP spoofing

In the Content Gateway manager:

Steps

- 1) Go to the Configure > Networking > ARM > General tab.
- 2) Under IP Spoofing, select Enabled.
- 3) Click Apply.
- 4) Go to the Configure > My Proxy > Basic > General tab and click Restart.
- Configure your network to ensure web traffic will be redirected back to the proxy.
 Contact your network equipment vendor or Technical Support for any needed assistance.



Warning

The ARM is a critical component of Content Gateway that should never be disabled. If it is disabled while IP spoofing is enabled, client requests receive a "Cannot display Web page" error and an error message is recorded in the /var/ log/messages directory.

For information about configuring WCCP routers, see Configuring WCCP v2 routers.

Related tasks

Configuring WCCP v2 routers on page 42

Configure range-based IP spoofing



Note

Range-based IP spoofing is not supported for IPv6.

When the proxy is configured to use range-based IP spoofing:

- Client IP address ranges and their corresponding spoofed IP address are specified in a table.
- The table is traversed top-down. The first match is applied.
- Requests from clients that do not match an IP address in the table are spoofed with their own IP address (basic IP spoofing).
- To create an entry that causes a set of IP addresses to appear to be coming from the proxy (as in ordinary forward proxy request handling), specify the desired client IP address range and then use the Spoofed IP Address field to specify the proxy's Internet-facing IP address.
- Create the smallest list that meets your needs. The list is traversed for every connection request. A very large list could contribute to latency. Use the Content Gateway manager performance charts (Monitor > Performance) to monitor proxy performance.

Procedure to create the range-based IP spoofing table

Steps

- 1) Go to the Configure > Networking > ARM > General page.
- 2) Under IP Spoofing, select Enabled. Basic IP spoofing must be enabled to enable range-based IP spoofing.
- 3) Under Range Based IP Spoofing, select Enabled.
- In the Client IP Addresses field, enter a comma separated list of individual IP addresses and/or IP address ranges.
 - In a range, the first IP address is separated from last with a hyphen. For example: 10.100.100.0-10.100.100.254
 - CIDR notation is allowed. Do not use spaces.
 - The Client IP Address list supports a maximum of:
 - 64 IPv4 addresses
 - 32 IPv4 address ranges
- 5) In the Specified IP Address field, enter a single IP address.
- 6) Click Apply to add the entry to the table.



Warning

If any of the formatting is invalid, all of the data in that row is cleared.

- 7) To add a new row to the table, click Add Row.
- 8) To put new entries into effect, click **Apply** and then restart Content Gateway.
- Configure your network to ensure web traffic will be redirected back to the proxy.
 Contact your network equipment vendor or Technical Support for any needed assistance.

Procedure to remove an entry from the IP spoofing table

Steps

- 1) Clear all the values in the row to be removed.
- 2) Click Apply.
- 3) To put the changes into effect, restart Content Gateway.

Content Gateway support for IPv6

Forcepoint security solutions, including Content Gateway, provide support for IPv6.

Important

In transparent proxy deployments, support requires WCCP v2.01. If you use a Cisco router, it must be version 15.4(1) or later.

Content Gateway support for IPv6 includes:

- IPv6 on dual IP stack Ethernet interfaces
- Support for these protocols: HTTP, HTTPS, FTP, DNS
- IPv6 traffic to the Internet, clients, and PAC file servers
- IPv6 virtual IP addresses (vaddrs.config)
- Authentication rules by client IPv6 address ranges
- Client IPv6 addresses and address ranges to allow or restrict access to the proxy (ip_allow.config)
- Client IPv6 addresses and address ranges to allow or restrict access to the Content Gateway manager (mgmt_allow.config)
- IPv6 Primary Destination value and Source IP values in proxy filtering rules (filter.config), and parent proxy servers in a chain (parent.config)
- IPv6 addresses in the SSL Incident List
- SNMP traps and counters for IPv6 data

Limits and restrictions

- IPv6-only internal networks are not supported
- IPv4 must be used to communicate among all Forcepoint components, including other members of a Content Gateway cluster (multicast address)
- With all user authentication, the domain controllers must be reachable on an IPv4 address
- Range-based IP Spoofing is not available for IPv6.
- SOCKS proxy is not supported
- IPv6 support is not available for FTP passive mode with the transparent proxy.
- IPv6 only clients do not display a block page correctly. The user is blocked from the site as expected but will receive a browser error rather than a block page. Dual-stack IPv6 clients receive the normal block page.

IPv6 proxy statistics

Content Gateway tracks IPv6 traffic. View statistics on the Monitor > Networking > System page.

Effect of IPv6 on Event logs

When IPv6 is enabled, Event log entries are normalized to IPv6 format. For example, "10.10.41.200" is logged as "::ffff:10.10.41.200".

To filter on a client at "10.10.41.200" in a custom log, requires the following filter:

```
<LogFilter>
<Name = "IPv6_Test_Machine"/>
<Condition = "chi MATCH ::ffff:10.10.41.200"/>
<Action = "ACCEPT"/>
</LogFilter>
```

IPv6 configuration summary

IPv6 support is disabled by default.

If Content Gateway is deployed on an appliance, first enable IPv6 on the appliance, then enable it for Content Gateway. See your Forcepoint appliance documentation.

To enable IPv6 support:

Steps

- 1) Log on to the Content Gateway manager.
- 2) Navigate to the Configure > My Proxy > Basic > General tab.
- 3) Under Networking, locate the IPv6 row and select On.

Next steps

Once IPv6 support is enabled, in any field that accepts an IPv6 address, the address can be entered in any format that conforms to the standard. For example:

- Leading zeros within a 16-bit value may be omitted
- One group of consecutive zeros may be replaced with a double colon

When IPv6 is disabled, IPv6 entry fields are hidden from view and IPv6 values are deleted from configuration files.

When the **DNS Resolver** is used, go to the **Configure > Networking > DNS Resolver** page to set an IPv4 or IPv6 preference. IPv4 is the default.

Chapter 6 Clusters

Contents

- Introduction on page 71
- Management clustering on page 72
- Changing clustering configuration on page 72
- Adding nodes to a cluster on page 73
- Deleting nodes from a cluster on page 75
- Virtual IP failover on page 76

Introduction

Content Gateway scales from a single node to a cluster of 2 or more nodes, with a maximum recommended limit of 16. This allows you to quickly increase capacity and improve system performance and reliability.



Note

For assistance with scaling your deployment, contact your Forcepoint account representative.

- Content Gateway detects the addition and deletion of nodes in the cluster and can detect when a node is down.
- You can add or delete a node from a cluster at any time.
- When you remove a node from the cluster, Content Gateway removes all references to the missing node.
- Restarting a node in the cluster causes all nodes in the cluster to restart.
- When the Virtual IP failoverfeature is enabled, the live nodes in a cluster can assume a failed node's traffic.
- Nodes in a cluster automatically share configuration information except for the following:
 - Filtering Service and Policy Service IP addresses are not propagated around the cluster.
 - In transparent proxy deployments with WCCP, the service group enabled/ disabled state and weight settings are not propagated. See Transparent interception with WCCP v2 devices.
 - When SSL support is enabled, the Dynamic Incident List is not propagated around the cluster.

Content Gateway uses a proprietary protocol for clustering, which is multicast for node discovery and heartbeat, and unicast for all data exchange within the cluster.



Important

It is recommended that a dedicated network interface be used for Content Gateway cluster communication, **except** when the host is a Forcepoint appliance, in which case the P1 interface is recommended.

In a proxy hierarchy, the nodes in the cluster cannot be a mixture of HTTP parents and children.

Related concepts Virtual IP failover on page 76 Transparent interception with WCCP v2 devices on page 38

Management clustering

In management clustering mode you can administer all Content Gateway nodes at the same time because cluster nodes share configuration information.

- Content Gateway uses a multicast management protocol to maintain a single system image of all nodes in the cluster.
- Information about cluster membership, configuration, and exceptions is shared across all nodes.
- The content_manager process propagates configuration changes to cluster nodes.
- When the HTTPS option is enabled (SSL support), its settings also propagate around the cluster, except for the Dynamic Incident List.

Changing clustering configuration

Clustering is usually configured when you install the proxy. You can, however, configure clustering afterward, or at any time, in the Content Gateway manager.

Steps

- 1) Go to the Configure > My Proxy > Basic > Clustering tab.
- 2) Under Cluster > Type:
 - Select Management Clustering to include this proxy in a cluster.
 - Select Single Node if this node is not part of a cluster.
- Under Interface, enter the name of the network interface. This is the interface used by Content Gateway to communicate with other nodes in the cluster.
 - It is recommended that you use a dedicated secondary interface.
 - Node configuration information is multicast, in plain text, to other Content Gateway nodes on the same subnet. Therefore, as a best practice, clients should be located on a separate subnet from Content Gateway nodes (multicast communications for clustering are not routed).
 - On Forcepoint appliances, P1 is the recommended interface. You may also use P2, however, if you are
 not using it for Internet egress traffic and want to isolate cluster management traffic.
4) In the Cluster Multicast Group Address area, enter the multicast group address that all members of the cluster share (224.0.1.37 by default).



Warning

Ensure that the multicast IP address does not conflict with the address used by any other application or service.

If there is a conflict and the Content Gateway node is allowed to restart, it will fail to initialize the interface and the Content Gateway instance will shut down. You can verify the condition by examining /var/log/ messages and looking for a message similar to:

[LocalManager::initCCom] Unable to find network interface eth2.#011 Exiting

To correct the problem, identify a unique multicast IP address that will work for all members of the cluster and do one of the following:

- If Content Gateway is on an appliance, see the Forcepoint Appliances CLI Guide.
- If Content Gateway is installed on a Linux server:
 - a) Log on to the server and go to /opt/WCG/config.
 - b) Edit (vi) records.config.
 - c) Find proxy.config.cluster.mc_group_addr and assign it the value of the multicast IP address.
 - d) Save and close the file.
 - e) Check each member of the cluster to ensure that they are all using the same multicast IP address.
 - Restart the node.
- 5) Click Apply.
- Select the General tab and click Restart.



Important

Content Gateway does not apply the clustering mode change to all of the nodes in the cluster. You must change the clustering mode on each node individually.

Adding nodes to a cluster

Content Gateway detects new Content Gateway nodes on your network and adds them to the cluster, propagating the latest configuration information to the newcomer. This provides a convenient way to bootstrap new machines.

To connect a node to a Content Gateway cluster, you need only install Content Gateway software on the new node, making sure during the process that the cluster name and port assignments are the same as those of the existing cluster. In this way, Content Gateway automatically recognizes the new node.



Important

The nodes in a cluster must be homogeneous; each node must be on the same hardware platform, each must be on the same operating system version, and Content Gateway must be installed in the same directory (/opt/WCG).

Steps

- 1) Install the appropriate hardware and connect it to your network.
- Install the Content Gateway software using the appropriate procedure for installing a cluster node. See the Forcepoint Web Security Installation Guide.
- 3) During the installation procedure, make sure that the following is true:
 - The cluster name that you assign to the new node is the same as the cluster name for the existing nodes.
 - The port assignments for the new node are the same as the port assignments used by the other nodes.
 - You have added multicast addresses and multicast route settings.
- Restart Content Gateway (/opt/WCG/WCGAdmin restart).

Procedure to add an existing Content Gateway installation to the cluster

- In the Content Gateway manager, go to the Configure > My Proxy > Basic > General tab and set Proxy Name to the name of the cluster.
- 2) Select the Clustering tab.
- 3) Set Interface to the interface used by the cluster. All members must use the same interface.
- 4) Set the Multicast Group Address to the address being used by the cluster.
- 5) In the Type area, select Management Clustering.
- 6) Click Apply.
- 7) Go back to the General tab and click Restart.

Procedure to add a node by editing variable values

You can also add a node by editing variable values in the **record.config** file of the node to be added.

Steps

- 1) On the node you want to add to the cluster, open the records.config file located in /opt/WCG/config.
- 2) Edit the following variables:

Variable	Description
proxy.local.cluster.type	Specify the clustering mode:
	2 = management mode
	3 = no clustering
proxy.config.proxy_name	Specify the name of the Content Gateway cluster. All nodes in a cluster must use the same name.
proxy.config.cluster.mc_group_addr	Specify the multicast address for cluster communications. All nodes in a cluster must use the same multicast address.
proxy.config.cluster.rsport	Specify the reliable service port. The reliable service port is used to send data between the nodes in the cluster. All nodes in a cluster must use the same reliable service port. The default value is 8087.
proxy.config.cluster.mcport	Specify the multicast port. The multicast port is used for node identification. All nodes in a cluster must use the same multicast port. The default port is 8088.
proxy.config.cluster.ethernet_interface	Specify the network interface for cluster traffic. All nodes in a cluster must use the same network interface.

- 3) Save and close the file.
- 4) Restart Content Gateway (/opt/WCG/WCGAdmin restart).

Deleting nodes from a cluster

On the node you want to remove from the cluster:

Steps

1) Log on to the Content Gateway manager and go to the **Configure > My Proxy > Basic > Clustering** tab.

- 2) Under Cluster Type, select Single Node.
- 3) Click Apply.
- If you are permanently removing the node from the cluster, it is a best practice to change the proxy name to a name other than the cluster name.
 Select the **Coneral** to hand change the **Proxy Name** to the system bestname or another meaningful value.
 - Select the General tab and change the Proxy Name to the system hostname or another meaningful value.
- 5) Restart the proxy.

Virtual IP failover

When virtual IP failover is enabled, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes in the cluster as necessary (see *What are virtual IP addresses*? for more information.) These addresses are virtual only in the sense that they are not tied to a specific machine; Content Gateway can assign them to any of its nodes. To the outside world, these virtual IP addresses are the addresses of Content Gateway servers.

Virtual IP failover assures that if a node in the cluster fails, other nodes can assume the failed node's responsibilities. Content Gateway handles virtual IP failover in the following ways:

- The content_manager process maintains cluster communication. Nodes automatically exchange statistics and configuration information through multicast communication. If multicast heartbeats are not received from one of the cluster nodes, the other nodes recognize it as unavailable.
- The content_manager process reassigns the IP addresses of the failed node to the remaining operational nodes within approximately 30 seconds, so that service can continue without interruption.
- The IP addresses are assigned to new network interfaces, and the new assignment is broadcast to the local network. The IP address reassignment is done through a process called **ARP rebinding**.

Related concepts What are virtual IP addresses? on page 77

Enabling or disabling virtual IP addressing

Steps

- 1) In the Content Gateway manager, navigate to the **Configure > My Proxy > Basic > General** tab.
- 2) Under Features > Networking, select On or Off for Virtual IP to enable or disable virtual IP addressing.
- 3) Click Apply.

This enables the Virtual IP page, used to add and edit virtual IP addresses.

Adding or editing virtual IP addresses

Virtual IP addresses must be pre-reserved, like all IP addresses, before they can be assigned to Content Gateway.



Warning

Incorrect IP addressing can disable your system. Make sure you understand how virtual IP addresses work before changing them.

Steps

- 1) In the Content Gateway manager, go to the Configure > Networking > Virtual IP page.
 - The page is available only after you have enabled the Virtual IP option.
 - The Virtual IP Addresses area displays the virtual IP addresses managed by Content Gateway.
- 2) Click Edit File to add new or edit existing virtual IP addresses.
 - To edit a virtual IP address, select it from the table at the top of the page, edit the fields provided, and then click Set.
 - To delete the selected IP address, click Clear Fields.
 - To add a virtual IP address, specify the virtual IP address, the Ethernet interface, and the Subinterface in the fields provided, and then click Add.
- 3) Click Apply, and then Close.

What are virtual IP addresses?

Virtual IP addresses are IP addresses that are not tethered to particular machines. Thus, they can rotate among nodes in a Content Gateway cluster.

It is common for a single machine to represent multiple IP addresses on the same subnet. This machine would have a primary or real IP address bound to its interface card and also serve many more virtual addresses.

You can set up your user base to use a DNS round-robin pointing at virtual IP addresses, as opposed to using the real IP addresses of the Content Gateway machines.

Because virtual IP addresses are not bound to machines, a Content Gateway cluster can take addresses from inactive nodes and distribute those addresses among the remaining live nodes.

Using a proprietary management protocol, Content Gateway nodes communicate their status with their peers. If a node fails, its peers notice the failure and negotiate which of the remaining nodes will mask the fault by taking over the failed node's virtual interface.

Chapter 7 Hierarchical Caching

Contents

- Introduction on page 79
- HTTP cache hierarchies on page 79

Introduction

Content Gateway can participate in HTTP cache hierarchies, in which requests not fulfilled in one cache can be routed to other regional caches, taking advantage of the contents and proximity of nearby caches.

A cache hierarchy consists of levels of caches that communicate with each other. Content Gateway supports several types of cache hierarchies. All cache hierarchies recognize the concept of *parent* and *child*. A parent cache is a cache higher up in the hierarchy, to which the proxy can forward requests. A child cache is a cache for which the proxy is a parent.

For more information, see:

Related concepts HTTP cache hierarchies on page 79

Related tasks

Configuring Content Gateway to use an HTTP parent cache on page 80

HTTP cache hierarchies

In an HTTP cache hierarchy, a Content Gateway node can search a parent cache—which itself can search other caches— before resorting to retrieving the object from the origin server. See *Configuring Content Gateway to use an HTTP parent cache*.

- You can configure a Content Gateway node to use one or more HTTP parent caches, so that if one parent is unavailable, another parent can service requests. This is called parent failover and is described below.
- If you do not want all requests to go to the parent cache, you can configure the proxy to route certain requests directly to the origin server (for example, requests that contain specific URLs) by setting parent proxy rules in the parent.config configuration file (described in *parent.config*).
- If the request is a cache miss on the parent, the parent retrieves the content from the origin server (or from another cache, depending on the parent's configuration). The parent caches the content and then sends a copy to the proxy (its child), (where it is served to the client).

Related concepts

parent.config on page 384

Related tasks

Configuring Content Gateway to use an HTTP parent cache on page 80

Parent failover

When you configure the proxy to use more than one parent cache, the proxy detects when a parent is not available and sends missed requests to another parent cache. If you specify more than two parent caches, the order in which the parent caches are queried depends upon the parent proxy rules configured in the parent configuration file described in *parent.config.* By default, the parent caches are queried in the order in which they are listed in the configuration file.

Related concepts parent.config on page 384

Configuring Content Gateway to use an HTTP parent cache

Steps

- 1) In the Content Gateway manager, navigate to the **Configure > Content Routing > Hierarchies** page.
- 2) The Parent Proxy option is **Enabled** by default.
- 3) Click Edit File to open the configuration file editor for the parent.configfile.
- 4) Enter information in the fields provided, and then click Add. All the fields are described in Hierarchies.
- 5) Click Apply, and then click Close.
- 6) On the **Parenting** tab, click **Apply** to save your configuration.



Important

Perform this procedure on the child proxy. Do not make any changes on the parent.

Related concepts parent.config on page 384

Related reference

Hierarchies on page 295

Chapter 8 Configuring the RAM Cache

Contents

- Introduction on page 83
- Configuring cache object size limit on page 83
- Changing the size of the RAM cache on page 84

Introduction

Content Gateway maintains a small RAM cache of popular objects which help to store, retrieve, and serve web pages and parts of web pages. This RAM cache serves the most popular objects as fast as possible providing bandwidth savings and reduced latency. You can configure the RAM cache size. See *Changing the size of the RAM cache*.

Related tasks

Changing the size of the RAM cache on page 84

Configuring cache object size limit

By default, Content Gateway allows objects of any size in the cache. You can change the default behavior and specify a size limit for objects in the cache.

Steps

- In the Content Gateway manager, go to the Configure > Subsystems > Cache > General Cache Configuration tab.
- In the Maximum Object Size field, enter the maximum size allowed (in bytes) for objects in the cache. Enter 0 (zero) if you do not want to have a size limit.
- 3) Click Apply.

When an object exceeds the size limit, the following message is entered in the system log file.

WARNING: Maximum document size exceeded

Changing the size of the RAM cache

Content Gateway provides a dedicated RAM cache for fast retrieval of popular small objects. You can increase the RAM cache size for better cache hit performance.



Warning

If you increase the size of the RAM cache and observe a decrease in Content Gateway performance (such as increased latencies), the operating system might require more memory for network resources. Return the RAM cache size to its previous value.

- In the Content Gateway manager, go to the Configure > Subsystems > Cache > General Cache Configuration tab.
- In the Ram Cache Size field, enter the amount of space (in bytes) you want to allocate to the RAM cache. Although the user interface will accept larger values, do not exceed 256 MB. The default size is 104857600 (100 MB).
- 3) Click Apply.
- 4) Go to the Configure > My Proxy > Basic > General tab and click Restart.

Chapter 9 DNS Proxy Caching

Contents

- Introduction on page 85
- Configuring DNS proxy caching on page 86

Introduction

Typically, clients send DNS requests to a DNS server to resolve hostnames. However, DNS servers are frequently overloaded or not located close to the client; therefore DNS lookups can be slow and can be a bottleneck to fulfilling requests.

The DNS proxy caching option allows Content Gateway to resolve DNS requests on behalf of clients. This option off-loads remote DNS servers and reduces response times for DNS lookups. See *Configuring DNS proxy caching*.



Important

You can use the DNS proxy caching option only with a layer 4 switch or a Cisco router running WCCP v2.

The following overview illustrates how Content Gateway serves a DNS request.

- 1) A client sends a DNS request. The request is intercepted by a router or L4 switch that is configured to redirect all DNS traffic on port 53 to Content Gateway.
- 2) The ARM examines the DNS packet. If the DNS request is **type A** (answer), the ARM forwards the request to Content Gateway. The ARM forwards all DNS requests that are not **type A** to the DNS server.
- 3) For type A requests, Content Gateway checks its DNS cache to see if it has the hostname to IP address mapping for the DNS request. If the mapping is in the DNS cache, Content Gateway sends the IP address to the client. If the mapping is not in the cache, Content Gateway contacts the DNS server to resolve the hostname. When Content Gateway receives the response from the DNS server, it caches the hostname to IP address mapping and sends the IP address to the client. If round-robin is used, Content Gateway sends the entire list of IP address mappings to the client and the round-robin order is strictly followed.



Note

If the hostname to IP address mapping is not in the DNS cache, Content Gateway contacts the DNS server specified in the **/etc/resolv.conf** file. Only the first entry in resolv.conf is used. This might not be the same DNS server for which the DNS request was originally intended.

The DNS cache is held in memory and backed up on disk. Content Gateway updates the data on disk every 60 seconds. The TTL (time-to-live) is strictly followed with every hostname to IP address mapping.

Related tasks

Configuring DNS proxy caching on page 86

Configuring DNS proxy caching

To configure Content Gateway as a DNS proxy cache:

- Add a remap rule in the **ipnat.conf** file.
- Enable the DNS proxy option and specify the port that Content Gateway will use for DNS proxy traffic.
- Configure your layer 4 switch or WCCP router to send DNS traffic on port 53 to Content Gateway.

You can use the DNS proxy caching option only with a layer 4 switch or a Cisco router running WCCP v2.

In the Content Gateway manager:

Important

- 1) Go to the Configure > Networking > ARM > General tab.
- 2) Under Redirection Rules, click Edit File to open the file editor for the ipnat.conf file.
- 3) Enter the following information:
 - a) Enter the Content Gateway Ethernet Interface to which client DNS requests are routed. For example, eth0.
 - b) In the Connection Type drop-down list, select udp.
 - c) In the Destination IP field, enter 0.0.0.0 to accept DNS requests from all clients.
 - d) (*Optional*) In the Destination CIDR field, enter the CIDR mask value. If you have specified 0.0.0.0 in the Destination IP field, enter **0** here.
 - e) In the Destination Port field, enter the port on which DNS requests are sent to Content Gateway (53, by default).
 - f) In the Redirected Destination IP field, enter the IP address of Content Gateway.
 - g) In the Redirected Destination Port field, enter the port that Content Gateway uses to communicate with the DNS server (5353, by default).
 - h) In the User Protocol drop-down list, select dns.

- 4) Click Add, then click Apply, and then click Close. Postpone the prompted restart until step 8.
- 5) Go to the My Proxy > Basic page.
- 6) Under Features > Networking, enable DNS Proxy and click Apply. Postpone the prompted restart until step 8.
- 7) Go to the **Networking > DNS Proxy** page.
- 8) Enter the **DNS Proxy Port** (5353, by default).
- 9) Click **Apply** and restart Content Gateway.
- **10)** Configure your layer 4 switch or WCCP v2 router to send DNS traffic to the Content Gateway DNS port (53, by default).

Chapter 10 Saving and Restoring Configurations

Contents

- Introduction on page 89
- Taking configuration snapshots on page 89
- Restoring configuration snapshots on page 90
- Deleting configuration snapshots on page 91

Introduction

The configuration snapshot feature lets you save all current configuration settings and restore them if needed. Content Gateway can store configuration snapshots on the node where they are taken, on an FTP server, and on portable media. Content Gateway restores a configuration snapshot on all the nodes in the cluster.



Note

It is recommended that you take a configuration snapshot before performing system maintenance or attempting to tune system performance. Taking a configuration snapshot takes only a few seconds.

This section describes how to perform the following tasks:

- Take a snapshot of the current configuration. See *Taking configuration snapshots*.
- Restore previously taken configuration snapshots. See Restoring configuration snapshots.
- Delete configuration snapshots stored on the Content Gateway node. See Deleting configuration snapshots.

Related concepts

Taking configuration snapshots on page 89 Restoring configuration snapshots on page 90

Related tasks

Deleting configuration snapshots on page 91

Taking configuration snapshots

You can save all of the current Content Gateway configuration settings via the Content Gateway manager.

Save a snapshot on the local system

Steps

- 1) Go to the Configure > My Proxy > Snapshots > File System tab.
- Use the Change Snapshot Directory field to find or edit the local directory for storing configuration snapshots (config/snapshots, by default).
 Relative paths are created in the Content Gateway config directory. To create a snapshot directory in another location, use the full path.
- 3) In the **Save Snapshot** field, type the name you want to use for the current configuration.
- 4) Click Apply.

Save a snapshot on an FTP server

Steps

- 1) Go to the **Configure > Snapshots > FTP Server** tab.
- 2) In the fields provided, enter the FTP server name, login, and password, and the remote directory where the FTP server stores configuration snapshots.
- Click Apply.
 After you have successfully logged on to the FTP server, the FTP Server page displays additional fields.
- 4) In the Save Snapshot to FTP Server field, enter a name for the configuration snapshot.
- 5) Click Apply.

Restoring configuration snapshots

Restore a saved configuration from the Content Gateway manager. If you are running a cluster of Content Gateway servers, the configuration is restored to all the nodes in the cluster.

Restore a configuration snapshot from the local system

Steps

- 1) Go to the Configure > Snapshots > File System tab.
- 2) Use the **Restore/Delete Snapshot** drop-down list to select the configuration snapshot that you want to restore.
- 3) Mark the Restore Snapshot from "directory_name" Directory box.
- Click Apply.
 The Content Gateway system or cluster uses the restored configuration.

Restore a configuration snapshot from an FTP server

Steps

- 1) Go to the **Configure > Snapshots > FTP Server** tab.
- 2) In the fields provided, enter the FTP server name, login, and password, and the remote directory in which the FTP server stores configuration snapshots.
- Click Apply.
 After you have successfully logged on to the FTP server, the FTP Server tab displays additional fields.
- 4) Use the Restore Snapshot drop-down list to select the configuration snapshot that you want to restore.
- Click Apply.
 The Content Gateway system or cluster uses the restored configuration.

Deleting configuration snapshots

- 1) In the Content Gateway manager, go to the Configure > Snapshots > File System tab.
- 2) From the **Restore > Delete a Snapshot** drop-down list, select the configuration snapshot you want to delete.
- 3) Mark the Delete Snapshot from "directory_name" directory box.

4) Click Apply.

The configuration snapshot is deleted.

Chapter 11 Monitoring Traffic

Contents

- Introduction on page 93
- Viewing statistics in the Content Gateway manager on page 94
- Viewing statistics from the command line on page 94
- Working with alarms on page 95
- Using Performance graphs on page 98
- Creating SSL certificate authorities reports on page 99
- Creating an SSL incidents report on page 100

Introduction

Content Gateway provides the following tools to monitor system performance and analyze network traffic:

- Statistics that show Content Gateway performance and network traffic information, available from the Content Gateway manager or the command line. See:
 - Viewing statistics in the Content Gateway manager
 - Viewing statistics from the command line.
- Alarms that signal detected failure conditions. See Working with alarms.
- Performance graphs that show historical Content Gateway performance and network traffic information. See Using Performance graphs.
- Reports for SSL traffic. See Creating SSL certificate authorities reports and Creating an SSL incidents report.

Related concepts

Viewing statistics in the Content Gateway manager on page 94 Working with alarms on page 95

Related tasks

Viewing statistics from the command line on page 94 Using Performance graphs on page 98 Creating SSL certificate authorities reports on page 99 Creating an SSL incidents report on page 100

Viewing statistics in the Content Gateway manager

Use the options on the Monitor tab of the Content Gateway manager to collect and interpret statistics about Content Gateway performance and web traffic.

Statistics are available regarding:

- My Proxy (the current Content Gateway instance, or nodes in the same cluster) See My Proxy and Working with alarms for details.
- Protocols (HTTP and FTP) See Protocols for details.
- Security (LDAP, NTLM, and IWA proxy authentication and SOCKS server connections) See Security for details.
- Subsystems (clustering and logging) See Subsystems for details.
- Networking (general network configuration, ARM, WCCP, DNS, virtual IP addressing, and client connections) See Networking for details.
- Performance See Performance and Using Performance graphs for details.
- SSL See SSL, Creating SSL certificate authorities reports, and Creating an SSL incidents report for details.

Related concepts

Working with alarms on page 95 My Proxy on page 243 Protocols on page 249 Security on page 251 Subsystems on page 255 Networking on page 256 Performance on page 261 SSL on page 263

Related tasks

Using Performance graphs on page 98 Creating SSL certificate authorities reports on page 99 Creating an SSL incidents report on page 100

Viewing statistics from the command line

You can use the command-line interface to view statistics about Content Gateway performance and web traffic.

To view specific information about a Content Gateway node or cluster, specify the variable that corresponds to the desired statistic.

Steps

1) Become root:

su

- 2) Log on to a Content Gateway node.
- 3) From the Content Gateway **bin** directory (/opt/WCG/bin), enter the following command:

```
./content_line -r <variable>
```

Here, *<variable>* is the variable that holds the information you want. For a list of the variables you can specify, see *Content Gateway variables*.

For example, the following command displays the open client connections for the node:

proxy.node.current_client_connections

Related concepts

Content Gateway variables on page 269

Working with alarms

Content Gateway signals an alarm when it detects a problem, for example if the space allocated to event logs is full, or if it cannot write to a configuration file. A general alarm message is displayed at the top of the content pane in the Content Gateway manager.



Not all alarms are critical. Some alarms report transient conditions. For example, a "Content Gateway subscription download failed: error connecting" alarm can be generated by a temporary disruption in Internet connectivity.

Content Gateway alarm messages provides a description of some of the alarm messages that Content Gateway generates.

Use the Monitor > My Proxy > Alarms page to see a listing of current alarms, as shown below.

immary	~	Alarm	! [2 pending]		
arms		Content Gate	way Alarms		
agnostics					
Protocols	~				Clea
Subsystems	~	Current Time: 1	Thu Jul 6 12:05:12 2023		
Networking	~	. Marda			Class
) Performance	~	localhost	[Thu Jul 6 12:00:18 2023] Content Gateway subscription download failed: error connecting	Alarm	Clear
SSL	~	localhost	[Wed Jun 28 11:50:44 2023] NIC virbr0 is disconnected.		



Note

Content Gateway also sends select alarms to the Web module of the Forcepoint Security Manager, where they are referred to as **alerts**. Summary alert messages are displayed on the System tab of the **Web > Status > Dashboard** page. Security Manager administrators can also configure SNMP and email alert notifications for Content Gateway alarms on the **Settings > Alerts** pages.

Related reference Content Gateway alarm messages on page 469

Clearing alarms

After you have addressed an alarm issue, click **Clear** in the alarm message window to dismiss the alarm.



Important

Clicking Clear only dismisses alarm messages; it does not resolve the cause of the alarms.

If the same alarm condition occurs a second time, it will not be logged if the first alarm has not been cleared.

Configuring Content Gateway to email alarm messages

Steps

- 1) In the Content Gateway manager, navigate to the Configure > My Proxy > Basic > General tab.
- 2) In the Alarm Email field, enter the email address to which you want to send alarms. Be sure to use the full mail address including @ notation, for example: receivername@example.com
- 3) Click Apply.

Using a script file for alarms

Alarm messages are built into Content Gateway; you cannot change them. However, you can write a script file to execute certain actions when an alarm is signaled.

A sample script file named **example_alarm_bin.sh** is provided in /opt/WCG/bin. You can modify this file.

Configuring SNMP alerting on Content Gateway (software)

Before configuring SNMP to monitor and report on Content Gateway processes, make sure you have installed Net-SNMP and performed a basic SNMP configuration.

Steps

- 1) Add the process names and MAX/MIN process values to the "Process checks" section of snmpd.conf. You also need to add the v2 trap specification.
- 2) Edit /etc/snmp/snmpd.conf and add the following lines in the "Process checks" area:

```
proc content_cop 1 1
proc content_gateway 1 1
proc content_manager 1 1
proc DownloadService 1 1
proc microdasys 2 1
proc microdasysws 1 1 # send v2 traps
trap2sink IP_address_of_SNMP_Manager:162 informsink IP_address_of_SNMP_Manager: 162 rwuser all
agentSecName all defaultMonitors yes
```

If Filtering Service is also running on the Content Gateway machine and you want to monitor it, add:

proc EIMServer 1 1

To verify that SNMP Agent is sending trap messages

Steps

- 1) On the SNMP Agent/Content Gateway machine, start a network packet analyzer and terminate the DownloadService process.
- In the packet capture data, look for an SNMPv2-Trap message for DownloadService going to the SNMP Manager. The trap message might be similar to:

Value: STRING: Too few DownloadService running (# = 0)

To verify that SNMP Manager is receiving trap messages

Steps

1) On the SNMP Agent/Content Gateway machine, terminate the DownloadService process. Note that it may take several minutes from the time the trap occurs until the trap is sent to the SNMP Manager.

2) On the SNMP Manager machine, check the SNMP trap log for an entry for DownloadService. The name and location of the log file is specified in the snmptrapd startup command (example provided above). Here is one way to find the message if it is being logged in /var/log/messages:

```
cat /var/log/messages | grep DownloadService
```

An entry might look like:

```
Nov 25 15:09:42 localhost snmptrapd[11980]: 10.10.10.10]: Trap,
DISPAN-EV = STRING , DISMAN-EVENT-MIB::mteHotOID = OID ,
DISMAN-EVENT-IB::prErrMessage.4 = STRING: Too few DownloadService
running (# = 0)
```

Grep for "snmptrapd" to see all log entries related to snmptrapd.

Use **nc** (netcat) to test basic UDP connectivity between the Agent and the Manager. For example, this command could be run on either side of the connection to test the designated UDP ports.

[root]# nc -u -v -z -w2 10.228.85.10 161-162

Here, "-u" indicates UPD, "-v" indicates verbose output, "-z" means to scan for listening daemons, and "-w2" indicates to wait 2 seconds before timing out.

Sample results:

10.228.85.10: inverse host lookup failed: Unknown host (UNKNOWN) [10.228.85.10]
161 (snmp) open

Using Performance graphs

The Performance graphing tool (Multi Router Traffic Grapher [MRTG]) allows you to monitor Content Gateway performance and analyze network traffic. Performance graphs show information about virtual memory usage, client connections, and so on. The information provided is recorded from the time that Content Gateway was started. Statistics are gathered at 5-minute intervals.

Use the Monitor > Performance page in the Content Gateway manager to access performance graphs.



Important

To run MRTG, you must have Perl v5.005 or later installed on your Content Gateway system.

- If your Content Gateway node is in a cluster, select the node whose statistics you want to view from the Monitor > My Proxy > Summary page.
- 2) Go to the Performance > Monitor page.
- 3) Select an option:
 - Click **Overview** to see a subset of available graphs.
 - Click Daily to see statistics for the current day.
 - Click Weekly to see statistics for the current week.
 - Click Monthly to see statistics for the current month.
 - Click Yearly to see statistics for the current year.

4) Wait at least 15 minutes after starting Content Gateway before looking at the graphs. It takes several 5minute sample intervals for the tool to initialize statistics.

To configure the tool

If MRTG has not been configured, the system displays a message indicating that it is not available.

Then configure the tool:

Steps

- 1) Make sure Perl 5.005 is installed on your system.
- 2) To ensure that the perl binary is in your PATH, open a command shell, navigate to the bin directory (/opt/ WCG/bin), and enter the following command:

perl ./pathfix.pl `which perl`

3) In the bin directory, use the following command to modify the MRTG update interval: ./update_mrtg;sleep 5;./update_mrtg;sleep 5;

By default, an MRTG update interval is set to 15 minutes. This command sets the update to 5 minutes.

- In the bin directory, start the MRTG cron updates: ./mrtgcron start
- 5) Wait about 15 minutes before accessing the performance graphs from the Content Gateway manager.



Note

To stop MRTG cron updates, use the following command:

./mrtgcron stop

Creating SSL certificate authorities reports

In the Content Gateway manager:

- 1) Navigate to the Monitor > SSL > Reports > Certificate Authorities tab.
- Select the format of the report: HTML or CSV (comma-separated values) If you select CSV, the report is created as an Excel spreadsheet.
- 3) Specify the time period the report will cover. The default is all records in the log.

- 4) Indicate the sort order for the report.
 - List authorities by date
 - List OCSP good responses first
 - List OCSP bad responses first

See Keeping revocation information up to date.

- 5) Click Generate Report. It may take several seconds for the report to be created.
 - HTML output is displayed in the content pane of the browser.
 - CSV output opens in Microsoft Excel, if it is present on the system.



Note

To delete the collected SSL log data, click Reset all collected data.

Related concepts Keeping revocation information up to date on page 134

Creating an SSL incidents report

In the Content Gateway manager:

Steps

- 1) Navigate to the Monitor > SSL > Reports > Incidents tab.
- Select the format of the report: HTML or CSV (comma-separated values) If you select CSV, the report is created as an Excel spreadsheet.
- Specify the time period the report should cover. You can specify a number of days or a date range, or all records since SSL support was enabled.
- 4) Indicate the sort order for the report.
 - List incidents by date
 - List incidents by URL
 - List the number of times each incident occurred See Managing HTTPS website access.
- 5) Click Generate Report. It may take several seconds for the report to be created.



Note

To delete the collected SSL log data, click **Reset all collected data**.

Related concepts

Managing HTTPS website access on page 137

Chapter 12 Working With Web DLP

Contents

- Introduction on page 101
- How the Forcepoint Web Security DLP Module works on page 101
- Deploying Content Gateway to work with Forcepoint DLP on page 102
- Registering Content Gateway with Forcepoint DLP on page 103
- Stopping and starting Forcepoint DLP processes on page 106
- Configuring the ICAP client on page 106

Introduction

When Forcepoint Web Security is deployed with the DLP Module:

- Organizations are protected from data loss over web channels (HTTP, HTTPS, FTP, and FTP over HTTP). A full Forcepoint DLP deployment can extend data loss prevention to include channels such as mobile devices, removable media, and printers.
- Forensics data appears in the Threats dashboard.
- Content Gateway records DLP Module transaction statistics.

To start using the DLP Module:

- 1) Install Forcepoint Web Security with the DLP Module, as described in the Installation Guide.
- 2) Configure Content Gateway to work with DLP Module components.

How the Forcepoint Web Security DLP Module works

When the DLP Module is enabled:

- 1) Content Gateway intercepts outbound content and provides that content to Forcepoint DLP.
- 2) Forcepoint DLP analyzes the content to determine if the web posting or FTP upload is allowed or blocked, based on the Web DLP policy.
 - Transactions over HTTP, HTTPS, FTP, and FTP over HTTP can be examined.
 - The disposition is communicated to the proxy.

- Forcepoint DLP logs the transaction.
- 3) The proxy acts on the Forcepoint DLP determination.
 - If the content is blocked, it is not transmitted to the remote host and Forcepoint DLP returns a block page to the sender.
 - If the content is allowed, it is forwarded to its destination.

Note

When a request is blocked and the DLP server sends a block page in response:

- Content Gateway forwards the block page to the sender in a 403 Forbidden message.
- The block page must be larger than 512 bytes or some browsers will substitute a generic error message.

In addition to applying Web DLP policies, the DLP Module can be used to enable data theft analysis for outbound traffic. Configure outbound security options in the Web Security module of the Forcepoint Security Manager on the **Scanning > Scanning Options** page.

Deploying Content Gateway to work with Forcepoint DLP

Content Gateway supports 2 methods of working with Forcepoint DLP:

- (Preferred) Some components are installed with Content Gateway.
- Over ICAP using Forcepoint DLP components located on a separate host. Only one method can be used at a time.

Forcepoint DLP components on the Content Gateway machine

When Forcepoint Web Security is deployed with the DLP Module or Forcepoint DLP, a small number of Forcepoint DLP components are typically installed on the Content Gateway machine. Content Gateway registers with Forcepoint DLP components when it's first configured and then checks the registration status whenever it's restarted, automatically re-registering if necessary.

For more information about Forcepoint DLP registration, see Registering Content Gateway with Forcepoint DLP.

Related tasks Registering Content Gateway with Forcepoint DLP on page 103

Forcepoint DLP over ICAP

When the Web DLP policy engine is located on a separate host, Content Gateway can communicate with Forcepoint DLP over ICAP v1.0. For configuration details, see *Configuring the ICAP client*. Note that integration with on-box components is the preferred deployment.

Related tasks Configuring the ICAP client on page 106

Registering Content Gateway with Forcepoint DLP

Content Gateway registers with on-box DLP Module components automatically once an administrator enables Web DLP integration.



Note

Automatic registration is not available with Forcepoint DLP Web Content Gateway. See *Registering Content Gateway with Forcepoint DLP manually*.

To enable Web DLP integration:

Steps

- 1) Make sure that:
 - a) The Forcepoint management server is running and accessible.
 - b) The Forcepoint management server includes both Forcepoint Web Security and Forcepoint DLP management components.
 - c) That the system clock on the Forcepoint management server and the Content Gateway machine are synchronized.
- 2) Go to the Configure > My Proxy > Basic > General tab in the Content Gateway manager.
- 3) Set Integration to On, then select the Web DLP (integrated on-box) option.



Note

To later disable the integration and unregister Content Gateway and Forcepoint DLP components, turn the Integration option to **Off** and restart Content Gateway.

4) Restart Content Gateway.

Next steps

Once the integration is enabled, Content Gateway registers with the Forcepoint management server, and Content Gateway queries the Forcepoint Security Manager for the presence of Forcepoint DLP.

Registration is tested and retried, if needed, every time Content Gateway is started. To perform registration, Content Gateway queries the Policy Broker for needed information, including IP address and cluster ID.

- Use the Monitor > Summary page in the Content Gateway manager to view registration status information. Click More Detail, then check the list at the bottom of the Subscription Details section.
- Registration success and failure information is logged in the /opt/WCG/logs/ dss_registration.log file.

If registration succeeds:

- Configure DLP Module integration on the Configure > Security > Web DLP page in the Content Gateway manager. See Web DLP configuration options for Content Gateway.
- Content Gateway uses the Forcepoint DLP policy engine for malware detection.
- Forensic reporting data for the Threats dashboard is collected automatically.
- DLP Module transaction statistics are displayed on the Monitor > Security > Web DLP page in the Content Gateway manager. For a complete list of statistics, see Web DLP.

If registration fails, an alarm displays. If this occurs, make sure that:

- Forcepoint Web Security and Forcepoint DLP management components reside on the same management server.
- The Content Gateway and management server system times are synchronized to within a few minutes.
- The ports used for communication between Forcepoint DLP components and Content Gateway are open in IPTables. See Forcepoint Ports and Configuring IPTables for Content Gateway.
- The server hosting software-based (non-appliance) instances of Content Gateway has an IPv4 address assigned to the eth0 network interface.

After registration, the IP address may move to another network interface on the system, but the IP address must remain available as long as the two modules are registered. The IP address is used for Web DLP policy configuration and deployment.

Related tasks

Registering Content Gateway with Forcepoint DLP manually on page 104 Web DLP configuration options for Content Gateway on page 105

Related reference Web DLP on page 255

Registering Content Gateway with Forcepoint DLP manually

If automatic registration between Content Gateway and Forcepoint DLP fails, administrators can attempt the following manual registration steps:

Steps

- 1) Ensure that the Content Gateway and Forcepoint management server systems are running and accessible, and that their system clocks are synchronized within a few minutes.
- Ensure that Web DLP (integrated on-box) option is enabled on the Configure > My Proxy > Basic > General tab in the Content Gateway manager.
- Next to Web DLP (integrated on-box), click the Not registered link to open the Configure > Security > Web DLP registration screen.
- 4) Enter the IP address of the management server.
- 5) Enter a user name and password for logging onto the Forcepoint Security Manager. The user must be an administrator with Data Security module Deploy Settings privileges.
- 6) Click **Register**. If registration is successful, a message confirms the result and prompts you to restart Content Gateway.

If registration fails, an error message indicates the cause of failure. Correct the problem and perform the registration process again.

Web DLP configuration options for Content Gateway

Once Content Gateway has registered with Forcepoint DLP, use the **Configure > Security > Web DLP** page in the Content Gateway manager to configure the following options:

- 1) If Content Gateway is configured to proxy FTP traffic, select **Analyze FTP Uploads** to send FTP uploads to Forcepoint DLP for analysis and policy enforcement.
- 2) If Content Gateway is configured to proxy HTTPS traffic, select **Analyze HTTPS Content** to send decrypted HTTPS posts to Forcepoint DLP for analysis and policy enforcement.
- 3) Click **Apply** to save your settings.
- 4) Go to the Configure > My Proxy > Basic > General tab and restart Content Gateway.

 Go to the Data Security module of the Forcepoint Security Manager to configure the Content Gateway module. See "Configuring the Web Content Gateway module" in Forcepoint DLP Help.



Note

A Content Gateway manager alarm is generated if:

- Web DLP is enabled but not registered.
- Web DLP is enabled and registered but not configured in the Data Security module of the Forcepoint Security Manager.

Stopping and starting Forcepoint DLP processes

When Content Gateway is registered with Forcepoint DLP and the Forcepoint DLP policy engine is running on the Content Gateway machine, 3 daemon processes are active on the Content Gateway machine:

- **PolicyEngine** handles transaction and data analysis.
- PAFPREP manages the Forcepoint DLP fingerprint repository.
- **mgmtd** handles configuration storage and replication.

These processes start automatically whenever the computer is started. You must have root privileges to stop or start the processes.

To stop or start **all** policy engine processes, on the command line enter:

/opt/websense/PolicyEngine/managePolicyEngine -command [stop|start]

To stop or start individual processes, on the command line enter:

```
service <service_name> [start|stop|restart]
```

Configuring the ICAP client

ICAP can be used with any version of Forcepoint DLP. **The direct interface is recommended**, however, when the policy engine is on the Content Gateway machine. See *Registering Content Gateway with Forcepoint DLP*.

Ę

Note

A secondary ICAP server can be specified as a failover should the primary server fail.

The primary and secondary can also be configured to perform load balancing.

See ICAP failover and load balancing, below.

To configure integration with ICAP:

Steps

1) Go to the **Configure > My Proxy > Basic > General** tab in the Content Gateway manager.

- 2) Under Networking > Integration, change Integration to On, then select ICAP.
- 3) Click Apply, and then click Restart.
- Navigate to the Configure > Networking > ICAP > General tab.
- 5) In the ICAP Service URI field, enter the Uniform Resource Identifier (URI) for the primary ICAP service, followed by a comma (no space) and the URI of the secondary ICAP service. A secondary ICAP service is optional.

Enter the URI in the following format:

icap://<hostname>:<port>/<path>

- <hostname> is the IP address or hostname of the Forcepoint DLP Protector appliance.
- The default ICAP port is 1344.
- *spath* is the path of the ICAP service on the host machine. For example:

icap://protector_app:1344/reqmod

You do not need to specify the port if you are using the default ICAP port.

- 6) Under Analyze HTTPS Content, indicate if decrypted traffic should be sent to Forcepoint DLP for analysis or sent directly to the destination. The HTTPS protocol option must be enabled to send HTTPS traffic to Forcepoint DLP. See Working With Encrypted Data.
- 7) Under Analyze FTP Uploads, select whether to send FTP upload requests to Forcepoint DLP for analysis. The FTP proxy feature must be enabled to send FTP traffic to Forcepoint DLP. See FTP.
- 8) Under Action for Communication Errors, select whether to permit traffic or send a block page if Content Gateway encounters an error while communicating with your data protection solution.
- 9) Under Action for Large Files, select whether to permit traffic or send a block page if a file larger than the size limit specified in your data protection solution is sent. The default size limit for Forcepoint DLP is 50 MB.
- 10) Click Apply.

Note

If you change the URI, you must restart Content Gateway. Other changes do not require a restart.

Related concepts

ICAP failover and load balancing on page 108

Related tasks

Registering Content Gateway with Forcepoint DLP on page 103

Related reference

FTP on page 294

Related information

Working With Encrypted Data on page 111

ICAP failover and load balancing

Content Gateway can be configured to failover to a backup ICAP server if the active ICAP server fails. The proxy detects the failure condition and sends traffic to the secondary server. If the secondary becomes unresponsive, the proxy uses the primary. If no ICAP servers are available, the proxy fails open.

Load balancing between 2 ICAP servers is also an option.

Time to failover

Content Gateway may experience temporary request-processing latency between the time the real failure occurs and the time the proxy marks the failed server as down. After the failed server is marked down, all new requests are sent to the second ICAP server. The time to failover is primarily limited by the connection timeout configuration.

Failure conditions

The following failure conditions lead to failover

- ICAP request failed due to layer-3 failure (twice for the same request)
- Failure to connect to a port within a given timeout
- Failure to send request (server resetting connection, and similar)

Content Gateway does not consider missing, invalid, or slow responses as failures.

Content Gateway does, however, verify that the ICAP server is valid at startup by verifying the response to the ICAP OPTIONS request.

Recovery conditions and actions

After the failed server is marked down, new requests are sent to the second server. No new ICAP requests are sent to the failed server until that server is detected to be active again, based on the recovery conditions below.

Content Gateway tests for the following recovery conditions for each down ICAP server at a specified interval:

- TCP connection success
- Successfully sent OPTIONS request
- Successfully received valid response to OPTIONS request

Upon server recovery (server comes back online and is marked as up):

- Load balancing ON: Requests start being distributed to the newly up server (round-robin)
- Load balancing OFF: If the primary server recovers, all requests start being sent to the primary. If the secondary server recovers, traffic continues to be sent to the primary, until the primary goes down.
Fail open

If all ICAP servers are down, a configuration option allows fail open or fail closed behavior. When all ICAP servers are down, the background thread continuously attempts to reestablish a new connection with each server.

Configuration settings

These ICAP failover parameters are set in the records.config file (defaults shown):

Configuration Variable	Data Type	Default Value	Description
proxy.config.icap. ICAPUri	STRING	(empty)	A comma-separated list of ICAP URIs. For example:
			icap://1.2.3.4:1344/ reqmod, icap://4.3.2.1:1344/ reqmod
proxy.config.icap. ActiveTimeout	INT	5	The read/response timeout in seconds. The activity is considered a failure if the timeout is exceeded.
proxy.config.icap. RetryTime	INT	5	The recovery interval, in seconds, to test whether a down server is back up
proxy.config.icap.	INT	1	Set to:
FailOpen			 1 to allow traffic when the ICAP servers are down
			 0 to send a block page if the ICAP servers are down
proxy.config.icap.	INT	1	Set to:
LoadBalance			 1 to distribute requests to all available servers
			 0 to distribute requests to only the primary server

Related concepts

records.config on page 387

Chapter 13 Working With Encrypted Data

Contents

- Introduction on page 111
- Enabling SSL support on page 113
- Initial SSL configuration tasks on page 114
- Certificates on page 115
- Internal Root CA on page 116
- Backing up your internal Root CA on page 123
- Managing certificates on page 124
- Decryption and Encryption on page 126
- Validating certificates on page 131
- Bypassing verification on page 133
- Keeping revocation information up to date on page 134
- Directing SSL traffic to Content Gateway via explicit proxy on page 135
- Managing HTTPS website access on page 137
- Client certificates on page 142
- Customizing SSL connection failure messages on page 144
- Custom certificate key on page 145
- SSL decryption port mirroring (appliance deployments) on page 146

Introduction

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are the industry standards for secure transmission of data on the Internet. They rely on data encryption and a system of trusted certificates issued by certificate authorities (CA) that are recognized by clients and servers. SSL/TLS requests made in a browser are easily identified by the "https" string that leads the URL.

In the topics that follow, for convenience and simplicity, SSL/TLS is referred to simply as SSL.

To establish an SSL connection, the client sends an SSL connection request to the server. If the server consents, the client and server use a standard handshake to negotiate an SSL connection.

Content Gateway offers 2 types of support for HTTPS traffic. Only one can be used at a time.

Simple connection management in which Content Gateway performs URL filtering and then allows the client to make the connection with the server.



Note

Even when HTTPS support is **not** enabled and HTTPS is not decrypted, Content Gateway performs a URL lookup and applies policy. In these circumstances:

In explicit proxy mode, Content Gateway performs URL filtering based on the hostname in the request. If the site is blocked, Content Gateway serves a block page. Some browsers do not support display of the block page.

To prevent this URL filtering, configure clients to not send HTTPS requests to the proxy.

In transparent proxy mode, if there is an SNI in the request, Content Gateway gets the hostname from the SNI and performs URL filtering based on the hostname. Otherwise, Content Gateway uses the Common Name in the certificate of the destination server. If the Common Name contains a wildcard (*), the lookup is performed on the destination IP address. If the site is blocked, the connection with the client is dropped; no block page is served.

To prevent this URL filtering with WCCP, do not create a service group for HTTPS.

- Advanced connection management in which Content Gateway:
 - Proxies requests
 - Decrypts content and performs real-time content and security analysis
 - Re-encrypts content for delivery to the client or origin server

When advanced connection management (HTTPS support *or* SSL support) is enabled, each HTTPS request consists of two separate sessions:

- One from the client browser to Content Gateway. This is the **inbound** connection.
- Another from Content Gateway to the origin server that will receive the secure data. This is the **outbound** connection.

Different certificates are required for each session.



Enabling SSL support

Steps

- 1) In the Content Gateway manager, go to the Configure > My Proxy > Basic > General tab.
- 2) Under Features > Protocols, set HTTPS to On.

Note

If you are deployed with the DLP Module and it is configured to inspect HTTPS traffic, you must enable HTTPS.

- 3) Click Apply and then Restart.
- 4) Enter the name of the SSL certificate file. See Creating a subordinate certificate authority.
- 5) Go to the Configure > Protocols > HTTPS page.
- 6) Specify the HTTPS Proxy Server Port used for client to Content Gateway connections (8080, by default). If traffic is transparent on 443, a default ARM redirection rule redirects the requests to 8080. See Configure > Networking > ARM > Redirection Rules.

7) To tunnel HTTPS requests when the SSL handshake results in an unknown protocol error, set Tunnel Unknown Protocols to **Enabled**.



Note

By default, Content Gateway does not try to tunnel non-SSL traffic. To change this, update the records.config file (in /opt/WCG/config, by default) as follows:

CONFIG proxy.config.ssl_decryption_bypas s.tunnel_non-ssl_traffic INT 1

Restart Content Gateway to implement the change.

Set the value to **0** to turn off tunneling of non-SSL traffic.



Warning

Tunneled connections are not decrypted or inspected.

When tunneling is enabled, Forcepoint Web Security behavior varies based on the type of proxy deployment.

- When Content Gateway is an explicit proxy, a URL lookup is performed and policy is applied before the SSL connection request is made. Transactions are logged as usual.
- When Content Gateway is a transparent proxy, if there is an SNI in the request, Content Gateway gets the hostname from the SNI and performs URL filtering based on the hostname. Otherwise, when Content Gateway sends the connect to the server, the unknown protocol error causes the request to be tunneled without the proxy being aware of it, and no transaction is logged.

Tunneling of WebSocket traffic over HTTPS (secure mode) is enabled by default.



Note

Client authentication may not work correctly for WebSocket traffic. To avoid the issue, it is recommended that a filtering rule be added for each WebSocket Primary Destination Type so client requests to those destinations are not authenticated. See *Configure > Security > Access Control > Filtering* for instructions.

Related concepts

Creating a subordinate certificate authority on page 118

Related reference

Configure > Security > Access Control > Filtering on page 304

Initial SSL configuration tasks

For inbound (client to Content Gateway) traffic, perform these steps to prepare for supporting HTTPS traffic through Content Gateway:

1) Create an internal root CA (certificate authority). In order to sign SSL traffic, Content Gateway requires an internal SSL Certificate Authority that has the ability to sign SSL certificates. This is for traffic between the browser and Content Gateway. See *Internal Root CA*.

- 2) Add this CA to the certificate tree. Servers, such as destination servers, check this tree to ensure that they can trust users because they have certificates from an authority listed here. The certificates listed on the certificate tree are certificate authorities you empower (trust) to verify the validity of individual websites. Any site signed by a certificate authority in the certificate tree with the "allow" status is allowed through Content Gateway. See *Managing certificates*
- 3) Customize pages that browser users will see. See *Customizing SSL connection failure messages*. Among the pages that can be customized are a connect failure and certificate verification failure page.

Related concepts Internal Root CA on page 116

Related tasks Managing certificates on page 124

Related reference

Customizing SSL connection failure messages on page 144

Certificates

HTTPS security revolves around certificates. A certificate must meet 3 criteria:

- It must be current (not expired or revoked). See Validating certificates.
- It must be issued by a trusted CA (certificate authority). See Managing certificates
- The URL and the certificate owner must match. See Configuring validation settings.

HTTPS connections between the client browser and Content Gateway require a certificate issued by an internal CA. See *Internal Root CA*.

Connections between Content Gateway and the origin server require a certificate signed by one of the certificate signing authorities listed in the Certificate Authority Tree on the **Configure > SSL > Certificates > Certificate Authorities** tab. See *Managing certificates*.

Related concepts

Validating certificates on page 131 Internal Root CA on page 116

Related tasks

Managing certificates on page 124 Configuring validation settings on page 131

Internal Root CA

The internal Root CA dynamically generates all certificates used between the client browser and Content Gateway.

- You must have an internal Root CA to complete an inbound connection.
- Only one internal Root CA can be active at a time.
- The internal Root CA is stored in the SSL configuration database.

Important

The default internal Root CA that is included with Content Gateway is not unique and should not be used in a production environment.

Replace the default internal Root CA with your organization's Root CA or create a new one.

There are three options for creating an internal Root CA:

- Leverage your organization's existing CA and import it into Content Gateway. See Importing your Root CA.
- Create a new Root CA and make that CA available to browsers. See *Creating a new Root CA*.
- Create a subordinate CA that leverages an existing CA, but can also be revoked by that CA. See Creating a subordinate certificate authority.



Important

Back up the existing internal Root CA before importing or creating a new one. This enables you to return to an earlier version, if necessary. See *Backing up your internal Root CA* for details.

Related concepts

Creating a subordinate certificate authority on page 118

Related tasks

Importing your Root CA on page 116 Creating a new Root CA on page 117 Backing up your internal Root CA on page 123

Importing your Root CA

If your organization already has a Root CA, or if you have created a certificate as described elsewhere in this document, you can import it into Content Gateway. The certificate must be trusted by all browsers in your organization.

Be sure to back up any new internal Root CA that you import. See Backing up your internal Root CA for details.

To import your Root CA:

Steps

1) In the Content Gateway manager, go to the **Configure > SSL > Internal Root CA > Import Root CA** tab.

- Click Choose File and browse to select the certificate. The certificate must be in X.509 format and base64encoded.
- Click Choose File and browse to select the private key. It must correspond to the certificate you selected in Step 2.
 - The certificate and private key format must match.
 - The private key format must match the format required by the importing node (unencrypted or encrypted). To verify the certificate and private key format, view the files in a text editor. Use **Backup Root CA** to export the CA from the database.



Note

For information about converting the private key format, see:

- Preparing an Internal Root CA for importing into a FIPS 140-2 enabled node
- Converting an RSA key type to a PKCS#8 key type
- Converting an encrypted private key to an RSA key
- 4) Enter and confirm the **Passphrase**.
- 5) Click Import Root CA. The imported CA is stored in the SSL configuration database.
- 6) Restart Content Gateway.

Related tasks

Backing up your internal Root CA on page 123

Creating a new Root CA

If you do not already have a Root CA, you can use the Content Gateway manager to create one. The process uses **openssl pkcs#8**.

Be sure to back up any new Root CAs that you create. See Backing up your internal Root CA for details.

- 1) In the Content Gateway manager, go to the **Configure > SSL > Internal Root CA > Create Root CA** tab.
- 2) Provide requested information in the fields, particularly noting the following:
 - The fields Organization, Organizational unit, and Common name comprise a distinguished name.
 - For **Organization**, enter the name of your company.
 - Optionally provide an **Organizational Unit** (for example, division, section, or department) name.
 - For **Common Name**, enter the name of your company certificate authority.
 - The comment becomes part of the certificate. The first line you enter can be seen by end users.
 - Enter, and then confirm, the passphrase. (A passphrase is similar to a password. Usually, however, it is longer to provide greater security. It is recommended that you use a strong passphrase, with a combination of numbers, characters, and upper- and lower-case letters.)

3) Click Generate and Deploy Certificate to deploy the certificate to the Content Gateway server.

Related tasks

Backing up your internal Root CA on page 123

Creating a subordinate certificate authority

Creating a subordinate certificate authority (sub CA) enables you to take advantage of all the information already existing for your Root CA. However, the Root CA can revoke the sub CA at any time.

Follow these steps to generate a sub CA using OpenSSL and the certificate services in Microsoft Windows.

Preparation

- If you are not the Enterprise domain administrator, you will need to work with that person to get the correct domain permissions to generate a sub CA.
- Install the OpenSSL toolkit (www.openssl.org) on a Windows or Linux machine.

Creating a Certificate Signing Request (CSR)

- 1) Log on to the Windows or Linux machines with root or Administrator permissions.
- 2) Open a Command Prompt or command shell.

3) Enter the following openssl command:

```
openssl req - sha256 -new -newkey rsa:2048 -keyout wcg.key - out wcg.csr
[root@ux81 ~]# openssl req -new -newkey rsa:2048 -keyout wcg.key -out wcg.csr
Generating a 2048 bit RSA private key
. . . . . . . . . . . . . +++
writing new private key to 'wcg.key'
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
Country Name (2 letter code) [XX]:US
State or Province Name (full name) []:Texas
Locality Name (eg, city) [Default City]:Austin
Organization Name (eg, company) [Default Company Ltd]:Forcepoint LLC
Organizational Unit Name (eg, section) []:IT
Common Name (eg, your name or your server's hostname) []:test.example.com
Email Address []:
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
[root@ux81 ~]#
```

4) There will be a series of questions. Answer each question and make note of the challenge password; it will be needed later in the process.

The openssl command generates 2 files:

- wcg.csr is the CSR that will be signed by the Certificate Authority to create the final certificate.
- wcg.key is the private key.
- If you created the CSR on a Linux system, copy it to your Windows host with WinSCP or some other file transfer utility.

Signing the request

To use Microsoft Certificate Services to sign the request:

Steps

 Open wcg.csr with WordPad (to preserve the formatting) and copy the contents onto the clipboard (Edit > Select all; Edit > Copy).



 In Internet Explorer, enter the following URL to go to the Microsoft CA server: http://<CA_server_IP_address>/certsrv/

The Certificate Services applet starts.



3) Under Select a task, click Request a certificate.



4) On the Request a Certificate page, click the link to submit an advanced certificate request.



5) On the Advanced Certificate Request screen, select the Submit a certificate request by using a base-64encoded CMC... link.

lo submit a sav renewal reques	/ed request to the CA, paste a base-64-encoded CMC or PKCS #10 certificate request or PKCS #7 It generated by an external source (such as a Web server) in the Saved Request box.
Saved Request:	
Zana 64 annadad	bd13geZX/EQ9zFHKqePKRU153bs32H3v0UJwVe2N
ertificate request	isdfdfhEQjHrhPgydnOoJsCwmMNJRgjy7d8ez2r6
CMC or	9eILRYus13zXqeWg3kU=
PKCS #10 or	END CERTIFICATE REQUEST
PKCS #7):	
	Browse for a file to insert.
Certificate Temp	late:
	Subordinate Certification Authority
Additional Attrib	utes:
Attributes:	<u>v</u>
	2

6) On the Submit a Certificate Request or Renewal Request screen, paste the content of the **wcg.csr** file (previously placed on the clipboard) in the field provided and click **Submit**.

7) The certificate is issued and the Certificate Issued screen displays.

If, instead, the Certificate Pending screen displays, you do not have sufficient privileges to create a sub CA. Contact your Enterprise domain administrator to complete the certificate creation process before proceeding.

Certificate Issued	•
The certificate you requested was issued to you.	
⊙ DER encoded or ○ Base 64 encoded	
Download certificate Download certificate chain	

- 8) Select the **Base 64 encoded** radio button, and then select **Download certificate**.
- 9) Save the certificate to your desktop. Later you will import it into Content Gateway. With the base 64 encoded certificate on your desktop, along with the private key created during the CSR generating process, you are ready to import both into Content Gateway. See *Importing your Root CA* for instructions.

Related tasks Importing your Root CA on page 116

Backing up your internal Root CA

Always back up the public and private keys of your internal Root CAs before importing or creating new ones. This enables you to return to an earlier version of the certificate, if necessary. In addition, back up any new Root CAs that you import or create.

- 1) In the Content Gateway manager, go to Configure > SSL > Internal Root CA > Backup Root CA tab.
- 2) Click Save Public CA Key to view or save the public CA key.
- 3) Click Save Private CA Key to view or save the private CA key. Depending on your browser settings, you may be prompted to open or save each key file, or the keys may automatically be saved to the browser's default downloads directory.

Managing certificates

Content Gateway initially populates its trusted certificate store, the Certificate Authority Tree (CA tree) with the list qualified by Mozilla for Firefox (see mozilla.org), by Microsoft for Internet Explorer, and by Apple for Safari. The CA tree appears on the **Configure > SSL > Certificates > Certificate Authorities** tab in the Content Gateway manager. Content Gateway trusts origin servers that offer these certificates.

In the CA tree, a small "i" appears before the names of certificates that can be validated via certificate revocation lists (CRL) or online certification status protocol (OCSP). Content Gateway checks the revocation status of certificates used for both inbound and outbound traffic. See *Keeping revocation information up to date* for information about checking the revocation status of a certificate.

To view, delete, or change the allow/deny status of a certificate:

Steps

- 1) In the Content Gateway manager, go to the **Configure > SSL > Certificates > Certificate Authorities** tab.
- 2) Select the name of an authority to open a small pop-up window with information about that authority.
- 3) Do one of the following:
 - To open or download the certificate for review, select Click to view certificate. Depending on your browser settings, you may be prompted to open or save the certificate file, or the file may automatically be saved to the browser's default downloads directory.
 - To delete a certificate, select Click to delete certificate, then confirm your choice. After deleting the certificate, verify that it no longer appears on the Certificate Authorities tab.
 - To allow or deny the certificate, select the Click to change status to option. Depending on the status of the certificate, your choice is allow or deny.
 - If you change the status to deny, a red X appears next to the name of the certificate authority in the certificate authority tree.
 - If you change the status to allow, a green circle appears next to the name of the certificate authority.

Related concepts

Keeping revocation information up to date on page 134

Adding new certificate authorities

Use the **Configure > SSL > Certificates > Add Root CA** tab to manually import additional certificate authorities. Certificates that you import manually have a default status of **allow**.



Important

Back up your current certificates before making any changes, such as adding or deleting certificates. See *Backing up certificates*. If you want to back up your entire Content Gateway configuration, see *Saving and Restoring Configurations*.

Steps

- 1) Browse to the certificate location. Look for files that have a ".cer" extension. The certificate must be in X.509 format and base64-encoded.
- 2) Click Add Certificate Authority.
- If the import was successful, check that the new certificate is listed on Configure > SSL > Certificates > Certificate Authorities.

New CAs are also added when users visit a site signed by that authority. These certificates may be allowed or denied.

Related tasks

Backing up certificates on page 125

Related information

Saving and Restoring Configurations on page 89

Backing up certificates

As a precaution, it is recommended that you back up the database containing the CA certificates whenever you make changes, such as adding or deleting a certificate. They can then be restored at a later date.

Steps

- 1) In the Content Gateway manager, go to the Configure > SSL > Certificates > Backup Certificates tab.
- Click Back Up Configuration to Database.
 To back up your entire Content Gateway configuration, see *Saving and RestoringConfigurations*.

Related information

Saving and Restoring Configurations on page 89

Restoring certificates

To restore saved certificate configuration information:

- 1) In the Content Gateway manager, go to the **Configure > SSL > Certificates > Restore Certificates** tab.
- Browse to the location of the backup certificate database.

 Click Restore. You receive a message telling you that the restore was successful and indicating where the previous certificate database was backed up.

The certificate database is propagated around the cluster.

If you are running multiple proxies, use this restore feature to ensure that all the proxies have the same configuration.

Automatic certificate updates

The information in the CA tree is automatically updated on a regular basis as well as each time Content Gateway is restarted. Updating the CA tree avoids the potential for using a root CA that has expired, is no longer a root CA, or if the certificate revocation list URL of the root CA has changed.

The update process inserts new trusted CAs and updates existing CAs that have updated certificate revocation lists, and at the same time removes expired CAs, any CA that is no longer a root CA, and non-trusted CAs.



Note

The update process maintains only Public certificates. Customers are responsible for maintaining Private certificates.

Enabled by default, the feature can be disabled by editing records.config using this command:

CONFIG proxy.config.ssl.catree_update INT 0

Restart Content Gateway after making this change. Reset the value to 1 to re-enable the updates.

To avoid file corruption, checks are in place to confirm the availability and health of each new update. Update attempts that fail generate an informational alarm. The existing set of certificates continues to be used until the next successful download.

This feature:

- Requires SSL decryption to be enabled.
- Does not check existing certificate revocation lists during the update process.
- Does not re-add CAs explicitly removed by a customer.
- When an update is in progress, provides a warning on the Configure > SSL > Certificates pages that changes made when the update is running are lost. The same message appears when a backup or restore is attempted

Decryption and Encryption

Use the **Configure > SSL > Decryption / Encryption** page in the Content Gateway manager to configure SSL and TLS settings and ciphers for inbound and outbound traffic.

For instructions, see:

Related tasks

SSL configuration settings for inbound traffic on page 127 SSL configuration settings for outbound traffic on page 129

SSL configuration settings for inbound traffic

To configure SSL and TLS settings and ciphers for inbound traffic:

Steps

- 1) In the Content Gateway manager, go to the **Configure > SSL > Decryption / Encryption > Inbound** tab.
- 2) Under Protocol Settings, mark the check box next to each protocol that you want Content Gateway to support. Supported protocols are:
 - TLSv1
 - TLSv1.1
 - TLSv1.2 (enabled by default)
 - TLSv1.3 (enabled by default)

Select the protocols that your organization's security policy has adopted and that your browsers support.

- You must select at least one protocol.
- These settings override the settings for these protocols in the users' browsers.
- You can select different protocols for outbound traffic.

- 3) Under Cipher Settings, select the appropriate Cipherlist for your deployment. The cipher list describes available algorithms and level of encryption between the client and Content Gateway. The Content Gateway DEFAULT cipher list matches the OpenSSL Default list, excluding those that Forcepoint experts believe provide the least security or encryption strength.
 - ADH
 - RC4
 - EXP
 - DES

Edit the variables defined in the records.config file to change the default list. See SSL Decryption.

The strongest cipher (providing the highest level of encryption) is applied first. This can be set to a different level of encryption than for outbound traffic.

Additional cipher settings are:

- HIGH encryption cipher suites are those with key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
- MEDIUM encryption cipher suites are the high cipher list plus additional cipher suites that use 128-bit encryption algorithms.
- **CUSTOM** allows the use of personalized cipher suites.

For inbound requests (clients connections to Content Gateway), consider using MEDIUM encryption to improve performance.

Similar to the OpenSSL string, in TLS1.3, the following five cipher suites are enabled for all settings except **CUSTOM**:

- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_GCM_SHA256
- TLS_AES_128_CCM_8_SHA256
- TLS_AES_128_CCM_SHA256

Regardless of the selected setting, specific insecure ciphers are disabled by default. Control this list via the **proxy.config.ssl.server.cipherlist_suffix** variable in the records.config file. See the information provided in the *SSL Decryption* section of *Content Gateway Configuration Files* for more information.

You can set the cipher lists to **CUSTOM** to specify your own cipher suites. If you choose this option, two additional input boxes will be available: one for TLS 1.2 and earlier cipher suites and another one for TLS 1.3 cipher suites. Provide your preferred cipher suites according to the formats recommended in the OpenSSL documentation.

For more information about ciphers, refer to www.openssl.org/docs.

- 4) Click Apply.
- 5) Go to the **Configure > My Proxy > Basic > General** tab and click **Restart**.

Related reference SSL Decryption on page 445

Related information

Content Gateway Configuration Files on page 363

SSL configuration settings for outbound traffic

Use **Configure > SSL > Decryption / Encryption > Outbound** to configure SSL and TLS settings and ciphers for outbound traffic (Content Gateway to the origin server).

Steps

- 1) Under **Protocol Settings**, indicate which protocols you want Content Gateway to support. Supported protocols are:
 - TLSv1
 - TLSv1.1
 - TLSv1.2 (enabled by default)
 - TLSv1.3 (enabled by default)

Select the protocols that your organization's security policy has adopted.

- You must select at least one protocol.
- You can select different protocols for inbound traffic.

 Under Cipher Settings, select the appropriate Cipherlist for your deployment. The cipher list describes available algorithms and level of encryption between the client and Content Gateway.

The Content Gateway **DEFAULT** cipher list matches the OpenSSL Default list, excluding those that Forcepoint experts believe provide the least security or encryption strength.

The strongest cipher (providing the highest level of encryption) is applied first. This can be set to a different level of encryption than for inbound traffic.

Additional cipher settings are:

- HIGH encryption cipher suites are those with key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
- MEDIUM encryption cipher suites include the high cipher list plus additional cipher suites that use 128-bit encryption algorithms.
- **CUSTOM** allows the use of personalized cipher suites.

For outbound requests, consider using HIGH to improve security.

Similar to the OpenSSL string, in TLS1.3, the following five cipher suites are enabled for all settings except **CUSTOM**:

- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_GCM_SHA256
- TLS_AES_128_CCM_8_SHA256
- TLS_AES_128_CCM_SHA256

Note that regardless of the selected setting, specific insecure ciphers are disabled by default. Control this list via the **proxy.config.ssl.client.cipherlist_suffix** variable in the records.config file. See the information provided in the *SSL Decryption* section of *Content Gateway Configuration Files* for more information.

You can set the cipher lists to **CUSTOM** to specify your own cipher suites. If you choose this option, two additional input boxes will be available: one for TLS 1.2 and earlier cipher suites and another one for TLS 1.3 cipher suites. Provide your preferred cipher suites according to the formats recommended in the OpenSSL documentation.

For more information about ciphers and cipher lists, refer to www.openssl.org/docs.

- 3) Click Apply.
- 4) Go to the Configure > My Proxy > Basic > General tab and click Restart.

Related reference

SSL Decryption on page 445

Related information

Content Gateway Configuration Files on page 363

Validating certificates

SSL certificate verification is an important component of SSL security. Through certificate exchange and verification, the client (Content Gateway) and the origin server verify that each is who it says it is.

Content Gateway performs this task with the certificate verification engine (CVE).

- Use the tabs on the Configure > My Proxy > SSL > Validation page to enable and configure the CVE.
- For information about options when verification fails and you prefer to trust the site, see Bypassing verification.
- For a comprehensive discussion of the use and best practices of the CVE, see SSL Certificate Verification Engine.

Related tasks

Bypassing verification on page 133

Configuring validation settings

Steps

- 1) In the Content Gateway manager, go to the **Configure > SSL > Validation > General** tab.
- 2) If it is not already selected, mark the Enable the certificate verification engine check box.
 - Certificate verification is enabled by default.
 - If this option is not selected, certificate validation does not occur.
- 3) Indicate whether or not to Deny certificates where the common name does not match the URL. When this option is selected, 2 checks are made:
 - The certificate's Common Name is checked for an exact match of the destination URL.
 - If the first check fails, the certificate's Subject Alternative Name (SAN) list is checked for an exact match of the destination URL.

Checks are case insensitive.

Because an exact match is required, there may be instances when a legitimate variation in the Common Name, or the absence of a matching variation in the SAN, may result in a block.

For example, using "https://cia.gov" to access "https://www.cia.gov" may result in a block. Additionally, a block may occur when users attempt to access a site by IP address.

4) If you have enabled the Deny certificates option, indicate whether or not to Allow wildcard certificates. When selected, this option allows matches with Common Names that include the "*" (wildcard) character in the name.

Some HTTPS servers use a wildcard in the Common Name so that a single certificate can cover an entire domain. For example, "*.example.com" could cover "email.example.com" and "stream.example.com", among others.

- Use of the wildcard means that individual servers within the domain are not verified; they are included as a result of the wildcard.
- Allowing wildcard certificates eases the strict matching burden when a Common Name match is required. It is also helpful for domains that have multiple subdomains like google.com or yahoo.com. It also introduces some risk that a fraudulent or undesirable variation of a domain may go unblocked.
- 5) Select the No expired or not yet valid certificates option to deny access to sites that offer an expired or not yet valid certificate. This is a basic check that is important because many malicious sites operate with expired certificates.

If this option is not selected, access to those sites is permitted.

- 6) Indicate whether or not to **Deny self-signed certificates**. By default, the option is enabled, and self-signed certificates (certificates without an official certificate authority) are considered invalid.
- 7) Indicate whether or not to Verify entire certificate chain. By default, this option is enabled, and Content Gateway verifies expiration and revocation status of all certificates between the site certificate and the root Certificate Authority as specified in the certification path of the certificate. This is an important check.
- 8) Indicate whether or not to Check certificate revocation by CRL. Certificate revocation lists (CRLs) are used to check a certificate's revocation status. CRLs list certificates that have been issued and subsequently revoked by the CA.

Verifying the revocation status is a basic check that is very important because certificates are revoked when they are improperly issued, have been compromised, have a false identity, or violate policies specified by the CA.

- If this option is enabled, verify that the daily CRL update feature is enabled on the Revocation Settings tab under CRL Settings.
- If this option is not used, disable the daily CRL update feature on the Revocation Settings tab under CRL Settings.
- 9) Indicate whether or not to Check certificate revocation by OCSP. Online Certificate Status Protocol (OCSP) is an alternate way to check a certificate's revocation status. While OCSP is beneficial, it is not used as widely as CRLs and therefore is not as reliable. Also, it is a real-time, Internet-hosted check that can introduce some request handling latency.



Note

It is recommended that you use OCSP in addition to, rather than instead of, CRLs. See *Keeping revocation information up to date* for more information about CRLs and OCSP.

- **10)** If you are using OCSP revocation checking, use the **Block certificates with Unknown OCSP state** option to determine whether to block certificates that return the "Unknown" status.
- 11) If both CRL and OCSP revocation checking are enabled, indicate your **Preferred method for revocation** check. The selected method (CRL, by default), is applied first.

- 12) If you have enabled CRL or OCSP checking (or both), use the Block certificates with no CRL URI and with no OCSP URI option to block certificates that do not have the expected, associated URIs. For example, if only CRL checking is enabled and the certificate doesn't have a CRL URI, if this option is enabled the connection is blocked. When both CRL and OCSP checking are enabled, the block occurs only if both CRL and OCSP lack a URI.
 - You can view URI information in the certificate when you select to view the certificate in your browser. See *Managing certificates* for details.
 - Because many certificates do not include CRL or OCSP information, this option can result in a high number of verification failures. Often the failures are reported as "Unknown revocation state" errors. This can result in a highly restrictive security policy, with many access denials.
 - As with all verification failures, you can allow for exceptions using the Incident List. See *Managing HTTPS website access*.

Related concepts

Keeping revocation information up to date on page 134 Managing HTTPS website access on page 137

Related tasks

Managing certificates on page 124

Bypassing verification

When verification bypass is enabled, users are allowed to access a website after they have been informed that the site has an invalid certificate.

It is recommended that organizations deploy initially with verification bypass enabled. Then, as the incident rate changes, administrators can use the Incident List to enforce policy. See *Managing HTTPS website access*.

Use the **Configure > SSL > Validation > Verification Bypass** tab in the Content Gateway manager to configure verification bypass settings.

- Select Permit users to visit sites with certificate failure after confirmation to enable verification bypass (default). If this check box is not selected, users do not have the option to browse to sites with an invalid certificate.
- 2) If verification bypass is enabled, use the Time before the user is notified again for the site field to specify a period of time, in minutes, that the user is allowed to visit a particular site without having to click through the warning again. The default is 6 minutes.
- 3) Select **Enable the SSL session cache for bypassed certificates** to store information about bypassed certificates in cache and reuse the connections.
 - If this option is selected, not all users are notified that they are trying to access a site where verification has failed.
 - If this option is not selected, all users are notified about sites that do not have valid certificates.

4) Click Apply.

Related concepts

Managing HTTPS website access on page 137

Keeping revocation information up to date

As a best practice, configure Content Gateway to check the status of any certificate before accepting it, to ensure that the certificate has not been revoked. There are 2 methods of doing this: through CRLs (see *Certificate revocation lists*) and through OCSP (see *Online certification status protocol*).

- CRLs may include information about thousands of certificates, and may therefore take some time to download and process.
- OCSP operates on a request/response basis for individual certificates, which may improve performance, but not all CAs provide OCSP responses.

Related concepts

Certificate revocation lists on page 134

Related tasks Online certification status protocol on page 135

Certificate revocation lists

Use the **Configure > SSL > Validation > Revocation Settings** tab to configure how Content Gateway keeps revocation information current, and to perform an immediate CRL update when needed.

By default, Content Gateway performs CRL downloads on a daily basis.

To configure a time for daily CRL downloads

- 1) Select Download the CRL at, then select a time.
- 2) Click Apply.

To perform an immediate CRL update

Steps

1) Click Update CRL Now to initiate the CRL download.



Note

Downloading CRL files can take some time and consume CPU resources. Download CRL updates at a time when Internet traffic on your system is light.

 Because the update process may take some time, click View CRL Update Progress to see the status of the update.

For more information about certificate revocation lists, see RFC 3280.

Online certification status protocol

With OCSP, when a site wants to verify the revocation status of a certificate, it sends a request to the CA about the status of the certificate. The CA then responds, confirming the validity (or revocation) of the certificate.

Because not all CAs provide responses, CRLs can provide information about the status of more certificates.

Content Gateway enables you to cache OCSP responses about the revocation state of a certificate. Caching responses may be useful in environments with high amounts of SSL traffic and where saving bandwidth is important.

Use the **Configure > SSL > Validation > Revocation Settings** tab to configure how Content Gateway keeps revocation information current.

Steps

- 1) Specify, in days, how long OCSP data should be cached. If you do not want to cache OCSP data, enter **0**. The maximum is 1000 days.
- 2) Click Apply.

For more information about OCSP, see RFC 2560.

Directing SSL traffic to Content Gateway via explicit proxy

Use an existing PAC file or create a new one to direct HTTPS traffic to Content Gateway.

Step 5, below, provides a script that can be used a basis for building a custom PAC file.

To configure Content Gateway to serve a PAC file:

Steps

- 1) In the Content Gateway manager, go to the **Configure > My Proxy > Basic > General** tab.
- Under Features > Protocols, make sure that, ensure that HTTPS is On. If HTTPS is disabled, set it to On, click Apply, and then click Restart.
- 3) Go to the Configure > Content Routing > Browser Auto-Config > PAC tab.
- 4) Specify an Auto-Configuration Port for the proxy to use to serve the PAC file (8083, by default).
- 5) Use the PAC Settings area to review or create the PAC file:
 - If an administrator has copied an existing PAC file into the Content Gateway config directory (as described in Using a PAC file), the contents of the file are displayed. Review and update the file as needed.
 - If no PAC file has been configured, the PAC Settings field is empty. To start creating a PAC file, copy and paste the following template into the PAC Settings field. Replace <host> with the IP address or hostname of the Content Gateway machine.

```
function FindProxyForURL(url, host)
{
    url = url.toLowerCase(); host = host.toLowerCase();
    if(url.substring(0, 5) == "http:"){ return "PROXY <host>:8080";
    else if(url.substring(0, 4) == "ftp:"){ return "PROXY <host>:2121";
    else if(url.substring(0, 6) == "https:"){ return "PROXY <host>:8080";
    else{
    return "DIRECT";
    }
}
```

The template is for basic testing only. Administrators should modify the file as needed to suit their organization's needs.

- 6) Click Apply.
- 7) Go to the Configure > My Proxy > Basic > General tab and click Restart.

Once the new PAC file is in place, configure users' browsers to use the PAC file. For example, if the PAC file is located on the proxy server with the hostname "proxy1" and Content Gateway uses the default port 8083 to serve the file, users' browsers must be configured to include the following URL in their proxy configuration settings:

http://proxy1.company.com:8083/proxy.pac

The procedures for specifying the PAC file location vary among browsers. See *Using a PAC file* for more information.

Related tasks

Using a PAC file on page 28

Managing HTTPS website access

Use the **Configure > SSL > Incident List** and **Add Website** tabs to manage access to websites and troubleshoot website access issues.

- When an end user receives an access denial message because a website does not comply with the organization's security policy, Content Gateway generates an incident. See *Viewing incidents*.
- To specify how Content Gateway treats a particular site, add it to the Incident List. See Adding websites to the Incident List.

Additional troubleshooting information can be found in SSL Certificate Verification Engine.

Related concepts

Viewing incidents on page 137

Related tasks Adding websites to the Incident List on page 141

Viewing incidents

Use the **Configure > SSL > Incidents > Incident List** tab to see a report of those times when clients received an access denial message.

Every node in a cluster has its own incident list.

- Incidents that are added or modified by the administrator are copied around the cluster (synchronized).
- Unexpected incidents that result in an access denial message are not synchronized in the cluster.

Use the fields in this report to specify how Content Gateway treats requested access to a site in the future.

- To view a specific incident in the local list, enter the ID number or URL and click Search Node. If the node is part of a cluster and you want see all instances of the ID or URL, in all lists, click Search Cluster.
- After performing a search, to restore the complete local list, click Show All in Node. When the list is very large, Show All displays only the first 2,500-3,000 records. Use the scroll bar to scroll through the list. Use the ">" and "<" buttons to view the next or previous page.</p>

The incident report

To sort on any column, click the small triangle next to the column heading. The incident report contains these fields:

Field	Description
Node	The name of the Content Gateway node on which the list entry is located.

Field	Description	
ID	The incident ID number assigned by the system, also called the Ticket ID. Help Desk can ask the user for the Ticket ID in the error message and quickly retrieve it from the Incident List.	
	The end user sees the Ticket ID and a denial message.	
Status	Determines how Content Gateway will treat this website in the future. Four conditions are possible:	
	 Allow Users can access the site even if the certificate is not valid. Traffic is decrypted, and certificate checking is disabled. 	
	 Blacklisted The site is completely blocked. Users cannot access this site even if the Verification Bypass is configured. 	
	 Block If certificate verification fails, access to the website is blocked, unless Verification Bypass is configured, in which case the block page includes a "Visit site anyway" button. See <i>Bypassing verification</i>. 	
	Tunnel The site is tunneled. Traffic is not decrypted and Content Gateway does not check the certificate. Tunneling can be used to bypass inspection of trusted sites and improve performance.	
	Note	
	Tunnel by URL does not work with all transparent proxy traffic. See <i>Adding</i> websites to the Incident List.	
	Use the drop-down list in the Action column to change the status of a site.	
Туре	Indicates whether the site was added based on its URL or its certificate. It is recommended that administrators add sites to the Incident List by certificate. See <i>Adding websites to the Incident List</i> .	
URL	The URL of a site whose certificate could not be validated.	
Message	Offers the option to edit the error message. See <i>Changing the text of a message</i> for information about customizing error messages. The pencil and the magnifying glass icon are both links. See <i>Viewing incident details</i> for details.	
Action	Offers the option to change the status of an incident and to delete the incident. See <i>Deleting an incident</i> .	

Related tasks

Bypassing verification on page 133 Adding websites to the Incident List on page 141 Changing the text of a message on page 140 Viewing incident details on page 140 Deleting an incident on page 139

Changing the status of an incident

When an administrator changes the status of an incident, that changes how Content Gateway treats the listed URL in the future.

Steps

- 1) In the Content Gateway manager, go to the Configure > SSL > Incidents > Incident List tab.
- 2) Identify the incident to update.
- 3) Select one of the following from the drop-down list in the Actions column. (See *The incident report* for an explanation of these options.)
 - Tunnel
 - Block
 - Blacklist
 - Allow
- 4) Click OK. The icon in the Status column changes to reflect the new status.

Related reference The incident report on page 137

Deleting an incident

- 1) In the Content Gateway manager, go to the Configure > SSL > Incidents > Incident List tab.
- 2) Select the incident to delete. If the incident is not visible, you can search by ID. See Viewing incidents.
- 3) In the Action column, select Delete from the Action drop-down list, and then click OK. If it is necessary or convenient, the entire Incident List can be deleted using a sqlite3 command: sqlite3 /opt/WCG/config/new_scip3.db "delete from certificate_acl;"

Related concepts

Viewing incidents on page 137

Changing the text of a message

Steps

- 1) In the Content Gateway manager, go to the Configure > SSL > Incidents > Incident List tab.
- 2) Identify the incident to update. See Viewing incidents.
- 3) Click the pencil icon to open a text editor window, then update the message. For example, an administrator could add more detail to an error message.
- 4) Click Submit to save the changes, or Close Window to close the text editor without saving.

Related concepts Viewing incidents on page 137

Viewing incident details

Steps

- 1) In the Content Gateway manager, go to the Configure > SSL > Incidents > Incident List tab.
- 2) Locate an incident to examine more closely. See *Viewing incidents*.
- 3) Click the magnifying glass icon to see additional details about the incident, such as:
 - The **Description** text that appears in the incident message.
 - The time the incident was **Created**.
 - The time the incident was **modified**.
 - The number of Access attempts (how many times users have attempted to access this site).

Related concepts Viewing incidents on page 137

Adding websites to the Incident List

Use the **Configure > SSL > Incidents > Add Website** tab to specify sites that you want to allow, blacklist, or tunnel. Sites that are added manually are assigned chronological Ticket IDs. These appear on the Incident List. See *Viewing incidents*.

Steps

1) Enter the URL of the site to add to the Incident List.



Note

When specifying an IPv6 address, enclose the address in square brackets ([]).

- 2) Select either By Certificate or By URL.
 - By Certificate provides greater security. When a site is added by certificate:
 - Clients cannot bypass the policy by using the IP address rather than the URL.
 - Content Gateway retrieves the server certificate and adds the site to the Incident List.
 If sites are blocked by certificates, wildcard certificates are not accepted, even if the common name is recognized.
 - Select By URL to tunnel, allow, or blacklist the site.
- 3) In the Action drop-down list, specify if the site should be added with Tunnel, Allow, or Blacklist status.
 - Tunnel: (Valid for By URL only) The site is tunneled. Traffic is not decrypted and Content Gateway does not check the certificate.



Important

Tunnel by URL does not work for all transparent proxy requests.

It works under these conditions:

- When the client application uses TLS and includes an SNI (server name indication), Content Gateway checks the Incident list for the hostname in the SNI.
- When there is no SNI, Content Gateway connects to the origin server to retrieve the certificate. If the Common Name is a unique FQDN, Content Gateway looks it up in the Incident list. If the Common Name contains a "*" (wildcard), or is not a unique FQDN, Content Gateway looks for the IP address in the Incident list.

Alternatively, use ARM Static bypass rules.

- Allow: Users can access the site even if the certificate is not valid. Traffic is decrypted, and certificate checking is disabled.
- Blacklist: The site is completely blocked. Users cannot access this site even if the Verification Bypass is configured.
- 4) Click Apply.

As a best practice, administrators should manually add sites to the Incident List after monitoring network traffic for a period of time with the CVE disabled. (See *Configuring validation settings*.) This enables administrators to improve performance by tunneling trusted sites and blocking those they know should not be accessed.

Related concepts Viewing incidents on page 137 Static bypass rules on page 60

Related tasks

Configuring validation settings on page 131

Client certificates

For security, the destination server may request a client certificate.

Responding to client certificate requests

Use the **Configure > SSL > Client Certificates > General** tab in the Content Gateway manager to configure how Content Gateway responds when the server requests a client certificate:

Steps

- 1) Under Action When Client Certificate Is Created:
 - Select Tunnel to always permit the request and provide the client certificate to the server.
 - Select Create incident to specify how Content Gateway should handle that certificate and site. This is the only way to specify a disposition other than tunnel. See *The incident report* for a list of possible dispositions.
- 2) Click Apply.

Related reference The incident report on page 137

Importing client certificates

Use the **Configure > SSL > Client Certificates > Import** tab in the Content Gateway Manager to import certificates from the organization represented by the client.

Note that a network administrator may need to provide the key and passphrase information needed to complete this configuration.



Important

Use only X.509-formatted, base64-encoded certificates.

Steps

- 1) Enter the name of the client certificate.
- 2) Browse to the public key for the certificate.
- 3) Browse to the private key for the certificate.
- 4) Enter, and then confirm, the passphrase. Use a strong passphrase, with a combination of numbers, characters, and upper- and lower-case letters.
- 5) Click Import.

When a client certificate is always required: the Hostlist

Use the **Configure > SSL > Client Certificates > Hostlist** tab in the Content Gateway manager to list destination servers that always require a client certificate.

Be sure to import the certificate before adding it to the Hostlist (see Importing client certificates).

Steps

- 1) Enter the IP address or hostname of the destination server that requires the client certificate.
- In the Client Certificate drop-down list, select the name of the client certificate. Only certificates you have already imported appear in this list.
- 3) Click Add.



Important

For browsers that don't send a Server Name Indicator (SNI), such as Internet Explorer version 8 and earlier, create an entry for both the destination IP address and the hostname.

Related tasks

Importing client certificates on page 142

Deleting client certificates

Use **Configure > SSL > Client Certificates > Manage Certificates** tab in the Content Gateway manager to delete imported client certificates.

Steps

- 1) Select the certificate you want to delete.
- 2) Click Delete.

Customizing SSL connection failure messages

Administrators can use the tabs of the **Configure > SSL > Customization** page in the Content Gateway manager as follows:

- **Certificate Failure**: customize the message users receive when they are trying to connect to a site that has an invalid certificate.
- **Connect Error**: customize the message users receive when Content Gateway is unable to connect to the destination web server.

The following variables may optionally be included in the message templates.

%P	Protocol (HTTP or HTTPS)
%о	The IP address of the host of the proxy that generated the message
%Н	Remote hostname of the request
%t	Time
%s	Name of the Content Gateway server
%u	Complete URL
\$\$DETAILS	Detailed error message
\$\$TICKETID	The ID number of the incident.

To customize the message

- 1) Select the appropriate tab (Certificate Failure or Connect Error).
- 2) Edit the HTML code in the window as needed.
3) Click Preview to see the changes.

Note

There is a known problem in Internet Explorer 10 that sometimes results in the wrong block page being displayed in the Preview pane. To work around the problem, click **Preview** repeatedly until the correct page is displayed, or disable TLS 1.0.

- 4) Repeat steps 1 and 2 until all changes have been made.
- 5) Click Apply to save the changes or Cancel to return to the original message.

Custom certificate key

Use the **Configure > SSL > Custom Certificate Key** options to specify your own base64-encoded key pair that Content Gateway will use, in place of a predefined pair, to generate the dynamic certificate that is sent to clients.

In environments with multiple clusters using the same load balancer, it is best to use the same custom certificate key. Use the backup option from in one cluster, and import the same key to other clusters. Optionally, import the same key to each cluster.



Note

The key format must be PKCS#8. See this article for information about converting to a PKCS#8 key type.

Importing a custom certificate key

If you have created a custom certificate key pair and a new passphrase, open the **Import Custom Certificate Key** tab provided from **Configure > SSL > Custom Certificate Key**.

Steps

- Use Choose File to browse and locate the key to be imported. The key must be in PKCS#8 format, be base-64 encoded, and 2048 or 4096 bits in length.
- Enter a Passphrase and then Confirm Passphrase. The passphrase and confirm passphrase entries must match and be the passphrase used when the key was created.
- 3) Click Import Custom Certificate Key.

Creating a custom certificate key

If you do not have a customer certificate key pair, use the options on the **Configure > SSL > Custom Certificate Key > Create Custom Certificate Key** tab.

Steps

- Enter the Key length.
 Valid vales are 2048 or 4096.
- Enter a Passphrase and then Confirm Passphrase.
 Enter a phrase that is 4-40 characters in length. The passphrase and confirm passphrase entries must match.
- Click Generate and Deploy Custom Certificate Key.
 Be sure to click this option only once. The process may take several seconds.
 Content Gateway must be restarted in order to use this new customer key.

Backing up a custom certificate key

Always back up your custom certificate key before importing or creating a new one. This allows you to return to an older key, if necessary.

To back up the current key:

Steps

- 1) Navigate to Configure > SSL > Custom Certificate Key > Back up Custom Certificate Key.
- 2) Click Save Custom Certificate Key to create a backup of the current custom key.

SSL decryption port mirroring (appliance deployments)

The Content Gateway proxy can be configured to decrypt HTTPS traffic for analysis. Port mirroring delivers all decrypted HTTPS traffic to a physical network interface. This allows a trusted service device to inspect and analyze the decrypted data for its own purpose. The trusted device, however, cannot modify the decrypted traffic and inject it back into the data stream.

SSL decryption port mirroring is available only when the proxy is hosted on a Forcepoint appliance. The feature can be enabled and configured using CLI commands.



Important

The mirror port interface should not be connected to a live network.

This feature is supported:

- If SSL decryption is enabled
- Using one of the interfaces on the Content Gateway appliance
- For both IPv4 and IPv6
- For both transparent and explicit proxy deployments

Only decrypted HTTPS traffic is delivered to the mirrored interface. The following SSL traffic is not delivered:

- Traffic that is set to bypass decryption
- Blocked traffic
- Tunneled traffic

See the Forcepoint Appliances CLI Guide for information about configuring port mirroring.

Chapter 14 Content Gateway Security

Contents

- Introduction on page 149
- Controlling client access to the proxy on page 150
- Controlling access to the Content Gateway manager on page 151
- FIPS 140-2 Mode on page 154
- Content Gateway filtering rules on page 155
- Configuring SOCKS firewall integration on page 160
- Setting SOCKS server bypass on page 163
- Using the Split DNS option on page 163
- Content Gateway user authentication on page 164

Introduction

Content Gateway allows administrators to establish secure communication between the proxy and other computers on the network. Administrators can:

- Control which clients are allowed to access the proxy. See Controlling client access to the proxy.
- Control access to the Content Gateway manager using:
 - Administrator accounts (see Setting the administrator ID and password and Creating a list of user accounts).
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access (see Using SSL for secure administration).
- Create filtering rules to control access to the Internet, specify special authentication requirements, and control other traffic transiting the proxy. See Content Gateway filtering rules.
- Configure Content Gateway integration into your firewall and control traffic through one or more SOCKS servers. See Configuring SOCKS firewall integration.
- Configure Content Gateway to use multiple DNS servers to match your site's security configuration. See Using the Split DNS option.
- Configure Content Gateway to perform user authentication. The proxy supports Integrated Windows Authentication (with Kerberos), legacy NTLM (NTLMSSP), LDAP, and RADIUS user authentication. There is also support for multiple authentication methods with multiple authentication realms. See Content Gateway user authentication.

Related concepts

Content Gateway filtering rules on page 155 Configuring SOCKS firewall integration on page 160 Content Gateway user authentication on page 164

Related tasks

Controlling client access to the proxy on page 150 Setting the administrator ID and password on page 151 Creating a list of user accounts on page 152 Using SSL for secure administration on page 153 Using the Split DNS option on page 163

Controlling client access to the proxy

Administrators can configure Content Gateway to allow only certain clients to use the proxy.

- When this configuration is in place, only clients whose IP address is included in the ip_allow.config file can access the proxy.

To restrict client access to the proxy:

Steps

- 1) In the Content Gateway manager, go to the **Configure > Security > Connection Control** page.
- 2) Click Edit File to open the configuration file editor for the ip_allow.config file.
- 3) To add a new rule:
 - a) Use the **IP Action** drop-down list to indicate whether to allow (ip_allow) or deny (ip_deny) requests from the specified IP address or range.
 - b) Enter a Source IP address or range.
 - c) Click Add.
- 4) To edit an existing rule:
 - a) Select a rule in the list at the top of the page.
 - b) Make changes as needed.
 - c) Click Set.
- 5) Click Apply to save the information, and then click Close.



Note

If an unauthorized client tries to access Content Gateway, a message is displayed in their browser, indicating that the requested content cannot be obtained.

Controlling access to the Content Gateway manager

Administrators can restrict access to the Content Gateway manager to ensure that only authenticated users can change configuration options and view performance and network traffic statistics.

Administrators can:

- Set the master administrator ID and password. A user who logs on to the Content Gateway manager with the administrator ID has access to all Content Gateway manager activities. See Setting the administrator ID and password.
- Create and maintain a list of user accounts that determines who can log on to the Content Gateway manager and which activities they can perform. See *Creating a list of user accounts*.
- Create an access control list of IP addresses that defines which machines can access the Content Gateway manager. See Controlling host access to the Content Gateway manager.
- Use SSL for secure administration. See Using SSL for secure administration.
- Require administrators to log on to the Forcepoint Security Manager, with or without two-factor authentication, and then use the Content Gateway Access page in the Security Manager to log on to the Content Gateway manager. See Accessing the Content Gateway manager.

Related concepts

Accessing the Content Gateway manager on page 16

Related tasks

Setting the administrator ID and password on page 151 Creating a list of user accounts on page 152 Controlling host access to the Content Gateway manager on page 153 Using SSL for secure administration on page 153

Setting the administrator ID and password

The administrator who installs Content Gateway sets a password that controls administrative access to the Content Gateway manager. A user who logs on to the Content Gateway manager using the correct ID and password can view all the statistics on the Monitor tab and change any configuration options on the Configure tab.

To change the administrator ID and password in the Content Gateway manager:

Steps

- 1) Navigate to the Configure > My Proxy > UI Setup > Login tab.
- 2) To change the current administrator ID, under Administrator > Login, type a new ID.

- 3) To change the current password, type the current password in the Old Password field. Type the new password in the New Password field, and then retype the new password in the New Password (Retype) field. Passwords must be 8 to 15 characters and include at least one:
 - Uppercase character
 - Lowercase character
 - Number
 - Special character

Supported characters include:

! # % & ' () * + , - . / ; < = > ? @ [] ^ { | }~

The following special characters are not supported:

Space \$:`\"

If you have forgotten the current administrator password, see Accessing the Content Gateway manager if you forget the master administrator password.

4) Click Apply.

Related tasks

Accessing the Content Gateway manager if you forget the master administrator password on page 19

Creating a list of user accounts

If a single administrator ID and password for the Content Gateway manager is not sufficient, an administrator can create a list of user accounts that define who has access to the Content Gateway manager and which activities they can perform.

Steps

- 1) In the Content Gateway manager, go to the Configure > My Proxy > UI Setup > Login tab.
- 2) Under Add New User, enter the name of the user allowed to access the Content Gateway manager.
- 3) Enter the password for the user, and then enter the password again in the New Password (Retype) field. Passwords must be 8 to 15 characters and include at least one:
 - Uppercase character
 - Lowercase character
 - Number
 - Special character

Supported characters include:

! # % & ' () * + , - . / ; < = > ? @ [] ^ { | }~

The following special characters are not supported:

```
Space $ : ` \ "
```

- Click Apply. 4)
- 5) In the Access drop-down list of the user table, select which Content Gateway manager activities the user can perform:
 - Select No Access to disable Content Gateway manager access for the user.
 - Select **Monitor Only** to allow the user to view statistics from the Monitor tab only.
 - Select Monitor and View Configuration to allow the user to view statistics from the Monitor tab and to view configuration options from the Configure tab.
 - Select Monitor and Modify Configuration to allow the user to view statistics from the Monitor tab and to change configuration options from the Configure tab.
- 6) Click Apply.
- Repeat this procedure for each user allowed to access the Content Gateway manager. 7)

Controlling host access to the Content Gateway manager

In addition to using an administrator ID and user accounts, it is possible to control which hosts have access to the Content Gateway manager.

Steps

- In the Content Gateway manager, go to the Configure > My Proxy > UI Setup > Access page. 1)
- In the Access Control area, click Edit File to open the configuration file editor for the mgmt_allow.config 2) file.
- Enter information in the fields provided, and then click Add. All the fields are described in UI Setup. 3)
- Click Apply, and then click Close. 4)

Related reference UI Setup on page 281

Using SSL for secure administration

Forcepoint Web Security uses the Secure Sockets Layer protocol (SSL) to protect administrator communication with the Content Gateway manager. SSL security provides authentication of both ends of a network connection using certificates, and provides privacy using encryption.

Administrators can optionally replace the Forcepoint-provided certificate with a custom certificate.

To do this:

Steps

- Obtain an SSL certificate from a recognized certificate authority (for example, VeriSign) or, if you use Active Directory Certificate Services, generate a certificate using Certificate Services and a script provided with your Content Gateway software. (See Creating an SSL Certificate for Content Gateway manager with Active Directory Certificate Services).
- Install the certificate in the Content Gateway config directory (/opt/WCG/bin).
 Either rename the certificate to the default filename (private_key.pem), or specify the name of the certificate in the Content Gateway manager.
- 3) If you have used a name other than the default, log on to the Content Gateway manager and navigate to the Configure > My Proxy > UI Setup > General tab. The HTTPS option is enabled by default.
- 4) In the Certificate File field, specify the filename of the SSL certificate.
- 5) Click Apply.

FIPS 140-2 Mode

FIPS (Federal Information Processing Standard) 140-2 is a U.S. government security standard for hardware and software cryptography modules. Modules validated against the standard assure government and other users that the cryptography in the system meets the standard.

The cryptographic libraries used in Forcepoint Web Security, including the Content Gateway component, have passed FIPS 140-2 Level 1 validation. For details on the cryptographic modules used and their FIPS certificates, see FIPS 140-2 with Forcepoint Appliances and Web Security.

By default, Content Gateway does not operate in FIPS 140-2 mode. Content Gateway still uses the FIPS-validated libraries, but it also allows cryptographic algorithms that are not supported by the FIPS 140-2 standard.

Administrators can configure Content Gateway to enforce FIPS 140-2 on HTTPS connections.

When FIPS is enabled:

- HTTPS connections use TLSv1.1 and TLSv1.2.
- HTTPS connections use FIPS 140-2 approved algorithms
- Content Gateway generates SHA-256 certificates in response to origin server certificate requests



Warning

Once the FIPS 140-2 option is enabled, it cannot be disabled without completely reinstalling Content Gateway. If Content Gateway is on an appliance, the appliance must be reimaged.

Important

FIPS 140-2 is not used for:

- Traffic that flows through the cloud (Hybrid Module).
- Traffic forwarded to Forcepoint Advanced Malware Detection.
- Forcepoint Mobile Security.
- IWA fallback to NTLM or Legacy NTLM user authentication.

To enable FIPS 140-2 on HTTPS connections

Steps

- 1) In the Content Gateway manager go to the Configure > Security > FIPS Security page.
- 2) Review the warning.
- 3) Select Enabled, then click Apply.
- 4) To enable FIPS, restart Content Gateway. Otherwise, select Disable and click Apply.



Important

After FIPS is enabled, you must re-install any hotfixes previously installed for the current version of Content Gateway.

Content Gateway filtering rules

Content Gateway supports the ability to create rules that inspect requests for certain parameters and, when matched, apply a specified action. Rules can be created to:

- Deny or allow URL requests
- Insert custom headers
- Allow specified applications, or requests to specified websites to bypass user authentication
- Keep or strip header information from client requests
- Prevent specified applications from transiting the proxy



Note

To create rules for IWA, NTLM, and LDAP user authentication, see *Rule-Based Authentication*. To get started with Content Gateway user authentication options, see *Content Gateway user authentication*.

Use the **Configure > Security > Access Control > Filtering** tab to create and modify filtering rules. Rules are stored in the **filter.config** file.

- Rules are applied in the order listed, top to bottom. Only the first match is applied. If no rule matches, the
 request proceeds.
- Secondary specifiers are optional. More than one secondary specifier can be used in a rule. You cannot, however, repeat a secondary specifier.
- Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 Forcepoint Advanced Malware Detection destinations.
- When Authentication bypass is enabled on the Web > Settings > Scanning > Bypass Settings page of the Forcepoint Security Manager, appropriate rules are added to filter.config.

After adding, deleting, or modifying a rule, restart Content Gateway. See *filter.config* for information about the structure of stored rules.

Related concepts

Content Gateway user authentication on page 164 Rule-Based Authentication on page 186 filter.config on page 372

Creating filtering rules

Steps

- 1) In the Content Gateway manager, go to the **Configure > Security > Access Control > Filtering** tab.
- 2) Click Edit File to open *filter.config*in the file editor.
- 3) Select a **Rule Type** from the drop down list. The Rule Type specifies the action the rule will apply. The supported options are:

allow: allows particular URL requests to bypass authentication; the proxy serves the requested content.

deny: denies requests for objects from specific destinations. When a request is denied, the client receives an access denied message.

keep_hdr: specifies which client request header information to keep.

strip_hdr: specifies which client request header information to strip.

add_hdr: causes a custom header-value pair to be inserted. Requires that **Custom Header** and **Header Value** are specified. Provides support for destination hosts that require a specific header-value pair. For an example, see *Creating an add_hdr rule to allow Google enterprise gmail* below.



Note

The "radius" rule type is **not** supported.

4) Select a Primary Destination Type and then enter a corresponding value in the Primary Destination Value field. Primary Destination Types include: dest_domain: a requested domain name. The value is a domain name.

dest_host: a requested hostname. The value is a hostname.

dest_ip: a requested IP address. The value is an IP address.

url_regex: a regular expression to be found in a URL. The value is a regular expression.

- 5) If the Primary Destination Type is **keep_hdr** or **strip_hdr**, select the type of information to keep or strip from the **Header Type** drop down list. Options include:
 - date
 - host
 - cookie
 - client_ip
- 6) If the rule applies to only inbound traffic on a specific port, enter a value for **Proxy Port**.
- 7) If the rule type is add_hdr, specify the Custom Header and Header Value. The Custom Header and Header Value must be values that the destination host expects. See the example for Google Business Gmail below.
- 8) Provide values for any required or desired **Secondary Specifiers**. They include:

Time: specifies a time range, such as 08:00-14:00. **Prefix**: specifies a prefix in the path part of a URL. **Suffix**: specifies a file suffix in the URL.

Source IP address: specifies a single client IP address, or an IP address range of clients.

Port: specifies the port in a requested URL.

Method: specifies a request URL method:

- get
- post
- put
- trace

Scheme: specifies the protocol of a requested URL. Options are:

- HTTP
- HTTPS
- FTP (for FTP over HTTP only)

User-Agent: specifies a request header User-Agent value. This is a regular expression (regex).

You can use the User-Agent field to create application filtering rules that:

- Allow applications that don't properly handle authentication challenges to bypass authentication
- Block particular client-based applications from accessing the Internet

See the knowledge base article titled "When authentication prevents devices, browsers, and custom applications from working with the proxy" for more information and several examples.

9) When you have finished defining the rule, click Add to add the rule and then Apply to save the rule.

10) When you are done adding rules, click Apply to save all the changes and then click Close to close the edit window.

Related concepts

filter.config on page 372

Related tasks

Creating an add_hdr rule to allow Google enterprise gmail on page 158

Editing a rule

Steps

- 1) In the Content Gateway manager, go to the Configure > Security > Access Control > Filtering tab.
- 2) Click Edit File to open *filter.config* in the file editor.
- 3) In the list, select the rule to be modified and change the values as desired.
- 4) Click Set to update the rule and click Apply to save the rule.
- 5) Click Close to close the edit window.

Related concepts

filter.config on page 372

Creating an add_hdr rule to allow Google enterprise gmail

Google provides a mechanism in the form of a custom header in the request, that allows Google to recognize and allow or block access to enterprise gmail and other Google Apps for Business.

To make Google's solution work for enterprise gmail:

Steps

- In the Web Security module of the Forcepoint Security Manager, permit the category Internet Communication > General Email.
- In the Content Gateway manager enable HTTPS (SSL decryption). If your site does not already use SSL support, acquaint yourself with the feature before enabling it.

 In the Content Gateway manager, on the Configure > Security > Access Control page, open filter.config and create an add_hdr rule.

The **add_hdr** rule type can be used with any site that uses a custom header-value pair to accomplish special handling.

a) Select add_hdr.

Note

- b) For Primary Destination Type select dest_domain.
- c) For Primary Destination Value specify "mail.google.com".
- d) In the Custom Header field, specify "X-GoogApps-Allowed-Domains".
- e) In the **Header Value** field, specify your domain, or a list of domains separated by commas. For example: www.example1.com, www.example2.com
- f) Optionally, in the Source IP field specify the source IP address or address range to which this rule will be applied. For example: 10.10.20.30 or 10.10.1.1-10.30.40.50.
- g) Click Add to add the rule.
- h) Click Apply to save all the changes, and then click Close to close the edit window.

When a user attempts to access Google services from an unauthorized account, Google displays a block page similar to this:

Google accounts

This service is not available

Gmail is not available for bob@gmail.com within this network. Gmail is only available for accounts in the following domains:

- example1.com
- example2.com

Please talk to your network administrator for more information.

Did you use this product with a different Google Account? <u>Sign out</u> of your current Google Account and then sign in to the account you want.

©2011 Google - Google Home - Terms of Service - Privacy Policy - Help

For Google's description of the filtering solution, see the article Block access to consumer accounts and services while allowing access to Google Apps for your organization.

Configuring SOCKS firewall integration

SOCKS is commonly used as a network firewall, allowing hosts behind a SOCKS server to gain full access to the Internet while preventing unauthorized access from the Internet to hosts inside the firewall.

In a SOCKS configuration, instead of accessing the origin server directly, the proxy goes through a SOCKS server. The SOCKS server authorizes communication between the proxy and the origin server and relays the data to the origin server. The origin server then sends the content back to the proxy through the SOCKS server.

- Content Gateway can act as a SOCKS client, where it receives and serves HTTP or FTP requests as usual.
- Content Gateway can act as a SOCKS proxy, relaying requests to and from the SOCKS server (usually on port 1080).
- When Content Gateway is installed on an appliance it can act as a SOCKS server, providing all of the services of a SOCKS server. (When Content Gateway is **not** installed on an appliance, it cannot act as a SOCKS server.)



Note

Content Gateway does not perform authentication with the client. However, Content Gateway can perform user name and password authentication with a SOCKS server running SOCKS version 5.

Configuring SOCKS servers

Content Gateway can be configured to work with one or more SOCKS servers in your network. When Content Gateway is installed on an appliance, a SOCKS server is included with the module.

Ę

When Content Gateway is **not** installed on an appliance, no SOCKS server is provided with Content Gateway.

To configure SOCKS servers:

Note

Steps

- 1) Enable the SOCKS feature.
 - a) Navigate to Configure > My Proxy > Basic > General.
 - b) In the Security section of the Features table, click SOCKS On, and click Apply.
 - c) Restart Content Gateway.
- 2) Specify the SOCKS version.
 - a) Go to Configure > Security > SOCKS > General.
 - b) Select the SOCKS version running on your SOCKS servers and click Apply.

- 3) To configure the on-appliance SOCKS server:
 - a) Select the Server tab.
 - b) In the **On-Appliance SOCKS Server** area, select **Enabled** and click **Apply**. An entry for the server is created in the **socks_server.config** file.
 - c) To change the default entry, in the SOCKS Server area click Edit File. In the editor, select the On-Appliance-SOCKS-Server rule.
 You can change the port, whether it will be the default SOCKS server, and whether server authentication is applied.

You cannot change the server name or the IP address, which is always the loopback address.

After you make the needed changes, click Set.

- 4) To configure use of other SOCKS servers in your network:
 - a) Select the Server tab and in the SOCKS Server area click Edit File.
 - b) Enter a SOCKS server name.
 - c) Enter the SOCKS server IP address or a domain name that is resolvable by the DNS server inside your network.
 - d) Select whether it will be the default SOCKS server.
 - e) If authentication will be used, provide a SOCKS user name and password.
 - f) Click Set to add the server to the list.
 You can always return to the editor, select the rule, make changes, and click Set to save them.
- 5) If there are multiple SOCKS servers, after they have been added, or while they are being added, you can arrange them in precedence-order by selecting an entry and moving it up or down the list with the up and down arrows.
- 6) Click Apply to accept your changes, and Close to close the editor.
- In the SOCKS Server Rules area you can create rules for specific routing and bypass by destination IP address. See, Setting SOCKS server bypass.
- 8) To review configuration options that apply to all SOCKS servers, select the **Options** tab.
 - a) Review and adjust the **Server Connection Timeout** value. It specifies how many seconds Content Gateway waits attempting to connect to a SOCKS server before timing out.
 - b) Review and adjust the **Connection Attempts Per Server** value. It specifies how many times Content Gateway attempts to connect to a given SOCKS server before marking the server as unavailable.
 - c) Review and adjust the Server Pool Connection Attempts value. It specifies how many times Content Gateway attempts to connect to a given SOCKS server in the pool before giving up.

9) When SOCKS server configuration is complete, click Apply and then go to Configure > My Proxy > General and restart Content Gateway.

Related tasks Setting SOCKS server bypass on page 163

To remove a server from the list

Steps

- 1) In the SOCKS Server area click Edit File.
- 2) In the list, select the entry you want to delete and click X, to the left of the list.
- 3) Click Apply and then Close, when you're ready to exit the editor.
- 4) When configuration is complete, go to Configure > My Proxy > General and restart Content Gateway.

Setting SOCKS proxy options

To configure Content Gateway as a SOCKS proxy, you must enable the SOCKS proxy option and specify the port on which Content Gateway accepts SOCKS traffic from SOCKS clients.

As a SOCKS proxy, Content Gateway can receive SOCKS packets (usually on port 1080) from the client and forward requests directly to the SOCKS server.



Note

You must set SOCKS proxy options in addition to enabling the SOCKS option and specifying SOCKS server information described in *Configuring SOCKS servers*.

Steps

- 1) Navigate to Configure > Security > SOCKS > Proxy.
- 2) Enable SOCKS Proxy.
- 3) Specify the port on which Content Gateway accepts SOCKS traffic. The default is port 1080.
- 4) Click Apply.
- 5) Click Restart on Configure > My Proxy > Basic > General.

Related tasks Configuring SOCKS servers on page 160

Social servers on page 100

Setting SOCKS server bypass

You can configure Content Gateway to bypass SOCKS servers and access certain origin servers directly.

Steps

- Navigate to Configure > Security > SOCKS > Server. In the SOCKS Server Rules area click Edit File to open socks.config.
- 2) To modify an existing rule, select it from the list, make your changes, and click Set.
- 3) To create a new rule, specify the parameters and click Add.
 - a) Select a Rule Type: Route through SOCKS server Do not route through SOCKS server
 - b) Specify a destination IP address or range of addresses. Never specify the all networks broadcast address: 255.255.255.255.255
 - c) Select the SOCKS servers to be used for the traffic.
 - d) Select whether the traffic will be distributed to the specified SOCKS servers in round robin fashion.
 - e) Click Add to add the rule.
- 4) Click Apply and then Close.
- 5) Click Restart on Configure > My Proxy > Basic > General.

Using the Split DNS option

You can configure Content Gateway to use multiple DNS servers, depending on your security requirements. For example, you can configure Content Gateway to look to one set of DNS servers to resolve host names on your internal network, while allowing DNS servers outside the firewall to resolve hosts on the Internet. This maintains the security of your intranet, while continuing to provide direct access to sites outside your organization.

To configure Split DNS, you must perform the following tasks:

- Specify the rules for performing DNS server selection based on the destination domain, the destination host, or a URL regular expression.
- Enable the Split DNS option.

In the Content Gateway manager:

Steps

- 1) Go to the Configure > Networking > DNS Resolver > Split DNS tab.
- 2) Enable the **Split DNS** option.
- 3) In the Default Domain field, enter the default domain for split DNS requests. Content Gateway appends this value automatically to a host name that does not include a domain before determining which DNS server to use.
- 4) In the DNS Servers Specification area, click Edit File to open the configuration file editor for the *splitdns.config* file.
- 5) Enter information in the fields provided, and then click Add. All the fields are described in splitdns.config.
- 6) Click Apply, and then click Close.
- 7) On the Split DNS tab, click Apply to save your configuration.
- 8) Click Restart on Configure > My Proxy > Basic > General.

Related concepts

splitdns.config on page 463

Content Gateway user authentication

Content Gateway supports several methods of authenticating users before their requests are allowed to proceed. These methods can be used together with Forcepoint Web Security user identification features to provide fallback should user authentication fail or become unavailable.



Important

Use the directory service used by Web Security.

In both explicit and transparent proxy modes, Content Gateway supports user authentication with:

- Integrated Windows Authentication (Kerberos with SPNEGO to NTLM)
- Legacy NTLM authentication (NTLMSSP)
- LDAP authentication
- RADIUS authentication

Content Gateway also supports combinations of Integrated Windows Authentication (IWA), Legacy NTLM, and LDAP using:

Rule-Based Authentication

Rule-Based Authentication summary

Rule-Based Authentication is an ordered list of authentication rules. When a request is processed, the list is traversed top to bottom and the first match is applied.

Rules specify:

- 1) How to match a client. By:
 - IP address
 - Inbound proxy port (explicit proxy only; do not use port 80)
 - User-Agent value
 - A combination of the above
- 2) The domain or ordered list of domains to authenticate against. With a list, the first successful authentication is remembered and used in subsequent authentications for that user.
- 3) Whether a customizable web portal page should be used for authentication.

Multiple Realm Networks

Rule-Based Authentication supports multiple realm network structures in which Windows Active Directory domains do not have mutual trust relationships and therefore require that each domain's members be authenticated by a domain controller within their domain. In this environment rules are created that specify:

- 1) Members of the realm (untrusted domain) by IP address or proxy port
- 2) The realm (domain) they belong to

Authenticating when domain membership is unknown

Some organizations do not always know what domain a user belongs to. For example, this can happen when organizations are rapidly acquiring new businesses. The unknown domain membership problem can be handled in rule-based authentication by creating a rule (or rules) for IP address lists or ranges that also specifies an ordered list of domains to attempt to authenticate against. The first successful authentication is remembered and used in later authentications.

Authentication based on User-Agent value

One or more User-Agent values can be specified in an authentication rule. Often this is a list of browsers. When the

User-Agent value matches a rule, authentication is performed against the specified domain(s). If the User-Agent value doesn't match any rule, and no rule matches based on other values, no authentication is performed (this is always true; if no rule matches, no authentication is performed).

Related concepts Integrated Windows Authentication on page 173 Legacy NTLM authentication on page 179 LDAP authentication on page 180 RADIUS authentication on page 184 Rule-Based Authentication on page 186

Selecting the authentication method

The authentication method is selected in the **Authentication** section of the **Configure > My Proxy > Basic** page. Configuring authentication for rule-based authentication begins with selecting **Rule-Based Authentication**.

Supported domain controllers and directories

- Windows NT domain controllers
- Windows 2008, 2008 R2, 2012, 2012 R2 Active Directory
- Novell eDirectory 8.5.1 or later (LDAP only)
- Oracle DSEE 11g (LDAP only)

Best practices when using Windows Active Directory

If you have only one Active Directory domain, or if all of your Active Directory domains share inbound and outbound trust relationships, the best option is to deploy Integrated Windows Authentication. However, if you want to control authentication based on User-Agent values, you must use Rule-Based Authentication.

If you have multiple domains or realms and user authentication is a requirement, you must use Rule-Based Authentication. For details, see *Rule-Based Authentication*.

If user identification is sufficient, you can use one of the Forcepoint Web Security user identification options. See the "User Identification" section of the Forcepoint Web Security Administrator Help.

Related concepts Rule-Based Authentication on page 186

Backup domain controllers

For Integrated Windows Authentication and Legacy NTLM, Content Gateway supports the specification of backup domain controllers for failover. If the primary domain controller (DC) does not respond to proxy requests, Content

Gateway contacts the next DC in the list (the backup domain controller). For the next request, the proxy tries to contact the primary DC again and then contacts the backup DC if the connection fails.

Transparent user authentication

Content Gateway supports both transparent (Single Sign-On) and interactive (prompted) authentication. Transparent authentication is supported with Integrated Windows Authentication and Legacy NTLM. Some browsers provide only limited

support. See Browser limitations.

On Windows networks, Single Sign-On allows users to sign on only once so that they can transparently access all authorized network resources. Therefore, if a user has already logged on to the Windows network successfully, the credentials specified during Windows logon are used for proxy authentication and the user is not prompted again for a username and password.

Interactive authentication is supported in networks that are not configured for Single Sign-On and for use with browsers that don't support Single Sign-On. With interactive authentication, users are prompted for credentials before they can access content through Content Gateway.



Browser limitations on page 167

Browser limitations

Not all web browsers support transparent user authentication.



Note

Please see the Certified Product Matrix for the most up-to- date information.

The following table indicates how a browser responds to an authentication request when Integrated Windows Authentication (IWA) is configured.

Browser/ Operating System	Internet Explorer (v10 and 11 tested)	Firefox	Chrome	Opera	Safari
Windows	Performs transparent authentication (v11 tested)	Performs transparent authentication (v53 tested)	Performs transparent authentication (v58 tested)	Performs transparent authentication	Falls back to NTLM and prompts for credentials
Mac OS X	Not applicable	Performs transparent authentication (v54 tested)	Falls back to NTLM and prompts for credentials (v55.0 tested)	Falls back to NTLM and prompts for credentials (v46 tested)	Performs transparent authentication (v10 tested)

Browser/ Operating System	Internet Explorer (v10 and 11 tested)	Firefox	Chrome	Opera	Safari
Red Hat Enterprise Linux, update 6	Not applicable	Performs transparent authentication (v45 tested)	Browser issue prevents IWA from working	Not tested.	Not applicable



Note

Each time Internet Explorer is accessed, authentication information is requested even if Integrated Windows Authentication (IWA) has been configured.

Global authentication options

Use the Configuration > Security > Access Control > Global Authentication Options page to configure:

- User authentication Fail Open/fail closed behavior
- Credential Caching options
- Redirect Options (required for transparent proxy deployments)
- Cookie Sharing options

These settings apply to all proxy user authentication configurations, within the parameters stated for each option below.

Whenever changes are made to any of these settings, click **Apply** to save your changes and then restart the proxy to put the changes into effect.

Related concepts

Fail Open on page 168 Credential Caching on page 169 Redirect Options on page 171 Cookie Sharing on page 172

Fail Open

Fail Open specifies whether requests are allowed to proceed for processing when user authentication fails.

When Fail Open is enabled and a Forcepoint Web Security transparent identification agent is configured, if authentication fails and the client is identified by the agent, user-based policy is applied. If the user cannot be identified and a policy is assigned to the client's IP address, that policy is applied. Otherwise, the Default policy is applied.



Important

The Fail Open setting does not apply when IWA is the authentication method and the client fails to retrieve a kerberos ticket from the domain controller (DC) because the DC is down.

The Fail Open setting does apply with IWA when IWA falls back to NTLM.

The Fail Open setting does not apply when using LDAP in explicit proxy mode.

Options include:

- Disabled specifies that requests do not proceed when authentication failures occur.
- Enabled only for critical service failures (default) specifies that requests proceed if authentication fails due to:
 - No response from the domain controller
 - The client is sending badly formatted messages
- Enabled for all authentication failures, including incorrect password specifies that requests proceed for all authentication failures, including password failures.



Important

When user authentication is rule-based with a domain list:

- If Enabled only for critical service failures is selected, when a critical service failure occurs fail open is not applied. An error always results in fail closed.
- If Enabled for all authentication failures, including incorrect password is selected, after trying basic credentials with every domain in the list, fail open is applied.

Credential Caching

Credential Caching options include:

- Credential Caching selections
- Caching Method
- Cache Time-To-Live (TTL), in minutes
- Cookie Expiration
- LDAP Specific Settings

Credential caching settings apply to all clients whether Content Gateway is in transparent or explicit mode. In explicit mode, caching exceptions can be configured.

Credential caching applies to:

- All authentication methods when Content Gateway is a transparent proxy
- When Content Gateway is an explicit proxy:
 - Integrated Windows Authentication (IWA)
 - Legacy NTLM

When IWA authenticates with Kerberos, Kerberos handles ticket (credential) caching.

Credential Caching selections

Cache all authentication credentials -- specifies that user credentials should be cached for all requests. This selection applies to both transparent and explicit proxy.

Cache all authentication credentials except from the specified IP addresses-- specifies that authentication credentials should not be cached when requests are made from the IP addresses added to the list provided. Exceptions listed here apply only when Content Gateway is in explicit proxy mode.

In the box provided, enter the IP addresses for which you do not want authentication information cached. Specify a comma separated list of IP addresses or IP address ranges that host multiple users.

Note

This feature should be used for clients for whom IP or cookie cache is not available, such as NATed IP addresses or IP addresses used for clients behind a terminal server. Since all requests from the clients listed are authenticated, performance may be impacted.

Caching Method options

Cache using IP address only – specifies that all credentials are cached with IP address surrogates. This is the recommended method when all clients have unique IP addresses.

Cache using Cookies only – specifies that all credentials are cached with cookie surrogates. This is recommended when all clients share IP addresses, as with multi-host servers such as Citrix servers, or when traffic is NATed by a device that is forwarding traffic to Content Gateway.

Cache using both IP addresses and Cookies – specifies to use cookie surrogates for the IP addresses listed in the cookie caching list, and to use IP address surrogates for all other IP addresses. This is recommended when the network has a mix of clients, some with unique IP addresses and some using multi-user hosts or that are subject to NATing.

The cookie caching list is a comma separated list that can contain up to:

- 64 IPv4 addresses
- 32 IPv4 address ranges
- 24 IPv6 addresses
- 12 IPv6 address ranges

For a description of surrogate credentials, see Surrogate credentials.

Cookie mode caching:

- Cookie mode caching does not work with applications that do not support cookies, or with browsers in which cookie support has been disabled.
- When the browser is Internet Explorer, the full proxy hostname in the form "http://host.domain.com" must be added to the Local intranet zone.
- When the browser is Chrome, it must be configured to allow third-party cookies or configured for an exception to allow cookies from the proxy hostname in the form "host.domain.com".
- When the IP address is set for cookie mode and the request method is CONNECT, no caching is performed.
- Cookie mode caching is not performed for FTP requests.
- Cookie mode caching is supported by Captive Portal and client certificate authentication.
- For explicit proxy, cookie-based authentication is not supported for HTTPS. IP-address authentication is used.



Note

The user interface setting to disable the NTLM cache for explicit proxy has been removed. Although not recommended, the cache can be disabled for explicit proxy traffic in records.config by setting the value of **proxy.config.ntlm.cache.enabled** to **0** (zero).

Related concepts

Surrogate credentials on page 173

Cache Time-To-Live

Cache Time-To-Live (TTL) specifies the duration, in minutes, that an entry in the cache is retained. When the TTL expires, the entry is removed and the next time that the user submits a request, the user is authenticated. If the authentication succeeds, an entry is placed in the cache.

The default TTL is 15 minutes. The range of valid values is 5 to 1440 minutes.

Cookie Expiration

Cookie caching allows a user to re-access the system without authentication until the cookie is no longer valid. Select **Delete cookies upon logout** to force cookies expire when the user ends a session.

This feature is recommended in deployments where multiple users share a machine.

LDAP Specific Settings

When enabled, **Purge LDAP cache on authentication failure** causes the proxy to delete the authorization record for the client from the LDAP cache when an LDAP user authentication failure occurs.

Redirect Options

Redirect Options include:

Redirect Hostname specifies an alternate hostname for the proxy.

By default, authenticating clients are redirected to the hostname of the Content Gateway machine. If clients are unable to resolve that hostname through DNS, or if an alternate DNS name for the proxy is defined, that hostname should be specified in the **Redirect Hostname** field.

To ensure that user authentication for transparent proxy occurs transparently (without prompting the user for credentials), the browser must be configured so that the Redirect Hostname is in its **Intranet Zone**. Typically, this is achieved by ensuring that the Redirect Hostname is in the same domain as the computer on which the browser is running. For example, if the client is **workstation.example.com** and the Redirect Hostname is **proxyhostname.example.com**, the browser allows authentication to occur transparently. Consult your browser documentation.

Content Gateway supports transparent authentication in proxy clusters that use WCCP load distribution. However, the **assignment method distribution attribute** must be the source IP address. For more information, see *WCCP load distribution*.

To ensure that user authentication for transparent proxy occurs transparently (without prompting the user for credentials), the browser must be configured so that the Redirect Hostname is in its **Intranet Zone**. Typically, this is achieved by ensuring that the Redirect Hostname is in the same domain as the computer on which the browser is running. For example, if the client is **workstation.example.com** and the Redirect Hostname is **proxyhostname.example.com**, the browser allows authentication to occur transparently. Consult your browser documentation.

Content Gateway supports transparent authentication in proxy clusters that use WCCP load distribution. However, the **assignment method distribution attribute** must be the source IP address. For more information see *WCCP load distribution*.

Related concepts WCCP load distribution on page 41

Cookie Sharing

Authentication credentials cached with cookie surrogates can be shared across all nodes in a cluster.

When cookie mode caching is enabled, after a user is authenticated the cookie for that user is used for subsequent authentication attempts by any of the proxies that are clustered with the proxy that did the initial authentication. This feature is especially useful in load balanced environments.

When either **Cache using Cookies only** or **Cache using both IP addresses and Cookies** is enabled, the Cookie Sharing option is automatically enabled.



Note

All proxies in the cluster must use the same caching method when cookie sharing is enabled.

- Select Choose File for both Public and Private keys to import your own keys for use with this feature. Browse to the file you want to use and select it. Files must be in PEM format. The same keys must be imported for each proxy in the cluster.
- After selecting each file, click Import Keys to import custom keys (recommended) and store them in the default location.

Note that default keys are provided and are added when the product is installed or upgraded. The default files are:

/opt/WCG/config/cookie_auth_public.pem

/opt/WCG/config/cookie_auth_private.pem

Select the files you wish to import. The custom keys are automatically copied to this folder and renamed to the default names.



Important

When custom keys are imported, the default files provided by Forcepoint are overwritten. You should backup the default keys prior to importing. See **Save Public Key** and **Save Private Key** below.

Keys must be PKCS#1 RSA public keys and are RSA 1024/2048/4096 bit public and private key pairs without a passphrase. Use the following commands to generate keys:

openssl genrsa -out cookie_auth_private.pem 1024

openssl rsa -in cookie_auth_private.pem -RSAPublicKey_out -out cookie_auth_public.pem

Change 1024 to 2048 or 4096 to generate 2048 or 4096 bit keys.

Select Save Public Key and Save Private Key to make a backup of the files. Select the location and filenames to use for the backup copy, keeping in mind that the default names are always used for the active keys.

Key files should be backed up prior to importing new keys.

When load balancing has been configured, all proxies must use the same setting for **Redirect Hostname**. The value must be the fully qualified domain name (FQDN) of the load balancer.

Important

Cookie sharing has the following limitations:

- Cookie caching limitations also apply to cookie sharing. Therefore, since cookie caching is not supported for CONNECT requests, cookie sharing is not supported.
- Custom keys must be imported manually. Custom Keys are not synchronized across the cluster.
- Cookie sharing is not supported with client certificate authentication.

Surrogate credentials

Surrogate credentials are entries placed in the credential cache after initial successful authentications.

- An IP address surrogate ties a credential to an IP address and assumes that the IP address is used by only one user at any given time.
- A cookie surrogate is tied to a cookie placed on the client's system and depends on client application support for cookies. This method is required when a client IP address is shared by more than one user at a time, as with multi-user hosts such as Citrix servers.

After the initial successful authentication, Content Gateway uses the surrogate credential to respond to subsequent authentication requests on behalf of the user, thus reducing latency and the load on domain controllers and directory services. Credential surrogate entries are deleted when the Time-To-Live expires.

Integrated Windows Authentication

Integrated Windows Authentication (IWA) is a robust method of authenticating users who belong to shared-trust Windows domains (one or many).

Integrated Windows Authentication:

- Uses Kerberos and SPNEGO
- Supports NTLM in both explicit and transparent proxy modes
- Supports NTLMv2 and NTLMv1 with Session Security
- Supports Windows Active Directory. (See this article for a list of supported versions.)
- Can be used with Rule-Based Authentication and Captive Portal Authentication.
- Supports Internet Explorer, Firefox, Google Chrome, Windows Safari, Safari on iPad iOS4, and Opera
- Supports UTF-8 user names
- Supports fall back to prompted authentication

Requires that:

- Clients be joined to the domain
- Client browsers specify the Fully Qualified Domain Name (FQDN) of Content Gateway as an intranet site or trusted site (HTTP://FQDN)
- When Redirect for HTTPSS Authentication is enabled on the Configure > Security > Access Control > Global Authentication page, Content Gateway will redirect over HTTPS. To avoid user prompts, HTTPS:// FQDN must also be specified as an intranet or trusted site in client browsers.



Note

Microsoft Edge does not support trusted sties. Intranet sites are required for clients using Edge.

In explicit proxy deployments, browsers must specify the FQDN of Content Gateway

If you are using IWA with rule-based authentication, see *Rule-Based Authentication* for configuration steps.

Related concepts Rule-Based Authentication on page 186

Integrated Windows Authentication: Configuration summary

Follow these steps to configure IWA as the user authentication method for your Content Gateway deployment:

- In the Content Gateway manager, enable Integrated Windows Authentication on the Configure > My Proxy > Basic page and click Apply.
- Configure Global authentication options.
- Join Content Gateway to the Windows domain. See Configuring Integrated Windows Authentication for a list of required conditions.

Related concepts

Global authentication options on page 168

Related tasks

Configuring Integrated Windows Authentication on page 174

Configuring Integrated Windows Authentication

Steps

- Go to Configure > My Proxy > Basic > General. In the Authentication section, click Integrated Windows Authentication On, and click Apply.
- 2) Configure the Global authentication options.

- 3) Join the Windows domain. To join the domain:
 - Content Gateway must be able to resolve the domain name.
 - Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
 - The correct domain Administrator name and password must be specified.
 - There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
 - If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services must be reachable by Content Gateway on the network.
 - a) In the **Domain Name** field, enter the fully qualified domain name.
 - b) In the Administrator Name field enter the Windows Administrator user name.
 - c) In the Administrator Password field enter the Windows Administrator password. The name and password are used only during the join and are not stored.
 - d) Select how to locate the domain controller:
 - Auto-detect using DNS
 - DC name or IP address

If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.

e) In the Content Gateway Hostname field, confirm that the hostname is the correct hostname and that it is no more than 15 characters (no more than 11 characters on appliances). If it is longer, it must be shortened if IWA is to be used. The length restriction results from the 15 character limit on NetBIOS hostnames.



Warning

Do not change the hostname after the domain is joined. If the hostname is changed, IWA immediately stops working and will not work again until the domain is unjoined and then rejoined with the new hostname.

f) Click Join Domain. If there is an error, ensure that the conditions outlined above are met and then see Failure to join the domain.

Important

All clients subject to authentication must be joined to the domain.

Browsers and other proxy clients must be configured to specify the FQDN of Content Gateway as an intranet site or trusted site.

g) Restart Content Gateway and run some test traffic through the proxy to verify that authentication is working as expected. If there is a problem, see *Troubleshooting Integrated Windows Authentication*.

Related concepts

Global authentication options on page 168 Failure to join the domain on page 177 Troubleshooting Integrated Windows Authentication on page 177

To unjoin the current domain and join a new domain

Steps

- Navigate to the Configure > Security > Access Control > Integrated Windows Authentication tab and click Unjoin.
- 2) To join a new domain, in the **Domain Name** field, enter the fully qualified domain name.
- 3) In the Administrator Name field enter the Windows Administrator user name.
- 4) In the Administrator Password field enter the Windows Administrator password. The name and password are used only during the join and are not stored.
- 5) Select how to locate the domain controller:
 - Auto-detect using DNS
 - DC name or IP address

If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.

6) Click Join Domain.

To change the way the domain controller is found

Steps

- 1) Navigate to the Configure > Security > Access Control > Integrated Windows Authentication tab.
- 2) In the Domain Controller section, select how to locate the domain controller:
 - Auto-detect using DNS
 - DC name or IP address

If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.

3) Click Apply.

Configuring Integrated Windows Authentication with a load balancer

Integrated Windows Authentication (IWA) with a load balancer is supported.



Important

After upgrade, check and, if necessary, rejoin IWA domains.

Transparent proxy deployments do not require any special configuration.

Explicit proxy deployments that are behind a load balancer require a custom configuration

With Content Gateway, IWA uses the Kerberos protocol, with NTLM fallback.

In a load-balanced environment:

- Clients explicitly point to the Content Gateway cluster via the FQDN, which, when a load balancer is used, must resolve to the load balancer's VIP.
- Kerberos then returns a ticket for the load balancer's VIP, which the client then sends to Content Gateway.
- Because the ticket is not issued for the proxy's IP address, but rather for the load balancer's VIP, Content Gateway cannot decrypt the ticket and authentication fails.

To restate the problem, it's not possible to configure clients to request Content Gateway's Kerberos ticket because the client's operating system handles the ticket request based on the FQDN of the proxy, which resolves to the VIP of the load balancer.

Normally, Content Gateway would be configured to share the hostname of the load balancer, but this is not possible when the load balancer requires hostname resolution (as with DNS-based load balancing).

Because it's not possible to stop clients from sending a load-balancer's Kerberos ticket to Content Gateway, the proxies must be configured to accept the load-balancer's ticket, making the Content Gateway nodes appear as the load-balancer within the scope of Kerberos.

Please contact this article for detailed, step-by-step configuration instructions.

Troubleshooting Integrated Windows Authentication

This section covers 2 common problems:

Related concepts

Failure to join the domain on page 177 Failure to authenticate clients on page 178

Failure to join the domain

These conditions are required for Content Gateway to join a domain:

- Content Gateway must be able to resolve the domain name.
- Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
- The correct domain Administrator name and password must be specified.
- There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
- If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services, must be reachable by Content Gateway on the network.
- If the Active Directory is configured with multiple Sites, ensure that the subnet that Content Gateway is on is added to one of them.

Troubleshooting

- Errors encountered in the join action are reported at the top of the screen (the Integrated Windows Authentication tab).
- The error message usually includes a link to the failure log where you can get more details.
- Join failures are logged to /opt/WCG/logs/smbadmin.join.log
- In most cases, the failure message in the log is a standard Samba and Kerberos error message that is easily found with an Internet search.

Failure to authenticate clients

These conditions are required to authenticate clients:

- Content Gateway clients must be a member of the same domain as that joined by Content Gateway.
- Client system time must be in sync with the domain controller and Content Gateway to plus or minus 1 minute.
- Explicit proxy clients must **not** be configured to send requests to the IP address of Content Gateway. Clients must use the Fully Qualified Domain Name (FQDN) of Content Gateway. If the IP address is used, NTLM authentication is always performed.
- The Content Gateway FQDN must be in DNS and resolvable by all proxy clients.
- Browsers and other client applications must specify the FQDN of Content Gateway as an intranet site or trusted site.
- When the Active Directory is configured with multiple Sites, the subnet that Content Gateway is on must be added to one of them. If it's not, the following alarm may be generated when Content Gateway is restarted:

Windows domain [domain name] unreachable or bad membership status

Troubleshooting

In the Content Gateway manager, use the **Diagnostic Test** function on the **Monitor > Security > Integrated Windows Authentication** tab. This Monitor page displays authentication request statistics and provides the diagnostic test function.

The **Diagnostic Test** function performs connectivity and authentication testing and reports errors. It also shows domain controller TCP port connectivity and latency.

Errors and messages are logged to:

- /var/log/messages
- content_gateway.out
- /opt/WCG/logs/smbadmin.log
- /opt/WCG/logs/smbadmin.join.log

Performance issues

- IWA (Kerberos): Authentication performance is bound by CPU. There is no communication to the domain controllers for Kerberos authentication.
- NTLM and Basic: Domain controller responsiveness effects performance. The Monitor > Security > Integrated Windows Authentication page shows average response time.

Legacy NTLM authentication

Content Gateway supports the NTLM (NT LAN Manager) authentication protocol as a method of ensuring that users in a Windows network are authenticated before they access the Internet.



Important

This implementation of NTLM support (Legacy NTLM) relies solely on the NTLMSSP protocol. Although it performs reliably as documented in this section, it is highly recommended that the *Integrated Windows Authentication* mode be used instead. It provides more robust and secure support for NTLM.

If rule-based authentication will be used, configure Legacy NTLM authentication through the *Rule-Based Authentication* option.

However, read this section to become familiar with Legacy NTLM features and restrictions.

When the Legacy NTLM option is enabled, the proxy challenges users who request content for proof of their credentials. The proxy then sends the proof of the user's credentials directly to the Windows domain controller to be validated. If the credentials are valid, the proxy serves the requested content and stores the credentials in the NTLM cache for future use. If the credentials are not valid, the proxy sends an *authentication failed* message.

Restrictions

- 1) WINS resolution is not supported. Domain controllers must have host names that can be resolved by a DNS server.
- 2) Extended security is not supported and cannot be enabled on the domain controller.
- NTLM2 session security is not supported and cannot be enabled on clients. In the Security Settings area of the Windows operating system, inspect the Network Security: Minimum session security settings.
- 4) NTLMv2 is not supported with Active Directory 2008. The required Network Security: LAN Manager Authentication setting is described in step 5 of *Configuring NTLM proxy authentication*, below.
- 5) Not all browsers support transparent NTLM authentication. See Browser limitations.

If you are using Legacy NTLM with rule-based authentication, see *Rule-Based Authentication*, for configuration steps.

Related concepts

Integrated Windows Authentication on page 173 Rule-Based Authentication on page 186

Related reference

Browser limitations on page 167

Configuring Legacy NTLM authentication

Steps

- 1) Go to Configure > My Proxy > Basic > General.
- 2) In the Authentication section, click Legacy NTLM On, and click Apply.
- 3) Configure the Global authentication options.
- 4) Go to Configure > Security > Access Control > Legacy NTLM.
- 5) In the **Domain Controller Hostnames** field, enter the hostname of the primary domain controller, followed, optionally, by a comma separated list of backup domain controllers. The format of the hostname must be: host_name[:port][%netbios_name]

or

IP_address[:port][%netbios_name]

Note



If you are using Active Directory 2008, you must include the netbios_name or use SMB port 445. If you **do not** use port 445, you must ensure that the Windows Network File Sharing service is running on the Active Directory server. See your Windows Server 2008

documentation for details.

If you are using Active Directory 2008, in the Windows **Network Security** configuration, **LAN Manager Authentication level** must be set to **Send NTLM response only**. See your Windows Server 2008 documentation for details.

6) Enable Load Balancing if you want the proxy to balance the load when sending authentication requests to multiple domain controllers.



Note

When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term failover provision, until such time that the primary domain controller can accept new connections.

7) Click Apply and restart Content Gateway (Configure > My Proxy > Basic > General). Optionally, you can configure Content Gateway to allow certain clients to access specific sites on the Internet without being authenticated by the NTLM server; See Access Control).

Related concepts

Global authentication options on page 168 Access Control on page 303

LDAP authentication

Content Gateway supports the LDAP option to ensure that users are authenticated with an LDAP server before accessing content through the proxy.


Important

If rule-based authentication will be used, configure LDAP authentication through the *Rule-Based Authentication* option. However, read this section to become familiar with LDAP features and restrictions.

When LDAP is enabled:

- Content Gateway acts as an LDAP client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the LDAP server to check that the credentials are correct.
- If the LDAP server accepts the username and password, the proxy serves the client the requested content and stores the username and password in the credential cache.
- Future authentication requests for that user are served from the cache until the cache entry expires (Time-To-Live value).
- If the LDAP server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

LDAP authentication supports both simple and anonymous bind.

LDAP user authentication can support passwords containing special characters. Configuration is made directly in the **records.config** file. The following parameter must be enabled, and the correct encoding name to which the special characters belong must be configured. Add these entries to **records.config**. Note that the default setting is 0 (feature disabled).

```
// To enable the feature specify 1.
CONFIG proxy.config.ldap.proc.encode_convert INT <1 or 0>
// Specify an encoding name here. For example,
// for German specify "ISO-8859-1".
CONFIG proxy.config.ldap.proc.encode_name STRING <encoding name>
```

Related concepts Rule-Based Authentication on page 186

Configuring Content Gateway to be an LDAP client

Steps

- 1) Go to Configure > My Proxy > Basic > General.
- 2) In the Authentication section, click LDAP On, and then click Apply.
- 3) Configure the *Global authentication options*.
- 4) Go to Configure > Security > Access Control > LDAP.
- 5) Enter the hostname of the LDAP server.

6) Enter the port on which Content Gateway communicates with the LDAP server. The default is port 389.

Note

When the LDAP directory service is Active Directory, requests from users located outside the global catalog's base domain will fail to authenticate. This is because the default port for LDAP is 389 and requests sent to 389 search for objects only within the global catalog's base domain. To authenticate users from outside the base domain, change the LDAP port to 3268. Requests sent to 3268 search for objects in the entire forest.

- 7) Enable **Secure LDAP** if you want the proxy to use secure communication with the LDAP server. Secure communication is performed on port 636 or 3269. Change the port value in the previous field, if necessary.
- 8) Select the type of directory service to set the filter for searching.
 - Microsoft Active Directory (sAMAccountName) sets the type to sAMAccountName (default).
 - Microsoft Active Directory (userPrincipalName) sets the type to userPrincipalName.
 - Other sets the type to uid for eDirectory or other directory services.
- 9) Enter the **Bind Distinguished Name** (fully qualified name) of a user in the LDAP-based directory service. For example:

CN=John Smith, CN=USERS, DC=MYCOMPANY, DC=COM

Enter a maximum of 128 characters in this field.

If no value is specified for this field, the proxy attempts to bind anonymously.

- **10)** Enter a password for the user specified in the previous step.
- 11) Enter the Base Distinguished Name (DN). Obtain this value from your LDAP administrator.
- 12) Click Apply.
- Click Restart on Configure > My Proxy > Basic > General. As optional steps, you can:
 - Change LDAP cache options. See Setting LDAP cache options.
 - Configure Content Gateway to allow certain clients to access specific sites on the Internet without being authenticated by the LDAP server. See *Access Control*).

Related concepts

Global authentication options on page 168 Access Control on page 303

Related tasks

Setting LDAP cache options on page 183

Setting LDAP cache options

By default, the LDAP cache is configured to store 5000 entries and each entry is considered fresh for 3000 minutes. Change these options by editing the **records.config** file.

Steps

- 1) Open the records.config file located in /opt/WCG/config.
- 2) Edit the following variables:

Variable	Description
proxy.config.ldap.cache.size	Specify the number of entries allowed in the LDAP cache.
	The default value is 5000. The minimum value is 256.
proxy.config.ldap.auth.ttl_value	Specify the number of minutes that Content Gateway can store username and password entries in the LDAP cache.
proxy.config.ldap.cache. storage_size	Specify the maximum amount of space (in bytes) that the LDAP cache can occupy on disk.
	When modifying this value, you must update the value of proxy.config.ldap.cache.size proportionally. For example, if you double the storage size, also double the cache size.
	Modifying this variable without modifying proxy.config.ldap.cache.size causes the LDAP subsystem to stop functioning.

- 3) Save and close the file.
- 4) From the Content Gateway **bin** directory (/opt/WCG/bin), run **content_line** -L to restart the proxy on the local node or **content_line** -M to restart the proxy on all the nodes in a cluster.

Configuring secure LDAP

By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology. You can enable LDAP over SSL (LDAPS) by installing a

properly formatted certificate from either a Microsoft certification authority (CA) or a non-Microsoft CA.

To use LDAPS with Content Gateway:

Steps

1) Open the records.config file located in /opt/WCG/config.

2) Add following entry to records.config:

CONFIG proxy.config.ldap.secure.bind.enabled INT 1

3) Navigate to Configure > Security > Access Control > LDAP and change the port to 3269.

_	

Note

The Directory Service must be configured to support LDAPS authentication. See to the documentation provided by the directory provider for instructions.

RADIUS authentication

Content Gateway supports the RADIUS option to ensure that users are authenticated with a RADIUS server before accessing content through the proxy.

When the RADIUS option is enabled:

- Content Gateway acts as a RADIUS client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the RADIUS server to check that the credentials are correct.
- If the RADIUS server accepts the username and password, the proxy serves the client with the requested content and stores the username and password entry in the RADIUS cache; all future authentication requests for that user are served from the RADIUS cache until the entry expires.
- If the RADIUS server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

Content Gateway supports a primary RADIUS server and a secondary RADIUS server for failover. If the primary server does not respond to the proxy request within the specified timeout (60 seconds by default), Content Gateway tries to check the username and password again. If a response from the primary RADIUS server is not received after the maximum number of retries (10 by default), the proxy contacts the secondary RADIUS server. If Content Gateway cannot contact the secondary RADIUS server, the user is prompted again for a username and password.

The RADIUS cache is held in memory and stored on disk. Content Gateway updates the data on disk every 60 seconds. In addition, Content Gateway stores username and password entries in the RADIUS cache for 60 minutes. If a password and username entry is expired in the RADIUS cache, Content Gateway contacts the RADIUS server to accept or reject the username and password.

To configure Content Gateway to be a RADIUS client:

- Enable the RADIUS option.
- Specify the hostname or IP address of the primary and secondary (optional) RADIUS servers, and the port and shared key that Content Gateway uses to communicate with the RADIUS servers.

See Configuring Content Gateway to be a RADIUS client.

Related tasks

Configuring Content Gateway to be a RADIUS client on page 185

Configuring Content Gateway to be a RADIUS client

Steps

- 1) Go to Configure > My Proxy > Basic > General.
- 2) In the Authentication section, click Radius On, and then click Apply.
- 3) Navigate to Configure > Security > Access Control > Radius.
- 4) Enter the hostname of your primary RADIUS server.
- 5) Enter the port number through which Content Gateway communicates with the primary RADIUS server.
- 6) Enter the key used for encoding.
- 7) If you are using a secondary RADIUS server, enter the hostname, port, and shared key in the appropriate fields of the **Secondary Radius Server (Optional)** area.
- 8) Click Apply.
- 9) Click Restart on Configure > My Proxy > Basic > General.



Note

In addition to performing these procedures, you must add the Content Gateway machine as a trusted client on the primary and secondary RADIUS servers and provide the shared key you want to use for the Content Gateway machine (the shared key must be the same one you specify in the procedure below). See your RADIUS server documentation.

Setting RADIUS cache and server timeout options

By default, the RADIUS cache and RADIUS server timeout options are configured as follows:

- The RADIUS cache is configured to store 1,000 entries and each entry is considered fresh for 60 minutes.
- Content Gateway can try to re-establish a connection to the RADIUS server if the connection remains idle for 10 seconds and can retry the connection a maximum of 10 times.

Change these default values by editing the records.config file.

Steps

1) Open the records.config file located in /opt/WCG/config.

2) Edit the following variables:

Variable	Description
proxy.config.radius.auth. min_timeout	Specify the amount of time in seconds that the Content Gateway connection to the RADIUS server remains idle before Content Gateway closes the connection.
proxy.config.radius.auth. max_retries	Specify the maximum number of times Content Gateway tries to connect to the RADIUS server.
proxy.config.radius.cache.size	Specify the number of entries allowed in the RADIUS cache.
	The minimum value is 256 entries. If you enter a value lower than 256, Content Gateway signals a SEGV.
proxy.config.radius.auth.ttl_value	Specify the number of minutes that Content Gateway can store username and password entries in the RADIUS cache.
proxy.config.radius.cache. storage_size	Specify the maximum amount of space that the RADIUS cache can occupy on disk.
	This value must be at least 100 times the number of entries. It is recommended that you provide the maximum amount of disk space possible.

- 3) Save and close the file.
- 4) From the Content Gateway **bin** directory (/opt/WCG/bin), run **content_line** -L to restart Content Gateway on the local node or **content_line** -M to restart WCG on all the nodes in a cluster.

Rule-Based Authentication

Using an ordered list of authentication rules, rule-based authentication provides support for multiple realm, multiple domain, and other special authentication requirements. When a request is processed, the rule list is traversed top to bottom, and the first match is applied.

Authentication rules specify:

- 1) How to match a user. By:
 - IP address
 - Inbound proxy port (explicit proxy only)
 - User-Agent value
 - A combination of the above
- 2) The domain or ordered list of domains to authenticate against. With a list of domains, the first successful authentication is cached and used in subsequent authentications. If IP address caching is configured, the IP address is cached. If Cookie Mode is configured, the cookie (user) is cached.

3) Whether a customizable web portal page should be used for authentication.

In rule-based authentication, only the first matching rule is tried. If authentication is unsuccessful, no further authentication is attempted.

Rule-based authentication is designed to meet these special requirements:

- Multiple realm networks: Rule-based authentication supports multiple realm networks in which domains do not share trust relationships and therefore require that each domain's members be authenticated by a domain controller within their domain. In this environment rules are created that specify:
 - Members of the realm (untrusted domain) by IP address or proxy port
 - The realm (domain) they belong to
- Authentication when domain membership is unknown: Some organizations do not always know what domain a user belongs to. For example, this can happen when organizations acquire new businesses and directory services are not mapped or consolidated. The unknown domain membership problem can be handled in rule-based authentication by creating a rule for IP address lists or ranges that specifies an ordered list of domains to attempt to authenticate against. The first successful authentication is remembered and used in later authentications. If authentication is not successful or the browser times out, no authentication is performed.
- Authentication based on User-Agent value: One or more User-Agent value can be specified in an authentication rule. Often this is a list of browsers. When the User-Agent value matches a rule, authentication is performed against the specified domain(s). If the User-Agent value doesn't match any rule and no rule matches based on other values, no authentication is performed (this is always true in rule-based authentication; if no rule matches, no authentication is performed).

For use case examples see Rule-based authentication use cases.



Note

If all the users in your network can be authenticated by domain controllers that share trust relationships, you probably don't need rule-based authentication.

However, the option is well suited to single domain environments that may benefit from multiple rules based on IP addresses, inbound proxy port (explicit proxy), and/ or User-Agent values.

Related information

Rule-based authentication use cases on page 200

Rule-based authentication structure and logic

Structure

A list of domains is created and maintained.
 When a domain is added to the list, the authentication method is specified: IWA, Legacy NTLM, or LDAP.
 RADIUS is not supported.

Only domains on the domain list can be specified in authentication rules.

The domain list is created and maintained on the **Configure > Security > Access Control > Domains** tab. The domain list is stored in the **auth_domains.config** file.

Authentication rules identify users (clients) by IP address, inbound proxy port (explicit proxy only), and/or User-Agent values, and attempt to authenticate the user against a specified domain or list of domains. Authentication rules are defined on the Configure > Security > Access Control > Authentication Rules tab. Rules are stored in the auth_rules.config file.



Note

Credential caching configuration is performed on the **Configure > Security > Access Control** > **Global Configuration Options** tab. On that page you specify IP address caching, cookie caching, or both. The setting applies to both transparent proxy and explicit proxy traffic. When both IP address caching and cookie caching are specified, the IP addresses that cookie caching is applied to must be specified.

See Credential Caching for more information.

Related concepts

Credential Caching on page 169

Logic

- One or more rules are defined for clients and domains (Configure > Security > Access Control > Authentication Rules).
- When a request for web content is received:
 - A top-down rule list traversal begins
 - The first match is applied
 - If the rule includes a list of domains, authentication proceeds as follows:
 - The proxy attempts to authenticate with the first domain using the method configured for that domain.
 For example, if the first domain is IWA, Content Gateway transparently negotiates with the browser for credentials (407 or 401).
 - If authentication fails and Content Gateway hasn't already challenged (prompted) for credentials, it then prompts for credentials.

Exception: When Content Gateway is an explicit proxy, the first and second domains are IWA, and the client has a ticket from the authentication domain, there is no prompt for basic credentials. Instead, Content Gateway uses the Kerberos ticket provided by the client to attempt to authenticate with the second domain. If the attempt fails and the fallback to NTLM authentication fails, the user is prompted for credentials.

When Content Gateway is a transparent proxy the standard behavior applies. This is because when the user is not a member of the first domain, the request for a Kerberos ticket fails because the client does not trust the FQDN sent with the request. The fallback to NTLM authentication also fails and the user is prompted for credentials.

- Content Gateway then uses the basic credentials with each domain, starting with the second, proceeding sequentially until authentication succeeds or the list is exhausted.
- Content Gateway then uses the basic credentials to attempt, again, to authenticate with the first domain.
- If authentication fails with all domains and the Fail Open (Configure > Security > Access Control > Global Configuration Options) setting is:

Enabled only for critical service failures, the proxy assumes that the user mis-entered their credentials, prompts again for basic credentials, and attempts, again, to authenticate sequentially against the list.

Enabled for all authentication failures, including incorrect password, fail open is applied.

- If no rule matches, no authentication is attempted
- Transactions are logged with the user name used by Filtering Service.

 Proxy authentication statistics are collected and reported individually for each authentication method. See Security (in the Statistics section).



Note

Content Gateway must be configured with a DNS server that can resolve the fully qualified domain name (FQDN) of Content Gateway for every realm used by IWA. If this isn't done, IWA fails to work. How to configure the DNS server is up to the network administrator. One option is to configure a DNS transfer zone (Sub Zone) between the primary DNS server of Content Gateway and the DNS server of each authentication realm (isolated domain).

Related concepts Security on page 251

Rule-based authentication configuration summary

Steps

 If Content Gateway is an explicit proxy and you want to bring traffic in on multiple ports, specify the ports on the Configure > Protocol > HTTP tab.



Important

You must also configure your clients to use the correct port.

- Configure Global authentication options (Configure > Security > Access Control > Global Authentication Options).
- 3) Create a domain list (Configure > Security > Access Control > Domains).
 - To specify a domain in a rule, it must be a member of the **Domain List**.
 - Active Directory domains used with IWA must be joined.

Handling of unknown users:

In rule-based authentication, Content Gateway may authenticate users that are outside the User Service primary domain. In these cases, Content Gateway can be configured to send an "alias" user name that User Service knows about. Or, you can send no name, in which case standard Filtering Service precedence is applied to determine the correct policy. (See Enforcement order in Administrator Help for the Web module.) This specification is made for each domain in the Domain list.

For more information, see Unknown users and the 'alias' option, below.

- 4) Create authentication rules (Configure > Security > Access Control > Authentication Rules).
- 5) Restart Content Gateway to make the new rules take effect.

Related concepts

Global authentication options on page 168 Unknown users and the 'alias' option on page 190

Rule-based authentication best practices

- If you don't need rules, don't use rule-based authentication. Deploying a single authentication method should provide the best performance.
- Use the fewest number of rules needed to satisfy your requirements.
- Do not use a domain list in a rule if it's not needed.

When a domain list is used

- If there is an IWA or NTLM domain, make it first in the list.
- If there is more than one IWA or NTLM domain, place the domain with the most active members first in the list. In other words, make the first domain the one that will most often authenticate users.
- Note that if an IWA domain is first in the list and the user is not joined to that domain, the user will be prompted for credentials.
- Note that if the first domain in the list is LDAP, every user who matches the rule will be prompted for credentials. The credentials provided will be offered to each successive domain.
- If client certificate authentication is enabled with Use the next selected authentication method if Client Certificate authentication fails option selected, the domain list cannot be empty.

Unknown users and the 'alias' option

In rule-based authentication it's possible for Content Gateway to authenticate a user who is not recognized by User Service because the name is not in the User Service directory.

When an authenticated user name is not found by User Service, standard Filtering Service precedence is used to determine correct policy. There are several ways to address this:

- Change the User Services configuration so that it can discover and add the names to its directory.
- Add the unrecognized names to the primary domain. The names must match exactly. Define policies for the new names.
- For users who match a particular authentication rule, pass an alias name and add the alias name to the primary domain. The names must match exactly. Define a policy for the alias name.
- Do nothing, or select to use a blank (empty) alias. This causes standard Filtering Service precedence to be applied to determine the correct policy. See Enforcement order in Administrator Help for the Web module.

For some illustrative use cases, see Rule-based authentication use cases.

Related information

Rule-based authentication use cases on page 200

Rule-based authentication Domain list

To use rule-based authentication, you create and maintain a **Domain List**. There must be at least one domain on the list before an authentication rule can be defined.

When a domain is added to the list, the authentication method is specified. When a rule is defined, the domain or domains are selected from the domain list. Supported domain types include:

- Active Directory (AD) domains to be used with IWA. These domains must be joined by Content Gateway, as well as by its members (users).
- Domain Controllers (DC) to be used with Legacy NTLM
- AD and uid domain controllers and directory servers to be used with LDAP

Domain specification configuration summary

Steps

- 1) Rule-based authentication must be enabled (Configure > My Proxy > General).
- 2) On Configure > Security > Access Control > Domains, click New Domain.
- 3) Select the authentication method.
- 4) Specify a unique name that will help you recognize the domain and its purpose.
- 5) Optionally, configure the Aliasing option.
- Specify the domain settings. These vary by authentication method. See:

Related tasks

Adding an Active Directory domain for use with IWA on page 191 Adding an NTLM domain controller for use with Legacy NTLM on page 193 Adding a domain (directory service) for use with LDAP on page 194

Adding an Active Directory domain for use with IWA

Active Directory (AD) domains to be used with IWA must be joined by both Content Gateway and directory members (clients).

If you are using IWA for the first time, see *Integrated Windows Authentication*, for a complete description of support and use.

To join a domain:

- Content Gateway must be able to resolve the domain name.
- Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
- The correct domain Administrator name and password must be specified.

- There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
- If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services, must be reachable by Content Gateway on the network.

To specify and join a domain:

Steps

- 1) Go to Configure > Security > Access Control > Domains and click New Domain.
- 2) Select Integrated Windows Authentication from the Authentication Method drop down box.
- 3) In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose.
- 4) Optionally, configure the Aliasing option. For information, see Unknown users and the 'alias' option.
- 5) In the **Domain Name** field, enter the fully qualified domain name. For example, ad1.example.com.
- 6) In the Administrator Name field enter the Windows Administrator user name.
- 7) In the Administrator Password field enter the Windows Administrator password. The name and password are used only during the join and are not stored.
- 8) Select how to locate the domain controller:
 - Auto-detect using DNS
 - DC name or IP address

If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.

9) Confirm the Content Gateway Hostname.



Warning

Do not change the hostname after the domain is joined. If it is changed, IWA immediately stops working and will not work again until the domain is unjoined and then re-joined with the new hostname.

10) Click Join Domain.

The Joined Domain Connections section of the Monitor > Security > Integrated Windows Authentication page displays a list of joined domains and connections, and provides a diagnostic test function.

For troubleshooting tips, see Failure to join the domain.

Related concepts

Integrated Windows Authentication on page 173 Unknown users and the 'alias' option on page 190 Failure to join the domain on page 177

To change the way the domain controller is found, and other attributes

Steps

- 1) On the Domains page, in the list select the domain you want to change and click Edit.
- 2) In the IWA Domain Details section, select how to locate the domain controller:
 - Auto-detect using DNS
 - DC name or IP address
 If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.
- 3) You can also change the Aliasing setting. See Unknown users and the 'alias' option.
- 4) Click Apply.

Related concepts

Unknown users and the 'alias' option on page 190

Adding an NTLM domain controller for use with Legacy NTLM

Support for Legacy NTLM has these restrictions:

- WINS resolution is not supported. Domain controllers must have hostnames that can be resolved by a DNS server.
- **Extended security** is not supported and cannot be enabled on the domain controller.
- NTLM2 session security is not supported and cannot be enabled on clients. In the Security Settings area of the Windows operating system, inspect the Network Security: Minimum session security settings.
- **NTLMv2** is not supported with Active Directory 2008.
- Not all browsers support transparent NTLM authentication. See Browser limitations.

For a complete description of support for Legacy NTLM, see Legacy NTLM authentication.

To add an NTLM domain for use in rule-based authentication:

Steps

- 1) Go to Configure > Security > Access Control > Domains and click New Domain.
- 2) Select Legacy NTLM from the Authentication Method drop down box.
- 3) In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose. After the domain is added, the name cannot be changed.
- 4) Optionally, configure the **Aliasing** option. For information see: *Unknown users and the 'alias' option*.

5) In the Legacy NTLM Domain Details section:

a) In the Domain Controller entry field enter the IP address and port number of the primary domain controller. If no port is specified, Content Gateway uses port 139.
 You can also specify secondary domain controllers in a comma-separated list. The supported formats are:

host_name[:port][%netbios_name]

IP_address[:port][%netbios_name]

The netbios_name is required with Active Directory 2008.

b) Specify whether load balancing should be applied among multiple DCs.



Note

Even if load balancing is **not** selected, if multiple domain controllers are specified and the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term fail over provision, until such time that the primary domain controller can accept new connections.

6) Click Add Domain.

Related concepts

Legacy NTLM authentication on page 179 Unknown users and the 'alias' option on page 190

Related reference

Browser limitations on page 167

Adding a domain (directory service) for use with LDAP

When LDAP is used:

- Content Gateway acts as an LDAP client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the LDAP server to check that the credentials are correct.
- If the LDAP server accepts the username and password, the proxy serves the client the requested content and stores the username and password in the credential cache.
- Future authentication requests for that user are served from the cache until the cache entry expires (Time-To-Live value).
- If the LDAP server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

LDAP authentication supports both simple and anonymous bind.

To add an LDAP domain to the Domains list:

Steps

- 1) Go to Configure > Security > Access Control > Domains and click New Domain.
- 2) Select LDAP from the Authentication Method drop down list.
- 3) In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose. After the domain is added, the name cannot be changed.
- 4) Optionally, configure the Aliasing option. For information see: Unknown users and the 'alias' option.
- 5) In the LDAP Domain Details section:
 - a) In the LDAP Server Name field, enter the fully qualified domain name or IP address of the LDAP server.
 - b) If the LDAP server port is other than the default (389), in the LDAP Server Port field, enter the LDAP server port.
 - c) Enter the LDAP Base Distinguished Name. Obtain this value from your LDAP administrator.
 - d) Select the LDAP Server Type from the drop down list.
 - Select sAMAccountName (MS AD) for Active Directory.
 - Select userPrincipalName (MS AD) for Active Directory.
 - Select uid (Other LDAP) for other directory services.
 - e) In the Bind Domain Name field, enter the bind distinguished name. This must be a Full Distinguished Name of a user in the LDAP directory service. For example: CN=John Smith,CN=USERS,DC=MYCOMPANY,DC=COM
 - f) In the Bind Password field, enter the password for the name given in the Bind Domain Name field.
 - g) Enable Secure LDAP if you want Content Gateway to use secure communication with the LDAP server. If enabled, set the LDAP port to 636 or 3269.
- 6) Click Add Domain.

Related concepts

Unknown users and the 'alias' option on page 190

To unjoin or remove a domain from the Domain List

On the **Domains** page, select the domain from the list and click **Unjoin** or **Delete**.

A confirmation dialogue displays. Confirm that you want to remove the domain from the list.

Warning

When a domain is removed, it is also removed from any authentication rules that specify it.

If it is the only domain specified in a rule, when the domain is removed the rule is made invalid and, therefore, the rule is removed.

Creating an authentication rule

Before you create an authentication rule you must:

- Enable Rule-Based Authentication on Configure > My Proxy > Basic > General.
- Configure Global authentication options
- Create a Rule-based authentication Domain list

You must also know:

- The name of the domain(s) to be specified in the rule. This is the unique name that was specified when the domain was added to the Domains list.
- How to match users. By:
 - IP address individual addresses or address ranges can be specified
 - Inbound proxy port (explicit proxy only)
 - User-Agent values
 - A combination of the above

To create a rule:

Note

In the Rule editor, after entering all specifiers, click **Add** before clicking **Apply**. If Apply is clicked first, or the edit window is closed, all entry fields are cleared.

The size of a rule cannot exceed 2048 characters.

Steps

- Go to Configure > Security > Access Control and review and adjust the Global Authentication Options and Domains list.
- 2) If AD domains are used with IWA, go to Monitor > Security > Integrated Windows Authentication and confirm that the IWA domains are joined and that connections are established.
- 3) Go to Configure > Security > Access Control > Authentication Rules. A list of existing authentication rules is displayed at the top of the page.
- 4) Click Edit File to open the rule editor.
- 5) If some rules have already been defined, note the order of the rules in the list at the top of the page.



Important

Rule order matters. The rule match traversal is performed top-to-bottom. Only the first match is applied.

- 6) Select **Enabled** next to Status if you want the rule to be active after the rule is added and Content Gateway is restarted.
- 7) Enter a unique **Rule Name** (required). A short, descriptive name will help you recognize the rule and its purpose. It is recommended that the name not exceed 50 characters.

8) If the rule applies to specific IP addresses, in the Source IP field, enter a comma- separated list of individual IP addresses and/or IP address ranges. Do not use spaces. For example:

10.4.1.1,10.12.1.1-10.12.254.254

The list can contain up to:

- 64 IPv4 addresses
- 32 IPv4 address ranges
- 24 IPv6 addresses
- 12 IPv6 address ranges

Source IP address ranges can overlap. Overlapping ranges may be useful as a quick way of identifying sub-groups in a large pool. In overlapping ranges, the first match is used.

If this field is empty (undefined), all IP addresses match.

9) If the rule applies to inbound traffic on a specific port, select the Proxy Port from the drop down list. This option is valid with explicit proxy only.

Inbound ports are specified on the **Configure > My Proxy > Protocols > HTTP > General** page in the **Secondary HTTP Proxy Server Ports** field. Client applications must be configured to send requests to the desired port.

If undefined, all ports match. Transparent proxy deployments should leave the field undefined.

10) To apply the rule to specific User-Agent values, enter POSIX-compliant regular expressions (regex) to match the desired values. To specify a common browser type, select a Predefined regex from the drop down list and click Include.

If undefined, all User-Agents match.

You can edit the field directly.

Use the "|" character (logical 'or') to separate regexes. The "^" regex operator is not supported.

The regex is validated when the rule is committed to the configuration file, which happens after clicking **Add** or **Set** and then **Apply**. **If the regex is not valid, the rule is deleted and must be recreated with a valid regex.**

For an extended description and examples, see Authentication based on User-Agent.

- Click Enabled next to Client Certificate to enable client certificate authentication. Click Disabled to disable the feature.
 - a) In the drop-down box next to Enabled, select a Client Certificate Authentication profile. See Client certificate authentication profiles.
 Only one profile is allowed.
 - b) Check the box next to Use the next selected authentication method if Client Certificate authentication fails to use one of the other authentication methods if certificate authentication fails for a user.

If this option is not selected, no further authentication is attempted for users who fail certificate authentication.

If the fallback option is enabled,

- The **Domain Sequence** list cannot be empty.
- Enable HTTP Authentication Page for Captive Portal is not supported and the option is disabled when the fallback option is selected.
- 12) Specify the domain(s) to authenticate against.
 - a) From the Domains drop down list, select the applicable domain and click Include. Only domains that have been added to the Domains list are available (Configure > Security > Access Control > Domains).
 - b) If an ordered list of domains will be used, select each domain one at a time and click **Include**. Then select domains in the list and use the up and down arrows to achieve the desired order.



Important

The Fail Open/fail closed setting is applied after every domain in the list is tried.

- 13) Next to Captive Portal, click:
 - Enabled for HTTPS Authentication page to redirect users to a customizable web portal page for authentication.

When this selection is enabled, the page will display using HTTPS.

When HTTPS is used, a server certification is generated based on the internal root CA. To use this feature, you must import the internal root CA to ensure there is no certificate error. See *Importing your Root CA* for details.

 Enabled for HTTP Authentication page to redirect users to a customizable web portal page for authentication.

With this selection, the page is displayed using the HTTP protocol.

Note that if client certificate authentication is enabled with **Use the next selected authentication method if Client Certificate authentication fails** option selected, this option is disabled.

This option is disabled if an IWA domain is included in the domains list.

If this option is enabled and an IWA domain is added to the domains list, an error message will display.

Note that when Content Gateway receives an unauthenticated POST request from a user who matches a Captive Portal rule, it redirects the user to the web portal authentication page and does not record the POST data. After successful authentication, the original POST data must be input again.

See Authentication using Captive Portal for additional details.

- 14) Click Add to add the rule.
- 15) At the top of the page, check and adjust the position of the rule in the rule list. The first rule matched is applied.
- 16) Click Apply and then restart Content Gateway to put the rule into effect.



Warning

If a rule has invalid values, a warning message displays that identifies the invalid rule. The rule in not written to the file.

Related concepts

Global authentication options on page 168 Rule-based authentication Domain list on page 191 Authentication based on User-Agent on page 203 Fail Open on page 168 Authentication using Captive Portal on page 205

Related tasks

Client certificate authentication profiles on page 208 Importing your Root CA on page 116

Working with existing authentication rules

Use the rule editor in the Content Gateway manager. Do not directly edit auth_rules.config.

Editing a rule

Steps

- 1) Go to Configure > Security > Access Control > Authentication Rules and click Edit File.
- 2) In the table of rules, click on the rule to be changed. Its values populate the fields in the definition area.
- 3) Make the desired changes, click Set and then click Apply.



Important

If a field value is not valid, the rule is not commit ed and the rule entry is discarded. To avoid difficulty in recreating a rule, separately record the field values so that it is easy to correct the bad field value and recreate the rule.

- Click Close to return to the Authentication Rules tab and click Refresh to see the updated list.
- Restart Content Gateway to put the changes into effect.

Reordering the list of rules

Authentication rules are matched top-down in the list. Only the first match is applied.

Steps

- 1) Go to Configure > Security > Access Control > Authentication Rules and click Edit File.
- In the table of rules, click on the rule that you want to reposition and then click the down or up arrow on the left to reposition the rule.
- 3) When the rules are in the desired order, click Apply.
- 4) Click Close to return to the Authentication Rules tab and click Refresh to see the updated list.
- 5) Restart Content Gateway to put the changes into effect.

Deleting a rule

Steps

- 1) Go to Configure > Security > Access Control > Authentication Rules and click Edit File.
- 2) In the table of rules, click on the rule to be deleted and click the "X" button on the left.
- 3) When you are done deleting rules, click **Apply**.
- 4) Click Close to return to the Authentication Rules tab and click Refresh to see the updated list.
- 5) Restart Content Gateway to put the changes into effect.

Rule-based authentication use cases

Related concepts

Authentication based on User-Agent on page 203

Related tasks

Multiple realm use case 1: Domain acquired; explicit proxy on page 201 Multiple realm use case 2: Internal domain added; explicit proxy on page 201 Multiple realm use case 3: Temporary domain added; transparent proxy on page 202

Multiple realm use case 1: Domain acquired; explicit proxy

This describes a common case in which a second domain is added to an existing, single-domain environment.Content Gateway is an explicit proxy; clients use a PAC file.

An organization—let's call them Quality Corp—uses a software installation of Content Gateway. They have one domain (QCORP), and one domain controller. They use NTLM to authenticate users.

Quality Corp acquires New Corp who has their own domain (NCORP) and domain controller. They use LDAP to authenticate users.

Quality Corp would like to manage the combined employees in a single domain, but they aren't ready to make the infrastructure changes. Until they are, they would like to

have a separate use policy for New Corp users (i.e., not use the "default" user on the QCORP domain).

Rule-based authentication makes this possible.

To configure the solution, Quality Corp would:

Steps

- 1) Enable Rule-Based Authentication.
- Add a second, non-default HTTP port (Configure > Protocols > HTTP > General). This port will be used by all members of NCORP.
- 3) Create a PAC file for members of NCORP that causes them to connect to Content Gateway on the new, second port.
- 4) Create authentication rules, one each for the QCORP and NCORP domains:
 - a) On Configure > Security > Access Control > Domains, add the QCORP and NCORP domains to the Domains list.
 When adding NCORP, use the Aliasing option to specify "NCorpUser" for use in policy determination.
 - b) On Configure > Security > Access Control > Authentication Rules, create an NCORP rule for connections on the second port. You must know the IP addresses/ranges of New Corp users, and specify the NCORP domain.
 - c) Define the QCORP rule to handle all other connections.
- 5) In the Web module of the Forcepoint Security Manager, add "NCorpUser" to the QCORP domain as a valid user and create policy for that user.

At this point, everyone connecting to Content Gateway from NCORP is authenticated against the NCORP domain controller and gets the group policy associated with NCorpUser. Note that no individual user-based policy or features, such as quota time, are possible in this scenario. Transactions are logged as NCorpUser. This is all performed with no effect on the authentication, policy, or logging of users on the QCORP domain.

Multiple realm use case 2: Internal domain added; explicit proxy

This describes a common case in which a second domain is added to an existing, single-domain environment. Content Gateway is an explicit proxy; clients use a PAC file.

An organization—let's call it BigStars—uses a software installation of Content Gateway. They have one domain (BIG), and one domain controller. They use NTLM to authenticate users.

A group in the company converts to Apple computers, which can't be authenticated with NTLM. The IT group installs an LDAP server and creates a new domain— BIGAPL—for the Apple users.

Because this group of users previously existed and was managed on the primary domain (BIG), the IT department expects that both user-based policy and logging still apply.

The Rule-Based Authentication feature makes this possible.

To configure the solution, BigStars would:

Steps

- 1) Verify that every user in BIGAPL is also in BIG with the exact same user name.
- 2) Enable Rule-Based Authentication.
- Add a second, non-default HTTP port (Configure > Protocols > HTTP). This port will be used by all members of BIGAPL.
- Create a PAC file for members of BIGAPL that causes them to connect to Content Gateway on the new, second port.
- 5) Create authentication rules, one each for the BIGAPL and BIG domains.
 - a) On Configure > Security > Access Control > Domains, add the BIGAPL and BIG domains to the Domains list.
 - b) On Configure > Security > Access Control > Authentication Rules, create a BIGAPL rule for connections on the second port.
 - c) Define the BIG rule to handle all other connections.

At this point, all members of BIGAPL are authenticated with LDAP, but maintain their individual policy as specified by their existing NTLM identities. Logs and reports also refer to that same user.

Multiple realm use case 3: Temporary domain added; transparent proxy

This describes a common case in which a second, special-purpose domain is added to an existing, single-domain environment. Content Gateway is a transparent proxy using WCCP v2.

An organization—let's call it Creative Corp—uses a software installation of Content Gateway. They have one domain (CCORP), and one domain controller. They use NTLM to authenticate users.

Creative Corp is about to launch a new product and wants to make a big splash. They decide to have an open house complete with kiosks, demonstrations, and presenters. The kiosks only need the default Internet policy to properly demonstrate the new product. The IT manager wants to keep the kiosk network as walled off from the corporate intranet as possible. In this scenario, logging individual users isn't a requirement.

The Rule-Based Authentication feature makes this possible.

To configure the solution, Creative Corp would:

Steps

- 1) Build a new, temporary network complete with its own domain controller. Let's call this domain CTEMP.
- 2) Add one or more users to CTEMP. They can either match one-to-one with existing users on the primary domain, or be one or more generic users for use by the presenters.
- 3) Redirect Internet traffic on CTEMP to Content Gateway with WCCP v2.
- 4) Enable Rule-Based Authentication.
- 5) Create authentication rules, one each for the CTEMP and CCORP domains:
 - a) On Configure > Security > Access Control > Domains, add the CTEMP domain, enable Aliasing and leave the name field blank. This will have the result of applying the Default policy to all users of CTEMP.
 - b) Add the CCORP domain to the Domains list.
 - c) On Configure > Security > Access Control > Authentication Rules, create a CTEMP rule to apply to all connections coming from the IP address range assigned to the CTEMP domain.
 - d) Define the CCORP rule to handle all other connections.

At this point, anyone using the Internet on one of the kiosks is authenticated against the CTEMP network and has the Default policy applied to their requests.

Authentication based on User-Agent

In an authentication rule, a Request header User-Agent value can be used to determine if user authentication will be performed. This is useful when you want to authenticate users using a known set of client applications, usually browsers, and allow other applications, often a set of applications that don't support authentication, to proceed without authentication. Such rules can also specify IP addresses and, if Content Gateway is an explicit proxy, inbound proxy port.

As with all authentication rules, the first matching rule is applied. (For a complete description of rule-based authentication, see *Rule-Based Authentication*.)

When the User-Agent field is used, the critical element is the regular expression (regex) that preforms the match.

- The regex must be POSIX-compliant. The "^" regex operator is not supported.
- Predefined regexes are provided for the most common browsers.
- When the field is empty, all User-Agent values match.
- You can create a custom regex by directly editing the field.
- Multiple regexes are allowed. They must be separated by a "|" ('or' operator).

When you click **Apply** (after Add or Set), the regex is parsed and validated. **If the regex is not valid, the rule is deleted and must be recreated with a valid regex.**

Following are a few examples of custom regexes. Microsoft Internet Explorer 7, 8, or 9:

MSIE $([7-9]{1}[\.0-9]{0}]$

Example User-Agent string:

Mozilla/5.0 (Windows; U; MSIE 9.0; Windows NT 9.0; en-US)

Microsoft Edge

Edge $([1]{1}[\.0-9]{0})$

Example User-Agent string:

Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/

537.36 (KHTML, like Gecko) Chrome/46.0.2486.0 Safari/

537.36 Edge/13.10586

Microsoft Internet Explorer Mobile, all versions:

IEMobile

Example User-Agent string:

Mozilla/5.0 (compatible; MSIE 9.0; Windows Phone OS 7.5; Trident/5.0; IEMobile/9.0)

Apple iPhone, all versions:

(iPhone) OS $(d+)_(d+)$?

Example User-Agent string:

Mozilla/5.0 (iPod; U; CPU iPhone OS 4_3_3 like Mac OS X; ja-jp) AppleWebKit/533.17.9 (KHTML, like Gecko) Version/

5.0.2 Mobile/8J2 Safari/6533.18.5

Apple iPad, all versions:

 $(iPad).+ OS (\d+)_(\d+)(?:_(\d+))?$

Example User-Agent string:

Mozilla/5.0 (iPad; CPU OS 6_0 like Mac OS X) AppleWebKit/

536.26 (KHTML, like Gecko) Version/6.0 Mobile/10A5355d Safari/8536.25

Search the Internet for lists of User-Agent strings, example regular expressions, regex checkers, and related resources.

Related concepts

Rule-Based Authentication on page 186

Use case

This describes a case in which an organization with a single domain wants to authenticate requests from 2 common web browsers. They also want to bypass authentication for web applications that do not support authentication.

An organization—let's call it Best Corp—uses Content Gateway. They have one domain (BCORP), and one domain controller. They use IWA to authenticate users.

Best Corp wants to ensure that:

- Requests from common web browsers are authenticated. They control which web browsers are allowed on their computers.
- Web applications that don't support authentication bypass authentication. The User-Agent feature of rulebased authentication makes this possible.

To configure the solution, Best Corp:

Steps

- 1) Enables Rule-Based Authentication.
- 2) Adds the BCORP domain to the Domains list.
- 3) Creates an IWA rule that:
 - a) Optionally, specifies the supported client IP address ranges.
 - b) Specifies, by User-Agent value, the web browsers to authenticate. In the User-Agent field, they use the Predefined drop down list to select and Add Internet Explorer and Firefox. The regex looks like:

MSIE.*|Firefox.*

That's it. With this configuration, all requests from Internet Explorer and Firefox, the only 2 browsers that can be installed on their computers, are subject to user authentication. All other requests, most particularly web applications, bypass authentication. To further customize the approach, Best Corp could create other authentication rules and/or add proxy filtering rules (filter.config) to deny or bypass specific applications by User-Agent value.

Authentication using Captive Portal

Content Gateway provides a Captive Portal option when adding an authentication rule. Captive Portal may be especially helpful in handling mobile and other personal devices brought in to your Forcepoint Web Security networks.

This feature:

- Redirects users to a web portal page for authentication.
- Supports captive, interactive (prompted) user authentication of IP addresses (users) that match the Captive Portal rule.
- Can be used with LDAP, Legacy NTLM, and IWA; RADIUS is not supported.
- Handles credential caching and expiration per the global configuration; cookie authentication and caching are also supported.

Note that most applications on mobile devices do not share cookies. For those applications, IP-based identification will be required. See the Credential Cashing section of *Global authentication options* for more information.

Also, for web applications that use Ajax, where Ajax is configured to prevent cookies, cookie-mode cannot support sites that include cross-origin requests (CORS) that rely on Ajax.

- Allows the authentication form (web portal page) to be customized to suit your needs.
- Supports only basic authentication.
- Provides the option to display the authentication page using either HTTP or HTTPS.

When adding an authentication rule (see *Creating an authentication rule*), an option is provided. Navigate to **Configure > Security > Access Control > Authentication Rules** and click **Enabled for HTTPS/HTTP Authentication page** next to Captive Portal to select the feature. Users who match the rule are redirected to the web portal authentication page.

Note that when Content Gateway receives an unauthenticated POST request from a user who matches a Captive Portal rule, it redirects the user to the web portal authentication page and does not record the POST data. After successful authentication, the original POST data must be input again.



If the requested URL is configured for tunneling or bypass, no user authentication is performed.

When a rule is added with the Captive Portal option enabled, users are reminded that they can customize the pre-defined web portal page. Go to the new Captive Portal Page Customization tab of **Configure > Security > Access Control**. Edit the text and HTML to suit your needs. For example, you may want to include your company logo in place of the default logo.

Related concepts

Note

Global authentication options on page 168

Related tasks

Creating an authentication rule on page 196

Customizing the web portal page

The web portal page is an HTML form that is presented to the user for interactive authentication.

Default contents are provided on the Captive Portal Page Customization tab of **Configure > Security > Access Control**. It is recommended that you customize the form to convey to users who see it that this logon portal is part of your network and organization. For example, you might:

- Replace the default logo with your organization's logo. To do that:
 - Edit the src tag and replace the png file name with your company logo file.
 - Copy your png file to /opt/WCG/config/ui_files/images.
- Include text to explain why the user is seeing this page

The form must be a valid HTML document, defined with valid HTML syntax.

The following variables are used in the document to ensure that it is delivered to the users properly. It is recommended that you do not change their placement or usage.

- %P is replaced with the protocol of the current transaction
- %h is replaced with "redirect_host:8080"
- %u is replaced with the URL request for the portal page
- \$\$DOMAIN is replaced with the basic authentication domain defined in the configuration variable proxy.config.proxy.authenticate.basic.realm. (See Authentication basic realm for more information.)

When you have entered all of the syntax, click **Preview** to preview the page you have created. When you are happy with the way the portal page looks, click **Apply** to save the content to a file. If you want to return to the default, pre-defined portal page syntax, click **Restore to Default Page**.

The customized Captive Portal page is saved to **auth_form.html**, which is stored in / opt/WCG/config. In addition, css and image files can be used to define the portal page. CSS files must be stored in /opt/WCG/ config/ui_files and image files must be store in /opt/WCG/config/ui_files/images, by default.

Ę

Note

The css and image files also reside in /opt/WCG/ui/ configure/auth_form and /opt/WCG/ui/ configure/ auth_form/images, respectively, for use by the **Preview** feature. Copy any new files to those directories to use **Preview**.

Add a variable to records.config to use a different name for the saved Captive Portal page or store the css and image files in a different directory.

Configuration Variable	Data Type	Default Value	Description
proxy.config.auth. form_filename	STRING	auth_form. html	Specifies the file that defines the Captive Portal authentication page.
			Changing this filename is not recommended.
proxy.config.internal. file.path	STRING	/config/ ui_files	Specifies the location of any css and image files used to define the Captive Portal authentication page. The full default path is /opt/ WCG/config/ ui_files. Image files are located in an /images sub-directory.

Related reference

Authentication basic realm on page 404

Using a custom certificate with Captive Portal

Upload the custom certificate and its key to the Content Gateway machine and update the following variable in records.config to force a custom certificate to work with Captive Portal.

CONFIG proxy.config.ssl.use_custom_cert_for_captive_portal

On the Content Gateway machine.

Steps

- 1) Copy the custom certificate and its key to the /opt/WCG/config directory
- 2) Make a backup of the /opt/WCG/config/ssl_multicert.config file
- 3) Open /opt/WCG/config/ssl_multicert.config with a text editor and make and:
 - a) Comment out the default entry
 - b) Add a new entry specifying the custom certificate and its key. For example: name=auth_server dest_ip=* ssl_cert_name=custom_cert.crt ssl_key_name=custom_key.key

Note that ssl_multicert.config file can only have one entry. It does not support multiple entries.

4) Make a backup of records.config. (Default location: /opt/WCG/config)

- 5) Run the following command to check the current value of the parameter. By default, the value is 0. /opt/WCG/bin/content_line -r proxy.config.ssl.use_custom_cert_for_captive_portal
- 6) Run the following command to enable the feature and make the custom certificate work with captive portal. /opt/WCG/bin/content_line -s proxy.config.ssl.use_custom_cert_for_captive_portal -v 1

7) Restart Content Gateway.



Important

Enable this parameter only if you have a custom certificate to be uploaded to the proxy. If the parameter is enabled but no custom certificate is entered in ssl_multicert.confg, the proxy can run into unexpected issues.

Client certificate authentication

Certificate authentication is available for use with mobile and other personal devices.

When client certificate authentication is enabled, unauthenticated users are redirected to an HTTPS page where they are prompted to select the certificate to send to Content Gateway. The user is considered authenticated if the certificate is signed by a trusted Certificate Authority (CA). The user name is extracted from the appropriate certificate field.

Client certificate authentication can also be configured to fall back to the domains list and Captive Portal for authentication. Users who cannot be authenticated using a certificate will then be authenticated using a different method.

Used with rule-based authentication, this feature is configured for each proxy and:

- Supports basic, LDAP, NTLM, and IWA authentication.
 If the fallback option is enabled, however, and Captive Portal is enabled for fallback, the Captive Portal limitations apply. See Authentication using Captive Portal.
- Supports credential and cookie caching.
- Requires a Client Certification Authentication Profile that explains where to extract user names from the certificates and includes a list of the CA Certificates valid for use by clients.
- Requires enabling SSL decryption.
 Access to HTTPS sites are not authenticated if HTTPS is not enabled on the Configure > My Proxy > Basic page.

Related concepts Authentication using Captive Portal on page 205

Client certificate authentication profiles

When client certificate authentication is enabled, a client certificate authentication profile must be selected. Configure client certificate authentication profiles on the **Client Cert Auth Profile** tab of the **Configure > Security > Access Control** page.



Note

You can have only one profile.

On the Client Certificate Authentication Profile page:

Steps

- 1) Enter a Profile Name. This name will appear in the drop-down list on the Authentication Rules page.
- Select an entry from the User Name Mapping drop-down.
 Valid selections are Common Name (CN), Distinguished Name, or Email. This entry tells the authentication process how to extract the user name from the certificate.
- 3) In the Certificate Authorities section, add, view, or delete certificates. The certificates used for authentication are manipulated the same way that SSL certificates are manipulated on the Configure > SSL > Certificates pages. Refer to the Adding new certificate authorities for assistance.
- 4) Click Apply to save your profile.

Related tasks Adding new certificate authorities on page 124

Troubleshooting authentication rules

In rule-based authentication, problems often present as:

- Users are not challenged for credentials when a challenge is expected
- Users are challenged for credentials when no challenge is expected
- User authentication is performed against the wrong domain

These problems occur in one of the following phases of user authentication processing:

- General user authentication logic (outlined below)
- Rule definition and matching
- User authentication protocol processing (IWA, NTLM, LDAP; for IWA troubleshooting, see *Troubleshooting* Integrated Windows Authentication.)

Related concepts

Troubleshooting Integrated Windows Authentication on page 177

Rule-based authentication logic

Rule-based authentication applies the following logic:

- The rules in filter.config are checked and applied. This action occurs first in every type of Content Gateway user authentication. If a filtering rule is matched, the rule is applied and user authentication processing stops. See Content Gateway filtering rules.
- 2) If no filtering rule matches, user authentication rule matching is performed.
 - a) The requestor's IP address is checked, top-down, against the rule set.

- b) If the IP address matches a rule, the source port is checked.
- c) If the IP address matches a rule, the User-Agent value is checked.
- d) The first rule matched is applied. If no rule matches, no authentication is attempted.
- If a rule is matched, the specified authentication protocol is applied against the specified domain. All rule configuration details are applied.
- 4) If the user is authenticated, the request proceeds or is denied per the assigned policy.
- 5) The transaction is logged.

To see how the logic is applied in a running environment, you can temporarily enable user authentication debug output. Among other details, the debug output shows the parsing of rules and matching. See *Enabling and disabling user authentication debug output*.

Related concepts Content Gateway filtering rules on page 155 Enabling and disabling user authentication debug output on page 211

Troubleshooting

When rule-based authentication doesn't produce the expected results, it is recommended that you troubleshoot the problem in the following order:

Steps

1) Check Redirection Rules

Confirm that there is no unexpected entries. In the Content Gateway manager, go to **Configure > Networking > ARM > General** and examine the **Redirection Rules**.

2) Check the rules in filter.config

Confirm that there is no unexpected matching of a **filter.config** rule. Among other purposes, filter.config rules can be used to bypass user authentication. See *Content Gateway filtering rules*.

3) Check rule matching

Using the IP address of a user who is or is not being challenged as expected, walk through each rule, top to bottom, examining the settings to find the first match. Be meticulous in your analysis. A common problem is that the IP address falls within a too-broad IP address range.

If the rule uses an alias, confirm that the alias is present in the User Service of the primary domain controller.

For explicit clients configured to send traffic to a specific port, check both the rule and the configuration of the client's browser.

4) Check the domain

If you are getting the match you expect, verify that the domain is reachable and that the user is a member of the domain. If yes, troubleshoot the problem at the authentication protocol level. For IWA, see *Troubleshooting Integrated Windows Authentication*.

5) When Content Gateway is in a proxy chain

If Content Gateway is a member of a proxy chain, verify that X-Forwarded-For headers are sent by the downstream proxy and read by Content Gateway.

- Use a packet sniffer to inspect inbound packets from the downstream proxy. Look for properly formed X-Forwarded-For headers.
- In the Content Gateway manager, go to Configure > My Proxy > Basic, scroll to the bottom of the page and verify that Read authentication from child proxy is enabled. If it's not, select On, click Apply, and then restart Content Gateway.

Related concepts

Content Gateway filtering rules on page 155 Troubleshooting Integrated Windows Authentication on page 177

Enabling and disabling user authentication debug output



Warning

Debug output should not be left enabled. Debug output slows proxy performance and can fill the file system with log output.

Debug log information is written to: /opt/WCG/logs/content_gateway.out

To enable user authentication debug information, edit: /opt/WCG/config/ records.config

(root)# vi /opt/WCG/config/records.config

Find and modify the following parameters and assign values as shown:

CONFIG proxy.config.diags.debug.enabled INT 1

CONFIG proxy.config.diags.debug.tags STRING

http_xauth.* | auth_* | winauth.* | ldap.* | ntlm.*

Save and close the file. Force Content Gateway to reread the file with the command:

(root)# /opt/WCG/bin/content_line -x

Follow the flow of debug information with the **tail -f** command:

(root)# tail -f /opt/WCG/logs/content_gateway.out

Use Ctrl+C to terminate the command.

When you have collected the debug output you want (after one or several user authentication processes is complete), disable debug output by editing **records.config** and modifying the parameter value as shown.

(root)# CONFIG proxy.config.diags.debug.enabled INT 0

Save and close the file. Force Content Gateway to reread the file with the command:

(root)# /opt/WCG/bin/content_line -x

Mac and iPhone/iPad authentication

Forcepoint Web Security solutions can be used to authenticate or identify Mac and iPhone/iPad users for user- or group-based filtering.

For Mac computers, see:

- Authentication for Mac computers
 - Enabling transparent identification of Mac users with DC Agent
 - Authenticating Mac users with Content Gateway
 - Typical steps for joining a Mac to an Active Directory domain

For iPhones/iPads, see:

Authentication for iPhones and iPads

For a list of Frequently Asked Questions regarding Mac and iPhone/iPad authentication, see this article.

Related concepts

Authentication for Mac computers on page 212 Enabling transparent identification of Mac users with DC Agent on page 213 Authenticating Mac users with Content Gateway on page 213 Authentication for iPhones and iPads on page 215

Related tasks

Typical steps for joining a Mac to an Active Directory domain on page 214

Authentication for Mac computers

Forcepoint Web Security solutions can be used to authenticate or identify Mac users for user- or group-based filtering. These restriction apply:

- Authentication and identification require that users belong to an Active Directory.
- Protocol block messages cannot be displayed on Macs.

If your organization uses DC Agent for transparent user identification, see *Enabling transparent identification of Mac users with DC Agent*.

If your organization uses Logon Agent for transparent user identification, see Deploying the logon application for Mac clients.

If your organization uses Content Gateway to authenticate users, see Authenticating Mac users with Content Gateway.

Manual (prompted) authentication can also be used to enable user and group-based filtering of Mac users.

Related concepts

Enabling transparent identification of Mac users with DC Agent on page 213 Authenticating Mac users with Content Gateway on page 213

Enabling transparent identification of Mac users with DC Agent

In order for DC Agent to identify the user on a Mac workstation, the Mac must mount a file share on the domain controller. This can be done by configuring the Mac to use a file share on the domain controller machine as the user's home directory, or by mounting another share with the domain controller.



Note

If the Mac only logs to the domain without mounting a file share, it will not be visible to DC Agent.

Configuration summary:

- Ensure that each participating Mac user is a member of a common Active Directory. See your Active Directory documentation.
- Create a home folder for each Mac user, and make sure that it is accessible to the user. See the first paragraph of this section.

When the user logs on to the properly configured Mac OS X system, the Mac mounts a network directory as the user's home directory, the DC Agent user map is populated, and user and group-based policies can be applied to user requests. When requests are blocked, browser-based block pages are displayed normally.

Authenticating Mac users with Content Gateway

Using the Integrated Windows Authentication (IWA) feature of Content Gateway, Mac users can be transparently authenticated when the user is a member of an Active Directory domain and the Mac computer is joined to the Active Directory domain. For more information see *Integrated Windows Authentication*.

Configuration summary:

- Ensure that each Mac computer is joined to the Active Directory domain. See Typical steps for joining a Mac to an Active Directory domain.
- Ensure that each participating Mac user is a member of a common Active Directory. See your Active Directory documentation.
- Ensure that Content Gateway is joined to the Active Directory domain.
 - If Content Gateway is not configured for IWA, see Integrated Windows Authentication and apply the configuration instructions.
 - If Content Gateway is already configured for IWA and your Mac users belong to the currently joined domain, there is nothing to do.
 - If Content Gateway is already configured for IWA and your Mac users belong to a different Active Directory domain, use the Rule-Based Authentication feature. See *Rule-Based Authentication* and follow the configuration instructions.
- When Content Gateway is an explicit proxy, configure participating Mac systems and browsers to send HTTP, HTTPS, and FTP requests to the Fully Qualified Domain Name (FQDN) of Content Gateway. Alternatively, specify the IP address of Content Gateway if NTLM is adequate.

If Content Gateway is a transparent proxy, no additional Mac system or browser configuration is required.



Important

Safari users may be prompted for credentials the first time they open a browser. The user should enter their credentials and check the "Remember password in keychain" check box.

Firefox users may receive an "Proxy Authentication Required" error message. This is a known issue in FireFox (http://support.mozilla.org/en-US/questions/926378) and is easily corrected by changing the browser configuration. In **About:Config** set the following options to **false**:

- network.automatic-ntlm-auth.allow-proxies
- network.negotiate-auth.allow-proxies

Related concepts

Rule-Based Authentication on page 186 Integrated Windows Authentication on page 173

Related tasks Typical steps for joining a Mac to an Active Directory domain on page 214

Typical steps for joining a Mac to an Active Directory domain

Steps

- 1) Using an account with Administrator privileges, log on to the Mac computer that you want to join to an Active Directory domain.
- 2) Open the Directory Utility. On OS X 10.6 (Snow Leopard), go to: /System/Library/CoreServices
- 3) If necessary, click the padlock icon and enter your password to unlock the Directory Utility.
- 4) Select the box next to Active Directory to enable Active Directory support.

5) Highlight Active Directory and click on the Pencil icon to configure the Active Directory connection.

	Select a service and click	the pencil icon to edit settings.
Enable	Name	Version
⊻	Active Directory	6.0
\checkmark	BSD Flat File and NIS	6.0
\checkmark	LDAPv3	6.0
\checkmark	Local	6.0

- Under Domain, enter the Fully Qualified Domain Name (FQDN).
- 7) Under Computer ID, enter the computer name.
- Click Bind. You are prompted for network credentials and a computer OU. Enter your OU admin account and password, and the computer OU location. For example:

ou=computers,ou=orgunits,dc=ad,dc=example,dc=com Your machine will be bound to the specified Active Directory.

Active Directory Forest:	- Automatic -
Active Directory Domain:	ad.example.com
Computer ID:	bsmith-mac
Show Advanced Options	Bind
	Cancel OK

9) Click Apply in the Directory Utility to save your changes and restart the machine.

Authentication for iPhones and iPads

Proxy-based user authentication is supported by the Content Gateway (proxy) component of Forcepoint Web Security, resulting in user- or group-based filtering.

User identification via DC Agent is not supported and, therefore, there is no user- or group-based filtering solution with Web Filter & Security or Forcepoint Web Security. Filtering can be provided to those devices based on IP address or network range.

Content Gateway user authentication has the following features and restrictions:

- Works with the authentication method configured in Content Gateway. Users must belong to the associated user directory.
- Supports the Safari browser. Other browsers may not work as expected.
- Transparent authentication is not supported. The user is always prompted for credentials.
- Works in transparent and explicit Content Gateway deployments.
- Many iPhone and iPad apps do not work well with Content Gateway (or any Web proxy) because they are not well programmed to handle proxy user authentication.

Explicit proxy settings can be configured in the iOS Network settings area.


Chapter 15 Working With Log Files

Contents

- Introduction on page 217
- Event log files on page 218
- Managing event log files on page 219
- Event log file formats on page 221
- Rolling event log files on page 227
- Splitting event log files on page 230
- Collating event log files on page 232
- Viewing logging statistics on page 235
- Viewing log files on page 235
- Example event log file entries on page 236

Introduction

Content Gateway keeps 3 types of log files:

System log files record system information, which includes messages about the state of Content Gateway and any errors or warnings that it produces. This information might include a note that event log files were rolled, a warning that cluster communication timed out, or an error indicating that Content Gateway was restarted. (Content Gateway posts alarms for error conditions in the Content Gateway manager; see *Working with alarms*, for details.)

All system information messages are logged with the system-wide logging facility **syslog** under the daemon facility. The **syslog.conf** configuration file (stored in the /etc directory) specifies where these messages are logged. A typical location is /var/log/messages.

The **syslog** process works on a system-wide basis, so it is the single repository for messages from all Content Gateway processes, including **content_gateway**, **content_manager**, and **content_cop**.

Each log entry in the log contains information about the date and time the error was logged, the hostname of the proxy server that reported the error, and a description of the error or warning.

See *Content Gateway Error Messages*, for a list of the system information messages that Content Gateway logs.

- Error log files record information about why a transaction was in error.
- Event log files (also called access log files) record information about the state of each transaction that Content Gateway processes.

Content Gateway creates both error and event log files and records system information in system log files. You can disable event logging and/or error logging. It is recommended that you log errors only or disable logging during peak usage hours.

On the Configure > Subsystems > Logging tab, select one of the following options:

- Log Transactions and Errors
- Log Transactions Only
- Log Errors Only
- Disabled

Related concepts

Working with alarms on page 95

Related information

Content Gateway Error Messages on page 467

Event log files

Event log files record information about every request that Content Gateway processes. By analyzing the log files, you can determine how many people use the proxy, how much information each person requested, what pages are most popular, and so on.

Content Gateway supports several standard log file formats, such as Squid and Netscape, and user-defined custom formats. You can analyze the standard format log files with off-the-shelf analysis packages. To help with log file analysis, you can separate log files so that they contain information specific to protocol or hosts. You can also configure Content Gateway to roll log files automatically at specific intervals during the day.

The following sections describe how to:

Manage your event log files

You can choose a central location for storing log files, set how much disk space to use for log files, and set how and when to roll log files. See *Managing event log files*.

Choose different event log file formats

You can choose which standard log file formats you want to use for traffic analysis (for example, Squid or Netscape). Alternatively, you can use the Content Gateway custom format, which is XML-based and enables you to institute more control over the type of information recorded in log files. See *Event log file formats*.

- Roll event log files automatically You can configure Content Gateway to roll event log files at specific intervals during the day so that you can identify and manipulate log files that are no longer active. See *Rolling event log files*.
- Separate log files according to hosts
 You can configure the proxy to create separate log files for different protocols based on the host. See Splitting event log files.
- Collate log files from different nodes You can designate one or more nodes on the network to serve as log collation servers. These servers, which might either be stand-alone or part of Content Gateway, enable you to keep all logged information in welldefined locations. See *Collating event log files*.
- View statistics about the logging system Content Gateway provides statistics about the logging system. You can access the statistics through the Content Gateway manager or through the command line interface. See *Viewing logging statistics*.
- View log files You can view the system, event, and error log files that Content Gateway creates. You can view an entire log file, a specified last number of lines in the log file, or all lines that contain a specified string.
- Interpret log file entries for the standard log file formats. See Example event log file entries.

Related concepts Managing event log files on page 219 Event log file formats on page 221 Rolling event log files on page 227 Splitting event log files on page 230 Collating event log files on page 232 Example event log file entries on page 236 Viewing logging statistics on page 235

Managing event log files

You can manage your event log files and control where they are located, how much space they can consume, and how low disk space in the logging directory is handled.

Choosing the logging directory

By default, Content Gateway writes all event log files in the /opt/WCG/logs directory, which is a subdirectory of the directory where you installed Content Gateway. To use a different directory, see *Setting log file management options*.

Related tasks

Setting log file management options on page 220

Controlling logging space

You can control the amount of disk space that the logging directory can consume. This allows the system to operate smoothly within a specified space window for a long period of time.

After you establish a space limit, Content Gateway continues to monitor the space in the logging directory. When the free space dwindles to the headroom limit (see *Setting log file management options*), Content Gateway enters a low space state and takes the following actions:

- If the autodelete option (discussed in *Rolling event log files*) is *enabled*, Content Gateway identifies previously rolled log files (log files with a .old extension) and starts deleting files one by one—beginning with the oldest file— until it emerges from the low state. Content Gateway logs a record of all files it deletes in the system error log.
- If the autodelete option is *disabled* or there are not enough old log files to delete for the system to emerge from its low space state, Content Gateway issues a warning and continues logging until space is exhausted. Content Gateway resumes event logging when enough space becomes available for it to exit its low space state. You can make space available by removing files from the logging directory or by increasing the logging space limit.

You can run a **cron** script in conjunction with Content Gateway to automatically remove old log files from the logging directory (before Content Gateway enters the low space state) and relocate them to a temporary partition.

Once the files are relocated, you can run log analysis scripts on them, and then you can compress the logs and move them to an archive location or delete them.

Related concepts

Rolling event log files on page 227

Related tasks

Setting log file management options on page 220

Setting log file management options

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > General** tab.
- In the Log Directory field, enter the path to the directory in which you want to store event log files. The default directory is /opt/WCG/logs, a subdirectory of the Content Gateway installation directory.



Note

The log directory you specify must already exist and must be /opt/WCG/logs or a subdirectory of it.

The user must have read/write permissions for the directory storing the log files.

 In the Limit field of the Log Space area, enter the maximum amount of space you want to allocate to the logging directory.

When Content Gateway is on an appliance, the size is set to 5120 (5 GB) and cannot be changed.

When Content Gateway is installed on a stand-alone server, the default size is 20480 (20 GB) and the size is configurable.



Note

All files in the logging directory contribute to the space used, even if they are not log files.

- 4) In the Headroom field, enter the tolerance for the log space limit. The default value is 100 MB. If the Auto-Delete Rolled Files option is enabled in the Log Rolling section, autodeletion is triggered when the amount of free space available in the logging directory is less than the headroom. For information about log file rolling, see Rolling event log files.
- 5) Click Apply.

Related concepts Rolling event log files on page 227

Event log file formats

Content Gateway supports the following log file formats:

- Standard formats, such as Squid or Netscape (see Using standard formats)
- Content Gateway custom format (see Custom format)

In addition to the standard and custom log file format, you must choose whether to save log files in *binary* or *ASCII*. See *Choosing binary* or *ASCII*.



Important

Event log files consume a large amount of disk space. Creating log entries in multiple formats at the same time can consume disk resources very quickly and affect proxy performance.

When IPv6 is enabled, Event log entries are normalized to IPv6 format.

For example, "10.10.41.200" is logged as "::ffff:10.10.41.200".

To filter on a client at "10.10.41.200" in a custom log, use:

```
<LogFilter>
<Name = "IPv6_Test_Machine"/>
<Condition =
"chi MATCH ::ffff:10.10.41.200"/>
<Action = "ACCEPT"/>
</LogFilter>
```

```
Related concepts
Using standard formats on page 221
Custom format on page 222
Choosing binary or ASCII on page 225
```

Using standard formats

The standard log formats include Squid, Netscape Common, Netscape Extended, and Netscape Extended-2.

The standard log file formats can be analyzed with a wide variety of off-the-shelf log- analysis packages. You should use one of the standard event log formats unless you need information that these formats do not provide. See *Custom format*.

By default, Content Gateway is configured to use the Netscape Extended log file format only.

```
Related concepts
Custom format on page 222
```

Setting standard log file format options

Steps

- 1) Navigate to Configure > Subsystems > Logging > Formats.
- 2) Enable the format you want to use.
- 3) Select the log file type (ASCII or binary).
- 4) In the Filename field, enter the name you want to use for your event log files.
- 5) In the **Header** field, enter a text header that appears at the top of the event log files. Leave this field blank if you do not want to use a text header.
- 6) Click Apply.
- 7) Click Restart on Configure > My Proxy > Basic > General.

Custom format

The XML-based custom log format is more flexible than the standard log file formats, giving you more control over the type of information in your log files. Create a custom log format if you need data for analysis that is not available in the standard formats. You can decide what information to record for each Content Gateway transaction and create filters to define which transactions to log.

The heart of the custom logging feature is an XML-based logging configuration file (**logs_xml.config**) that enables you to create modular descriptions of logging objects. The **logs_xml.config** file uses three types of objects to create custom log files:

- The LogFormat defines the content of the log file using printf-style format strings.
- The LogFilter defines a filter so that you include or exclude certain information from the log file.
- The **LogObject** specifies all the information needed to produce a log file. For example:
 - The name of the log file (required).
 - The format to be used (required). This can be a standard format (Squid or Netscape) or a previously defined custom format (a previously defined LogFormat object).
 - The file mode (ASCII, Binary, or ASCII_PIPE). The default is ASCII. The ASCII_PIPE mode writes log entries to a UNIX named pipe (a buffer in memory). Other processes can then read the data using standard I/O functions. The advantage of using this option is that Content Gateway does not have to write to disk, freeing disk space and bandwidth for other tasks.

Note

When the buffer is full, Content Gateway drops log entries and issues an error message indicating how many entries were dropped. Content Gateway writes only complete log entries to the pipe; therefore, only full records are dropped.

- Any filters you want to use (previously defined **LogFilter** objects).
- The collation servers that are to receive the log files.
- The protocols you want to log (if the protocols tag is used, Content Gateway logs only transactions from the protocols listed; otherwise, all transactions for all protocols are logged).

- The origin servers you want to log (if the servers tag is used, Content Gateway logs only transactions for the origin servers listed; otherwise, transactions for all origin servers are logged).
- The header text you want the log files to contain. The header text appears at the beginning of the log file, just before the first record.
- The log file rolling options.



Note

To generate a custom log format, you must specify at least one **LogObject** definition. One log file is produced for each **LogObject** definition. You can create a custom log format in the Content Gateway manager or by editing a configuration file.

How to Configure Custom Logging in Content Gateway

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > Custom** tab.
- 2) Enable the Custom Logging option.
- 3) The Custom Log File Definitions area displays the logs_xml.config file. Add LogFormat, LogFilter, and LogObject specifications to the configuration file. For detailed information about the logs_xml.config file and associated object specifications, see logs_xml.config.
- 4) Click Apply.

Related concepts

logs_xml.config on page 377

Creating summary log files

Content Gateway performs several hundred operations per second; therefore, event log files can grow quite large. Using SQL-like aggregate operators, you can configure Content Gateway to create summary log files that summarize a set of log entries over a specified period of time. This can reduce the size of the log files generated.

You generate a summary log file by creating a **LogFormat** object in the XML-based logging configuration file (**logs_xml.config**) using the following SQL-like aggregate operators:

- COUNT
- SUM
- AVERAGE
- FIRST
- LAST

You can apply each of these operators to specific fields, requesting it to operate over a specified interval.

Summary log files represent a trade-off between convenience and information granularity. Since you must specify a time interval during which only a single record is generated, you can lose information. If you want the convenience of summary logs and need the detail of a conventional log file, consider creating and enabling two custom log formats—one using aggregate operators and the other not using aggregate operators.

To create a summary log file format

Steps

- In the Content Gateway manager, go to the Configure > Subsystems > Logging > Custom tab to display the logs_xml.config file.
- 2) Define the format of the log file as follows:

```
<LogFormat>
<Name = "summary"/>
<Format = "%<operator(field)> : %<operator(field)>"/>
<Interval = "n"/>
</LogFormat>
```

Here:

- "operator" is one of the five aggregate operators (COUNT, SUM, AVERAGE, FIRST, LAST). You can specify more than one operator in the format line.
- "field" is the logging field that you want to aggregate.
- "n" is the interval in seconds between summary log entries. For more information, see logs_xml.config.

For example, the following format generates one entry every 10 seconds, with each entry summarizing the time stamp of the last entry of the interval, a count of the number of entries seen within that 10-second interval, and the sum of all bytes sent to the client:

```
<LogFormat>
<Name = "summary"/>
<Format = "%<LAST(cqts)> : %<COUNT(*)> :
%<SUM(psql)>"/>
<Interval = "10"/>
</LogFormat>
```



Important

You cannot create a format specification that contains both aggregate operators and regular fields. For example, the following specification would be invalid:

```
<Format = "%<LAST(cqts)> : %<COUNT(*)> :
```

```
%<SUM(psql)> : %<cqu>"/>
```

- Define a LogObject that uses this format.
- 4) Click Apply.

Related concepts logs xml.config on page 377

Applying logs_xml.config file changes to all nodes in a cluster

After modifying the **logs_xml.config** file on one Content Gateway node, enter the following command from the Content Gateway **bin** directory (/opt/WCG/bin):

content_line -x

Content Gateway applies the changes to all nodes in the cluster. The changes take effect immediately.

Choosing binary or ASCII

You can configure Content Gateway to create event log files in either of the following:

- ASCII: these files can be processed using standard, off-the-shelf log-analysis tools. However, Content Gateway must perform additional processing to create the files in ASCII, resulting in an increase in overhead. Also, ASCII files tend to be larger than the equivalent binary files. ASCII log files have a .log filename extension by default.
- Binary: these files generate lower system overhead, as well as generally occupying less space on the disk, depending on the type of information being logged. You must, however, use a converter application before you can read or analyze these files using standard tools. Binary log files use a .blog filename extension by default.

While binary log files typically require less disk space, this is not always the case. For example, the value 0 (zero) requires only one byte to store in ASCII but requires four bytes when stored as a binary integer. If you define a custom format that logs IP addresses, a binary log file would require only four bytes of storage per 32-bit address. However, the same IP address stored in dot notation would require around 15 characters (bytes) in an ASCII log file.

For standard log formats, you select **Binary** or **ASCII** on the **Configure > Subsystems > Logging > Formats** tab in the Content Gateway manager. See *Setting standard log file format options*. For the custom log format, you specify

ASCII or Binary mode in the LogObject. Refer to Custom format.



Note

For custom log files, in addition to the ASCII and Binary options, you can also write log entries to a UNIX named pipe (a buffer in memory). Other processes can then read the data using standard I/ O functions. The advantage of using this option is that Content Gateway does not have to write to disk, freeing disk space and bandwidth for other tasks. In addition, writing to a pipe does not stop when logging space is exhausted because the pipe does not use disk space. See *logs_xml.config*, for more information about the ASCII_PIPE option.

Before selecting ASCII versus binary for your log files, consider the type of data that will be logged. Try logging for one day using ASCII and then one day using binary. Assuming that the number of requests is roughly the same for both days, you can calculate a rough metric comparing the two formats.

Related concepts Custom format on page 222 logs_xml.config on page 377

Related tasks

Setting standard log file format options on page 221

Using logcat to convert binary logs to ASCII

You must convert a binary log file to ASCII before you can analyze it using standard tools.

Steps

- 1) Change to the directory containing the binary log file.
- 2) Make sure that the logcat utility is in your path.

3) Enter the following command:

logcat <options> <input_filename>

The command-line options are:

Option	Description
-o output_file	Specifies where the command output is directed.
-a	Automatically generates the output filename based on the input filename. If the input is from stdin , this option is ignored.
	For example:
	logcat -a squid-1.blog squid-2.blog squid-3.blog
	generates:
	squid-1.log, squid-2.log, squid-3.log
-S	Attempts to transform the input to Squid format, if possible.
-C	Attempts to transform the input to Netscape Common format, if possible.
-E	Attempts to transform the input to Netscape Extended format, if possible.
-2	Attempt to transform the input to Netscape Extended-2 format, if possible.

Ę

Note

Use only one of the following options at any given time:

-S, -C, -E, or -2.

If no input files are specified, **logcat** reads from the standard input (**stdin**). If you do not specify an output file, **logcat** writes to the standard output (**stdout**).

For example, to convert a binary log file to an ASCII file, you can use the logcat

command with either of the following options:

logcat binary_file > ascii_file

logcat -o ascii_file binary_file

The binary log file is not modified by this command.

Rolling event log files

Content Gateway provides automatic log file rolling. This means that at specific intervals during the day, Content Gateway closes its current set of log files and opens new log files.

Log file rolling offers the following benefits:

It defines an interval over which log analysis can be performed.

- It keeps any single log file from becoming too large and assists in keeping the logging system within the specified space limits.
- It provides an easy way to identify files that are no longer being used so that an automated script can clean the logging directory and run log analysis programs.

You should roll log files several times a day. Rolling every six hours is a good guideline to follow.

Rolled log filename format

Content Gateway provides a consistent name format for rolled log files that allows you to identify log files.

When Content Gateway rolls a log file, it saves and closes the old file and starts a new file. Content Gateway renames the old file to include the following information:

- The format of the file (for example, squid.log).
- The hostname of the Content Gateway server that generated the log file.
- Two timestamps separated by a hyphen (-). The first time stamp is a lower bound for the time stamp of the first record in the log file. The lower bound is the time when the new buffer for log records is created. Under low load, the first time stamp in the filename can be different from the timestamp of the first entry. Under normal load, the first time stamp in the filename and the time stamp of the first entry are similar. The second time stamp is an upper bound for the time stamp of the last record in the log file (this is normally the rolling time).
- The suffix .old, which makes it easy for automated scripts to find rolled log files.

The timestamps have the following format:

%Y%M%D.%Hh%Mm%Ss-%Y%M%D.%Hh%Mm%Ss

The following table describes the format:

Code	Definition	Example
%Y	The year in four-digit format	2000
%M	The month in two-digit format, from 01-12	07
%D	The day in two-digit format, from 01-31	19
%Н	The hour in two-digit format, from 00-23	21
%M	The minute in two-digit format, from 00-59	52
%S	The second in two-digit format, from 00-59	36

The following is an example of a rolled log filename:

squid.log.mymachine.20000912.12h00m00s- 20000913.12h00m00s.old

In this example, the file is squid log format and the host machine is mymachine. The first time stamp indicates a date and time of year 2000, month September, and day 12 at 12:00 noon. The second time stamp indicates a date and time of year 2000, month September, and day 13 at 12:00 noon. At the end, the file has a .old suffix.

The logging system buffers log records before writing them to disk. When a log file is rolled, the log buffer might be partially full. If so, the first entry in the new log file will have a time stamp earlier than the time of rolling.

When the new log file is rolled, its first time stamp will be a lower bound for the time stamp of the first entry. For example, suppose logs are rolled every three hours, and the first rolled log file is:

squid.log.mymachine.19980912.12h00m00s- 19980912.03h00m00s.old

If the lower bound for the first entry in the log buffer at 3:00:00 is 2:59:47, the next log file, when rolled, will have the following time stamp:

squid.log.mymachine.19980912.02h59m47s- 19980912.06h00m00s.old

The contents of a log file are always between the two timestamps. Log files do not contain overlapping entries, even if successive timestamps appear to overlap.

Rolling intervals

Log files are rolled at specific intervals relative to a given hour of the day. Two options control when log files are rolled:

- The offset hour, which is an hour between 0 (midnight) and 23
- The rolling interval

Both the offset hour and the rolling interval determine when log file rolling starts. Rolling occurs every rolling interval *and* at the offset hour.

For example, if the rolling interval is six hours and the offset hour is 0 (midnight), the logs roll at midnight (00:00), 06:00, 12:00, and 18:00 each day. If the rolling interval is 12 hours and the offset hour is 3, logs roll at 03:00 and 15:00 each day.

Setting log file rolling options

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > General** tab.
- 2) In the Log Rolling section, ensure the Log Rolling option is enabled (the default).
- In the Offset Hour field, enter a specific time each day you want log file rolling to take place. Content Gateway forces the log file to be rolled at the offset hour each day. You can enter any hour in the range 0 (midnight) to 23.
- In the Interval field, enter the amount of time Content Gateway enters data in the log files before rotation takes place.

The minimum value is 300 seconds (five minutes). The maximum value is 86400 seconds (one day).



Note

If you start Content Gateway within a few minutes of the next rolling time, rolling may not occur until the following rolling time.

- 5) Ensure the Auto-Delete Rolled Files option is enabled (the default). This enables auto deletion of rolled log files when available space in the log directory is low. Auto deletion is triggered when the amount of free space available in the log directory is less than the headroom.
- 6) Click Apply.



Note

You can fine tune log file rolling settings for a custom log file in the **LogObject** specification in the **logs_xml.config** file. The custom log file uses the rolling settings in its **LogObject**, which override the default settings you specify in the Content Gateway manager or the **records.config** file described above.

Splitting event log files

By default, Content Gateway uses standard log formats and generates log files that contain HTTP and FTP transactions in the same file. However, you can enable host log splitting if you prefer to log transactions for different origin servers in separate log files.

HTTP host log splitting

HTTP host log splitting enables you to record HTTP and FTP transactions for different origin servers in separate log files. When HTTP host log splitting is enabled, Content Gateway creates a separate log file for each origin server listed in the **log_hosts.config** file (see *Editing the log_hosts.config file*).

When HTTP host log splitting is enabled, Content Gateway generates separate log files for HTTP/FTP transactions, based on the origin server.

For example, if the **log_hosts.config** file contains the two origin servers **uni.edu** and **company.com**, and the Squid format is enabled, Content Gateway generates the following log files:

Log Filename	Description
squid-uni.edu.log	All HTTP and FTP transactions for uni.edu
squid-company.com.log	All HTTP and FTP transactions for
	company.com
squid.log	All HTTP and FTP transactions for other hosts

Content Gateway also enables you to create XML-based custom log formats that offer even greater control over log file generation based on protocol and host name. See *Custom format*.

Related concepts	
Custom format on page 222	

Related tasks

Editing the log_hosts.config file on page 231

Setting log splitting options

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > Splitting** tab.
- 2) Enable the Split Host Logs option to record all HTTP and FTP transactions for each origin server listed in the log_hosts.config file in a separate log file. Disable the Split Host Logs option to record all HTTP and FTP transactions for each origin server listed in the log_hosts.config file in the same log file.
- 3) Click Apply.

Editing the log_hosts.config file

The default **log_hosts.config** file is located in /opt/WCG/config. To record HTTP and FTP transactions for different origin servers in separate log files, you must specify each origin server's hostname on a separate line in the file.



Note

You can specify keywords in the **log_hosts.config** file to record in a separate log file all transactions from origin servers that contain the specified keyword in their names. For example, if you specify the keyword **sports**, Content Gateway records all HTTP and FTP transactions from **sports.yahoo.com** and www.foxsports.com in a log file called **squid-sports.log** (if the Squid format is enabled).

If Content Gateway is clustered and if you enable log file collation, it is recommended that you use the same **log_hosts.config** file on every node in the cluster.

Steps

- 1) Open the log_hosts.config file located in /opt/WCG/config.
- 2) Enter the hostname of each origin server on a separate line in the file. For example:

webserverl

webserver2

webserver3

- 3) Save and close the file.
- 4) To apply the changes, run the following command from the Content Gateway bin directory (/opt/WCG/bin): ./content_line -x

Collating event log files

You can use the log file collation feature to keep all logged information in one place. This allows you to analyze Content Gateway as a whole rather than as individual nodes and to use a large disk that might only be located on one of the nodes in a cluster.

Content Gateway collates log files by using one or more nodes as log collation servers and all remaining nodes as log collation clients. When a node generates a buffer of

event log entries, it determines whether it is the collation server or a collation client. The collation server node simply writes all log buffers to its local disk, just as it would if log collation were not enabled.

The collation client nodes prepare their log buffers for transfer across the network and send the buffers to the log collation server. When the log collation server receives a log buffer from a client, it writes it to its own log file as if it were generated locally.

If log clients cannot contact their log collation server, they write their log buffers to their local disks, into *orphan* log files. Orphan log files require manual collation.

Log collation servers can be stand-alone or they can be part of a node running Content Gateway.

Ę

Note

Log collation can have an impact on network performance. Because all nodes are forwarding their log data buffers to the single collation server, a bottleneck might occur in the network, where the amount of data being sent to a single node in the network exceeds the node's ability to process it quickly.

Collated log files contain time-stamp information for each entry, but entries do not appear in the files in strict chronological order. You can sort collated log files before doing analysis.

Configuring Content Gateway to be a collation server

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > Collation** page.
- 2) In the Collation Mode section, enable the Be A Collation Server option.
- In the Log Collation Port field, enter the port number used for communication with collation clients. The default port number is 8085.
- 4) In the Log Collation Secret field, enter the password used to validate logging data and prevent the exchange of arbitrary information.



Note

All collation clients must use this same secret.

5) Click Apply.



Important

If you modify the collation port or secret after connections between the collation server and collation clients have been established, you must restart Content Gateway.

Configuring Content Gateway to be a collation client

Steps

- 1) In the Content Gateway manager, go to the **Configure > Subsystems > Logging > Collation** tab.
- 2) In the Collation Mode section, enable the Be a Collation Client option to set the Content Gateway node as a collation client and send the active standard formatted log entries (such as Squid and Netscape) to the log collation server.



Note

To send custom XML-based formatted log entries to the collation server, you must add a log object specification to the **logs_xml.config** file. See *Custom format*.

- 3) In the **To Collation Server** field, enter the hostname of the collation server. This could be the Content Gateway collation server or a stand-alone collation server.
- In the Log Collation Port field, enter the port number used for communication with the collation server. The default port number is 8085.
- 5) In the Log Collation Secret field, enter the password used to validate logging data and prevent the exchange of arbitrary information. This must be the same secret you set on the collation server.
- 6) Enable the Log Collation Host Tagged option if you want to preserve the origin of log entries in the collated log files.
- 7) In the Log Collation Orphan Space field, enter the maximum amount of space (in megabytes) you want to allocate to the logging directory on the collation client for storing orphan log files. (Orphan log files are created when the log collation server cannot be contacted). The default value is 25 MB.
- 8) Click Apply.



Important

If you modify the collation port or secret after connections between the collation clients and collation server have been established, you must restart Content Gateway.

Related concepts

Custom format on page 222

Using a stand-alone collator

If you do not want the log collation server to be a Content Gateway node, you can install and configure a standalone collator (SAC) which can dedicate more of its power to collecting, processing, and writing log files.



Note

The stand-alone collator is currently available for the Linux platform only.

Steps

- 1) Configure your Content Gateway nodes as log collation clients. See *Configuring Content Gateway to be a collation client*.
- Copy the sac binary from the Content Gateway bin directory (/opt/WCG/bin) to the machine serving as the stand-alone collator.
- 3) Create a directory called **config** in the directory that contains the **sac** binary.
- 4) Create a directory called **internal** in the **config** directory you created in *Step 3*. This directory will be used internally by the stand-alone collator to store lock files.
- 5) Copy the records.config file (/opt/WCG/config) from a Content Gateway node configured to be a log collation client to the config directory you created in *Step 3* on the stand-alone collator. The records.config file contains the log collation secret and port you specified when configuring nodes to be collation clients. The collation port and secret must be the same for all collation clients and servers.
- 6) Open the **records.config** file on the stand-alone collator and edit the **proxy.config.log2.logfile_dir** variable to specify the directory where you want to store log files.
 - You can specify an absolute path to the directory or a path relative to the directory from which the sac binary is executed.
 - The directory must already exist on the machine serving as the stand-alone collator.
- 7) Save and close the file.
- Enter the following command:

```
sac -c config
```

Related tasks

Configuring Content Gateway to be a collation client on page 233

Viewing logging statistics

Content Gateway generates statistics about the logging system that help you see the following information:

- How many log files (formats) are currently being written.
- The current amount of space being used by the logging directory, which contains all of the event and error logs.
- The number of access events that have been written to log files since Content Gateway installation. This counter represents one entry in one file. If multiple formats are being written, a single event will create multiple event log entries.
- The number of access events skipped (because they were filtered out) since Content Gateway installation.
- The number of access events that have been written to the event error log since Content Gateway installation.

You can view the statistics from the Monitor tab in the Content Gateway manager or retrieve them through the command-line interface. See *Monitoring Traffic*.

Related information Monitoring Traffic on page 93

Viewing log files

You can view the system, event, and error log files that Content Gateway creates from the Content Gateway manager. You can view an entire log file, a specified last number of lines in the log file, or all lines that contain a specified string.

You can also delete a log file or copy it to your local system.



Note

You must have the correct user permissions to copy and delete log files.

Content Gateway displays only the first 1 MB of data in the log file. If the log file you select is larger than 1 MB, Content Gateway truncates the file and displays a warning message indicating that the file is too big.

You can now access log files through the Content Gateway manager.

Steps

- Navigate to the Configure > My Proxy > Logs > System tab.
- To view, copy, or delete a system log file, go to Step 3.
 To view, copy, or delete an event or error log file, select the Access tab.

3) In the Log File drop-down list, select the log file you want to view, copy, or delete.

Content Gateway lists the system log files logged with the system-wide logging facility **syslog** under the daemon facility.

Content Gateway lists the event log files located in the directory specified in the **Logging Directory** field in the **Configure > Subsystems > Logging > General** tab or by the configuration variable **proxy.config.log2.logfile_dir** in the **records.config** file. The default directory is **logs** in the Content Gateway installation directory.

- 4) In the Action area, select one of the following options:
 - Display the selected log file to view the entire log file. If the file is larger than 1 MB, only the first MB of data is displayed.
 - Display last lines of the selected file to view the last lines of the log file. Enter the number of lines you want to view in the field provided.
 - **Display lines that match in the selected log file** to view all the lines in the log file that match a particular string. Enter the string in the field provided.
 - Remove the selected log file to delete the selected log file from the Content Gateway system.
 - Save the selected log file in local filesystem to save a copy of the selected log file on your local system.
- 5) Click Apply.

If you selected to view the log file, Content Gateway displays the file at the end of the page.

If you selected to delete the log file, Content Gateway deletes the file. You are not prompted to confirm the deletion.

If you selected to save the log file, you are prompted for the location where you want to save the file on your local system.

Example event log file entries

This section shows examples of a log file entry in each of the standard log formats supported by Content Gateway:

Related concepts Netscape Extended format on page 238

Related reference

Squid format on page 237 Netscape Extended-2 format on page 238

Related information Netscape examples on page 238

Squid format

The following figure shows a sample log entry in a **squid.log** file. The table below describes each field.



Field	Description
1	The client request time stamp in Squid format; the time of the client request in seconds since January 1, 1970 UTC (with millisecond resolution).
2	The time the proxy spent processing the client request; the number of milliseconds between the time that the client established the connection with the proxy and the time that the proxy sent the last byte of the response back to the client.
3	The IP address of the client's host machine.
4	The cache result code; how the cache responded to the request: HIT, MISS, and so on. Cache result codes are described in <i>Cache result codes in Squid- and</i> <i>Netscape-format log files</i> .
	The proxy response status code (the HTTP response status code from Content Gateway to client).
5	The length of the Content Gateway response to the client in bytes, including headers and content.
6	The client request method: GET, POST, and so on.
7	The client request canonical URL; blanks and other characters that might not be parsed by log analysis tools are replaced by escape sequences. The escape sequence is a percentage sign followed by the ASCII code number of the replaced character in hex.
8	The authenticated client's user name. A hyphen (-) means that no authentication was required.
9	The proxy hierarchy route; the route Content Gateway used to retrieve the object. The proxy request server name; the name of the server that fulfilled the request. If the request was a cache hit, this field contains a hyphen (-).
10	The proxy response content type; the object content type taken from the Content Gateway response header.

Related reference

Cache result codes in Squid- and Netscape-format log files on page 240

Netscape examples

Netscape Common format

The following figure shows a sample log entry in a **common.log** file. The table below describes each field.



Netscape Extended format

The following figure shows a sample log entry in an extended.log file. The table below describes each field.



Netscape Extended-2 format

The following figure shows a sample log entry in an extended2.log file. The table below describes each field.



Field	Description
	Netscape Common
1	The IP address of the client's host machine.

Field	Description
2	This hyphen (-) is always present in Netscape log entries.
3	The authenticated client user name. A hyphen (-) means no authentication was required.
4	The date and time of the client's request, enclosed in brackets.
5	The request line, enclosed in quotes.
6	The proxy response status code (HTTP reply code).
7	The length of the Content Gateway response to the client in bytes.
	Netscape Extended
8	The origin server's response status code.
9	The server response transfer length; the body length in the origin server's response to the proxy, in bytes.
10	The client request transfer length; the body length in the client's request to the proxy, in bytes.
11	The proxy request transfer length; the body length in the proxy request to the origin server.
12	The client request header length; the header length in the client's request to the proxy.
13	The proxy response header length; the header length in the proxy response to the client.
14	The proxy request header length; the header length in the proxy request to the origin server.
15	The server response header length; the header length in the origin server's response to the proxy.
16	The time Content Gateway spent processing the client request; the number of seconds between the time that the client established the connection with the proxy and the time that the proxy sent the last byte of the response back to the client.
	Netscape Extended-2
17	The proxy hierarchy route; the route Content Gateway used to retrieve the object.
18	The client finish status code: FIN if the client request completed successfully or INTR if the client request was interrupted.
19	The proxy finish status code: FIN if the Content Gateway request to the origin server completed successfully or INTR if the request was interrupted.

Field	Description
20	The cache result code; how the Content Gateway cache responded to the request: HIT, MISS, and so on. Cache result codes are described in <i>Cache result codes in Squid- and Netscape-format log files</i> .

Related reference

Cache result codes in Squid- and Netscape-format log files on page 240

Cache result codes in Squid- and Netscapeformat log files

Cache result codes in the Squid and Netscape log files:

Cache Result Code	Description
TCP_HIT	Indicates that a valid copy of the requested object was in the cache and that the proxy sent the object to the client.
TCP_MISS	Indicates that the requested object was not in the cache and that the proxy retrieved the object from the origin server or from a parent proxy and sent it to the client.
TCP_REFRESH_HIT	Indicates that the object was in the cache but was stale. Content Gateway made an if-modified-since request to the origin server and the origin server sent a 304 not-modified response. The proxy sent the cached object to the client.
TCP_REF_FAIL_HIT	Indicates that the object was in the cache but was stale. Content Gateway made an if-modified-since request to the origin server but the server did not respond. The proxy sent the cached object to the client.
TCP_REFRESH_MISS	Indicates that the object was in the cache but was stale. Content Gateway made an if-modified-since request to the origin server and the server returned a new object. The proxy served the new object to the client.
TCP_CLIENT_REFRESH	Indicates that the client issued a request with a no- cache header. The proxy obtained the requested object from the origin server and sent a copy to the client. Content Gateway refreshes any previous copy of the object in the cache.

Cache Result Code	Description
TCP_IMS_HIT	Indicates that the client issued an if-modified-since request and the object was in the cache and fresher than the IMS date, or an if-modified-since to the origin server found that the cache object was fresh. The proxy served the cached object to the client.
TCP_IMS_MISS	Indicates that the client issued an if-modified-since request and the object was either not in cache or was stale in cache. The proxy sent an if-modified-since request to the origin server and received the new object. The proxy sent the updated object to the client.
TCP_SWAPFAIL	Indicates that the object was in the cache but could not be accessed. The client did not receive the object.
ERR_CLIENT_ABORT	Indicates that the client disconnected before the complete object was sent.
ERR_CONNECT_FAIL	Indicates that Content Gateway could not reach the origin server.
ERR_DNS_FAIL	Indicates that the Domain Name Server could not resolve the origin server name, or that no Domain Name Server could be reached.
ERR_INVALID_REQ	Indicates that the client HTTP request was invalid. Content Gateway forwards requests with unknown methods to the origin server.
ERR_READ_TIMEOUT	Indicates that the origin server did not respond to the Content Gateway request within the timeout interval.
ERR_PROXY_DENIED	Indicates that client service was denied by access control configuration.
ERR_UNKNOWN	Indicates that the client connected but subsequently disconnected without sending a request.

Appendix A Statistics

Contents

- My Proxy on page 243
- Protocols on page 249
- Security on page 251
- Subsystems on page 255
- Networking on page 256
- Performance on page 261
- SSL on page 263

This section describes the following statistics accessed on the Content Gateway manager Monitor tab:

Related concepts My Proxy on page 243 Protocols on page 249 Security on page 251 Subsystems on page 255 Networking on page 256 Performance on page 261 SSL on page 263

My Proxy

My Proxy statistics are divided into the following categories:

Related concepts Graphs on page 246 Alarms on page 247 Diagnostics on page 247

Related reference

Summary on page 244 Node on page 245

Summary

Statistic/Field	Description
	Subscription Details
Feature	Lists available features, such as analytic options, threat detection, and the file sandbox.
Purchased Status	Indicates if a feature has been purchased.
Expiration Date	If a feature has been purchased, displays the expiration date of the subscription.
	More Detail
Subscription key	Displays the subscription key. See <i>Entering your</i> subscription key.
Last successful subscription download time	Displays the time of the last successful validation of the subscription key. The check is made once a day.
Connection status	Displays the Content Gateway connection status to Policy Server, Policy Broker, and Filtering Service.
Registration status	Displays the Content Gateway registration status with the Forensics Repository.
	Scanning Data Files
Engine Name	Displays the name of each scanning engine.
Engine Version	Displays the version number of the scanning engine.
Data File Version	Displays the version number of the data file currently in use by the scanning engine.
Last update	Displays the time and date when Content Gateway last successfully loaded that analytics data files, settings, and policies.
Last time Content Gateway loaded data	Displays the time and date when Content Gateway last successfully loaded databases, settings, and policies.
Last time Content Gateway checked for updates	Displays the time and date when Content Gateway last successfully communicated with the download server to check for data file updates.
	Node Details
Node	Name of the Content Gateway node or cluster.
On/Off	Indicates if the proxy and manager services are running.
Objects Served	The total number of objects served by the node.
Ops/Sec	The number of operations per second processed by the node.

Statistic/Field	Description
Throughput (Mbit/sec)	The number of megabits per second passing through the node (and cluster).
	The proxy updates the throughput statistic after it transfers an entire object. For larger files, the byte count increases sharply at the end of a transfer. The complete number of bytes transferred is attributed to the last 10- second interval, although it can take several minutes to transfer the object. This transient inaccuracy is more noticeable with a light load.
	More Detail
errors	The percentage of requests that end in early hangups.
aborts	The percentage of aborted requests.
active clients	The current number of open client connections.
active servers	The current number of open origin server connections.
node IP address	The IP address assigned to the node. If virtual IP addressing is enabled, several virtual IP addresses could be assigned to this node.
HostDB hit rate	The ratio of host database hits to total host database lookups, averaged over a 10-second period.

Related tasks

Entering your subscription key on page 20

Node

Browser limitations require configuring a specific port in order for these graphs to display properly. The Node and Graphs options are disabled until a port is specified in records.config (in /opt/WCG/config, by default).

 Update this variable to enable the Node and Graphs pages: proxy.config.admin.overseer_port INT ##

where ## is a valid port number.

2) Restart Content Gateway.

If the node is part of a cluster, two sets of statistics are shown:

- Information about the single node and
- Information showing an average value for all nodes in the cluster. Click the name of a statistic to display the information in graphical format.

Statistic	Description
	Node Summary

Statistic	Description
Status	Indicates if Content Gateway is running on this node (active or inactive).
Up Since	Date and time Content Gateway was started.
Clustering	Indicates if clustering is on or off on this node.
	In Progress
Open Server Connections	Number of currently open origin server connections.
Open Client Connections	Number of currently open client connections.
	Network
Client Throughput (Mbit/Sec)	Number of megabits per second passing through the node (and cluster).
Transactions per Second	Number of HTTP transactions per second.
	Name Resolution
Host Database Hit Rate	Ratio of host database hits to total host database lookups, averaged over 10 seconds. This value is refreshed every 10 seconds.
DNS Lookups per Second	Number of DNS lookups per second.

Graphs

Browser limitations require configuring a specific port in order for these graphs to display properly. The Node and Graphs options are disabled until a port is specified in records.config (in /opt/WCG/config, by default).

 Update this variable to enable the Node and Graphs pages: proxy.config.admin.overseer_port INT ##

where ## is a valid port number.

2) Restart Content Gateway.

The Graphs page displays the same statistics listed on the *Node* page (current connections and transfers, network, and name resolution) but in graphical format. You can choose the statistics you want to present in a graph. See *Viewing statistics in the Content Gateway manager*.



Important

The graph is displayed in your browser using a Java applet. You should have the latest version of Java installed on your PC (at least version 1.7). To validate your access to Content Gateway statistics, you will be prompted for Content Gateway logon credentials.

Related concepts

Viewing statistics in the Content Gateway manager on page 94

Related reference

Node on page 245

Alarms

Content Gateway signals an alarm when it detects a problem (for example, if the space allocated to event logs is full or if Content Gateway cannot write to a configuration file) and displays a description of the alarm in the alarm message window. In addition, the **Alarm! [pending]** bar at the top of the Content Gateway manager display indicates when alarms are detected and how many alarms exist.

After you have read an alarm message, click **Clear** in the alarm message window to dismiss the alarm. Clicking **Clear** only dismisses alarm messages; it does not actually resolve the cause of the alarms.

For information about working with alarms, see Working with alarms.

Related concepts Working with alarms on page 95

Diagnostics

Use the tools provided to help diagnose communication or connection issues, trace network packets, or capture network packets.

Related concepts

Automatic diagnostic tests on page 247 Manual diagnostic tests on page 248

Automatic diagnostic tests

By default, the page opens to the **Automatic** tab. Click **Run Diagnostics** to execute all of the tests listed in the table. Connectivity is tested from the Content Gateway host machine to each of the servers listed under **Test**. In addition, the availability of the DNS servers is confirmed.

- The IPv4 default gateway
- The IPv6 default gateway
- Your primary DNS server
- Your secondary DNS server
- download.forcepoint.com (a download server)
- ddsdom.websense.com (a download server)
- ddsint.websense.com (a download server)
- my.websense.com (customer account portal)

Once the diagnostics are run, additional information is provided:

- Result indicates whether the test is running, passed, failed, or could not complete.
- Latency provides the round-trip latency of the Ping command used to test the connection. The value, reported in milliseconds, is the amount of time between the command being sent and the response being received from the server.

An empty latency value does not necessarily indicate a problem. Rather, it indicates either (a) that the test passed, but the packet that holds the value was banned by something in the network, or (b) that the test failed, and thus no latency value could be obtained.

If the value seems high (a full 10 seconds, for example) when compared to other latency values, it may indicate a problem in the network.

Details offers additional information for any test that failed or could not complete.

Below the table, the Last update information reflects the date and time the connections were last tested. Each time you access the page, the results of the last test will display.

Manual diagnostic tests

The Manual tab offers 4 commands typically run from the Linux command line.

- Ping, used to determine if a remote device can be reached across the network.
- **Traceroute**, used to determine the path network packets take and measure delays across the network.
- NSlookup, used to obtain domain name or IP address mapping.
- **TCPDump**, used to analyze network packets.

Click the radio button next to the command you want to execute and enter parameters for the command in the entry field provided.

- Enter a server name or IP address for Ping or Traceroute.
- Enter a server name for **NSlookup**.
- Enter valid parameters for TCPDump. For additional information about using TCPDump with Content Gateway, see this article.

Click the **Run** button next to your selected command to execute the test. The results for Ping, Traceroute, and NSlookup display in the **Test Results** section at the bottom of the pane.

Test results for TCPDump are typically too long to easily display and review in the Test Results window. When TCPDump is run, the Test Results window simply indicates the success or failure of the command.

As TCPDump runs, output is written to /opt/WCG/logs/tcpdump.pcap. This file is overwritten each time TCPDump is executed. When a test is successful, a link is provided so that you can download and view or save a copy of the most recent file.

To avoid disk space problems, tcpdump.pcap is limited to 10,000 packets. Once that limit is reached, no additional output is written to the file.



Important

TCPDump uses a lot of system resources. Try to avoid using it during peak hours when the system is busy.

As each command executes, the Run button becomes a Stop button. Click Stop to abort the command.

Protocols

Protocol statistics are divided into the following categories:

- HTTP
- FTP

For SSL statistics, click the SSL button at the bottom of the Monitor tab.

Related concepts SSL on page 263

Related reference HTTP on page 249 FTP on page 250

HTTP

Statistic	Description
	General
Client	
Total Document Bytes	Total amount of HTTP data served to clients since installation.
Total Header Bytes	Total amount of HTTP header data served to clients since installation.
Total Connections	Total number of HTTP client connections since installation.
Current Connections	Current number of HTTP client connections
Transactions in Progress	Total number of HTTP client transactions in progress.
Server	
Total Document Bytes	Total amount of HTTP data received from origin servers since installation.
Total Header Bytes	Total amount of HTTP header data received from origin servers since installation.
Total Connections	Total number of HTTP server connections since installation.
Current Connections	Current number of HTTP server connections
Transactions in Progress	Total number of HTTP server connections currently in progress.
	Transaction

Statistic	Description
Errors	
Connection Failures	Percentage of connect errors and their average transaction times.
Other Errors	Percentage of other errors and their average transaction times.
Aborted Transactions	
Client Aborts	Percentage of client-aborted transactions and their average transaction times.
Questionable Client Aborts	Percentage of transactions that could possibly be client aborted and their average transaction times.
Partial Request Hangups	Percentage of early hangups (after partial requests) and their average transaction times.
Pre-Request Hangups	Percentage of pre-request hangups and their average transaction times.
Pre-Connect Hangups	Percentage of pre-connect hangups and their average transaction times.
Other Transactions	1
Unclassified	Percentage of unclassified transactions and their average transaction times.
	FTP over HTTP
Connections	
Open Server Connections	Number of open connections to the FTP server.
Successful PASV Connections	Number of successful PASV connections since installation.
Failed PASV Connections	Number of failed PASV connections since installation.
Successful PORT Connections	Number of successful PORT connections since installation.
Failed PORT Connections	Number of failed PORT connections since installation.

FTP

Statistic	Description
	Client
Open Connections	Number of client connections currently open.
Bytes Read	Number of client request bytes read since installation.
Bytes Written	Number of client request bytes written since installation.

Statistic	Description
	Server
Open Connections	Number of FTP server connections currently open.
Bytes Read	The number of bytes read from FTP servers since installation.
Bytes Written	Number of bytes written since installation.

Security

Security statistics are divided into the following categories:



Note

Even when multiple authentication rules are used, Content Gateway reports authentication statistics discreetly for each authentication method (IWA, LDAP, Legacy NTLM).

Related reference

Integrated Windows Authentication on page 251 LDAP on page 253 Legacy NTLM on page 254 SOCKS on page 255 Web DLP on page 255

Integrated Windows Authentication

Statistic	Description
	Diagnostic Test
	This function runs diagnostic tests on the Kerberos connection to the selected domain. Results are displayed on screen and written to /opt/WCG/ logs/ content_gateway.out and /opt/WCG/logs/ smbadmin.log.
Domain drop down box	Select a joined domain. Unless Rule-Based Authentication is configured, there will only be 1 joined domain.
Run Test button	Click to initiate a test.

Statistic	Description
	Active Directory Joined Domains list
	Lists all joined AD domains.
	The Content Gateway Hostname DNS is the name that clients must specify in their browser proxy settings for Kerberos authentication to occur.
	Kerberos request counters
Total Kerberos requests	The total number of Kerberos authentication requests.
Authentication succeeded	The number of Kerberos authentication requests that resulted in successful authentication.
Authentication failed	The number of Kerberos authentication requests that resulted in authentication failure.
Kerberos errors	The number of Kerberos process errors.
	NTLM request counters
Total NTLM requests	The total number of NTLM authentication requests.
Authentication succeeded	The number of NTLM authentication requests that resulted in successful authentication.
Authentication failed	The number of NTLM authentication requests that resulted in authentication failure.
NTLM request errors	The number of NTLM process errors.
NTLM within negotiate requests	The number of NTLM requests encapsulated in Negotiate requests.
	Basic authentication request counters
Total basic authentication requests	The total number of basic authentication requests.
Authentication succeeded	The number of basic authentication requests that resulted in successful authentication.
Authentication failed.	The number of basic authentication requests that resulted in authentication failure.
Basic authentication request errors	The number of basic authentication process errors.
	Performance counters
Kerberos - Average time per transaction	The average time, in milliseconds, to complete a Kerberos transaction.
NTLM - Average time per transaction	The average time, in milliseconds, to complete a NTLM transaction.
Basic - Average time per transaction	The average time, in milliseconds, to complete a basic transaction.
Average helper latency per transaction	The average time for Samba to process an authentication request.
Statistic	Description
---	---
Time authentication spent offline	The time, in seconds, that Content Gateway was unable to perform NTLM authentication due to service or connectivity failures. (This measure does not apply to Kerberos because no communication with the DC is needed.)
	If the Fail Open option is enabled (<i>Global authentication options</i>), proxy requests may proceed without authentication.
	The counter is incremented when connectivity is reestablished after a failure.
Number of times authentication servers or services went offline	The number of times that connectivity with authentication servers or services has been lost.
	Top lists counters
	These user authentication lists provide a view into which User-Agent values and client IP addresses are most active. Four counters tally the top 20 User-Agent and client IP addresses that are passing or failing user authentication.
Button: Reset Top Lists to Zero	Resets all Top Lists counters to zero.
Top User-Agents passing authentication	Lists the top 20 User-Agent matches by number of authentication attempts that pass authentication.
Top User-Agents failing authentication	Lists the top 20 User-Agent matches by number of authentication attempts that fail authentication.
Top Client IP addresses passing authentication	Lists the top 20 client IP addresses by number of authentication attempts that pass authentication.
Top Client IP addresses failing authentication	Lists the top 20 client IP addresses by number of authentication attempts that fail authentication.

Related concepts

Global authentication options on page 168

LDAP

Statistic	Description
	Cache
Hits	Number of hits in the LDAP cache.
Misses	Number of misses in the LDAP cache.
	Errors
Server	Number of LDAP server errors.

Statistic	Description
	Successful Authentications
Authentication Succeeded	Number of times authentication was successful.
	Unsuccessful Authentications
Authentication Denied	Number of times the LDAP Server denied authentication.
Authentication Timeouts	Number of times authentication timed out.
Authentication Cancelled	Number of times authentication was terminated after LDAP authentication was started and before it was completed.
	Note: This does not count the number of times that an authentication request was cancelled by the client by clicking "Cancel" in the dialog box that prompts for credentials.

Legacy NTLM

Statistic	Description
	Cache
Hits	Number of hits in the NTLM cache.
Misses	Number of misses in the NTLM cache.
	Errors
Server	Number of NTLM server errors.
	Successful Authentications
Authentication Succeeded	Number of times authentication was successful.
	Unsuccessful Authentications
Authentication Denied	Number of times the NTLM server denied authentication.
Authentication Cancelled	Number of times authentication was cancelled.
Authentication Rejected	Number of times authentication failed because the queue was full.
	Queue Size
Authentication Queued	Number of requests that are currently queued because all of the domain controllers are busy.

SOCKS

Statistic	Description
On-Appliance SOCKS Server (when Content Gateway is on an appliance)	Indicates whether the on-appliance SOCKS server is on (enabled) or off (disabled).
Unsuccessful Connections	Number of unsuccessful connections to the SOCKS server since Content Gateway was started.
Successful Connections	Number of successful connections to the SOCKS server since Content Gateway was started.
Connections in Progress	Number of connections to the SOCKS server currently in progress.

Web DLP

Statistic	Description
Total Posts	Total number of posts sent to Web DLP.
Total Analyzed	Total number of posts analyzed by Web DLP.
FTP Analyzed	Total number of FTP requests analyzed by DLP.
Blocked Requests	Total number of requests blocked after analysis and policy enforcement.
Allowed Requests	Total number of requests allowed after analysis and policy enforcement.
Failed Requests	Total number of posts sent to Web DLP that timed out or otherwise failed to complete.
Huge Requests	Total number of requests that exceeded the maximum transaction size.
Tiny Requests	Total number of requests that were smaller than the minimum transaction size.
Decrypted Requests	Total number of SSL requests decrypted and sent to Web DLP.
Total Bytes Scanned	Total number of bytes scanned by Web DLP.
Average Response Time	Average time needed to by Web DLP to complete a scan since the last time Content Gateway was started.

Subsystems

Subsystems statistics are divided into the following categories:

Related reference Clustering on page 256 Logging on page 256

Clustering



Note

Cache statistics may be non-zero even if all content sent to Content Gateway is not cacheable. Content Gateway performs a cache-read even if the client sends a no-cache control header.

Statistic	Description
Clustering Nodes	Number of clustering nodes.

Logging

Statistic	Description
Currently Open Log Files	Number of event log files (formats) that are currently being written.
Space Used for Log Files	Current amount of space being used by the logging directory, which contains all of the event and error logs.
Number of Access Events Logged	Number of access events that have been written to log files since Content Gateway installation. This counter represents one entry in one file. If multiple formats are being written, a single access creates multiple event log entries.
Number of Access Events Skipped	Number of access events skipped (because they were filtered out) since Content Gateway installation.
Number of Error Events Logged	Number of access events that have been written to the event error log since Content Gateway installation.

Networking

Networking statistics are divided into the following categories:

Related concepts	
Virtual IP on page 261	

Related reference	
System on page 257	
ARM on page 258	
ICAP on page 259	
WCCP on page 259	
DNS Resolver on page 260	

System

Statistic/Field	Description
	General
Hostname	The hostname assigned to this Content Gateway machine.
Search Domain	Search domain that this Content Gateway machine uses.
IPv4 or IPv6	
Default Gateway	IP address of the default gateway used to forward packets from this Content Gateway machine to other networks or subnets.
Primary DNS	IP address of the primary DNS server that this Content Gateway machine uses to resolve host names.
Secondary DNS	Secondary DNS server that this Content Gateway machine uses to resolve host names.
Tertiary DNS	Third DNS server that this Content Gateway machine uses to resolve host names.
	NIC <interface_name></interface_name>
Status	Indicates whether the NIC is up or down.
Start on Boot	Indicates whether the NIC is configured to start on boot.
IPv4 or IPv6	
IP address	The assigned IP address of the NIC.
Netmask	The netmask that goes with the IP address.
Gateway	The configured default gateway IP address for the NIC.

ARM

Statistic	Description
	Network Address Translation (NAT) Statistics
Client Connections Natted	Number of client connections redirected transparently by the ARM.
Client Connections in Progress	Number of client connections currently in progress with the ARM.
Total Packets Natted	Number of packets translated by the ARM.
DNS Packets Natted	Number of DNS packets translated by the ARM.
	Bypass Statistics
Total Packets Bypassed	Total number of packets bypassed by the ARM.
Packets Dynamically Bypassed	Total number of packets dynamically bypassed. See
	Dynamic bypass rules.
DNS Packets Bypassed	Number of DNS packets bypassed by the ARM.
Packets Shed	Total number of packets shed.
	HTTP Bypass Statistics
Bypass on Bad Client Request	Number of requests forwarded directly to the origin server because Content Gateway encountered non-HTTP traffic on port 80.
Bypass on 400	Number of requests forwarded directly to the origin server because an origin server returned a 400 error.
Bypass on 401	Number of requests forwarded directly to the origin server because an origin server returned a 401 error.
Bypass on 403	Number of requests forwarded directly to the origin server because an origin server returned a 403 error.
Bypass on 405	Number of requests forwarded directly to the origin server because an origin server returned a 405 error.
Bypass on 406	Number of requests forwarded directly to the origin server because an origin server returned a 406 error.
Bypass on 408	Number of requests forwarded directly to the origin server because an origin server returned a 408 error.
Bypass on 500	Number of requests forwarded directly to the origin server because an origin server returned a 500 error.

Re	ated	concepts	
1.0	acca	concepto	

Dynamic bypass rules on page 59

ICAP

Statistic	Description
Total Posts	Total number of posts sent to Forcepoint DLP.
Total Analyzed	Total number of posts analyzed by Forcepoint DLP.
FTP Analyzed	Total number of FTP requests analyzed by Forcepoint DLP.
Blocked Requests	Total number of requests blocked after analysis and policy enforcement.
Allowed Requests	Total number of requests allowed after analysis and policy enforcement.
Failed Requests	Total number of posts sent to Forcepoint DLP that timed out or otherwise failed to complete.
Huge Requests	Total number of requests that exceeded the maximum transaction size.
Decrypted Requests	Total number of SSL requests decrypted and sent to Forcepoint DLP.

WCCP

WCCP v2 statistics are displayed only if WCCP version v2 is enabled.

Statistic/Field	Description	
	WCCP v2.0 Statistics	
WCCP Fragmentation		
Total Fragments	Total number of WCCP fragments.	
Fragmentation Table Entries	Number of entries in the fragmentation table.	
Out of Order Fragments	Number of fragments out of order.	
Matches	Number of fragments that match a fragment in the fragmentation table.	
Service group name		
Service Group ID	Service Group ID for the protocol being serviced.	
Configured mode	The forward, return and assignment settings.	
IP Address	IP address to which the router is sending traffic.	
Leader's IP Address	IP address of the leader in the WCCP cache farm.	
Number of Buckets Assigned	Number of buckets assigned to this Content Gateway node. Determined by the value of Weight and the current active nodes.	

Statistic/Field	Description	
Number of Caches	The number of caches in the WCCP cache farm.	
Number of Routers	The number of routers sending traffic to this Content Gateway node.	
Router IP Address	IP address of the WCCP router sending traffic to Content Gateway.	
	Note: If the WCCP router is configured with multiple IP addresses, as for example when the router is configured to support multiple VLANs, the IP address reported in Monitor > Networking > WCCP statistics, and in packet captures, may differ from the IP address configured here. This is because the router always reports traffic on the highest active IP address.	
	One way to get the router to always report the same IP address is to set the router's loopback address to a value higher than the router's highest IP address, then the loopback address is always reported as the router's IP address. This is the recommended configuration.	
Router ID Received	The number of times that Content Gateway has received WCCP protocol messages from the router(s).	
Router Negotiated mode	The return, forward, and assignment modes negotiated with the router.	

DNS Proxy

Statistic	Description
Total Requests	Total number of DNS requests received from clients.
Hits	Number of DNS cache hits.
Misses	Number of DNS cache misses.

DNS Resolver

Statistic	Description
	DNS Resolver
Total Lookups	Total number of DNS lookups (queries to name servers) since installation.
Successes	Total number of successful DNS lookups since installation.
Average Lookup Time (ms)	Average DNS lookup time.

Statistic	Description
	Host Database
Total Lookups	Total number of lookups in the Content Gateway host database since installation.
Total Hits	Total number of host database lookup hits since installation.
Average TTL (min)	Average time to live in minutes.

Virtual IP

The Virtual IP table displays the virtual IP addresses that are managed by the proxies in the cluster.

Client Connection Status

Statistic	Description
	Clients Connections
Current Unique Clients Connected	
Total Unique Clients that have Connected	Total since Content Gateway last started.
Total Clients that have Exceeded the Limits	Total clients that exceeded the connection limits since Content Gateway last started. See Configure > Connection Management > Client Connection Control.
Total Clients for which Connections were Closed	Total since Content Gateway last started.

Performance

Performance graphs allow you to monitor Content Gateway performance and analyze network traffic. Performance graphs also provide information about virtual memory usage, client connections, and so on.

Performance graphs are created by the Multi Router Traffic Grapher tool (MRTG). MRTG uses 5-minute intervals to accumulate statistics.

Performance graphs provide the following information.

Statistic	Description
Overview	Displays a subset of the graphs available.
Daily	Displays graphs that provide historical information for the current day.
Weekly	Displays graphs that provide historical information for the current week.

Statistic	Description
Monthly	Displays graphs that provide historical information for the current month.
Yearly	Displays graphs that provide historical information for the current year.



Important

To run the Multi Router Traffic Grapher tool in Linux, you must have Perl version 5.005 or later installed on your Content Gateway system.

A description is given adjacent to each graph. Click on a graph to get the daily, weekly, monthly, and yearly on a single screen.

These graphs are available (sorted alphabetically):

- Active Client Connections
- Active Native FTP Client Connections
- Active Origin Server Connections
- Active Parent Proxy Connections
- Analytic Response Latency
- Bandwidth Savings
- Completed Client Transactions Per Second
- Content Gateway Manager Memory Usage
- Content Gateway Uptime
- CPU Available
- CPU Busy
- Web DLP Module Memory Usage
- DNS Cache Usage
- DNS Lookup Latency
- HTTP Abort Latency
- HTTP and HTTPS Transactions Per Second
- HTTP Connection Errors & Aborts (Count)
- HTTP Connection Errors & Aborts (Percentage)
- HTTP Error Latency
- HTTP Hits & Misses (Count)
- HTTP Hits & Misses (Percentage)
- HTTP POST and FTP PUT Transactions Per Second
- IWA Basic & NTLM Latency
- IWA Negotiate Latency
- Microsoft Internet Explorer Browser Requests (Percentage)
- MRTG Runtime
- Network Reads
- Network Writes
- Origin Server Connection Latency

- Outbound Analysis Latency
- RAM Cache Read I/O Hit Rate
- RAM Cache Usage
- System Memory
- TCP CLOSE_WAIT Connections
- TCP Connect Rate
- TCP ESTABLISHED Connections
- TCP FIN_WAIT_1 Connections
- TCP FIN_WAIT_2 Connections
- TCP LAST_ACK Connections
- TCP Segments Transmitted
- TCP Throughput
- TCP TIME_WAIT Connections
- Throughput in Bytes
- Throughput in Error and Dropped Packets
- Throughput in Packets
- Transaction Buffer Memory Usage
- URL Policy Lookup Latency
- WCCP Exceptional Input Fragments
- WCCP Fragment Table Size
- WCCP Input Fragments
- Scanned Transactions (Percentage)
- Slow Scanned Transactions
- Slow Transactions
- Content Gateway Memory Usage

SSL

The following tabs monitor and report on SSL traffic.

Related concepts

Reports on page 265

Related reference

SSL Key Data on page 264 CRL Statistics on page 264

SSL Key Data

These fields	provide information	about SSL	connections	and activity.
--------------	---------------------	-----------	-------------	---------------

Statistic/Field	Description
	SSL Inbound Key Data
Is alive	Online indicates that SSL support is enabled.
Current SSL connections	The number of active inbound SSL requests (browser to Content Gateway).
Total SSL server connections	The number of browser requests.
Total finished SSL server connections	The number of browser requests that resulted in decryption.
Total SSL renegotiation requests sent by Content Gateway as a server	The number of browser requests renegotiated due to handshake failures or invalid certificates between the browser and Content Gateway.
	SSL Outbound Key Data
Is alive	Online indicates that SSL support is enabled.
Current SSL connections	The number of active outbound SSL requests (Content Gateway to origin server).
Total SSL client connections	The number of Content Gateway requests to origin servers.
Total finished SSL client connections	The number of requests where data went from Content Gateway to the origin server.
Total SSL renegotiation requests sent by Content Gateway as a client	The number of requests that were renegotiated due to handshake failures or invalid certificates between Content Gateway and the origin server
Total SSL session cache hits	The number of times that a request was validated by a key in the session cache.
Total SSL session cache misses	The number of times that a request could not be validated by a key in the session cache.
Total SSL session cache timeouts	The number of times that keys were removed from the session cache because the timeout period expired.

CRL Statistics

These fields provide information about certificate status.

Statistic/Field	Description
	CRL Statistics

Statistic/Field	Description
CRL list count	The number of certificates on the Certificate Revocation List. This list is downloaded every night. See <i>Keeping revocation information up to date</i> .
	OCSP Statistics
OCSP good count	The number of responses that certificates are valid.
OCSP unknown count	The number of OCSP responses where the certificate cannot be verified.
OCSP revoked count	The number of certificates found to have been revoked.

Related concepts

Keeping revocation information up to date on page 134

Reports

See *Creating SSL certificate authorities reports* and *Creating an SSL incidents report* for information about creating reports on certificate authorities or incidents.

Related tasks

Creating SSL certificate authorities reports on page 99 Creating an SSL incidents report on page 100

Appendix B Commands and Variables

Contents

- Content Gateway commands on page 267
- Content Gateway variables on page 269

Content Gateway commands

Use the command line to execute individual commands and when scripting multiple commands in a shell.

Run commands as 'root'.

Execute Content Gateway commands from the Content Gateway bin directory.



Note

If the Content Gateway bin directory is not in your path, prepend the command with:

./

For example:

./content_line -p

Command	Description
WCGAdminstart	Starts the Content Gateway service
WCGAdminstop	Stops the Content Gateway service
WCGAdminrestart	Stops the Content Gateway service and then starts it again
WCGAdminstatus	Displays the status (running or not running) of the Content Gateway services: Content Cop, Content Gateway, Content Gateway Manager, and Analytics Server.
WCGAdminhelp	Displays a list of the WCGAdmin commands
content_line -h	Displays the list of Content Gateway commands.

Command	Description
content_line-p socket_path	Specifies the location (directory and path) of the file used for Content Gateway command line and Content Gateway manager communication. The default path is install_dir/config/cli
content_line-r variable	Displays specific performance statistics or a current configuration setting. For a list of the variables you can specify, see <i>Content</i> <i>Gateway variables</i> .
content_line -s variable -v value	Sets configuration variables. <i>variable</i> is the configuration variable you want to change and <i>value</i> is the value you want to set. See <i>records.config</i> , for a list of the configuration variables you can specify.
content_line-x	Initiates a Content Gateway configuration file reread. Executing this command is similar to clicking Apply in the Content Gateway manager.
content_line-y	Clears Forcepoint dynamically signed certificates from the cache and the SSL sqlite database.
content_linedb_clear -y	Clears Forcepoint dynamically signed certificates from the SSL sqlite database.
content_line-M	Restarts the content_manager process and the content_gateway process on all the nodes in a cluster.
content_line-L	Restarts the content_manager process and the content_gateway process on the local node.
content_line-S	Shuts down Content Gateway on the local node.
content_line-U	Starts Content Gateway on the local node.
content_line-B	Bounces Content Gateway cluster- wide. Bouncing Content Gateway shuts down and immediately restarts the proxy node-by-node.

Command	Description
content_line-b	Bounces Content Gateway on the local node. Bouncing Content Gateway shuts down and immediately restarts the proxy on the local node.
content_line-W	Enables WCCP router communication.
content_line-w	Disables WCCP router communication. After changing the Content Gateway WCCP configuration, or the router WCCP configuration, force WCCP communication down for 60 seconds to force WCCP to negotiate a new connection.
content_line -N snapshot_name	Perform a Content Gateway snapshot (backup). See <i>Taking</i> configuration snapshots.
content_line -n snapshot_name	Restore a Content Gateway snapshot. See <i>Restoring</i> <i>configuration snapshots.</i>

Related conceptsContent Gateway variables on page 269records.config on page 387Taking configuration snapshots on page 89Restoring configuration snapshots on page 90

Content Gateway variables

You can change the value of a specific configuration variable on the command line with the content_line -s command. The variables that can be set are described in *records.config*.

You can view statistics related to specific variables on the command line with the content_line -r command. See below for a list of variables.

See, also, Viewing statistics from the command line and Using the command-line interface.

Related concepts

records.config on page 387

Related tasks

Viewing statistics from the command line on page 94 Using the command-line interface on page 23

Statistics

The following table lists the variables you can specify on the command line to view individual statistics. See *Statistics* for additional information.

To view a statistic, at the prompt enter:

content_line -r <variable>

Variable
Summary
proxy.node.hostname
proxy.node.user_agents_total_documents
_served
proxy.node.user_agent_xacts_per_second
Node
proxy.node.current_server_connections
proxy.cluster.current_server_connections
proxy.node.current_client_connections
proxy.cluster.current_client_connections
proxy.node.client_throughput_out
proxy.cluster.client_throughput_out
<pre>proxy.node.http.user_agent_xacts_per</pre>
_second
proxy.cluster.http.user_agent_xacts_per
_second
proxy.node.dns.lookups_per_second
proxy.cluster.dns.lookups_per_second
proxy.node.hostdb.hit_ratio_avg_10s
proxy.cluster.hostdb.hit_ratio_avg_10s
нттр
proxy.process.http.user_agent_response
_document_total_size
proxy.process.http.user_agent_response
_header_total_size
proxy.process.http.current_client
_connections

Statistic	Variable
Current unique clients connected	proxy.process.http.client.unique _clients.active
Total unique clients that have connected	proxy.process.http.client.unique _clients.total
Total clients that exceeded limit	proxy.process.http.client.exceeding_limit
Total clients for which connections were closed	proxy.process.http.client.closed _connections
Open HTTP client connections	proxy.process.http.current_active_http _client_connections
Open HTTPS client connections	proxy.node.process.http.current_active_ https_client_connections
Client Requests (IPv4 +IPv6)	proxy.process.http.real_client_requests
Client IPv6 Requests	proxy.process.http.real_client _ipv6_requests
Client transactions in progress	proxy.process.http.current_client _transactions
Total document bytes from origin server	proxy.process.http.origin_server_response _document_total_size
Total header bytes from origin server	proxy.process.http.origin_server_response _header_total_size
Total connections to origin server	proxy.process.http.current_server _connections
Origin server transactions in progress	proxy.process.http.current_server _transactions
	FTP
Currently open FTP connections	proxy.process.ftp.connections _currently_open
Successful PASV connections	proxy.process.ftp.connections_successful _pasv
Unsuccessful PASV connections	proxy.process.ftp.connections_failed_pasv

Successful PORT connections proxy.process.ftp.connections_successful unsuccessful PORT connections proxy.process.ftp.connections_failed_port WCCP model Enabled proxy.config.wccp.enabled WCCP interface proxy.process.hostdb.total_interface Total lookups proxy.process.hostdb.total_lookups Total hits proxy.process.hostdb.total_hits Time TTL (min) proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.dns.total_dns_lookups DNS successes proxy.process.cluster.read_bytes Bytes read proxy.process.cluster.read_bytes Bytes written proxy.process.cluster.onnections_open Total operations proxy.process.cluster.net_backup Cluster proxy.process.cluster.net_backup Bytes written proxy.process.cluster.connections_open Total operations proxy.process.cluster.net_backup Clustering nodes proxy.process.socks.connections_open Successful connections proxy.process.cluster.net_backup Clustering nodes proxy.process.cluster.net_backup Cuscessful connections proxy.process.cluster.net_backup Successful connections pr	Statistic	Variable
_port Unsuccessful PORT connections proxy.process.ftp.connections_failed_port WCCP Enabled proxy.config.wccp.enabled WCCP interface proxy.local.wccp2.ethernet_interface WCCP proxy.process.hostdb.total_lockups Total lookups proxy.process.hostdb.total_hits Total hits proxy.process.hostdb.total_hits Time TTL (min) proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.cluster.read_bytes Bytes read proxy.process.cluster.onnections_open Total operations proxy.process.cluster.connections_open Total operations proxy.process.cluster.net_backup Clustering nodes proxy.process.socks.connections_open Successful connections proxy.process.socks.connections_currently Successful connections proxy.process.socks.connections_currently Questions in progress proxy.process.socks.connections_currently _open proxy.process.socks.connections_currently	Successful PORT connections	proxy.process.ftp.connections_successful
Unsuccessful PORT connections proxy.process.ftp.connections_failed_port WCCP Enabled proxy.config.wccp.enabled WCCP interface proxy.local.wccp2.ethernet_interface Mott DB Proxy.process.hostdb.total_lookups Total lookups proxy.process.hostdb.total_hits Time TTL (min) proxy.process.hostdb.total_hits DNS DNS DNS total lookups proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.cluster.read_bytes Bytes read proxy.process.cluster.onnections_open Total operations proxy.process.cluster.connections_open Total operations proxy.process.cluster.net_backup Clustering nodes proxy.process.cluster.net_backup Socks proxy.process.cluster.net_backup Unsuccessful connections proxy.process.socks.connections_open Successful connections proxy.process.socks.connections_open Connections in progress proxy.process.socks.connections_open		_port
WCCPEnabledproxy.config.wccp.enabledWCCP interfaceproxy.local.wccp2.ethernet_interfaceHost DBTotal lookupsproxy.process.hostdb.total_lookupsTotal hitsproxy.process.hostdb.total_hitsTime TTL (min)proxy.process.hostdb.total_dns_lookupsDNS total lookupsproxy.process.dns.total_dns_lookup_avg_timeDNS successesproxy.process.dns.lookup_avg_timeDNS successesproxy.process.cluster.read_bytesBytes readproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.net_backupUnsuccessful connectionsproxy.process.cluster.net_backupConnections in progressproxy.process.socks.connections_currently _openLoogingproxy.process.socks.connections_currently	Unsuccessful PORT connections	proxy.process.ftp.connections_failed_port
Enabledproxy.config.wccp.enabledWCCP interfaceproxy.local.wccp2.ethernet_interfaceWCCP interfaceHost DBTotal lookupsproxy.process.hostdb.total_lookupsTotal hitsproxy.process.hostdb.total_hitsTime TTL (min)proxy.process.hostdb.total_dns_lookupsDNSproxy.process.dns.total_dns_lookupsAverage lookup time (ms)proxy.process.dns.lookup_avg_timeDNS successesproxy.process.cluster.read_bytesBytes readproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_opendTotal operationsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.net_backupUnsuccessful connectionsproxy.process.cluster.net_backupSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open		WCCP
WCCP interfaceproxy.local.wccp2.ethernet_interfaceHost DBTotal lookupsproxy.process.hostdb.total_lookupsTotal hitsproxy.process.hostdb.total_hitsTime TTL (min)proxy.process.hostdb.total_hitsDNSproxy.process.dns.total_dns_lookupsAverage lookup time (ms)proxy.process.dns.lookup_avg_timeDNS successesproxy.process.cluster.read_bytesBytes readproxy.process.cluster.write_bytesConnections openproxy.process.cluster.onnections_opendTotal operationsproxy.process.cluster.net_backupClustering nodesproxy.process.socks.connections_ unsuccessfulSuccessful connectionsproxy.process.socks.connections _unsuccessfulSuccessful connections in progressproxy.process.socks.connections_currently _open	Enabled	proxy.config.wccp.enabled
Host DBTotal lookupsproxy.process.hostdb.total_lookupsTotal hitsproxy.process.hostdb.total_hitsTime TTL (min)proxy.process.hostdb.total_dns_lookupsDNSproxy.process.dns.total_dns_lookupsAverage lookup time (ms)proxy.process.dns.lookup_avg_timeDNS successesproxy.process.dns.lookup_successesBytes readproxy.process.cluster.read_bytesBytes writtenproxy.process.cluster.onnections_openTotal operationsproxy.process.cluster.net_bockupNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.net_backupUnsuccessful connectionsproxy.process.socks.connections_openSuccessful connectionsproxy.process.socks.connections_openSuccessful connectionsproxy.process.socks.connections_openConnectionsproxy.process.socks.connections_openCustering nodesproxy.process.socks.connections_openCustering nodesproxy.process.socks.connections_openCustering nodesproxy.process.socks.connections_openSuccessful connectionsproxy.process.socks.connections_openCustering nodesproxy.process.socks.connections_openSuccessful connectionsproxy.process.socks.connections_openConnections in progressproxy.process.socks.connections_openConnections in progressproxy.process.socks.connections_open	WCCP interface	proxy.local.wccp2.ethernet_interface
Total lookups proxy.process.hostdb.total_lookups Total hits proxy.process.hostdb.total_hits Time TTL (min) proxy.process.hostdb.ttl DNS proxy.process.dns.total_dns_lookups DNS total lookups proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.dns.lookup_avg_time DNS successes proxy.process.cluster.read_bytes Bytes read proxy.process.cluster.vrite_bytes Connections open proxy.process.cluster.connections_open Total operations proxy.process.cluster.net_backup Clustering nodes proxy.process.cluster.net_backup Unsuccessful connections proxy.process.socks.connections_open Successful connections proxy.process.socks.connections_open Successful connections proxy.process.socks.connections_open Successful connections proxy.process.socks.connections_open Connections proxy.process.socks.connections		Host DB
Total hitsproxy.process.hostdb.total_hitsTime TTL (min)proxy.process.hostdb.ttlDNSDNSDNS total lookupsproxy.process.dns.total_dns_lookupsAverage lookup time (ms)proxy.process.dns.lookup_avg_timeDNS successesproxy.process.dns.lookup_successesBytes readproxy.process.cluster.read_bytesBytes writtenproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_opendTotal operationsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connections_opendConnections in progressproxy.process.socks.connections_currently _open	Total lookups	proxy.process.hostdb.total_lookups
Time TTL (min)proxy.process.hostdb.ttlIme TTL (min)proxy.process.hostdb.ttlIme TTL (min)proxy.process.dbs.total_dbs_lookupsDNS total lookupsproxy.process.dbs.total_dbs_lookupsAverage lookup time (ms)proxy.process.dbs.lookup_avg_timeDNS successesproxy.process.dbs.lookup_successesDNS successesproxy.process.cluster.read_bytesBytes readproxy.process.cluster.write_bytesBytes writtenproxy.process.cluster.connections_openTotal operationsproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.nodesClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open	Total hits	proxy.process.hostdb.total_hits
Image:	Time TTL (min)	proxy.process.hostdb.ttl
DNS total lookups proxy.process.dns.total_dns_lookups Average lookup time (ms) proxy.process.dns.lookup_avg_time DNS successes proxy.process.dns.lookup_successes DNS successes Cluster Bytes read proxy.process.cluster.read_bytes Bytes written proxy.process.cluster.write_bytes Connections open proxy.process.cluster.connections_opend Total operations proxy.process.cluster.net_backup Network backups proxy.process.cluster.nodes Clustering nodes proxy.process.socks.connections_opend Unsuccessful connections proxy.process.cluster.nodes Successful connections proxy.process.socks.connections_opend Successful connections proxy.process.socks.connections_opend Connections in progress proxy.process.socks.connections_opend Open proxy.process.socks.connections_opend		DNS
Average lookup time (ms)proxy.process.dns.lookup_avg_timeDNS successesproxy.process.dns.lookup_successesBytes readproxy.process.cluster.read_bytesBytes writtenproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_opendTotal operationsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open	DNS total lookups	proxy.process.dns.total_dns_lookups
DNS successes proxy.process.dns.lookup_successes Gluster Bytes read proxy.process.cluster.read_bytes Bytes written proxy.process.cluster.write_bytes Connections open proxy.process.cluster.connections_opend Total operations proxy.process.cluster.connections_opend Network backups proxy.process.cluster.net_backup Clustering nodes proxy.process.cluster.nodes Unsuccessful connections proxy.process.socks.connections_ _unsuccessful Successful connections proxy.process.socks.connections_ _unsuccessful Connections in progress proxy.process.socks.connections_ _uopen	Average lookup time (ms)	proxy.process.dns.lookup_avg_time
ClusterBytes readproxy.process.cluster.read_bytesBytes writtenproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_opendTotal operationsproxy.process.cluster.net_backupNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open	DNS successes	proxy.process.dns.lookup_successes
Bytes readproxy.process.cluster.read_bytesBytes writtenproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_openTotal operationsproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.net_backupUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open		Cluster
Bytes writtenproxy.process.cluster.write_bytesConnections openproxy.process.cluster.connections_opendTotal operationsproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open	Bytes read	proxy.process.cluster.read_bytes
Connections openproxy.process.cluster.connections_openTotal operationsproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsSuccessful connectionsproxy.process.socks.connectionsConnections in progressproxy.process.socks.connections_currently _open	Bytes written	proxy.process.cluster.write_bytes
Total operationsproxy.process.cluster.connections_openedNetwork backupsproxy.process.cluster.net_backupClustering nodesproxy.process.cluster.nodesUnsuccessful connectionsproxy.process.socks.connections _unsuccessfulSuccessful connectionsproxy.process.socks.connections _successfulConnections in progressproxy.process.socks.connections_currently _openLoggingLogging	Connections open	proxy.process.cluster.connections_open
Network backups proxy.process.cluster.net_backup Clustering nodes proxy.process.cluster.nodes SOCKS SOCKS Unsuccessful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections _unsuccessful proxy.process.socks.connections _open _open	Total operations	proxy.process.cluster.connections_opened
Clustering nodes proxy.process.cluster.nodes Image: Clustering nodes SOCKS Unsuccessful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections Connections in progress proxy.process.socks.connections_currently _open Logging	Network backups	proxy.process.cluster.net_backup
SOCKSUnsuccessful connectionsproxy.process.socks.connections _unsuccessfulSuccessful connectionsproxy.process.socks.connections _successfulConnections in progressproxy.process.socks.connections_currently _openLoggingLogging	Clustering nodes	proxy.process.cluster.nodes
Unsuccessful connections proxy.process.socks.connections Successful connections proxy.process.socks.connections Successful proxy.process.socks.connections Connections in progress proxy.process.socks.connections_currently _open Logging		SOCKS
_unsuccessful Successful connections proxy.process.socks.connections _successful Connections in progress proxy.process.socks.connections_currently _open Logging	Unsuccessful connections	proxy.process.socks.connections
Successful connections proxy.process.socks.connections _successful Connections in progress proxy.process.socks.connections_currently _open Logging		_unsuccessful
_successful Connections in progress proxy.process.socks.connections_currently _open Logging	Successful connections	proxy.process.socks.connections
Connections in progress proxy.process.socks.connections_currently _open		_successful
open	Connections in progress	proxy.process.socks.connections_currently
Logging		_open
		Logging
Currently open log files proxy.process.log2.log_files_open	Currently open log files	proxy.process.log2.log_files_open
Space used for log files proxy.process.log2.log_files_space_used	Space used for log files	proxy.process.log2.log_files_space_used
Number of access events logged proxy.process.log2.event_log_access	Number of access events logged	proxy.process.log2.event_log_access
Number of access events skipped proxy.process.log2.event_log_access_skip	Number of access events skipped	proxy.process.log2.event_log_access_skip

Statistic	Variable
Number of error events logged	proxy.process.log2.event_log_error

Related information	
Statistics on page 243	

Appendix C Configuration Options

Contents

- My Proxy on page 275
- Protocols on page 287
- Content Routing on page 295
- Security on page 300
- Subsystems on page 324
- Networking on page 329
- SSL on page 353

Options are grouped as follows on the left side of the Configure pane:

Related concepts

My Proxy on page 275 Protocols on page 287 Content Routing on page 295 Security on page 300 Subsystems on page 324 Networking on page 329

My Proxy

The My Proxy options are:

Related reference	
Basic on page 275	
Subscription on page 280	
UI Setup on page 281	
Snapshots on page 284	
Logs	

Basic

Configure > My Proxy > Basic > General

Restart	Restarts the proxy and manager services (the content_gateway and content_manager processes). You must restart the proxy and manager services after modifying certain configuration options. A message is displayed in the manager when a restart is required. Important In a cluster configuration, the Restart button restarts the proxy and manager services on all nodes in the cluster.
Proxy Name	Specifies the name of your Content Gateway node. By default, this is the hostname of the machine running Content Gateway.
	name of the Content Gateway cluster. In a cluster, all nodes must share the same name.
	Valid characters for Proxy Name are: A-Z, a-z,0-9 and
Alarm email	Specifies the email address to which Content Gateway sends alarm notifications.
Features	
Protocols: FTP	When this option is enabled, Content Gateway accepts FTP requests from FTP clients.
	If this option is changed you must restart Content Gateway.
Protocols: HTTPS	Enables/disables Content Gateway HTTPS traffic management and security analysis. After selecting HTTPS On , you must provide additional information about the Configure > Protocols > HTTPS page and on the Configure > SSL pages. See <i>Working With</i> <i>Encrypted Data</i> .
Networking: WCCP	Enable this option to use a WCCP v2-enabled router for transparent redirection to Content Gateway.
	WCCP v1 is not supported.
	See Transparent interception with WCCP v2 devices.
	If you change this option, you must restart Content Gateway.
Networking: DNS Proxy	When this option is enabled, Content Gateway resolves DNS requests on behalf of clients. This option offloads remote DNS servers and reduces response time for DNS lookups. See <i>DNS Proxy Caching</i> .
Networking: Virtual IP	When this option is enabled, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes in a cluster as necessary. See <i>Virtual IP failover</i> .

Networking: IPv6	When this option is enabled, Content Gateway provides support for IPv6.
	IPv6 addresses can be used on any dual stack Ethernet interface that services client and/or Internet traffic.
	IPv4 addresses must be used to communicate with all Forcepoint components.
	To see a complete description of the feature and an important list of restrictions, see <i>Content Gateway support for IPv6</i> .
Networking: Web DLP	Enables a connection to Forcepoint DLP. There are 2 options:
	 Automatic registration through the Forcepoint management server
	 ICAP communication to a remote Forcepoint DLP deployment (not recommended)
	See Working With Web DLP.
	If you change this option, you must restart Content Gateway.
Networking: Integration > Web DLP (integrated on-box)	Enables registration with the on-box Web DLP components and the Forcepoint management server. See <i>Registering Content Gateway with Forcepoint</i> <i>DLP</i> .
Networking: Web DLP: ICAP	Enables ICAP for use with Forcepoint DLP. See Configuring the ICAP client.
Security: SOCKS	When SOCKS is enabled, Content Gateway communicates with your SOCKS servers. See <i>Configuring SOCKS firewall integration</i> . If you change this option, you must restart Content Gateway.
Authentication: None	Content Gateway supports several types of user authentication. When this option is selected, the proxy does not perform user authentication. This is the default setting.
Authentication: Integrated Windows Authentication	When Integrated Windows Authentication (IWA) is enabled, users are authenticated by IWA before they are allowed access to content.
	See Integrated Windows Authentication.
	If you change this option, you must restart Content Gateway.

Authentication: LDAP	When LDAP is enabled, users are authenticated by an LDAP server before they are allowed access to content.See <i>LDAP authentication</i> .
	If you change this option, you must restart Content Gateway.
Authentication: Radius	When RADIUS is enabled, users are authenticated by a RADIUS server before they are allowed access to content. See <i>RADIUS authentication</i> .
	If you change this option, you must restart Content Gateway.
Authentication: Legacy NTLM	When legacy NTLM (NTLMSSP) is enabled, users in a Windows network are authenticated by a Domain Controller before they are allowed access to content.
	See Legacy NTLM authentication.
	If you change this option, you must restart Content Gateway.
Authentication: Rule-Based Authentication	When Rule-Based Authentication is enabled, users are authenticated based on the parameters of the rule that they match. Rule-based authentication supports multiple realm, multiple domain, and other user authentication scenarios. See <i>Rule-Based</i> <i>Authentication</i> .
	If you change this option, you must restart Content Gateway.
Authentication: Read authentication from child proxy	Enables or disables the reading of X-Authenticated- User and X-Forwarded-For header values in incoming requests. This option is disabled by default.
	Enable this option when Content Gateway is the parent (upstream) proxy in a chain and the child (downstream) proxy is sending X-Authenticated- User and X-Forwarded-For header values to facilitate authentication.
Authentication: Send authentication to parent proxy	Enables or disables the insertion of X-Authenticated- User header values in outgoing requests. This option is disabled by default.
	Enable this option when Content Gateway is the child (downstream) proxy in a chain and the parent (upstream) proxy wants X-Authenticated-User values to facilitate authentication.
	If this option is enabled, the user name will be sent only to a configured parent proxy. To send user names to all outbound requests, enable proxy.config.http.insert_xua_to_external.

Configure > My Proxy > Basic > Clustering

Cluster: Type	Specifies the clustering mode:
	Select Single Node to run this Content Gateway server as a single node. This node will not be part of a cluster.
	Select Management Clustering to activate management clustering mode. The nodes in the cluster share configuration information and you can administer all the nodes at the same time.
	For complete information about clustering, see
	Clusters.
	If you change this option, you must restart Content Gateway.
Cluster: Interface	Specifies the interface on which Content Gateway communicates with other nodes in the cluster. For example, eth1 .
	It is recommended that you use a dedicated secondary interface.
	Node configuration information is multicast, in plain text, to other Content Gateway nodes on the same subnet. Therefore, as a best practice, clients should be located on a separate subnet from Content Gateway nodes (multicast communications for clustering are not routed).
	On appliances, P1 is the recommended interface. However, you may also use P2 if you are not using it for Internet egress traffic and want to isolate cluster management traffic.
	See Changing clustering configuration.
	If you change this option, you must restart Content Gateway.
Cluster: Multicast Group Address	Specifies the multicast group address on which Content Gateway communicates with its cluster peers.
	See Changing clustering configuration.

Related concepts

Transparent interception with WCCP v2 devices on page 38

- Virtual IP failover on page 76
- Content Gateway support for IPv6 on page 69
- Configuring SOCKS firewall integration on page 160
- Integrated Windows Authentication on page 173
- LDAP authentication on page 180
- **RADIUS** authentication on page 184
- Legacy NTLM authentication on page 179
- Rule-Based Authentication on page 186

Related tasks

Registering Content Gateway with Forcepoint DLP on page 103 Configuring the ICAP client on page 106 Changing clustering configuration on page 72 Providing system information on page 21

Related information

Working With Encrypted Data on page 111 DNS Proxy Caching on page 85 Working With Web DLP on page 101 Clusters on page 71

Subscription

Configure > My Proxy > Subscription > Subscription Management

Subscription Key	Displays the subscription key you received from Forcepoint LLC.
	If Content Gateway is used with Forcepoint Web Security, this is the subscription key you entered in the Web Security module of the Forcepoint Security Manager.
	If Content Gateway is deployed with only Forcepoint DLP, you must enter your Content Gateway subscription key in this field.

Configure > My Proxy > Subscription > Scanning

Policy Server	
IP address	The IP address of the Policy Server. This value is specified when Content Gateway is installed.
Port	The port used by Policy Server. The default port is 55806.
Filtering Service	
IP address	Specify the IP address of the Filtering Service. This value is specified when Content Gateway is installed.
Port	Specify the port used by Filtering Service. The default port is 15868.
Communication Timeout	Specifies the timeout, in milliseconds, in which Policy Server and Filtering Service must respond before a communication timeout condition occurs and the Action for Communication Errors setting is applied. The default value is 5000 ms (5 seconds).

Action for Communication Errors	
Permit traffic	Permits all traffic if communication with Policy Server or Filtering Service fails.
Block traffic	Blocks all traffic if communication with Policy Server or Filtering Service fails.
Scanning Data Files Update	
Delay time	Specifies the length of time scanning data file downloads are delayed. The default value is No delay.
	See the Scanning Data Files Opdate section of
	Providing system information.

Related tasks

Providing system information on page 21

UI Setup

Configure > My Proxy > UI Setup > General

UI Port	Specifies the port on which browsers can connect to the Content Gateway manager. The default port is 8081. If you change this setting, you must restart Content Gateway.
HTTPS: Enable/Disable	Enables or disables support for SSL connections to the Content Gateway manager (enabled by default). SSL provides protection for remote administrative monitoring and configuration. To use SSL for Content Gateway manager connections, you must install an SSL certificate on the Content Gateway server machine. For more information, see <i>Using SSL for</i> <i>secure administration</i> .
HTTPS: Certificate File	Specifies the name of the SSL certificate file used to authenticate users who want to access the Content Gateway manager.
Monitor Refresh Rate	Specifies how often Content Gateway manager refreshes the statistics on the Monitor pane. The default value is 30 seconds.
Default Help Language	Specifies the language that Content Gateway Manager Help displays by default. If a page is not available in the default language, another language may be substituted.

Configure > My Proxy > UI Setup > Login

Administrator: Login	Specifies the administrator login. The default is
	The administrator login is the master login that has access to both Configure and Monitor mode in the Content Gateway manager.
Administrator: Password	Lets you change the administrator password that controls access to the Content Gateway manager.
	Enter the current password in the Old Password field. Enter the new password in the New Password field, re- enter it in the New Password (Retype) field, and then click Apply .
	Passwords must be 8 to 15 characters and include at least one:
	 Uppercase character
	Lowercase character
	Number
	Special character
	Supported characters include:
	! # % & ' () * + , / ; < = > ? @ [] ^_{
	}~
	The following special characters are not supported:
	Space \$: ` \ "
	During installation, you select the administrator password. The installer automatically encrypts the password and stores the encryptions in the records.config file so that no one can read them. Each time you change the password in the Content Gateway manager, Content Gateway updates the records.config file. If you forget the administrator password and cannot access the Content Gateway manager, see Accessing the Content Gateway manager if you forget the master administrator password.

Additional Users	Lists the current user accounts and lets you add new user accounts. User accounts determine who has access the Content Gateway manager and which activities they can perform. You can create a list of user accounts if a single administrator login and password is not sufficient security for your needs.
	To create a new account, enter the user login in the New User field, and then enter the user password in the New Password field. Retype the user password in the New Password (Retype) field, and then click Apply .
	Passwords must be 8 to 15 characters and include at least one:
	 Uppercase character
	 Lowercase character
	Number
	 Special character
	Supported characters include:
	! # % & ' () * + , / ; < = > ? @ [] ^ {
	}~
	The following special characters are not supported:
	Space \$: ` \ "
	Information for the new user is displayed in the table. From the Access drop-down list in the table, select the activities that the new user can perform (Monitor , Monitor and View Configuration , or Monitor and Modify Configuration). For more information about user accounts, see <i>Creating a list of user accounts</i> .

Configure > My Proxy > UI Setup > Access

Access Control	Displays a table listing the rules in the <i>mgmt_allow.config</i> file. Rules specify the remote hosts allowed to access the Content Gateway manager. The entries in this file ensure that only authenticated users can change configuration options and view performance and network traffic statistics. Note: By default, all remote hosts are allowed to access the Content Gateway manager.
Refresh	Updates the table to display the most up-to-date rules in the mgmt_allow.config file.
Edit File	Opens the configuration file editor so that you can edit and add rules to the mgmt_allow.config file.
	mgmt_allow.config Configuration File Editor

rule display box	Lists the mgmt_allow.config file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. Content Gateway applies the rules in the order listed, starting from the top.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
IP Action	Lists the type of rules you can add.
	An ip_allow rule allows the remote hosts specified in the Source IP field to access the Content Gateway manager.
	An ip_deny rule denies the remote hosts specified in the Source IP field access to the Content Gateway manager.
Source IP	Specifies the IP addresses that are allowed or denied access to the Content Gateway manager. You can enter a single IP address (111.111.11.1) or a range of IP addresses (0.0.0.255.255.255.255).
Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Related concepts

mgmt_allow.config on page 384

Related tasks

Using SSL for secure administration on page 153 Accessing the Content Gateway manager if you forget the master administrator password on page 19 Creating a list of user accounts on page 152

Snapshots

Configure > My Proxy > Snapshots > File System

Change Snapshot Directory	Specifies the directory in which snapshots are stored
	on this Content Gateway hode.

Snapshots: Save Snapshot	Specifies the name of the configuration snapshot you want to take. Click Apply to save the configuration on the local node. Content Gateway saves the configuration snapshot in the directory specified in the Change Snapshot Directory field.
	It is recommended that you take a snapshot before performing system maintenance or attempting to tune system performance. Taking a snapshot takes only a few seconds and can save you hours of correcting configuration mistakes.
Snapshots: Restore/Delete Snapshot	Lists the snapshots that are stored on this node. Select the snapshot that you want to restore or delete from the drop-down list.
Snapshots: Restore Snapshot from "directory_name" Directory	Restores the snapshot selected in the Restore/ Delete Snapshot drop-down box. In a cluster configuration, snapshots are restored on all nodes in the cluster
Snapshots: Delete Snapshot from "directory_name" Directory	Deletes the snapshot selected in the Restore/ Delete Snapshot drop-down box.

Configure > My Proxy > Snapshots > FTP server

FTP Server	Specifies the name of the FTP server from which you want to restore a configuration snapshot or to which you want to save a configuration snapshot.
Login	Specifies the login needed to access the FTP server.
Password	Specifies the password needed to access the FTP server.
Remote Directory	Specifies the directory on the FTP server from which you want restore, or in which you want to save a configuration snapshot.
Restore Snapshot	Lists the configuration snapshots on the FTP server that you can restore. This field appears after you have logged on to the FTP server successfully.
Save Snapshot to FTP Server	Specifies the name of the configuration snapshot you want to take and save on the FTP server. This field appears after you have logged on to the FTP server successfully.

Logs

Configure > My Proxy > Logs > System

Log File	Lists the system log files you can view, delete or copy to your local system. Content Gateway lists the system log files logged with the system-wide logging facility syslog under the daemon facility.
Action: Display the selected log file	When this option is enabled, Content Gateway displays the first MB of the system log file selected in the Log File drop-down list.
	To view the entire file, select "Save the selected log file in local filesystem" and view the file with a local viewer.
Action: Display last lines of the selected file	When this option is enabled, Content Gateway displays the last specified number of lines in the selected system log file.
Action: Display lines that match in the selected log file	When this option is enabled, Content Gateway displays all the lines in the selected system log file that match the specified string.
Action: Remove the selected log file	When this option is enabled, Content Gateway deletes the selected log file.
Action: Save the selected log file in local filesystem	When this option is enabled, Content Gateway saves the selected log file on the local system in a location you specify.

Configure > My Proxy > Logs > Access

Log File	Lists the event or error log files you can view, delete, or copy to your local system. Content Gateway lists the event log files located in the directory specified in the Logging Directory field under Subsystems/Logging and by the configuration variable proxy.config.log2.logfile_dir in the records.config file. The default directory is logs in the Content Gateway installation directory.
Action: Display the selected log file	When this option is enabled, Content Gateway displays the first MB of the event or error log file selected in the Log File drop-down list. To view the entire file, select "Save the selected log file in local filesystem" and view the file with a local viewer.
Action: Display last lines of the selected file	When this option is enabled, Content Gateway displays the last specified number of lines in the event or error log file selected from the Log File drop-down list.
Action: Display lines that match in the selected log file	When this option is enabled, Content Gateway displays all the lines in the selected event or error log file that match the specified string.
Remove the selected log file	When this option is enabled, Content Gateway deletes the selected log file.
Action: Save the selected log file in local filesystem	When this option is enabled, Content Gateway saves the selected log file on the local system in a location you specify.

Protocols

The Protocol configuration options are divided into the following categories:

Related reference	
HTTP on page 287	
HTTP Responses on page 292	
HTTPS on page 293	
FTP on page 294	

HTTP

Configure > Protocols > HTTP > General

HTTP Proxy Server Port	Specifies the port that Content Gateway uses when acting as a Web proxy server for HTTP traffic or when serving HTTP requests transparently. The default port is 8080. If you change this option, you must restart Content Gateway.
Secondary HTTP Proxy Server Ports	For explicit proxy configurations only, specifies additional ports on which Content Gateway listens for HTTP traffic. Transparent proxy configurations always send all HTTP traffic to port 8080.
Unqualified Domain Name Expansion	Enables or disables .com name expansion. When this option is enabled, Content Gateway attempts to resolve unqualified hostnames by redirecting them to the expanded address, prepended with www. and appended with .com . For example, if a client makes a request to <i>company</i> , Content Gateway redirects the request to www.company.com If local domain expansion is enabled (see <i>DNS</i> <i>Resolver</i>), Content Gateway attempts local domain expansion before .com domain expansion; Content Gateway tries .com domain expansion only if local domain expansion fails.
Send HTTP 1.1 by Default	Enables the sending of HTTP 1.1 as the first request to the origin server (the default). If the origin server replies with HTTP 1.0, Content Gateway switches to HTTP 1.0 (most origin servers use HTTP 1.1). When disabled, HTTP 1.0 is used in the first request to the origin server. If the origin server replies with HTTP 1.1, Content Gateway switches to HTTP 1.1.

Reverse DNS	Enables reverse DNS lookup when the URL has an IP address (instead of a hostname) and there are rules in filter.config or parent.config . This is necessary when rules are based on destination hostname and domain name.
Tunnel Ports	Specifies the ports on which Content Gateway allows tunneling. This is a space separated list that also accepts port ranges (e.g. 1-65535).
	When SSL is not enabled, all traffic destined for the specified ports is allowed to tunnel to an origin server.
	When SSL is enabled, traffic to any port that is also listed in the HTTPS ports field is not tunneled, but is decrypted and filtering policy is applied.
HTTPS ports	When SSL support is enabled, specifies ports on which HTTPS traffic is decrypted and policy is applied. Note that Content Gateway receives HTTPS traffic on the port specified in Configure > Protocols > HTTP > HTTPS Proxy: Server Port .
	When SSL support is disabled, traffic to these ports is not decrypted. However, filtering policy is applied based on:
	 Explicit proxy: the server hostname in the CONNECT request.
	 Transparent proxy: the SNI hostname or the server hostname in the server's certificate. If the hostname in the server's certificate includes a wildcard (*), the lookup is performed on the destination IP address.
FTP over HTTP: Anonymous Password	Specifies the anonymous password Content Gateway must use for FTP server connections that require a password. This option affects FTP requests from HTTP clients.
FTP over HTTP: Data Connection Mode	An FTP transfer requires two connections: a control connection to inform the FTP server of a request for data and a data connection to send the data. Content Gateway always initiates the control connection. FTP mode determines whether Content Gateway or the FTP server initiates the data connection.
-------------------------------------	---
	Select PASV then PORT for Content Gateway to attempt PASV connection mode first. If PASV mode fails, Content Gateway tries PORT mode and initiates the data connection. If successful, the FTP server accepts the data connection.
	Select PASV only for Content Gateway to initiate the data connection to the FTP server. This mode is firewall friendly, but some FTP servers do not support it.
	Select PORT only for the FTP server to initiate the data connection and for Content Gateway to accept the connection.
	The default value is PASV then PORT .

Configure > Protocols > HTTP > Privacy

Insert Headers: Client-IP	When enabled, Content Gateway inserts the Client-IP header into outgoing requests to retain the client's IP address.
	This option is mutually exclusive with the Remove Headers: Client-IP option. When Insert Headers: Client-IP is enabled the Remove Headers: Client-IP option is automatically disabled.
	Insert Headers: Client-IP and Remove Headers: Client-IP can both be disabled.
Insert Headers: Via	When enabled, Content Gateway inserts a Via header into the outgoing request. The Via header informs the destination server of proxies through which the request was sent.
Insert Headers: X-Forwarded-For	When enabled, Content Gateway inserts an X- Forwarded-For header into the outgoing request. The X-Forwarded-For value contains the originating IP address
	If enabled, header information is sent only to a configured parent proxy. To send header values for all outbound requests, enable proxy.config.http.insert_xff_to_external.

	л. — — — — — — — — — — — — — — — — — — —
Remove Headers: Client-IP	When this option is enabled, Content Gateway removes the Client-IP header from outgoing requests to protect the privacy of your users.
	This option is mutually exclusive with the Insert Headers: Client-IP option. When Remove Headers: Client-IP is enabled the Insert Headers: Client-IP option is automatically disabled.
	Remove Headers: Client-IP and Insert Headers: Client-IP can both be disabled.
Remove Headers: Cookie	When this option is enabled, Content Gateway removes the Cookie header from outgoing requests to protect the privacy of your users. The Cookie header often identifies the user that makes a request.
Remove Headers: From	When this option is enabled, Content Gateway removes the From header from outgoing requests to protect the privacy of your users. The From header identifies the client's email address.
Remove Headers: Referer	When this option is enabled, Content Gateway removes the Referer header from outgoing requests to protect the privacy of your users. The Referer header identifies the Web link that the client selects.
Remove Headers: User-Agent	When this option is enabled, Content Gateway removes the User-Agent header from outgoing requests to protect the privacy of your users. The User-Agent header identifies the agent that is making the request, usually a browser.
Remove Headers: Remove Others	Specifies headers other than From , Referer , User- Agent , and Cookie , that you want to remove from outgoing requests to protect the privacy of your users.
	Use a comma separated list for multiple entries.

Configure > Protocols > HTTP > Timeouts

See this knowledge base article for a discussion of HTTP timeout options.

Keep-Alive Timeouts: Client	Specifies (in seconds) how long Content Gateway keeps connections to clients open for a subsequent request after a transaction ends. Each time Content Gateway opens a connection to accept a client request, it handles the request and then keeps the connection alive for the specified timeout period. If the client does not make another request before the timeout expires, Content Gateway closes the connection. If the client does make another request, the timeout period starts again. The client can close the connection at any time.
-----------------------------	--

Keep-Alive Timeouts: Origin Server	Specifies (in seconds) how long Content Gateway keeps connections to origin servers open for a subsequent transfer of data after a transaction ends. Each time Content Gateway opens a connection to download data from an origin server, it downloads the data and then keeps the connection alive for the specified timeout period. If Content Gateway does not need to make a subsequent request for data before the timeout expires, it closes the connection. If it does, the timeout period starts again. The origin server can close the connection at any time.
Inactivity Timeouts: Client	Specifies how long Content Gateway keeps connections to clients open if a transaction stalls. If Content Gateway stops receiving data from a client or the client stops reading the data, Content Gateway closes the connection when this timeout expires. The client can close the connection at any time.
Inactivity Timeouts: Origin Server	Specifies how long Content Gateway keeps connections to origin servers open if the transaction stalls. If Content Gateway stops receiving data from an origin server, it does not close the connection until this timeout has expired. The origin server can close the connection at any time.
Active Timeouts: Client	Specifies how long Content Gateway remains connected to a client. If the client does not finish making a request (reading and writing data) before this timeout expires, Content Gateway closes the connection. The default value of 0 (zero) specifies that there is no timeout. The client can close the connection at any time.
Active Timeouts: Origin Server Request	Specifies how long Content Gateway waits for fulfillment of a connection request to an origin server. If Content Gateway does not establish connection to an origin server before the timeout expires, Content Gateway terminates the connection request. The default value of 0 (zero) specifies that there is no timeout. The origin server can close the connection at any time.
Active Timeouts: Origin Server Response	Specifies how long Content Gateway waits for a response from the origin server.
FTP Control Connection Timeout	Specifies how long Content Gateway waits for a response from an FTP server. If the FTP server does not respond within the specified time, Content Gateway abandons the client's request for data. This option affects FTP requests from HTTP clients only. The default value is 300.

Related concepts

DNS Resolver on page 344

HTTP Responses

Configure > Protocols > HTTP Responses > General

Response Suppression Mode	If Content Gateway detects an HTTP problem with a particular client transaction (such as unavailable origin servers, authentication requirements, and protocol errors), it sends an HTML response to the client browser. Content Gateway has a set of hard-coded default response pages that explain each HTTP error in detail to the client.
	Select Always Suppressed if you do not want to send HTTP responses to clients.
	Select Intercepted Traffic Only if you want to send HTTP responses to nontransparent traffic only.
	Select Never Suppressed if you want to send HTTP responses to all clients.
	If you change this option, you must restart Content Gateway.

Configure > Protocols > HTTP Responses > Custom

Custom Responses	You can customize the responses Content Gateway sends to clients. By default, the responses you can customize are located in the Content Gateway config/ body_factory/ default directory.
	Select Enabled Language-Targeted Response to send your custom responses to clients in the language specified in the Accept-Language header.
	Select Enabled in "default" Directory Only to send the custom responses located in the default directory to clients.
	Select Disabled to disable the custom responses. If Never Suppressed or Intercepted Traffic Only is selected for the Response Suppression Mode option, Content Gateway sends the hard-coded default responses.
	If you change this option, you must restart Content Gateway.

Custom Response Logging	When enabled, Content Gateway sends a message to the error log each time custom responses are used or modified.If you change this option, you must restart Content Gateway.
Custom Response Template Directory	Specifies the directory where the custom responses are located. The default location is the Content Gateway config/body_factory directory. If you change this option, you must restart Content Gateway.

Incorporating images, animated gifs, and Java applets on the response page

Content Gateway can respond to clients with only a single text or HTML document.

However, you can provide references on your custom response pages to images, animated gifs, Java applets, or objects other than text that are located on a Web server.

Add links in the **body_factory** template files in the same way you would for any image in an HTML document, with the full URL in the SRC attribute.

It is recommended that you do not run the Web server and Content Gateway on the same system, to prevent both programs from trying to serve documents on the same port number.

HTTPS

Configure > Protocols > HTTPS

This page is displayed only when HTTPS is enabled on Configure > My Proxy > Basic > General

HTTPS Proxy Server Port	Specifies the port that Content Gateway uses when acting as a Web proxy server for HTTPS traffic. The default value is 8080.
	See also, <i>Configure > Protocols > HTTP > General</i> : HTTPS Ports.

Tunnel Unknown Protocols	Enables and disables tunneling of HTTPS requests when the SSL handshake results in an unknown protocol error.
	Tunneled connections are not decrypted or inspected.
	When Content Gateway is an explicit proxy, a URL lookup is performed and policy is applied before the SSL connection request is made with the server. Therefore, tunneled transactions appear in the Forcepoint Web Security transaction log.
	When Content Gateway is a transparent proxy, if there is an SNI a URL lookup is done on the hostname in the SNI. Otherwise no URL lookup is possible and tunneled transactions are not logged. This is because an initial connection with the server is required to get the Common Name from the SSL certificate. It is used for the URL lookup. If the connection handshake fails and this option is enabled, the connection is tunneled without the proxy being aware of it.
	Important: This setting persists after the HTTPS feature is disabled (on Configure > My Proxy > Basic > General . Therefore, disable this option before disabling HTTPS support.

Related reference

HTTP on page 287

FTP



Note

The FTP configuration options appear on the Configure pane only if you have enabled FTP processing in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Protocols > FTP > General

FTP Proxy Server Port	Specifies the port that Content Gateway uses to accept FTP requests. The default port is 2121.
Listening Port Configuration	Specifies how FTP opens a listening port for a data transfer.
	Select Default Settings to let the operating system choose an available port. Content Gateway sends 0 and retrieves the new port number if the listen succeeds.
	Select Specify Range if you want the listening port to be determined by the range of ports specified in the Listening Port (Max) and Listening Port (Min) fields.

Default Data Connection Method	Specifies the default method used to set up data connections with the FTP server.
	Select Proxy Sends PASV to send a PASV to the FTP server and let the FTP server open a listening port.
	Select Proxy Sends PORT to set up a listening port on the Content Gateway side of the connection first.
Shared Server Connections	When enabled, server control connections can be shared between multiple anonymous FTP clients.

Configure > Protocols > FTP > Timeouts

Keep-Alive Timeout: Server Control	Specifies the timeout value when the FTP server control connection is not used by any FTP clients. The default value is 90 seconds.
Inactivity Timeouts: Client Control	Specifies how long FTP client control connections can remain idle. The default value is 900 seconds.
Inactivity Timeouts: Server Control	Specifies how long the FTP server control connection can remain idle. The default value is 120 seconds.
Active Timeouts: Client Control	Specifies the how long FTP client control connections can remain open. The default value is 14400 seconds.
Active Timeouts: Server Control	Specifies how long the FTP server control connection can remain open. The default value is 14400 seconds.

Content Routing

The Content Routing configuration options are divided into the following categories:

Related reference Hierarchies on page 295 Mapping and Redirection on page 298 Browser Auto-Config on page 300

Hierarchies

Configure > Content Routing > Hierarchies

Parent Proxy	Enables or disables the HTTP parent caching option. When this option is enabled, Content Gateway can participate in an HTTP cache hierarchy. You can point your Content Gateway server at a parent network cache to form a cache hierarchy where a child cache relies upon a parent cache in fulfilling client requests. See <i>HTTP cache hierarchies</i> . This setting must be enabled when Protected Cloud Apps is enabled and configured in the Forcepoint Security Manager
No DNS and Just Forward to Parent	When enabled, and if HTTP parent caching is enabled, Content Gateway does no DNS lookups on requested hostnames. If rules in the parent.config file are set so that only selected requests are sent to a parent proxy, Content Gateway skips name resolution only for requests that are going to the parent proxy. Name resolution is performed as usual for requests that are not sent to a parent proxy. If the parent proxy is down and the child proxy can go directly to origin servers, the child performs DNS resolution.
Uncacheable Requests Bypass Parent	When enabled, and if parent caching is enabled, Content Gateway bypasses the parent proxy for uncacheable requests.
HTTPS Requests Bypass Parent	When enabled, Content Gateway bypasses the parent proxy for HTTPS requests.
Tunnel Requests Bypass Parent	When enabled, Content Gateway bypasses parent proxy for non-HTTPS tunnel requests.
Parent Proxy Cache Rules	Displays a table listing the rules in the <i>parent.config</i> file that identify the HTTP parent proxies used in an HTTP cache hierarchy and configure selected URL requests to bypass parent proxies. Rules are applied from the list top-down; the first match is applied.
Refresh	Updates the table to display the most up-to-date rules in the parent.config file.
Edit File	Opens the configuration file editor so that you can edit and add rules to the parent.config file.
	parent.config Configuration File Editor
rule display box	Lists the <i>parent.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.

Primary Destination Type	Lists the primary destination types: dest_domain is a requested domain name. dest_host is a requested hostname. dest_ip is a requested IP address.
	url_regex is a regular expression to be found in a URL.
Primary Destination Value	Specifies the value of the primary destination type. For example:
	If the primary destination is dest_domain , a value for this field can be yahoo.com
	If the primary destination type is dest_ip , the value for this field can be 123.456.78.9.
	If the primary destination is url_regex , a value for this field can be politics.
Parent Proxies	Specifies the IP addresses or hostnames of the parent proxies and the port numbers used for communication. Parent proxies are queried in the order specified in the list. If the request cannot be handled by the last parent server in the list, it is routed to the origin server. Separate each entry with a semicolon; for example: parent1:8080; parent2:8080
Round Robin	Select true for the proxy to go through the parent cache list in a round-robin based on client IP address.
	Select strict for the proxy to serve requests strictly in turn. For example, machine proxy1 serves the first request, proxy2 serves the second request, and so on.
	Select false if you do not want round-robin selection to occur.
Go direct	Select true for requests to bypass parent hierarchies and go directly to the origin server.
	Select false if you do not want requests to bypass parent hierarchies.
Secondary Specifiers: Time	Specifies a time range, using a 24-hour clock, such as 08:00-14:00. If the range crosses midnight, enter this as two comma-separated ranges. For example, if a range extends from 6:00 in the evening until 8:00 in the morning, enter the following:
	18:00 - 23:59, 0:00 - 8:00
Secondary Specifiers: Prefix	Specifies a prefix in the path part of a URL.
Secondary Specifiers: Suffix	Specifies a file suffix in the URL, such as .htm or .gif.
Secondary Specifiers: Source IP	Specifies the IP address or range of IP addresses of the clients.
Secondary Specifiers: Port	Specifies the port in a requested URL.
Secondary Specifiers: Method	Specifies a request URL method. For example:
	get post put trace

Secondary Specifiers: Scheme	Specifies the protocol of a requested URL. This must be either HTTP or FTP.
Secondary Specifiers: User-Agent	Specifies a request header User-Agent value.

Related concepts HTTP cache hierarchies on page 79 parent.config on page 384

Mapping and Redirection

Configure > Content Routing > Mapping and Redirection

Serve Mapped Hosts Only	Select Required if you want the proxy to serve requests only to origin servers listed in the mapping rules of the remap.config file. If a request does not match a rule in the remap.config file, the browser receives an error. This option provides added security for your Content Gateway system.
Retain Client Host Header	When this option is enabled, Content Gateway retains the client host header in a request (it does not include the client host header in the mapping translation).
Redirect No-Host Header to URL	Specifies the alternate URL to which to direct incoming requests from older clients that do not provide a Host: header.
	It is recommended that you set this option to a page that explains the situation to the user and advises a browser upgrade or provides a link directly to the origin server, bypassing the proxy. Alternatively, you can specify a map rule that maps requests without Host: headers to a particular server.
URL Remapping Rules	Displays a table listing the mapping rules in the remap.config file so that you can redirect HTTP requests permanently or temporarily without the proxy having to contact any origin servers.
	Note: Mapping a URL to another URL in the same domain requires that a "/" be specified in From Path Prefix field. See the example following this table.
Refresh	Updates the table to display the most up-to-date rules in the remap.config file.
Edit File	Opens the configuration file editor so that you can edit and add rules to the remap.config file.
	remap.config Configuration File Editor
rule display box	Lists the remap.config file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.

Add	Adds a the cor	new rule to the rule display box at the top of figuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.	
Rule Type	Lists the type of rules you can add to the remap.cor file:	
	map pi redirec	ovides the same function as redirect. Use of t is recommended.
	redired without redirec returnin browse	et redirects HTTP requests permanently thaving to contact the origin server. Permanent ts notify the browser of the URL change (by an HTTP status code 301) so that the er can update bookmarks.
	redired tempor Tempo change HTTP s	ct_temporary redirects HTTP requests arily without having to contact the origin server. rary redirects notify the browser of the URL of or the current request only (by returning an status code 307).
	revers	e_map is not supported.
From Scheme	Specifi mms	es the protocol of the mapping rule. rtsp and
	are not	supported.
		Note
		Mapping a URL of one protocol (scheme) to a different protocol (scheme) is not supported.
From Host	Specifi	es the hostname of the URL to map from.
From Port (Optional)	Specifi	es the port number in the URL to map from.
From Path Prefix (Optional)	Specifies the path prefix of the URL to map from. Sometimes it is desirable to redirect a URL to a sub- page in the same domain. For example, to redirect "www.cnn.com" to "www.cnn.com/tech". To make this rule work you must specify "/" in the From Path Prefix field. If it is not specified, the redirection results in a URL that recursively adds the page specifier to the URL. For example, "www.example.com/tech" becomes "www.example.com/tech/tech/tech/tech/tech/tech/tech/tech	
From Query (Optional)	Specifi	es the query of the URL to map from.
To Scheme	Must m	natch From Scheme.
To Host	Specifi	es the hostname of the URL to map to.
To Port (Optional)	Specifies the port number of the URL to map to.	
To Path Prefix (Optional)	Specifies the path prefix of the URL to map to.	

To Query (Optional)	Specifies the query of the URL to map to.
{undefined}	Specifies the media protocol type of the mapping rule. Not supported.

Browser Auto-Config

Configure > Content Routing > Browser Auto-Config > PAC

Auto-Configuration Port	Specifies the port Content Gateway uses to download the auto-configuration file to browsers. The port cannot be assigned to any other process. The default port is 8083.
	If you change this option, you must restart Content Gateway.
PAC Settings	Lets you edit the PAC file (proxy.pac). See <i>Using a PAC file</i> .

Configure > Content Routing > Browser Auto-Config > WPAD

WPAD Settings	Lets you edit the wpad.dat file. See <i>Using WPAD</i> .

Related concepts

Using WPAD on page 30

Related tasks

Using a PAC file on page 28

Security

The Security configuration options are divided into the following categories:

Related concepts

FIPS Security on page 301 Access Control on page 303 SOCKS on page 321

Related reference

Connection Control on page 301 Web DLP on page 302

Connection Control

Configure > Security > Connection Control

Option	Description
	Proxy Access
Access Control	Displays the rules in the <i>ip_allow.config</i> file that control which clients can access Content Gateway.
	By default, all remote hosts are allowed to access the proxy.
Refresh	Updates the table to display the most up-to-date rules in the ip_allow.config file.
Edit File	Opens the configuration file editor for to the ip_allow.config file.
	ip_allow.config Configuration File Editor
rule display box	Lists the <i>ip_allow.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
IP Action	Lists the type of rules you can add.
	An ip_allow rule allows the clients listed in the Source IP field to access the proxy.
	An ip_deny rule denies the clients listed in the Source IP field access to the proxy.
Source IP	Specifies the IP address or range of IP addresses of the clients.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Related concepts

ip_allow.config on page 375

FIPS Security

Configure > Security > FIPS

Important

After FIPS is enabled, you must re-install any hotfixes previously installed for the current version of Content Gateway.

When FIPS mode is enabled:

- HTTPS connections use only TLSv1 or higher
- HTTPS connections use FIPS 140-2 approved algorithms
- Content Gateway generates SHA-256 certificates in response to origin server certificate requests



Warning

Once enabled, FIPS 140-2 mode cannot be disabled without reinstalling Content Gateway. If Content Gateway is on an appliance, the appliance must be reimaged.

A

Important

Due to a system limitation, FIPS 140-2 mode cannot be used with NTLM user authentication (IWA fallback to NTLM or Legacy NTLM).

For complete information, see FIPS 140-2 Mode.

Option	Description
FIPS Enable/Disable radio buttons	By default, Content Gateway is installed in non-FIPS 140-2 mode.
	To switch to FIPS 140-2 mode, select the Enabled radio button, click Apply , and restart Content Gateway.
	Warning
	Once enabled, FIPS 140-2 mode cannot be disabled without reinstalling Content Gateway. For appliance installations, reinstallation requires reimaging the system.

Related concepts FIPS 140-2 Mode on page 154

Web DLP

Note

The Web DLP configuration options appear on the Configure menu only if you have enabled Web DLP (integrated on-box) on the Configure > My Proxy > Basic > General tab and selected Integration > Web DLP (integrated on-box) in the Features table.

Configure > Security > Web DLP

Option	Description
Forcepoint management server IP address	Specifies the IP address of the Forcepoint management server. Configure Web DLP policy in the Data Security module of the Forcepoint Security Manager.
Analyze HTTPS Content	Specifies whether decrypted traffic should be sent to Forcepoint DLP for analysis, or sent directly to the destination.
Analyze FTP Uploads	Specifies whether to send FTP upload requests to Forcepoint DLP for analysis. The FTP proxy feature must be enabled. See <i>FTP</i> .

Registration screen fields:

Option	Description
Forcepoint management server IP	Specifies the IP address of the Forcepoint management server. This is where data security policy configuration and management is performed.
Administrator user name	Specifies the account name of a Forcepoint DLP administrator. The administrator must have Deploy Settings privileges.
Administrator password	Specifies the password of the Forcepoint DLP administrator.
Register button	Initiate the registration action. This button is enabled only after data is entered in all of the fields.

Related reference

FTP on page 294

Access Control

Use the Access Control tabs to:

- Create custom filtering rules
- Configure proxy user authentication

The *Filtering* tab is always available on the **Access Control** page.

Other tabs are dynamic based on the authentication method selected in the **Authentication** section of **Configure** > **My Proxy** > **Basic**.

If an authentication method is enabled, the Global Configuration Options tab is always displayed.

If Integrated Windows Authentication is selected, these tabs display:

- Integrated Windows Authentication
- Global Configuration Options

If LDAP is selected, these tabs display:

- LDAP
- Global Configuration Options

If Radius is selected, these tabs display:

- Radius
- Global Configuration Options

If NTLM is selected, these tabs display:

- NTLM
- Global Configuration Options

If Rule-Based Authentication is selected, these tabs display:

- Domains
- Authentication Rules
- Global Configuration Options

The tables below describe the purpose of each field on each tab. Use your browser's Search feature to find the field that you're looking for.

For a complete description of Content Gateway user authentication features, see *Content Gateway user authentication*.

Related concepts

Content Gateway user authentication on page 164

Related reference

Configure > Security > Access Control > Filtering on page 304 Configure > Security > Access Control > Global Configuration Options on page 307 Configure > Security > Access Control > IWA on page 311 Configure > Security > Access Control > LDAP on page 312 Configure > Security > Access Control > Radius on page 313 Configure > Security > Access Control > NTLM on page 313 Configure > Security > Access Control > NTLM on page 314 Configure > Security > Access Control > Domains on page 314

Configure > Security > Access Control > Filtering

Filtering rules can be used to:

- Deny or allow URL requests
- Insert custom headers
- Allow specified applications, or requests to specified websites to bypass user authentication
- Keep or strip header information from client requests
- Prevent specified applications from transiting the proxy

Rules are ordered checked prior to user authentication (if configured). Rules are applied based on first match in a top-down traversal of the list. If no rule matches, the request is allowed to proceed.

Rules are stored in *filter.config*.

After adding, deleting, or modifying a rule, restart Content Gateway.

For complete information about filtering rules, see Content Gateway filtering rules.

Filtering	Displays an ordered list of filtering rules.
	Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 file sandbox destinations.
Refresh	Updates the table to display the most up-to-date rules in the filter.config file.
Edit File	Opens the configuration file editor for the filter.config file.
	filter.config Configuration File Editor
rule display box	Lists the rules currently stored in <i>filter.config</i> . Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page. Click Add after selecting or entering values for the rule.
Set	Updates the rule display box at the top of the configuration file editor page.
Rule Type	Specifies the rule type:
	Select allow to allow particular URL requests to bypass authentication.
	Select deny to deny requests for objects from specific destinations. When a request is denied, the client receives an access denied message.
	Select keep_hdr to specify which client request header information you want to keep.
	Select strip_hdr to specify which client request header information you want to strip.
	Select add_hdr to cause a custom header to be added to the request. This rule type requires that values be defined for Custom Header and Header Value . Add custom headers to satisfy specific requirements of a destination domain. See <i>Content Gateway filtering</i> <i>rules</i> .
	The radius rule type is not supported.
Primary Destination Type	Lists the primary destination types: dest_domain is a requested domain name. dest_host is a requested host name. dest_ip is a requested IP address.
	url_regex is a regular expression to be found in a URL.

Primary Destination Value	Specifies the value of the Primary Destination Type. For example, if the Primary Destination Type is dest_ip , the value for this field might be 123.456.78.9.
Additional Specifiers: Header Type	Specifies the client request header information that you want to keep or strip.
	This option applies to only keep_hdr or strip_hdr
	rule types.
Additional Specifiers: Realm (optional)	Not supported.
Additional Specifiers: Proxy Port (optional)	Specifies the proxy port to match for this rule.
Additional Specifiers: Custom Header (optional)	For use when the rule type is add_hdr . Specifies the custom header name that the destination domain expects to find in the request.
Additional Specifiers: Header Value (optional)	For use when the rule type is add_hdr . Specifies the custom header value that the destination domain expects to be paired with the custom header.
Secondary Specifiers: Time	Specifies a time range, such as 08:00-14:00.
Secondary Specifiers: Prefix	Specifies a prefix in the path part of a URL.
Secondary Specifiers: Suffix	Specifies a file suffix in the URL.
Secondary Specifiers: Source IP	Specifies the IP address of the client.
Secondary Specifiers: Port	Specifies the port in a requested URL.
Secondary Specifiers: Method	Specifies a request URL method:
	■ get
	■ post
	= put
	• trace
Secondary Specifiers: Scheme	Specifies the protocol of a requested URL. Options are:
	HTTP
	HTTPS
	FTP (for FTP over HTTP only)
	rtsp and mms are not supported.
Secondary Specifiers: User-Agent	Specifies the Request header User-Agent value.
	Use this field to create application filtering rules that:
	 Allow applications that don't properly handle authentication challenges to bypass authentication
	 Block specified client-based applications from accessing the Internet
Apply	Applies the configuration changes.

Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Related concepts filter.config on page 372 Content Gateway filtering rules on page 155

Configure > Security > Access Control > Global Configuration Options

Use this page to specify global options for:

- The fail open/fail closed action to take when user authentication fails
- Credential caching
- For transparent proxy, an alternate hostname for the proxy that all clients on the network can resolve. Required.
- Cookie sharing

For more information, see Global authentication options.



Note

The user interface setting to disable the NTLM cache for explicit proxy has been removed. Although not recommended, the cache can be disabled for explicit proxy traffic in records.config by setting the value of **proxy.config.ntlm.cache.enabled** to **0** (zero).

Global Configuration Options

Fail Open	Disabled – Prevents requests from proceeding to the Internet when an authentication failure occurs.
	Enabled only for critical service failures (default) – Allows requests to proceed if authentication fails because there is no response from the domain controller or because the client is sending badly formatted messages.
	Enabled for all authentication failures – Allows requests to proceed for all authentication failures, including password failures.
	When a fail open setting is enabled, if a Forcepoint Web Security transparent user identification agent is configured an attempt is made to identify the requester and apply user-based policy. Otherwise, if a policy has been assigned to the client's IP address, that policy is applied. Otherwise, the Default policy is applied.
	Important
	When user authentication is rule-based with a domain list:
	If Enabled only for critical service failures is selected, when a critical service failure occurs fail open is not applied. An error always results in fail closed.
	If Enabled for all authentication failures, including incorrect password is selected, after trying basic credentials with every domain in the list, fail open is applied.
	The Fail Open setting does not apply when IWA is the authentication method and the client fails to retrieve a kerberos ticket from the domain controller (DC) because the DC is down. The Fail Open setting does apply with IWA when IWA falls back to NTLM and authentication fails.

Credential Caching: Caching Method	Cache using IP address only – specifies that all credentials are cached with IP address surrogates. This is the recommended method when all clients have unique IP addresses.
	Cache using Cookies only – specifies that all credentials are cached with cookie surrogates. This is recommended when all clients share IP addresses, as with multi-host servers such as Citrix servers, or when traffic is NATed by a device that is forwarding traffic to Content Gateway.
	Cache using both IP addresses and Cookies – specifies to use cookie surrogates for the IP addresses listed in the cookie caching list, and to use IP address surrogates for all other IP addresses. This is recommended when the network has a mix of clients, some with unique IP addresses and some using multi- user hosts or that are subject to NATing.
	The cookie caching list is a comma separated list that can contain up to:
	64 IPv4 addresses
	32 IPv4 address ranges
	■ 24 IPv6 addresses
	12 IPv6 address ranges
	For a description of surrogate credentials, see <i>Surrogate credentials</i> .
	 Cookie mode caching does not work with applications that do not support cookies, or with browsers in which cookie support has been disabled.
	 When the browser is Internet Explorer, the full proxy hostname in the form "http:// host.domain.com" must be added to the Local intranet zone.
	When the browser is Chrome, it must be configured to allow third-party cookies or configured for an exception to allow cookies from the proxy hostname in the form "host.domain.com".
	 When the IP address is set for cookie mode and the request method is CONNECT, no caching is performed.
	 Cookie mode caching is not performed for FTP requests.
	 Cookie mode caching is supported with Captive Portal and client certificate authentication.
	 For explicit proxy, cookie-based authentication is not supported for HTTPS. IP-address authentication is used.

	1
Credential Caching: Time-To-Live	Specifies the duration, in minutes, that an entry in the cache is retained. When the TTL expires, the entry is removed and the next time that the user submits a request, the user is authenticated. If the authentication succeeds, an entry is placed in the cache.
Cookie Expiration	Specifies whether a user is allowed to re-access the system without authentication until the cookie is no longer valid. When enabled, cookies expire when the user ends a session
Purge LDAP cache on authentication failure	Specifies that when an LDAP user authentication failure occurs, Content Gateway will delete the authorization record for that client from the LDAP cache.
Redirect Hostname	For transparent proxy, specifies an alternate hostname for the proxy that all clients on the network can resolve. Required.
	Valid characters for Redirect HostName are: A-Z, a-z,0-9 and
	For complete information see <i>Redirect Options</i> .
Cookie Sharing	When cookie caching is enabled, cookie surrogates can be shared across all nodes in a cluster.
	Select and import both private and public keys and then make a backup of them.
	Used with load balancing, the entry in Redirect Hostname must be the FQDN of the load balancer.
	Note:
	 Cookie caching limitations also apply to cookie sharing. Therefore, since cookie caching is not supported for CONNECT requests, cookie sharing is not supported.
	 Custom keys must be imported manually. Custom Keys are not synchronized across the cluster.
	 Cookie sharing is not supported with client certificate authentication.
	Keys must be PKCS#1 RSA public keys.
	For more information, see <i>Cookie Sharing</i> .

Related concepts

Global authentication options on page 168 Surrogate credentials on page 173 Redirect Options on page 171 Cookie Sharing on page 172

Configure > Security > Access Control > IWA

The Integrated Windows Authentication (IWA) page appears only if you have enabled IWA in the Features table on the **Configure > My Proxy > Basic > General** tab.

Use this page to join or unjoin the Windows domain. When a domain has been joined, the page provides a summary of the domain attributes and an Unjoin button.

For a complete description, see Integrated Windows Authentication.

Integrated Windows Authentication

Domain Name	Specifies the fully qualified Windows domain name.
Administrator Name	Specifies the Windows Administrator user name.
Administrator Password	Specifies the Windows Administrator password.
	Note
	The name and password are used only during the join and are not stored.
Domain Controller	Specifies how to locate the domain controller:
	 Auto-detect using DNS
	 DC name or IP address
	If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list.
Contont Cotoway Hastnama	
Content Galeway Hostname	Specifies the Content Gateway hostname.
Content Galeway Hostname	Specifies the Content Gateway hostname. Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on Forcepoint appliances (which add 4 characters to the hostname to ensure that the hostname is unique across modules (Doms).
Content Galeway Hostname	Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on Forcepoint appliances (which add 4 characters to the hostname to ensure that the hostname is unique across modules (Doms).
Content Galeway Hostname	Specifies the Content Gateway hostname. Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on Forcepoint appliances (which add 4 characters to the hostname to ensure that the hostname is unique across modules (Doms). Important Once the domain is joined the hostname cannot be changed. If it is, IWA will immediately stop working until the domain is unjoined and then rejoined with the new hostname.

Related concepts

Integrated Windows Authentication on page 173

Configure > Security > Access Control > LDAP

The LDAP configuration options appear on the Configure pane only if you have enabled LDAP in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information about configuring LDAP see *LDAP authentication*.

LDAP	
LDAP Server: Hostname	Specifies the hostname of the LDAP server.
	If you change this option, you must restart Content Gateway.
LDAP Server: Port	Specifies the port used for LDAP communication. The default port number is 389.
	To use the default Global Catalog server port, specify port 3268.
	If Secure LDAP is enabled, set the port to 636 or 3269 (the secure LDAP ports).
	If you change this option, you must restart Content Gateway.
LDAP Server: Secure LDAP	Specifies whether Content Gateway will use secure communication with the LDAP server. If enabled, set the LDAP Port field (above) to 636 or 3269 (the secure LDAP ports).
LDAP Server: Server Type	Specifies the search filter. Select either a Microsoft Active Directory option or other directory services.
LDAP Server: Bind Distinguished Name	Specifies the Full Distinguished Name (fully qualified name) of a user in the LDAP-based directory service. For example:
	CN=John Smith,CN=USERS,DC=MYCOMPANY, DC=COM
	Enter a maximum of 128 characters in this field.
	If you do not specify a value for this field, the proxy attempts to bind anonymously.
LDAP Server: Password	Specifies a password for the user identified in the
	Bind_DN field.
LDAP Server: Base Distinguished Name	Specifies the base Distinguished Name (DN). You can obtain this value from your LDAP administrator.
	You must specify a correct base DN; otherwise LDAP authentication will fail to operate.
	If you change this option, you must restart Content Gateway.

Related concepts	
LDAP authentication on page	180

Configure > Security > Access Control > Radius

The Radius configuration options appear on the Configure pane only if you have enabled Radius in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information about configuring Radius, see RADIUS authentication.

Radius	
Primary Radius Server: Hostname	Specifies the hostname or IP address of the primary RADIUS authentication server.
	If you change this option, you must restart Content Gateway.
Primary Radius Server: Port	Specifies the port that Content Gateway uses to communicate with the primary RADIUS authentication server. The default port is 1812.
	If you change this option, you must restart Content Gateway.
Primary Radius Server: Shared Key	Specifies the key to use for encoding.
	If you change this option, you must restart Content Gateway.
Secondary Radius Server (optional): Hostname	Specifies the hostname or IP address of the secondary RADIUS authentication server.
	If you change this option, you must restart Content Gateway.
Secondary Radius Server (optional): Port	Specifies the port that Content Gateway uses to communicate with the secondary RADIUS authentication server. The default port is 1812.
	If you change this option, you must restart Content Gateway.
Secondary Radius Server (optional): Shared Key	Specifies the key to use for encoding.
	If you change this option, you must restart Content Gateway.

Related concepts RADIUS authentication on page 184

Configure > Security > Access Control > NTLM

The NTLM configuration options appear on the Configure pane only if you have enabled NTLM in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information about configuring NTLM, see Legacy NTLM authentication.

NTLM

Domain Controller Hostnames	Specifies the hostnames of the domain controllers in a comma separated list. The format is:
	host_name[:port][%netbios_name] or IP_address[:port] [%netbios_name]
	If you are using Active Directory 2008, you must include the netbios_name or use SMB port 445.
	If you change this option, you must restart Content Gateway.
Load Balancing	Enables or disables load balancing. When enabled, Content Gateway balances the load when sending authentication requests to the domain controllers.
	Note
	When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term failover provision, until such time that the primary domain controller can accept new connections.
	If you change this option, you must restart Content Gateway.

Related concepts

Legacy NTLM authentication on page 179

Configure > Security > Access Control > Domains

The Domains tab appears in the Access Control list only if you have enabled **Rule-Based Authentication** in the Features table on **Configure > My Proxy > Basic > General**.

Use this tab to create and maintain a list of domains that can be specified in authentication rules. Use the Authentication Rules tab to define authentication rules.

Be sure to set the Global authentication options.



Important

You must configure the Domains list before you configure authentication rules.

If you have never configured rule-based authentication, see *Rule-Based Authentication*, for complete information.

Domains

	•
Domain List	An unordered list of domains that have been identified for use in authentication rules.
	Use the Edit button to change some attributes associated with the domain.
	Use the Delete or Unjoin button to remove a domain from the list.
	The domain list is stored in auth_domains.config.
Domain list: New Domain button	Use the New Domain button to add a domain to the Domains list. The screen is expanded to allow for specification of the domain.
	New Domain action
Domain Details: Domain Identifier	Specify a unique name for the domain. The name is used only by Content Gateway; it does not change any attribute of the actual domain or directory.
	Important
	You cannot change the domain identifier after it has been added to the list. To change the name, delete the entry from the list and re-add it with the new name.
Domain Details: Authentication Method	Specify the authentication method: IWA, Legacy NTLM, or LDAP. Radius is not supported.
	When you select an authentication method, configuration options specific to that method are added to the page.
	Important
	You cannot change the authentication method after you add the domain to the list. To change the authentication method, delete the entry from the list and re-add the domain specifying the new authentication method.
Domain Details: Aliasing	Specify an alias to send to the filtering service for all users who match this rule (optional). The alias must be static. It can be empty (blank). The alias must exist in the primary domain controller (the DC visible to the filtering service). See Unknown users and the 'alias' option.
IWA Domain Details	These options are presented when IWA is specified as the authentication method.
Domain Name	Specify the fully qualified domain name. For example: corp-domain.example.com
Administrator Name	Specify a Windows Active Directory domain administrator user name.

Administrator Password	Specify the corresponding domain administrator password.
	Note
	The name and password are used only during the join and are not stored.
Domain Controller	Specify how to locate the domain controller:
	 Auto-detect using DNS
	DC name or IP address
	If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list.
Content Gateway Hostname	Specify the Content Gateway hostname.
	Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on Forcepoint appliances (which add 4 characters to the hostname to ensure that the hostname is unique across modules (Doms).
	Warning
	Once the domain is joined the hostname cannot be changed. If it is, IWA will immediately stop working until the domain is unjoined and then rejoined with the new hostname.
Join Domain	Click Join Domain to join the domain.
Legacy NTLM Domain Details	
Domain Controller	Specify the IP address and port number of the primary domain controller (if no port is specified, Content Gateway uses port 139), followed by a comma separated list of secondary domain controllers to be used for load balancing and failover.

Load Balance	Select the check box to balance the load across multiple NTLM DCs.
	Note
	When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term failover provision, until such time that the primary domain controller can accept new connections.
LDAP Domain Details	
LDAP Server Name	Specify the LDAP server name.
LDAP Server Port	Specify the LDAP Server Port (optional) The default is 389.
LDAP Base Distinguished Name	Specify the LDAP Base Distinguished Name.
LDAP Server Type	Set the search filter to "sAMAccountName (MS AD)" or "userPrincipalName (MS AD)" for Active Directory, or "uid" for other directory services.
Bind Domain Name	Specify the LDAP bind account distinguished name. For example:
	CN=John Smith,CN=USERS,DC=MYCOMPANY, DC=COM
	The field length is limited to 128 characters.
	If no value is specified, Content Gateway attempts to bind anonymously.
Bind Password	Specify the LDAP bind account password.
Secure LDAP	Specify whether Content Gateway will use secure communication with the LDAP server. If enabled, you must set the LDAP port to one of the secure ports: 636 or 3269.

Related concepts	
Global authentication options on page 168	
Rule-Based Authentication on page 186	
Unknown users and the 'alias' option on page 190	

Configure > Security > Access Control > Authentication Rules

The Authentication Rules tab appears in the Access Control list only if you have enabled **Rule-Based Authentication** in the Features table on the **Configure > My Proxy > Basic > General** tab.

Use this tab to create and maintain authentication rules. Use the **Domains** tab to build and maintain a list of domains that can be used in authentication rules. You must configure the Domains list before you define authentication rules.

Be sure to set the Global authentication options.



Important

If you have never configured rule-based authentication, see *Rule-Based Authentication*, for complete information.

Authentication Rules	
Authentication Rule List	Displays a table of the ordered list of rules defined for user authentication. Rules are defined for sets of clients to be authenticated against one or more IWA, LDAP and NTLM domains. See <i>Rule-Based</i> <i>Authentication</i> .
Refresh	Updates the table to display the current rules in the auth_rules.config file.
Edit File	Opens the authentication rule editor.
	Warning
	Do not edit rules directly in the configuration file.
	auth_rules.config Configuration File Editor
rule display box	Lists, in order, the current rule set. When user authentication is performed, the list is traversed, top- down and the first match is applied.
	Select a rule to edit it.
	The arrows to the left of the box allow you to move the selected rule up or down in the list.
	The "X" button deletes the selected rule. Rules cannot be more than 2048 characters.
Add	Adds a new rule.
Set	Updates the selected rule with the current values.

Status	Specifies whether the rule is enabled (active) or disabled after the rule is saved and Content Gateway is restarted.
	elements of your network are ready to support it.
Rule Name	Specifies a unique, descriptive name for the rule. It is recommended that the name not exceed 50 characters.
Source IP	Specifies IP addresses or IP address ranges for this rule (must be entered without any spaces).
	Example: 10.1.1.1 or 0.0.0.0-255.255.255.255 or
	10.1.1.1,20.2.2.2,3.0.0.0-3.255.255.255
	The comma separated list can contain up to:
	64 IPv4 addresses
	 32 IPv4 address ranges
	24 IPv6 addresses
	 12 IPv6 address ranges
Proxy Port	Specifies the inbound port for traffic when Content Gateway is deployed as an explicit proxy. If undefined, all ports match, as configured on Configure > Protocols > HTTP > General .
	Transparent proxy deployment should leave this field undefined.
User-Agent	Specifies 1 or more regular expressions used to match text in the User-Agent string, for example to match common browsers.
	Regexes must be POSIX-compliant. The "^" operator is not supported.
	When the field is empty, all User-Agent values match. You can edit the field directly.
	To insert a predefined regex for a common browser, select it from the drop down list and click Add .
	Multiple regexes can be specified. Use the " " character to separate entries (logical 'or').
	For more information, including regex examples, see <i>Authentication based on User-Agent</i> .
Client Certificate	Click Enabled to enable client certificate authentication.
	Select Use the next selected authentication method if Client Certificate authentication fails to use one of the other authentication methods if certificate authentication fails for a user.
	See Client certificate authentication for details.

Auth Sequence	Specifies 1 or more domains to use for authentication.
	Select a domain from the Domains drop down list (populated from the Domains List), and click Include to add it to the list.
	If you add more than one domain, you can set the order by selecting an entry and using the up and down arrows. You can delete a selected domain with the "X" button.
	Best practice: If you know what domain a set of users belongs to, create a rule just for that group.
	Best practice: Place the rule with the largest number of users authenticating with known domain membership at the top of the list. These are the fastest authentications.
	Best practice: If you don't know what domain a set of users belongs to, specify the fewest number of domains needed to authenticate the users in the set.
	Best practice: It is always better to create targeted rules because attempting to authenticate against a large set of domains can introduce noticeable latency.
	When user authentication is rule-based with a domain list:
	 For each user, the first successful authentication is cached and used in subsequent authentications. If IP address caching is configured, an IP address surrogate is cached. If Cookie Mode is configured, a cookie surrogate is cached.
	For Fail Open:
	If Enabled only for critical service failures is selected, the fail open setting is not applied. The user continues to be prompted for credentials until there is a timeout.
	 If Enabled for all authentication failures, including incorrect password is selected, after trying basic credentials with every domain in the list, fail open is applied.
Captive Portal	Click Enabled for HTTPS/HTTP Authentication page to redirect users to a customizable web portal page for authentication.
	See Authentication using Captive Portal for details.
Apply	Applies the configuration changes.
	If the rule specifies a regex for User- Agent, the regex is validated when Apply is clicked. If the regex is not valid, the rule is deleted and must be recreated.

Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Related concepts Global authentication options on page 168 Rule-Based Authentication on page 186 Authentication based on User-Agent on page 203 Client certificate authentication on page 208 Authentication using Captive Portal on page 205

SOCKS

Note

For more information about Content Gateway support for SOCKS, see Configuring SOCKS firewall integration.

Ę

The SOCKS configuration options appear on the Configure pane only if you have enabled SOCKS in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Security > SOCKS > General

SOCKS Version	Specifies the version of SOCKS used on your SOCKS server. Content Gateway supports SOCKS version 4 and version 5.
	If you change this option, you must restart Content Gateway.

Configure > Security > SOCKS > Proxy

SOCKS Proxy	Enables or disables the SOCKS Proxy option. As a SOCKS proxy, Content Gateway can receive SOCKS packets (usually on port 1080) from the client, and forward requests directly to the SOCKS server.
	For more information about the SOCKS Proxy option, see
	Configuring SOCKS firewall integration.
	If you change this option, you must restart Content Gateway.
SOCKS Proxy Port	Specifies the port on which Content Gateway accepts SOCKS traffic. This is usually port 1080.
	If you change this option, you must restart Content Gateway.

Configure > Security > SOCKS > Server

	* 2
On-Appliance SOCKS server	Displays only when Content Gateway is on an appliance. Enables or disables the on-appliance SOCKS server.
	The SOCKS proxy option must be enabled to route client requests through the SOCKS server.
	You can configure Content Gateway to use other SOCKS servers in your network by editing socks_server.config . See the next entry.
Socks Servers table	Displays a table of configured SOCKS servers. For information about adding and configuring SOCKS servers, see <i>Configuring SOCKS servers</i> .
Refresh	Updates the table to display the current entries in socks_server.config .
Edit File	Opens the configuration file editor for socks_server.config .
	socks_server.config Configuration File Editor
entry display box	Lists the SOCKS servers that have been configured for use with Content Gateway. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected entry up or down in the list.
Add	Adds an entry to the server list.
Set	Updates the selected entry. Select a server from the list; modify the settings; click Set to update the entry.
Clear Fields	Clears all fields for the selected server.
SOCKS Server Name	Specify a name that helps distinguish this SOCKS server from other SOCKS servers.
SOCKS Server Host	Specify the SOCKS server IP address, or a hostname that is resolvable by your internal DNS service.
SOCKS Port	Specify the port on which the SOCKS server listens.
Default SOCKS Server	Select this option to make this SOCKS server the default SOCKS server.
SOCKS User Name	When SOCKS authentication is used, specify the SOCKS user name with which to authenticate.
SOCKS Password	When SOCKS authentication is used, specify the password that goes with the specified user.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes are lost.

Socks Server Rules	Displays a table listing the rules in the socks.config file that specify the SOCKS servers that Content Gateway must go through to access specific origin servers, and the order in which Content Gateway goes through the SOCKS server list.
	You can also specify the origin servers that you want the proxy to access directly, without going through a SOCKS server.
	Do not route through SOCKS server Rule Type does not support non-HTTP traffic.
Refresh	Updates the table to display the current rules in the socks.config file.
Edit File	Opens the configuration file editor for the socks.config file.
	socks.config Configuration File Editor
rule display box	Lists the <i>socks.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
Rule Type	Select Route through SOCKS server to specify the origin servers that you want the proxy to route through a SOCKS server.
	Select Do not route through SOCKS server to specify the origin servers that you want the proxy to access directly, bypassing the SOCKS server(s).
	Do not route through SOCKS server Rule Type does not support non-HTTP traffic.
Destination IP	For Route through SOCKS server , specify either a single IP address <i>or</i> a range of IP addresses of origin servers for which Content Gateway must use the SOCKS servers specified in the SOCKS Servers field below.
	For Do not route through SOCKS server , specify the IP addresses of the origin servers that you want the proxy to access directly (without going through the SOCKS server). You can enter a single IP address, a range of IP addresses, or a list of IP addresses. Separate each entry in the list with a comma. Do not specify the all networks broadcast address: 255.255.255.255.
SOCKS Server	For a Route through SOCKS server rule , select the SOCKS server(s) through which to route requests.
Round Robin	Specifies how strictly Content Gateway will follow round robin. You can select strict , or false .

Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Configure > Security > SOCKS > Options

Server Connection Timeout	Specifies how many seconds Content Gateway waits attempting to connect to a SOCKS server before timing out.
Connection Attempts Per Server	Specifies how many times Content Gateway attempts to connect to a given SOCKS server before marking the server as unavailable.
Server Pool Connection Attempts	Specifies how many times Content Gateway attempts to connect to a given SOCKS server in the pool before giving up.

Related concepts

Configuring SOCKS firewall integration on page 160 socks.config on page 461

Related tasks Configuring SOCKS servers on page 160

Subsystems

The Subsystems configuration options are divided into the following categories:

Related concepts Logging on page 325 Networking on page 329

Related reference

Cache on page 324

Cache

Configure > Subsystems > Cache > General Cache Configuration
Ram Cache Size	Specifies the size of the RAM cache, in bytes. The default size is 104857600 (100 MB). If you change this option, you must restart Content Gateway.
Maximum Object Size	Specifies the maximum size allowed for objects in the cache. A value of 0 (zero) means that there is no size restriction.

Logging

Configure > Subsystems > Logging > General

Logging	Enables or disables event logging so that transactions are recorded into event log files and/or error log files.
	Select Log Transactions and Errors to log transactions into your selected event log files and errors in the error log files.
	Select Log Transactions Only to log transactions into your selected event log files only. Content Gateway does not log errors in the error log files.
	Select Log Errors Only to log errors in the error log files only. Content Gateway does not log transactions into your selected event log files.
	Select Disabled to turn off logging.
Log Directory	Specifies the path of the directory in which Content Gateway stores event logs. The path of this directory must be the same on every node in the Content Gateway cluster failover group. The default is: /opt/ WCG/logs

Log Space: Limit	Specifie megaby log files	es the maximum amount of space (in /tes) allocated to the logging directory for the s.
	When C set to 5	Content Gateway is on an appliance, the size is 120 (5 GB) and cannot be changed.
	When C server, is config	Content Gateway is installed on a stand-alone the default size is 20480 (20 GB) and the size gurable.
	Ę	Note
		Transaction logs can consume a lot of space. Make sure that this limit is smaller than the actual space available on the partition that contains the logging directory.
Log Space: Headroom	Specifie the Aut autodel space a headroo	es the tolerance for the log space limit. If o-Delete Rolled Files option is enabled, etion is triggered when the amount of free available in the logging directory is less than the om.
Log Rolling: Enable/ Disable	Enables or disables log file rolling. To keep log files down to manageable sizes, you can roll them at regular intervals. See <i>Rolling event log files</i> .	
Log Rolling: Offset Hour	Specifies the hour when log rolling takes place. You can set a time of the day in the range 0 to 23. For example, if the offset hour is 0 (midnight) and the roll interval is 6, the log files are rolled at 00:00, 06:00, noon, and 18:00.	
Log Rolling: Interval	Specifies the amount of time Content Gateway enters data in log files before rolling them to .old files. The minimum value is 300 seconds (five minutes). The default value is 21600 seconds (6 hours). The maximum value is 86400 (1 day).	
Log Rolling: Auto- Delete Rolled Files	Enables space in triggere in the lo Headro	s autodeletion of rolled log files when available in the log directory is low. Autodeletion is ad when the amount of free space available og directory is less than the Log Space bom .
Reverse DNS lookup for Threat Tracking	Enables inclusio dashbo Forcepo	s or disables reverse DNS lookups to facilitate n of the client host name in the Threats ard in the Web Security module of the pint Security Manager, and in logs and reports.
		CAUTION
		To achieve the expected results and avoid unexpected network behaviors, before enabling this option be sure that reverse DNS is configured in your network.

Related concepts

Rolling event log files on page 227

Configure > Subsystems > Logging > Formats

Squid Format: Enable/ Disable	Enables or disables the Squid log format
	Chables of disables the oquid log format.
Squid Format: ASCII/ Binary	Select ASCII or Binary as the type of log files to be created.
Squid Format: Filename	Specifies the name used for Squid log files. The default filename is squid.log .
Squid Format: Header	Specifies the text header you want Squid log files to contain.
Netscape Common Format: Enable/ Disable	Enables or disables the Netscape Common log format.
Netscape Common Format: ASCII/ Binary	Select ASCII or Binary as the type of log file to be created.
Netscape Common Format: Filename	Specifies the name used for Netscape Common log files. The default filename is common.log .
Netscape Common Format: Header	Specifies the text header you want Netscape Common log files to contain.
Netscape Extended Format: Enable/ Disable	Enables or disables the Netscape Extended log format.
Netscape Extended Format: ASCII/ Binary	Select ASCII or Binary as the type of log file to be created.
Netscape Extended Format: Filename	Specifies the name used for Netscape Extended log files. The default filename is extended.log .
Netscape Extended Format: Header	Specifies the text header you want Netscape Extended log files to contain.
Netscape Extended 2 Format: Enable/Disable	Enables or disables the Netscape Extended-2 log format.
Netscape Extended 2 Format: ASCII/Binary	Select ASCII or Binary as the type of log file to be created.
Netscape Extended 2 Format: Filename	Specifies the name used for Netscape Extended-2 log files. The default filename is extended2.log .
Netscape Extended 2 Format: Header	Specifies the text header you want Netscape Extended-2 log files to contain.

Configure > Subsystems > Logging > Splitting

Split ICP Logs	When enabled, Content Gateway records ICP transactions in a separate log file. When disabled, Content Gateway records ICP transactions in the same log file with HTTP and FTP entries.
Split Host Logs	When enabled, Content Gateway creates a separate log file for each of the hosts listed in the log_hosts.config file. When disabled, Content Gateway records transactions for all hosts in the same log file.

Configure > Subsystems > Logging > Collation

	4
Collation Mode	Specifies the log collation mode for this Content Gateway node. You can use the log file collation feature to keep all logged information in one place. For more information about log file collation, see <i>Collating</i> <i>event log files</i> .
	Select Collation Disabled to disable log collation on this Content Gateway node.
	Select Be a Collation Server to configure this Conter Gateway node to be the collation server.
	Select Be a Collation Client to configure this Content Gateway server to be a collation client. A Content Gateway server configured as a collation client sends only the active standard log files, such as Squid, Netscape Common, and so on, to the collation server. If you select this option, enter the hostname of the collation server for your cluster in the Log Collation Server field.
	Note
	When logs are collated, the source of the log entry—its node of origin—is lost unless you turn on the Log collation host tagged option (described below).
	Log collation consumes cluster bandwidth in sending all log entries to a single node. It can therefore affect the performance of the cluster.
	If you want Content Gateway as a collation client to send custom (XML-based) log files, you must specify LogObject in the logs_xml.config file.
Log Collation Server	Specifies the hostname of the log collation server to which you want to send log files.

Log Collation Port	Specifies the port used for communication between the collation server and client. You must specify a port number in all cases, except when log collation is inactive. The default port number is 8085. Note: Do not change the port number unless there is a conflict with another service already using the port.
Log Collation Secret	Specifies the password for the log collation server and the other nodes in the cluster. This password is used to validate logging data and prevent the exchange of arbitrary information.
Log Collation Host Tagged	When this option is enabled, Content Gateway adds the hostname of the node that generated the log entry to end of the entry in the collated log file.
Log Collation Orphan Space	Specifies the maximum amount of space (in megabytes) allocated to the logging directory for storing orphan log files on the Content Gateway node. Content Gateway creates orphan log entries when it cannot contact the log collation server.

Related concepts

Collating event log files on page 232

Configure > Subsystems > Logging > Custom

Custom Logging	Enables or disables custom logging.
Custom Log File Definitions	Displays the <i>logs_xml.config</i> file so that you can configure custom (XML-based) logging options.

Related concepts

logs_xml.config on page 377

Networking

The Networking configuration options are divided into the following categories:

Related concepts ARM on page 331 DNS Proxy on page 344 DNS Resolver on page 344 Health Check URLs on page 350

Related reference Connection Management on page 330 WCCP on page 339 ICAP on page 349 Virtual IP on page 350

Connection Management

The options on the Connection Management pages allow you to tune several important properties of proxy behavior, including connection throttling and load shedding, and individual client connection limits and rates.

By default, Content Gateway accepts 60,000 connections. A connection throttle event occurs when client or origin server connections reach 90% of half the configured limit (27,000 by default). When a connection throttle event occurs, Content Gateway continues processing all existing connections and queues new client connection requests until the connection count falls below the limit.

If you think that Content Gateway is hitting the connection limits, you should monitor the Performance graphs to get an accurate reading of connection activity. In particular, check the **Active Client Connections** and **TCP ESTABLISHED Connections** graphs. You can also check error messages in the system log file, error log file, or event log files.

Throttling Net Connections	Specifies the maximum number of network connections that Content Gateway accepts. The default value is 60,000.
	Setting a Content Gateway throttle limit helps to prevent system overload when traffic bottlenecks develop. When network connections reach this limit, Content Gateway queues new connections until existing connections close.
	Do not set this variable below the minimum value of 100.

Configure > Networking > Connection Management > Throttling

Configure > Networking > Connection Management > Load Shedding

Maximum Connections	Specifies the maximum number of client connections allowed before the ARM starts forwarding incoming requests directly to the origin server. The default value is 1 million connections.
	If you change this option, you must restart Content Gateway.

Configure > Networking > Connection Management > Client Connection Control

Specifies:

- Client concurrent connection limits
- Client connection rate limits
- Proxy response when a limit is exceeded
- A list of clients excepted from the limits

Concurrent Connection Limit: Maximum concurrent connections	Specifies the maximum number of concurrent HTTP/ HTTPS connections a client is allowed. The default is 1000. The supported range is: 1 - 45000
Concurrent Connection Limit: Alert when limit exceeded	When enabled, causes Content Gateway to generate an alert when a client exceeds the maximum concurrent connection limit. In addition to displaying the alert in the Content Gateway manager, it is also logged in /var/log/ messages and content_gateway.out.
Concurrent Connection Limit: Close excessive connections when limit exceeded	When enabled, causes Content Gateway to close excessive connections when the limit is exceeded.
Connection Rate Limit: Maximum connection rate	Specifies the maximum connections per second, averaged over a minute, that a client can make. The default is 100. The supported range is: 1 - 1000
Connection Rate Limit: Alert when limit exceeded	When enabled, causes Content Gateway to generate an alert when a client exceeds the maximum connection rate limit.
	In addition to displaying the alert in the Content Gateway manager, it is also logged in /var/log/ messages and content_gateway.out.
Connection Rate Limit: Close excessive connections when limit exceeded	When enabled, causes Content Gateway to close excessive connections when the limit is exceeded.
Exceptions	 Specifies IP addresses and/or IP address ranges to which connection limits are not applied. IP addresses can be IPv4 or IPv6 (IPv6 support must be enabled). Multiple addresses or ranges can be specified in a comma-separated list that can contain up to: 64 IPv4 addresses 32 IPv4 address ranges 24 IPv6 addresses
	12 IPv6 address ranges

ARM

The Adaptive Redirection Module (ARM) performs several essential functions including sending device notifications for cluster communication interface failover and inspection of incoming packets before a routing decision is made and redirecting the packets to Content Gateway for processing.

The ARM is always active. For more information, see The Content Gateway ARM.

Related tasks

The Content Gateway ARM on page 36

Configure > Networking > ARM > General

Redirection Rules	Displays the redirection rules in the <i>ipnat.conf</i> file that specify how incoming packets are redirected when the proxy is serving traffic transparently. During installation, Content Gateway creates a small number of default rules. These rules can be added to and modified. IPv4 and IPv6 addresses are supported. During operation, Content Gateway traverses the list top down and applies the first matching rule. Updates the table to display the most up-to-date rules
	in the ipnat.conf file.
Edit File	Opens the configuration file editor for the ipnat.conf file.
	ipnat.conf Configuration File Editor
rule display box	Lists the <i>ipnat.conf</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
Ethernet Interface	Specifies the Ethernet interface that traffic will use to access the Content Gateway machine: for example, eth0 on Linux.
Connection Type	Specifies the connection type that applies for the rule: TCP or UDP.
Destination IP	Specifies the IP address from which traffic is sent.
	0.0.0.0 or :: match all IP addresses.
Destination CIDR	Specifies the IP address in CIDR (Classless Inter- Domain Routing) format, such as 1.1.1.0/24. Entering a value in this field is optional.
Destination Port	Specifies the traffic destination port: for example, 80 for HTTP traffic.
Redirected Destination IP	Specifies the IP address of your Content Gateway server.
Redirected Destination Port	Specifies the proxy port: for example, 8080 for HTTP traffic.
User Protocol (Optional)	When dns is selected, the ARM redirects DNS traffic to Content Gateway: otherwise, DNS traffic is bypassed.
Apply	Applies the configuration changes.

Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes are discarded.
IP Spoofing: Enabled/ Disabled	Enables or disables the IP spoofing option, which configures Content Gateway to establish connections to origin servers with the client IP address instead of the Content Gateway IP address. For more information, see <i>Content Gateway IP spoofing</i> .
	Note IP spoofing requires precise control of the routing paths on your network, overriding the normal routing process for traffic running on TCP port 80 and 443.
	ipnat.conf Configuration File Editor
rule display box	Lists the <i>ipnat.conf</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
Ethernet Interface	Specifies the Ethernet interface that traffic will use to access the Content Gateway machine: for example, eth0 on Linux.
Connection Type	Specifies the connection type that applies for the rule: TCP or UDP.
Destination IP	Specifies the IP address from which traffic is sent. 0.0.0.0 or :: match all IP addresses.
Destination CIDR	Specifies the IP address in CIDR (Classless Inter- Domain Routing) format, such as 1.1.1.0/24. Entering a value in this field is optional.
Destination Port	Specifies the traffic destination port: for example, 80 for HTTP traffic.
Redirected Destination IP	Specifies the IP address of your Content Gateway server.
Redirected Destination Port	Specifies the proxy port: for example, 8080 for HTTP traffic.
User Protocol (Optional)	When dns is selected, the ARM redirects DNS traffic to Content Gateway: otherwise, DNS traffic is bypassed.
Apply	Applies the configuration changes.

Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes are discarded.
IP Spoofing: Enabled/ Disabled	Enables or disables the IP spoofing option, which configures Content Gateway to establish connections to origin servers with the client IP address instead of the Content Gateway IP address. For more information, see <i>Content Gateway IP spoofing</i> , page 77.
	Warning
	IP spoofing requires precise control of the routing paths on your network, overriding the normal routing process for traffic running on TCP port 80 and 443.
Range Based IP Spoofing: Enabled/ Disabled	Enables or disables the range-based IP spoofing extension. This extension supports the specification of IP addresses and ranges of addresses that are mapped to specified IP addresses for spoofing.
	Many groups can be specified. However, use this feature judiciously because list traversal adds overhead to every connection request. The larger the list, the more overhead.
	The list is traversed in order (as displayed). The first match is applied.
	Clients that don't match a grouping are spoofed with their own IP address (basic IP spoofing).
	For more information, see <i>Content Gateway IP</i> spoofing.

Range Based IP Spoofing: Address table	In the Client IP Addresses field, enter a comma separated list of individual IP addresses and/or IP address ranges. Do not use spaces.
	You can use:
	A simple IP address, such as 123.45.67.8
	 CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24.
	 A range separated by a dash, such as 1.1.1.1-2.2.2.2
	 Any combination of the above, separated by commas, such as:
	1.1.1.0/24,25.25.25.25,123.1.23.1-123.1.23.123
	 A maximum of 64 IPv4 addresses or 32 IPv4 address ranges.
	In the Spoofed IP Address field, enter the IP address to use with matching clients. This is the spoofed IP address.
	To add a row to the table, click Add Row.
	To remove a row from the table, delete the contents of the cells. When you click Apply the empty row(s) is removed
	The table always has a minimum of 5 rows.
	Restart Content Gateway to put changes into effect.

Related concepts ipnat.conf on page 376 Content Gateway IP spoofing on page 63

Configure > Networking > ARM > Static Bypass

Static bypass rules route requests around the proxy (bypass). Rules can be defined for clients (sources), origin servers (destinations), or both (pairs). See *Static bypass rules*.



Important

This feature is for transparent proxy deployments only.

Static Bypass table	Lists the configured static bypass rules. When Content Gateway is serving transparent traffic, the proxy uses these rules to determine whether to bypass incoming client requests or attempt to serve them transparently. Rules are stored in <i>bypass.config</i>
Refresh	Updates the table to display the most up-to-date rules in the bypass.config file.

Edit File	Opens the configuration file editor for the bypass.config file.
	bypass.config Configuration File Editor
rule display box	Lists the <i>bypass.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
Rule Type	Specifies the rule type:
	A bypass rule bypasses specified incoming requests.
	A deny_dyn_bypass rule prevents the proxy from bypassing specified incoming client requests dynamically (a deny bypass rule can prevent Content Gateway from bypassing itself).
Source IP	Specifies the source IP address in incoming requests that the proxy must bypass or deny bypass. The IP address can be one of the following:
	A simple IP address, such as 123.45.67.8
	In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24.
	A range separated by a dash, such as 1.1.1.1-2.2.2.2 Any combination of the above, separated by commas,
	such as 1.1.1.0/24, 25.25.25.25, 123.1.23.1-123.1.23.123
Destination IP	Specifies the destination IP address of incoming requests that the proxy must bypass or deny bypass. The IP address can be one of the following:
	A simple IP address, such as 123.45.67.8
	In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24
	A range separated by a dash, such as 1.1.1.1-2.2.2.2 Any combination of the above, separated by commas,
	such as 1.1.1.0/24, 25.25.25.25, 123.1.23.1-123.1.23.123
Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Related concepts Static bypass rules on page 60 bypass.config on page 370

Configure > Networking > ARM > Dynamic Bypass

Dynamic Bypass	Enables or disables the dynamic bypass option to bypass the proxy and go directly to the origin server when clients or servers cause problems. Dynamic bypass rules are deleted when you stop Content Gateway.
Behavior: Non-HTTP, Port 80	Select Enabled to enable dynamic bypass when Content Gateway encounters non-HTTP traffic on port 80.
	Select Disabled to disable dynamic bypass when Content Gateway encounters non-HTTP traffic on port 80.
	Select Source-Destination to enable dynamic source/ destination bypass when Content Gateway encounters non-HTTP traffic on port 80.
	Select Destination Only to enable dynamic destination bypass when Content Gateway encounters non-HTTP traffic on port 80.
Behavior: HTTP 400	Select Enabled to enable dynamic bypass when an origin server returns a 400 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 400 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 400 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 400 error.
Behavior: HTTP 401	Select Enabled to enable dynamic bypass when an origin server returns a 401 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 401 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 401 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 401 error.

Behavior: HTTP 403	Select Enabled to enable dynamic bypass when an origin server returns a 403 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 403 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 403 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 403 error.
Behavior: HTTP 405	Select Enabled to enable dynamic bypass when an origin server returns a 405 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 405 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 405 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 405 error.
Behavior: HTTP 406	Select Enabled to enable dynamic bypass when an origin server returns a 406 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 406 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 406 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 406 error.
Behavior: HTTP 408	Select Enabled to enable dynamic bypass when an origin server returns a 408 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 408 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 408 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 408 error.

Behavior: HTTP 500	Select Enabled to enable dynamic bypass when an origin server returns a 500 error.
	Select Disabled to disable dynamic bypass when an origin server returns a 500 error.
	Select Source-Destination to enable dynamic source/ destination bypass when an origin server returns a 500 error.
	Select Destination Only to enable dynamic destination bypass when an origin server returns a 500 error.

WCCP



Note

The WCCP configuration options appear on the Configure pane only if you have enabled WCCP in the Features table on the **Configure > My Proxy > Basic > General** tab.

The options defined in the **wccp.config** configuration file control the use of WCCP with Content Gateway. Entries should be defined and maintained using the editor provided on **Configure > Networking > WCCP**.

Administrators should have a good working knowledge of WCCP. Only WCCP v2 is supported.

It is recommended that you consult the documentation and the manufacturer's support site for information regarding optimal configuration and performance of your

WCCP v2 device. Most devices should be configured to take best advantage of hardware-based redirection. With Cisco devices, the most recent version of IOS is usually best.

For every active WCCP service group, there must be a corresponding ARM redirection rule. See ARM.

For a complete description of Content Gateway support for WCCP v2, see *Transparent interception with WCCP v2 devices*.

Option	Description
WCCP Service Groups	Displays a table of the service groups defined in the wccp.config file. WCCP service group configuration defines WCCP behavior. Column fields are explained in the Configuration Editor entries below.
Refresh	Refreshes the table to display the current definitions in the wccp.config file.
Edit File	Opens wccp.config in the configuration file editor.

Option	Description
Synchronize in the Cluster	When there are several Content Gateway nodes in a cluster:
	Enable this option to cause the WCCP configuration (wccp.config) to be synchronized in the cluster. This allows configuration changes to be made on any node in the cluster.
	Disable this option to cause the WCCP configuration to not be synchronized in the cluster. This requires that changes to the WCCP configuration be made individually on each node. A common use case for this is to control which service groups are enabled/ disabled on each node, and to use proportional load distribution with weight .
	If after being disabled this option is enabled, the configuration on the node on which the option is enabled is used to initially synchronize the cluster.
	wccp.config Configuration File Editor
Service group display box	Lists the WCCP service group definitions. Select an entry in the list to edit it.
	Use the "X" button to delete the selection.
	List order has no meaning; therefore, the up and down arrows can be ignored.
Add	Adds a new service group definition. After Add is clicked, the new definition is displayed in the box at the top of the page.
Set	Accepts modifications to the selected service group definition, displaying the new values in the box at the top of the page.
	Service Group Information
Service Group Status	Enables or disables the service group.
	If you change this option, you must restart Content Gateway.
Service Group Name	Specifies a unique service group name. This is as an aid to administration.
Service Group ID	Specifies a service group ID between 0-255. This ID must also be configured on the router(s).
	If the specified number is already in use, an error is displayed when Add or Set is clicked.
Protocol	Specifies the protocol, TCP or UDP, that applies to this service group.

Option	Description
Ports	Specifies the ports the service group will use.
	Specify ports can be used to list up to 8 ports in a comma-separated list.
	All ports can be selected to redirect traffic from all ports.
Network Interface	Specifies the Ethernet interface on this Content Gateway host system to use with this service group. On Forcepoint appliances, use the CLI command 'show interface info' to view the logical name to physical interface bindings.
	Mode Negotiation
Special Device Profile	Select ASA Firewall to specify that traffic is routed to the proxy by a Cisco ASA firewall. When this option is selected, GRE is automatically selected as the Packet Forward Method and Packet Return Method. These settings are required and cannot be changed.
Packet Forward Method	Specifies the preferred encapsulation method used by the WCCP router to transmit intercepted traffic to the proxy. If the router supports GRE and L2, the method specified here is used.
	Important
	GRE and Multicast are incompatible.
	If you change the forward or return method configuration while there is an active connection with the WCCP device, in order to re- negotiated the method you must force the current connection to terminate. Typically, this means turning off the service group on the WCCP device for 60 seconds. See the documentation for your WCCP device.
Packet Return Method	Specifies the preferred packet encapsulation method
	used to return rejected or declined traffic to the WCCP router.
	Note
	If Content Gateway is configured with a Forward/Return method that the router does not support, the proxy attempts to negotiate a method supported by the router.
	Selecting L2 requires that the router or switch be Layer 2-adjacent (in the same subnet) as Content Gateway.

Option	Description
	Advanced Settings
Assignment Method	Specifies the method that the router will use to distribute intercepted traffic across multiple proxy servers. Choices are HASH and MASK .
	The MASK value is applied up to 6 significant bits (in a cluster, a total of 64 buckets are created).
	See your WCCP documentation for more information about assignment method. Use the value recommended in the manufacturer's documentation for your device.
Distribution attribute(s)	Specifies the attribute that the assignment method uses to determine which requests are distributed to which proxy servers.
	If the assignment method is HASH, select one or more distribution attributes.
	If the assignment method is MASK, select one distribution attribute.
Weight	This option is only useful when Synchronize in the Cluster is disabled .
	Specifies the distribution of requests to servers in a cluster by proportional weighting. Set weight to a value that is the desired proportion of the total flow of traffic.
	When all cluster members have a value of 0 (the default), distribution is equal. If any member has a non-zero value, distribution is proportional, relative to the weight values of other members. Members that continue to have a value of zero, receive no traffic.
	See WCCP load distribution.
Reverse Service Group ID	For use when IP spoofing is enabled.
	When IP spoofing is enabled, the proxy advertises a reverse service group for each enabled WCCP forward service group. The reverse service group must be applied along the return path of origin server responses to the proxy.
	Router Information
Security (optional)	Enables or disables security so that the router and Content Gateway can authenticate each other.
	If you enable security in Content Gateway, you must also enable security on the router. See your router documentation.
	If you change this option, you must restart Content Gateway.

Option	Description
Security:Password	Specifies the password used for authentication. The password must be the same password as that configured on the router for the associated service group ID and can be a maximum of eight characters long. If you change this option, you must restart Content Gateway.
Multicast (optional)	Enables or disables WCCP multicast mode.
	Important
	Cannot be used with GRE packet Forward/Return method.
	If you change this option, you must restart Content Gateway.
Multicast: IP Address	Specifies the multicast IP address.
	If you change this option, you must restart Content Gateway.
WCCP Routers: Router IP Address	Specifies the IP addresses of up to 10 WCCP v2- enabled routers.
	If ASA_Firewall was selected as the Service Device Profile, entries should include both the router IP Address and the WCCP router ID, separated by /.
	A total of 24 WCCP routers across all service groups is supported if the Packet Forward Method or Packet Return Method is GRE. An IP address for the local GRE tunnel endpoint for each router must also be provided.
	If you change this option, you must restart Content Gateway.

Option	Description
WCCP Routers: Local GRE Tunnel Endpoint IP Address	If GRE is selected for Packet Return Method, also specify Local GRE Tunnel Endpoint IP Addresses, except when the device is an ASA firewall.
	These are Content Gateway tunnel endpoints for the associated Router IP Addresses.
	A Local GRE Tunnel Endpoint IP Address:
	 Must be unique for every router in the table
	 Must not be assigned to any other device
	 Must be a routable IP address
	Should reside on the same subnet as the proxy. If it is not, you must define a route for it.
	 Is not intended to be a client-facing proxy IP address
	 Is bound to the physical interface specified for the service group (on Forcepoint appliances, use the CLI command 'show interface info' to view the logical name to physical interface bindings.
WCCP Routers: GRE Tunnel Next Hop Router IP Address	Specify a GRE Tunnel Next Hop Router IP Address (must be in IPv4 format) when GRE Packet Return Method is configured and Content Gateway does not have a route back to the WCCP router. You can use "ping" to test connectivity to the router.

Related concepts ARM on page 331 Transparent interception with WCCP v2 devices on page 38 WCCP load distribution on page 41

DNS Proxy

Note

Ę

The DNS Proxy configuration options appear on the Configure pane only if you have enabled DNS Proxy in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Networking > DNS Proxy

DNS Proxy Port	Specifies the port that Content Gateway uses for DNS
	traffic. The default port is 5353.

DNS Resolver

Configure > Networking > DNS Resolver > Resolver

Local Domain Expansion	Enables or disables local domain expansion so that Content Gateway can attempt to resolve unqualified hostnames by expanding to the local domain. For example, if a client makes a request to an unqualified host named hostx, and if the WCG local domain is y.com , Content Gateway expands the hostname to hostx.y.com .
DNS Preference	Specifies the IP version preference when IPv6 support is enabled in Content Gateway and a web server supports both IPv4 and IPv6.
	Select IPv4 to cause the proxy to prefer IPv4. Select IPv6 to cause the proxy to prefer IPv6.
	The DNS Preference is not applied to FTP requests made in transparent proxy mode. The proxy uses the IP address sent with the request.
DNS Preference Exceptions	List IPv4 / IPv6 preference rules for specific origin servers.
Refresh	Updates the table to display the most up-to-date rules. Click this button after you have added or modified rules with the configuration file editor.
Edit File	Opens the configuration file editor.
	dns_prefer_exception.config File Editor
rule display box	dns_prefer_exception.config File Editor Displays an ordered list of the
rule display box	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.
rule display box	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of thebox allow you to delete or move the selected rule up ordown in the list.
rule display box Add	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.
rule display box Add Set	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.Updates the selected rule with the values in the entry fields.
rule display box Add Set Name	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.Updates the selected rule with the values in the entry fields.Specify a unique name to aid in administering rules.
rule display box Add Set Name Destination Host	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.Updates the selected rule with the values in the entry fields.Specify a unique name to aid in administering rules.Specify the destination hostname.
rule display box Add Set Name Destination Host Preferred Format	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this
rule display box Add Set Name Destination Host Preferred Format Apply	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.Updates the selected rule with the values in the entry fields.Specify a unique name to aid in administering rules.Specify the destination hostname.Specify the preferred IP version, IPv4 or IPv6.Applies the configuration changes.
rule display box Add Set Name Destination Host Preferred Format Apply Close	dns_prefer_exception.config File EditorDisplays an ordered list of thedns_prefer_exception.config file rules.Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button.Updates the selected rule with the values in the entry fields.Specify a unique name to aid in administering rules.Specify the destination hostname.Specify the preferred IP version, IPv4 or IPv6.Applies the configuration changes.Exits the configuration file editor.

Configure > Networking > DNS Resolver > Host Database

These setting pertain to all DNS name resolution performed by Content Gateway, including DNS Proxy.

DNS Lookup Timeout		Specifies the maximum number of seconds the proxy can wait for a lookup response from the DNS server.	
		es how long, in seconds, the proxy will wait making a second DNS request if there is no se to the first request. The value is stored in config.hostdb.lookup_timeout". The default s 120 seconds.	
		Important	
		This setting is not used. Instead the records.config entry "proxy.config.dns.lookup_timeout" is used. The default value is 20 seconds. proxy.config.dns.lookup_timeout specifies how long the proxy will wait for the DNS response after sending the request.	
Foreground Timeout	Specifie databas is used zero (th to-live (es how long DNS entries remain in the host se before they are flagged as stale. This setting only when "proxy.config.hostdb.ttl_mode" is not be default value is 0, which means use the time- ttl) value set by the DNS server. See <i>HostDB</i> .	
	For exa request hours o serving	ample, if this timeout is 24 hours and a client as an entry that has been in the database for 24 or longer, the proxy refreshes the entry before it.	
	The def	fault is 86400 seconds (144 minutes).	
		CAUTION	
		Setting the foreground timeout too low might slow response time. Setting it too high risks accumulation of incorrect information.	

Failed DNS Timeout	Specifies how long, in seconds, that a hostname is retained in the failed DNS lookup cache (default = 60). When the timeout expires, the hostname is removed from the cache and the next request for that hostname is sent to the DNS server.
	A DNS lookup failure is considered to have occurred when:
	There is no DNS response
	 There is a DNS response error code, including NXDOMAIN
	 There is an error parsing the DNS response code (there is a malformed response).
	Zero (0) is not a legal value.

Related reference

HostDB on page 436

Configure > Networking > DNS Resolver > Split DNS

Split DNS	Enables or disables the Split DNS option. When enabled, Content Gateway can use multiple DNS servers, depending on your security requirements. For example, you can configure the proxy to look to one set of DNS servers to resolve hostnames on your internal network, while allowing DNS servers outside the firewall to resolve hosts on the Internet. For information about using Split DNS, see <i>Using the Split</i> <i>DNS option</i> .
Default Domain	Specifies the default domain used for split DNS requests. If a hostname does not include a domain, Content Gateway appends the default domain name to the hostname before choosing which DNS server to use.
DNS Servers Specification	Displays a table listing the rules in the <i>splitdns.config</i> file that control which DNS server the proxy uses for resolving hosts under specific conditions.
Refresh	Updates the table to display the most up-to-date rules in the splitdns.config file. Click this button after you have added or modified rules with the configuration file editor.
Edit File	Opens the configuration file editor so that you can edit and add rules to the splitdns.config file. The configuration file editor page is described below.

	splitdns.config Configuration File Editor
rule display box	Lists the <i>splitdns.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list.
Add	Adds a new rule to the rule display box at the top of the configuration file editor page. Enter information in the fields provided before you click this button.
Set	Updates the rule display box at the top of the configuration file editor page. Select a rule and change its properties before you click this button.
Primary Destination Type	Specifies that DNS server selection is based on the destination domain (dest_domain), destination host (dest_host), or on a regular expression (url_regex).
Primary Destination Value	Specifies the value of the primary destination. Place the symbol "!" at the beginning of the value to specify the NOT logical operator.
	The NOT logical operator applies only if the number of rules does not exceed the value set in proxy.config.dns.splitdns.file_match.count defined in records.config .
DNS Server IP	Specifies the DNS server to use with the primary destination specifier. You can specify a port using a colon (:). If you do not specify a port, 53 is used. You can specify multiple DNS servers separated by spaces or by semicolons (;).
Default Domain Name (Optional)	Specifies the default domain name to use for resolving hosts. Only one entry is allowed. If you do not provide the default domain, the system determines its value from /etc/resolv.conf .
Domain Search List (Optional)	Specifies the domain search order. You can specify multiple domains separated by spaces or by semicolons (;). If you do not provide the search list, the system determines the value from /etc/resolv.conf .
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes are lost.

Related concepts

splitdns.config on page 463

Related tasks

Using the Split DNS option on page 163

ICAP



Note

The ICAP configuration option appears on the Configure pane only if you have enabled **ICAP** in the **Features** table on the **Configure > My Proxy > Basic > General** tab.

ICAP provides an alternate interface to Forcepoint DLP, and other data security services that are ICAPconversant. A primary and backup URI can be specified, and failover and load balancing can be configured. See *Configuring the ICAP client* and the subsection for *ICAP failover and load balancing*.

ICAP Service URI	Specifies the Uniform Resource Identifier for the ICAP service. The format is: icap://hostname:port/path
	For example:
	icap://ICAP_machine:1344/reqmod The default ICAP port is 1344. If you are using the default port, you need not specify it in the URI.
	An optional secondary URI service can be specified immediately after the first by adding a comma and the URI of the second service, no spaces.
Analyze HTTPS Content	Select whether decrypted traffic should be sent to the data protection software for analysis or sent directly to the destination.
Analyze FTP Uploads	Select whether to send FTP upload requests to the data protection software for analysis. The FTP proxy feature must be enabled. See <i>FTP</i> .
Action for Communication Errors	Select whether to allow traffic or send a block page if Content Gateway receives an error while communication with the data protection software.
Action for Large files	Select whether to allow traffic or send a block page if a file larger than the size limit specified in the data protection software is sent. The default size limit in Forcepoint DLP is 50 MB.

Configure > Networking > ICAP

Related concepts

ICAP failover and load balancing on page 108

Related tasks

Configuring the ICAP client on page 106

Related reference

FTP on page 294

Virtual IP



Note

The Virtual IP configuration options appear on the Configure pane only if you have enabled Virtual IP in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Networking > Virtual IP

Virtual IP Addresses	Displays a table listing the virtual IP addresses managed by Content Gateway.
Refresh	Updates the table to display the most up-to-date list of virtual IP addresses. Click this button after you have added to or modified the list of virtual IP addresses with the configuration file editor.
Edit File	Opens the configuration file editor so that you can edit and add to the list of virtual IP addresses.
	vaddrs.config Configuration File Editor
rule display box	Lists the virtual IP addresses. Select a virtual IP address to edit it. The buttons on the left of the box allow you to delete or move the selected virtual IP address up or down in the list.
Add	Adds a new virtual IP address to the rule display box at the top of the configuration file editor page.
Set	Updates the rule display box at the top of the configuration file editor page.
Virtual IP Address	Specifies the virtual IP address managed by Content Gateway.
Ethernet Interface	Specifies the network interface assigned to the virtual IP address.
Sub-Interface	Specifies the subinterface ID. This is a number between 1 and 255 that the interface uses for the address.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor.
	Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Health Check URLs

Content Gateway includes 3 URLs that return proxy health and performance information in the HTTP response. These URLs are designed to help load balancers optimize performance by acquiring and adjusting for real-time state information of each proxy node.

The default port for health check URLs is 8083. The value can be changed in records.config by assigning the desired value to **proxy.config.admin.autoconf_port**

Configure > Networking > Health Check URLs

Force Health Checks to Report Proxy Down	
Enable/Disable	When enabled, all health check URLs sent to this proxy report WSDOWN.
	The URL response will be similar to:
	HTTP/1.0 503 Service Unavailable
	Server: Content Gateway Manager 7.7.0
	Date: Thu, 26 Jul 2012 20:26:14 GMT
	Cache-Control: no-store
	Pragma: no-cache
	Content-type: text/plain
	Content-length: 6
	WSDOWN
Health Check URLs	The load balancer should consider the service down if the URL request fails for the following reasons:
	No TCP connection proxy down
	 Response too slow proxy deadlocked or not responsive
	Invalid response
http://[Content Gateway IP address]: 8083/health.basic	Checks connectivity with Content Gateway and responds with WSUP or WSDOWN.
http://[Content Gateway IP address]: 8083/ health.app.filtering	Checks the health of Filtering Service responses to Content Gateway requests and reports WSUP or WSDOWN.

http://[Content Gateway IP address]: 8083/health.load	If the health.basic URL reports WSDOWN, this URL also reports WSDOWN
	Otherwise, health.load returns:
	 CPU usage (operating system load average)
	 Connection usage (number of open connections)
	 Bandwidth usage
	How these values are calculated and how they can be customized is described below.
	The default response will look similar to:
	HTTP/1.0 200 OK
	Server: Content Gateway Manager 7.7.0
	Date: Thu, 26 Jul 2012 20:26:14 GMT
	Cache-Control: no-store
	Pragma: no-cache
	Content-type: text/plain
	Content-length: xx
	Load=2253
	Conns=5150
	Mbps=6.42

A format file, /opt/WCG/config/health.load.template, allows for customization of the response format.

Format specifiers are:

%L = Load (integer)

```
%C = Connections integer)
```

%B = Bandwidth in Mbps (double)

```
%% = %
```

The default health.load.template file is:

Load=%L Conns=%C Mbps=%B

Here is health.load.template modified to respond with an xml-like format:

```
<load>
```

<item name="Load" value="%L" />

```
<item name="Conns" value="%C" />
```

```
<item name="Mbps" value="%B" />
```

</load>

How the values are calculated

The **Load** value, **%L**, is derived from the LINUX system load average. To make the value comparable across machines with varying numbers of cores, the number is divided by the number of cores on the system.

The calculation is:

```
// load avg values are 0.00 precision double avgs[3];
// get load averages for 1, 5, and 15 minutes getloadavg(avgs, 3);
// 5 minute_load_average * 10000 / number_of_cores Load = avgs[1] * 10000 / get_nprocs();
```

The **Connection** value, **%C**, is the sum of proxy.process.http.current_server_connections and proxy.process.http.current_client_connections.

The **Bandwidth** value, **%B**, is the value of proxy.node.client_throughput_out.



Note

HTTP connection and bandwidth information can be viewed in the Content Gateway manager on the **Monitor > Protocols > HTTP** page.

SSL

The SSL configuration options are divided into the following categories:

- Certificates (see Managing certificates)
- Decryption/Encryption (see SSL configuration settings for inbound traffic and SSL configuration settings for outbound traffic)
- Validation (see Validating certificates)
- Incidents (see Managing HTTPS website access)
- Client certificates (see Client certificates)
- Customization (see Customizing SSL connection failure messages)
- Internal Root CA (see Internal Root CA)

Related concepts

Validating certificates on page 131 Managing HTTPS website access on page 137 Client certificates on page 142 Internal Root CA on page 116

Related tasks

Managing certificates on page 124 SSL configuration settings for inbound traffic on page 127 SSL configuration settings for outbound traffic on page 129

Related reference

Customizing SSL connection failure messages on page 144

Appendix D Event Logging Formats

Contents

- Custom logging fields on page 355
- Logging format cross-reference on page 358

Custom logging fields

% <field symbol=""></field>	Description
{HTTP header field}cqh	Logs the information in the requested field of the client request HTTP header; for example, to log the Accept- Language field in client request headers, use:
	<{Accept-Language}cqh>
	This field cannot be used in custom log filters.
{HTTP header field}cqhua	Logs the information in the requested field of the client request HTTP header; for example, to log the User- Agent field in client request headers, use:
	<{User-Agent}cqhua>
{HTTP header field}pqh	Logs the information in the requested field of the proxy request HTTP header; for example, to log the Authorization field in proxy request headers, use:
	<{Authorization}pqh>
	This field cannot be used in custom log filters.
{HTTP header field}psh	Logs the information in the requested field of the proxy response HTTP header; for example, to log the Retry- After field in proxy response headers, use:
	<{Retry-After}psh>
	This field cannot be used in custom log filters.
{HTTP header field}ssh	Logs the information in the requested field of the server response HTTP header; for example, to log the Age field in server response headers, use:
	%<{Age}ssh>
	This field cannot be used in custom log filters.
caun	The client authenticated user name; result of the RFC931/ident lookup of the client user name.

% <field symbol=""></field>	Description
cfsc	The client finish status code; specifies whether the client request to the proxy was successfully completed (FIN) or interrupted (INTR).
chi	The client host IP; the IP address of the client's host machine.
cqbl	The client request transfer length; the body length in the client's request to Content Gateway in bytes.
cqhl	The client request header length; the header length in the client's request to Content Gateway.
cqhm	The HTTP method in the client request to Content Gateway: GET, POST, and so on (subset of cqtx).
cqhv	The client request HTTP version.
cqtd	The client request time stamp; specifies the date of the client request in the format yyyy-mm-dd , where yyyy is the 4-digit year, mm is the 2-digit month, and dd is the 2-digit day.
cqtn	The client request time stamp; date and time of the client's request (in the Netscape time stamp format).
cqtq	The client request time stamp with millisecond resolution.
cqts	The client request time stamp in Squid format; the time of the client request in seconds since January 1, 1970.
cqtt	The client request time stamp; the time of the client request in the format hh:mm:ss , where hh is the 2-digit hour in 24- hour format, mm is the 2-digit minutes, and ss is the 2-digit seconds. For example, 16:01:19.
cqtx	The full HTTP client request text, minus headers. For example: GET http://www.company.com HTTP/1.0
cqu	The client request URI; universal resource identifier (URI) of the request from client to Content Gateway (subset of cqtx).
cquc	The client request canonical URL; differs from cqu in that blanks (and other characters that might not be parsed by log analysis tools) are replaced by escape sequences. The escape sequence is a percentage sign followed by the ASCII code number in hex.
cqup	The client request URL path; specifies the argument portion of the URL (everything after the host). For example, if the URL is http://www.company.com/ images/x.gif, this field displays /images/x.gif.
cqus	The client request URL scheme (HTTP, FTP, etc.).
crc	The cache result code; specifies how the cache responded to the request (HIT, MISS, and so on).

% <field symbol=""></field>	Description
pfsc	The proxy finish status code; specifies whether the Content Gateway request to the origin server was successfully completed (FIN) or interrupted (INTR).
phn	The host name of the Content Gateway server that generated the log entry in collated log files.
phr	The proxy hierarchy route; the route that Content Gateway used to retrieve the object.
pqbl	The proxy request transfer length; the body length in the Content Gateway request to the origin server.
pqhl	The proxy request header length; the header length in the Content Gateway request to the origin server.
pqsi	The proxy request server IP address (0 on cache hits and parent-ip for requests to parent proxies).
pqsn	The proxy request server name; the name of the server that fulfilled the request.
pscl	The proxy response transfer length; the length of the Content Gateway response to the client in bytes.
psct	The proxy response content type; content type of the document (for example, img/gif) from server response header.
pshl	The proxy response header length; the header length in the Content Gateway response to the client.
psql	The proxy response transfer length in Squid format (includes header and content length).
pssc	The proxy response status code; the HTTP response status code from Content Gateway to the client.
shi	The IP address resolved from the DNS name lookup of the host in the request. For hosts with multiple IP addresses, this field records the IP address resolved from that particular DNS lookup. This can be misleading for cached documents.
	For example, if the first request was a cache miss and came from IP1 for server S and the second request for server S resolved to IP2 but came from the cache, the log entry for the second request will show IP2.
shn	The host name of the origin server.
sscl	The server response transfer length; response length, in bytes, from origin server to Content Gateway.
sshl	The server response header length; the header length in the origin server's response to Content Gateway in bytes.
sshv	The server response HTTP version (1.0, 1.1, and so on).

% <field symbol=""></field>	Description
SSSC	The server response status code; the HTTP response status code from origin server to Content Gateway.
ttms	The time Content Gateway spends processing the client request; the number of milliseconds between the time that the client establishes the connection with Content Gateway and the time that Content Gateway sends the last byte of the response back to the client.
ttmsf	The time Content Gateway spends processing the client request as a fractional number of seconds; specifies the time in millisecond resolution, but instead of formatting the output as an integer (as with ttms), the display is formatted as a floating-point number representing a fractional number of seconds. For example, if the time is 1500 milliseconds, this field displays 1.5 while the ttms field displays 1500 and the tts field displays 1.
tts	The time Content Gateway spends processing the client request; the number of seconds between the time that the client establishes the connection with the proxy and the time that the proxy sends the last byte of the response back to the client.
wc	The predefined or custom category of the URL for the data being scanned. For example, "News and Media".
wct	The content type of the web page. For example, "text/ html; charset=UTF-8".
wsds	The scan disposition string. For example: CATEGORY_BLOCKED, PERMIT_ALL, FILTERED_AND_PASSED.
wsr	The scan recommended bit ("true" or "false"). The URL database identifies and recommends data that should be analyzed further. Depending on the policy used, the data may or may not be analyzed further.
wstms	The scan time in milliseconds that it took to scan a downloaded file or page.
wui	The authenticated user's ID used to select the policy for scanning data of the client request.

Logging format cross-reference

The following sections illustrate the correspondence between Content Gateway logging fields and standard logging fields for the Squid and Netscape formats.

Squid logging formats

Squid	Content Gateway
time	cqts
elapsed	ttms
client	chi
action/code	crc/pssc
size	psql
method	cqhm
url	cquc
ident	caun
hierarchy/from	phr/pqsn
content	psct

For example, if you want to create a custom format called **short_sq** based on the first three Squid fields, enter a line in the **logs.config** file as follows:

format:enabled:1:short_sq:%<cqts> %<ttms>

%<chi>:short_sq:ASCII:none

See Custom format for more information about defining custom log files.

Related concepts

Custom format on page 222

Netscape Common logging formats

Netscape Common	Content Gateway
host	chi
usr	caun
[time]	[cqtn]
"req"	"cqtx"
s1	pssc
c1	pscl

Netscape Extended logging formats

Netscape Extended	Content Gateway
host	chi
usr	caun
[time]	[cqtn]
"req"	"cqtx"
s1	pssc
c1	pscl
s2	SSSC
c2	sscl
b1	cqbl
b2	pqbl
h1	cqhl
h2	pshl
h3	pqhl
h4	sshl
xt	tts

Netscape Extended-2 logging formats

Netscape Extended-2	Content Gateway
host	chi
usr	caun
[time]	[cqtn]
"req"	"cqtx"
s1	pssc
c1	pscl
s2	SSSC
c2	sscl
b1	cqbl
b2	pqbl
h1	cqhl
h2	pshl
Netscape Extended-2	Content Gateway
---------------------	-----------------
h3	pqhl
h4	sshl
xt	tts
route	phr
pfs	cfsc
SS	pfsc
crc	crc

Appendix E Content Gateway Configuration Files

Contents

- Specifying URL regular expressions (url_regex) on page 364
- auth_domains.config on page 365
- auth_rules.config on page 368
- bypass.config on page 370
- filter.config on page 372
- ip_allow.config on page 375
- ipnat.conf on page 376
- log_hosts.config on page 376
- logs_xml.config on page 377
- mgmt_allow.config on page 384
- parent.config on page 384
- records.config on page 387
- remap.config on page 460
- socks.config on page 461
- socks_server.config on page 463
- splitdns.config on page 463
- wccp.config on page 465

Content Gateway contains the following configuration files that you can edit to customize the proxy.

Related concepts

auth_domains.config on page 365 auth_rules.config on page 368 bypass.config on page 370 filter.config on page 372 ip_allow.config on page 375 ipnat.conf on page 376 log_hosts.config on page 376 logs_xml.config on page 377 mgmt_allow.config on page 384 parent.config on page 384 records.config on page 387 socks.config on page 461 socks_server.config on page 463 splitdns.config on page 463

Related reference remap.config on page 460

Specifying URL regular expressions (url_regex)

Entries of type **url_regex** within the configuration files use regular expressions to perform a match.

The following table offers examples to illustrate how to create a valid **url_regex**.

Value	Description
x	Matches the character x.
	Match any character.
^	Specifies beginning of line.
\$	Specifies end of line.
[xyz]	A <i>character class</i> . In this case, the pattern matches either x, y, or z.
[abj-oZ]	A <i>character class</i> with a range. This pattern matches a, b, any letter from j through o, or Z.
[^A-Z]	A <i>negated character class</i> . For example, this pattern matches any character except those in the class.
r*	Zero or more r's, where r is any regular expression.
r+	One or more r's, where r is any regular expression.

Value	Description
r?	Zero or one r, where r is any regular expression.
r{2,5}	From two to five r's, where r is any regular expression.
r{2,}	Two or more r's, where r is any regular expression.
r{4}	Exactly 4 r's, where r is any regular expression.
"[xyz]\"images"	The literal string [xyz]"images"
١X	If X is a, b, f, n, r, t, or v, then the ANSI-C interpretation of \x; Otherwise, a literal X. This is used to escape operators such as *.
\0	A NULL character.
\123	The character with octal value 123.
\x2a	The character with hexadecimal value 2a.
(r)	Matches an r; where r is any regular expression. You can use parentheses to override precedence.
rs	The regular expression r, followed by the regular expression s.
rls	Either an r or an s.
# <n>#</n>	Inserts an <i>end</i> node causing regular expression matching to stop when reached. The value n is returned.

Examples

To match any host in mydomain.com, specify:

dest_domain=mydomain.com

Likewise, to match any request, you can specify:

dest_domain=.

auth_domains.config

The **auth_domains.config** file stores the list of domains that have been identified for use with *Rule-Based Authentication*.

Domains must be identified (added to this file) using the interface in the Content Gateway manager on the **Configure > Security > Access Control > Domains** tab. Do not edit this configuration file.

Related concepts Rule-Based Authentication on page 186

Format

Domains must be identified (added to this file) using the interface in the Content Gateway manager on the **Configure > Security > Access Control > Domains** tab. Do not edit this configuration file.

Each line in **auth_domains.config** consists of a set of tags; each tag is followed by its value. For example:

type=<auth_method> name=<unique_name> use_alias=<0 or 1> <additional tags> The set of tags varies depending on the selected authentication method.

The following table lists all of the tags.

Тад	Allowed value
type	Specifies the authentication method: IWA, NTLM, LDAP
name	Specifies a unique name for the domain. This is not the actual domain name, but rather a name that is unique to the proxy and rule-based authentication.
use_alias	 Specifies the user name sent to filtering service if authentication is successful. 0 = send actual authenticated user name (default).
	 1 = send a blank username 2 = send the string specified in auth_name_string
alias	Only active if use_alias=2. Specifies the static string to send as the user name for all successful authentications using this rule.

The following table lists the additional tags used with IWA domains.

IWA Tag	Allowed Value
winauth_realm	Specifies the joined Windows domain to use with the rule. Content Gateway must be joined and active in that domain.

The following table lists the additional tags used with NTLM domains.

NTLM Tag	Allowed Value
dc_list	Takes the IP address and port number of the primary domain controller (if no port is specified, Content Gateway uses port 139), followed by a comma separated list of secondary domain controllers to be used for load balancing and failover.

NTLM Tag	Allowed Value
NTLM Tag dc_load_balance (optional)	Allowed Value Specifies whether load balancing is used: 0 = disabled 1 = enabled Note When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term failover provision, until such time that the primary
	domain controller can accept new connections.

The following table lists the additional tags used with LDAP domains.

LDAP Tag	Allowed Value
server_name	Specifies the fully qualified domain name of the LDAP server.
server_port (optional)	Specifies the LDAP server port. The default is 389.
	To use the default Global Catalog server port, specify port 3268.
	If Secure LDAP is enabled, set the port to 636 or 3269 (the secure LDAP ports).
base_dn (optional)	Specifies the LDAP base distinguished name.
uid_filter (optional)	Specifies the type of service, if different from that configured on the LDAP tab. Enter sAMAccountName (MS AD) or userPrincipalName (MS AD) for Active Directory, or uid for any other service.
bind_dn (optional)	Specifies the bind distinguished name. This must be a Full Distinguished Name of a user in the LDAP directory service. For example:
	CN=John Smith,CN=USERS,DC=MYCOMPANY, DC=COM
bind_pwd (optional)	Specifies the password for the bind distinguished name.
sec_bind	Specifies whether Content Gateway will use secure communication with the LDAP server.
	• 0 = disabled
	1 = enabled
	If enabled, set the LDAP port to 636 or 3269 (secure LDAP ports).

auth_rules.config

The **auth_rules.config** file stores rules that direct specified IP addresses and IP address ranges, and/or traffic on specified inbound ports (explicit proxy only), and/or matching Request header User-Agent values to authenticate with distinct domain controllers. One or more domain controllers can be specified in an ordered list. This feature is called *Rule-Based Authentication*.

Rule-based authentication rules must be defined in the Content Gateway manager on the **Configure > Security > Access Control > Authentication Rules** tab. Do not edit this configuration file.

- Rule-based authentication is supported for Integrated Windows Authentication (IWA), legacy NTLM, and LDAP authentication only.
- Each authentication rule can specify source IP addresses, inbound port (explicit proxy only), and/or a User-Agent regex
- Each authentication rule can specify one or more domains in an ordered list. Domains are identified on the Configure > Security > Access Control > Authentication Rules tab. That process includes specifying the authentication method (IWA, Legacy NTLM, LDAP).
- When a rule matches, authentication is performed against one or more domains in the ordered list. The first successful authentication ends domain list traversal and the authenticating domain is cached for later use.
- Authentication rules are applied from the list top-down; only the first match is applied. If no rule matches, no user authentication is performed.



Note

If all the users in your network can be authenticated by domain controllers that share trust relationships, you probably don't need rule-based authentication.

However, rule-based authentication can be useful in any deployment that needs to perform special authentication handling based on IP address, inbound proxy port (explicit proxy), and/or User-Agent values.

Related concepts

Rule-Based Authentication on page 186

Format

Each line in **auth_rules.config** contains an authentication rule that consists of a set of tags, each followed by its value. Authentication rules have the format:

rule_name=<name> src_ip=<IP addresses> user_agent=<regex> <additional tags>

The following table lists all of the tags.

Tags	Allowed value
rule_name	A short, unique name.
enabled	Specifies whether the rule will be active:
	• 0 = disabled
	1 = enabled

Tags	Allowed value
src_ip	Takes a comma separated list of IP addresses and IP address ranges. No spaces. If this field is empty, all IP addresses match. The list can contain up to:
	64 IPv4 addresses
	 32 IPv4 address ranges
	 24 IPv6 addresses
	12 IPv6 address ranges
user_agent (optional)	Takes a regular expression that is applied to the user- agent string. See <i>Specifying URL regular expressions</i> <i>(url_regex)</i> for information about using regular expressions.
proxy_port (optional)	Takes a port number. Valid with explicit proxy only. Client applications must be configured to send requests to the correct port.
domain_list	An ordered, comma separated list of domains the Content Gateway will attempt to authenticate a matching user with.
use_captive_portal	Specifies whether Captive Portal is used.
	• 0 = disabled
	1 = enabled using HTTP
	2 = enabled using HTTPS
use_clientcert_auth	Specifies whether Client Certificate Authentication is used.
	0 = disabled
	■ 1 = enabled
clientcert_profile	Takes a text string. The name of the Client Certificate Authentication profile to be used with the authentication rule.
clientcert_fallback	Specifies whether the next selected authentication method should be used if Client Certificate Authentication fails.
	• 0 = disabled
	1 = enabled

Related reference

Specifying URL regular expressions (url_regex) on page 364

bypass.config

The **bypass.config** file contains *static* bypass rules that Content Gateway uses in transparent proxy mode. Static bypass rules instruct Content Gateway to bypass certain incoming client requests so that they are served by the origin server.

The bypass.config file also accepts dynamic deny bypass rules. See Dynamic deny bypass rules.

You can configure three types of static bypass rules:

- Source bypass rules configure the proxy to bypass a particular source IP address or range of IP addresses.
- Destination bypass rules configure the proxy to bypass a particular destination IP address or range of IP addresses. For example, you can bypass origin servers that use IP authentication based on the client's real IP address.
- Source/destination pair bypass rules configure the proxy to bypass requests that originate from the specified source to the specified destination. For example, you can route around specific client-server pairs that experience broken IP authentication. Source/ destination bypass rules can be preferable to destination rules because they block a destination server only for users that experience problems.

Related concepts

Dynamic deny bypass rules on page 371

Format

Bypass rules have the following format:

bypass src ipaddress dst ipaddress src ipaddress AND dst ipaddress	
Option	Description
src ipaddress	Specifies the source (client) IP address in incoming requests that the proxy must bypass.
	<i>ipaddress</i> can be one of the following: A simple IP address, such as 123.45.67.8
	 In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24
	 A range separated by a dash, such as 1.1.1.1-2.2.2.2
	 Any combination of the above, separated by commas, such as 1.1.1.0/24, 25.25.25.25, 123.1.23.1-123.1.23.123

Option	Description
dst <i>ipaddress</i>	Specifies the destination (origin server) IP address in incoming requests that the proxy must bypass.
	<i>ipaddress</i> can be one of the following: A simple IP address, such as 123.45.67.8
	 In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24
	 A range separated by a dash, such as 1.1.1.1-2.2.2.2
	 Any combination of the above, separated by commas, such as 1.1.1.0/24, 25.25.25.25, 123.1.23.1-123.1.23.123
src ipaddress AND dst ipaddress	Specifies the source and destination IP address pair that the proxy must bypass.
	<i>ipaddress</i> can be a single IP address, an IP address range, or a combination of both separated by commas

Dynamic deny bypass rules

In addition to static bypass rules, the bypass.config file also accepts dynamic deny

bypass rules.

Deny bypass rules prevent the proxy from bypassing certain incoming client requests dynamically (a deny bypass rule can prevent the proxy from bypassing itself).

Dynamic deny bypass rules can be source, destination, or source/destination and have the following format:

deny_dyn_bypass src ipaddress | dst ipaddress | src ipaddress AND dst ipaddress

For a description of the options, see the table in *Format*.

For the dynamic deny bypass rules to work, you must either:

- Enable the **Dynamic Bypass** option in the Content Gateway manager.
- Set proxy.config.arm.bypass_dynamic_enabled to 1 in the records.config file.



Important

Static bypass rules overwrite dynamic deny bypass rules. Therefore, if a static bypass rule and a dynamic bypass rule contain the same IP address, the dynamic deny bypass rule is ignored.

Related reference Format on page 370

Examples

The following example shows source, destination, and source/destination bypass rules:

bypass src 1.1.1.0/24, 25.25.25, 128.252.11.11-128.252.11.255 bypass dst 24.24.24.0/24 bypass src 25.25.25.25 AND dst 24.24.24.0

The following example shows source, destination, and source/destination dynamic deny bypass rules:

```
deny_dyn_bypass src 128.252.11.11-128.252.11.255
deny_dyn_bypass dst 111.111.11.1
deny_dyn_bypass src 111.11.11.1 AND dst 111.11.1
```

filter.config

Filtering rules stored in filter.config allow you to:

- Deny or allow URL requests
- Keep or strip header information from client requests
- Insert custom headers
- Allow specified applications or requests to specified web sites to bypass authentication
- Prevent specified applications from transiting the proxy

Filtering rules should be defined in the Content Gateway manager on the **Configure > Security > Access Control > Filtering** tab. See *Creating filtering rules*.



Important

After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Related tasks

Creating filtering rules on page 156

Format

Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 file sandbox destinations.

Each line in **filter.config** is a filtering rule. Content Gateway applies the rules in the order listed, starting at the top of the file. If no rule matches, the request is allowed to proceed.

Content Gateway recognizes three space-delimited tags:

primary_destination=value secondary_specifier=value action=value

The following table lists the possible primary destination types.

Primary Destination Type	Allowed Value
dest_domain	A requested domain name
dest_host	A requested hostname
dest_ip	A requested IP address
url_regex	A regular expression to be found in a URL. See Specifying URL regular expressions (url_regex) for information about using regular expressions.

Secondary specifiers are optional. The following table lists the possible secondary specifiers and their purpose.



Note

You can use more than one secondary specifier in a rule. However, you cannot repeat a secondary specifier.

Secondary Specifier	Allowed Value
time	A time range, such as 08:00-14:00
prefix	A prefix in the path part of a URL
suffix	A file suffix in the URL
src_ip	A single client IP address, or a client IP address range.
port	A requested URL port
method	A request URL method; one of the following:
	 get post put trace
scheme	A request URL protocol. You can specify one of the following: HTTP HTTPS FTP (for FTP over HTTP only)
user_agent	A request header User-Agent value. Takes a regular expression that is applied to the user- agent string. See <i>Specifying URL regular expressions</i> <i>(url_regex)</i> for information about using regular expressions.

The following table lists the possible actions and their allowed values.

Action	Allowed Value
action	Specify one of the following:
	 allow - to allow particular URL requests to bypass authentication. The proxy serves the requested content.
	 deny - to deny requests for HTTP or FTP objects from specific destinations. When a request is denied, the client receives an access denied message.
	 radius - not supported.
keep_hdr	The client request header information that you want to keep. You can specify the following options:
	 date
	host
	cookie
	<pre>client_ip</pre>
strip_hdr	The client request header information that you want to strip. You can specify the same options as with keep_hdr .
add_hdr	The custom header value you want to add. Requires specification of the custom header and a header value. For example:
	add_hdr=" <i>header_name</i> : <i>header_value</i> "

Related reference

Specifying URL regular expressions (url_regex) on page 364

Examples

The following example configures Content Gateway to deny all FTP document requests to the IP address 112.12.12.12:

dest_ip=112.12.12.12 scheme=ftp action=deny

The following example configures Content Gateway to keep the client IP address header for URL requests that contain the regular expression politics and whose path prefix is

/viewpoint:

url_regex=politics prefix=/viewpoint keep_hdr=client_ip

The following example configures Content Gateway to strip all cookies from client requests destined for the origin server www.server1.com:

dest_host=www.server1.com strip_hdr=cookie

The following example configures Content Gateway to disallow **puts** to the origin server www.server2.com:

dest_host=www.server2.com method=put action=deny

Content Gateway applies the rules in the order listed in the file. For example, the following sample **filter.config** file configures Content Gateway to do the following:

- Allow all users (except those trying to access internal.com) to access server1.com
- Deny all users access to notthatsite.com dest_host=server1.com action=allow

dest_host=notthatsite.com action=deny

ip_allow.config

The **ip_allow.config** file controls client access to the proxy. You can specify ranges of IP addresses that are allowed to use Content Gateway.



Note

After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Format

Each line in the **ip_allow.config** file must have the following format:

src_ip=ipaddress action=ip_allow | ip_deny

where *ipaddress* is the IP address or range of IP addresses of the clients allowed to access the proxy.

The action ip_allow allows the specified clients to access the proxy. The action ip_deny denies the specified clients to access the proxy.

By default, the **ip_allow.config** file contains the following line, which allows all clients to access the proxy. Comment out or delete this line before adding rules to restrict access.

src_ip=0.0.0.0-255.255.255.255 action=ip_allow

Examples

The following example allows all clients to access the proxy:

src_ip=0.0.0.0-255.255.255.255 action=ip_allow

The following example allows all clients on a specific subnet to access the proxy:

src_ip=123.12.3.000-123.12.3.123 action=ip_allow

The following example denies all clients on a specific subnet to access the proxy:

src_ip=123.45.6.0-123.45.6.123 action=ip_deny

ipnat.conf

The **ipnat.conf** file contains redirection rules that specify how incoming packets are readdressed when the proxy is serving traffic transparently. Content Gateway creates the redirection rules during installation. You can modify these rules.



Important

After you modify this file, you must restart the proxy.

Format

Each line in the ipnat.conf file must have the following format:

```
rdr interface 0.0.0.0/0 port dest -> ipaddress port proxy tcp|udp
```

where:

interface is the Ethernet interface that traffic will use to access the Content Gateway machine (for example, eth0 on Linux).

dest is the traffic destination port (for example, 80 for HTTP traffic).

ipaddress is the IP address of your Content Gateway server.

proxy is the Content Gateway proxy port (usually 8080 for HTTP traffic).

Examples

The following example configures the ARM to redirect all incoming HTTP traffic to the Content Gateway IP address (111.111.11.1) on the Content Gateway proxy port 8080:

rdr hme0 0.0.0.0/0 port 80 -> 111.111.11.1 port 8080 tcp

log_hosts.config

To record HTTP/FTP transactions for different origin servers in separate log files, you must list each origin server's hostname in the **log_hosts.config** file. In addition, you must enable the HTTP host splitting option (see *HTTP host log splitting*).



Note

It is recommended that you use the same **log_hosts.config** file on every Content Gateway node in your cluster.



Important

After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Related concepts

HTTP host log splitting on page 230

Format

Each line in the log_hosts.config file has the following format:

hostname

where *hostname* is the hostname of the origin server.



You can specify keywords in the **log_hosts.config** file to record all transactions from origin servers with the specified keyword in their names in a separate log file. See the example below.

Examples

Note

The following example configures Content Gateway to create separate log files containing all HTTP/ FTP transactions for the origin servers webserver1, webserver2, and webserver3.

webserver1

webserver2

webserver3

The following example records all HTTP and FTP transactions from origin servers that contain sports in their names (for example, sports.yahoo.com and www.foxsports. com) in a log file called **squid-sport.log** (the Squid format is enabled):

sports

logs_xml.config

The **logs_xml.config** file defines the custom log file formats, filters, and processing options. The format of this file is modeled after XML, the Extensible Markup Language.

Format

The logs_xml.config file contains the following specifications:

- LogFormat specifies the fields to be gathered from each protocol event access. See LogFormat.
- LogFilter specifies the filters that are used to include or exclude certain entries being logged based on the value of a field within that entry. See LogFilter.
- LogObject specifies an object that contains a particular format, a local filename, filters, and collation servers. See LogObject.



Note

The logs_xml.config file ignores extra white space, blank lines, and all comments.

Related reference	
LogFormat on page 378	
LogFilter on page 379	
LogObject on page 380	

LogFormat

The following table lists the LogFormat specifications.

Field	Allowed Inputs
<name "valid_format_name"="" ==""></name>	Required. Valid format names include any name except squid, common, extended, or extended2, which are predefined formats. There is no default for this tag.
<format "valid_format_specification"="" ==""></format>	Required. A valid format specification is a printf-style string describing each log entry when formatted for ASCII output. Use "% <field>" as placeholders for valid field names. For more information, see <i>Custom logging</i> <i>fields</i>. The specified field can be of two types:</field>
	Simple: for example, % <cqu></cqu>
	A field within a container, such as an HTTP header or a Content Gateway statistic. Fields of this type have the following syntax:
	<{field}container>

Field	Allowed Inputs
<interval "aggregate_interval_secs"="" ==""></interval>	Use this tag when the format contains aggregate operators. The value "aggregate_interval_secs" represents the number of seconds between individual aggregate values being produced. The valid set of aggregate operators are:
	 COUNT
	SUM
	AVG
	■ FIRST
	LAST

Related reference

Custom logging fields on page 355

LogFilter

The following table lists the LogFilter specifications.

Field	Allowed Inputs
<name "valid_filter_name"="" ==""></name>	Required. All filters must be uniquely named.

Field	Allowed Inputs
<condition "valid_log_field="" =="" operator<br="" valid_="">valid_comparison_value"/></condition>	Required. This field contains the following elements:
	 valid_log_field is the field that will be compared against the given value. For more information, see Logging format cross-reference.
	 valid_operator_field is any one of the following: MATCH, CASE_INSENSITIVE_MATCH, CONTAIN, CASE_INSENSITIVE_CONTAIN. MATCH is true if the field and value are identical (case sensitive). CASE_INSENSITIVE_MATCH is similar to MATCH, only case insensitive.
	CONTAIN is true if the field contains the value (the value is a substring of the field). CASE_INSENSITIVE_
	CONTAIN is a case-insensitive version of CONTAIN.
	 valid_comparison_value - any string or integer matching the field type. For integer values, all of the operators are equivalent and mean that the field must be equal to the specified value.
	Note
	There are no negative comparison operators. If you want to specify a negative condition, use the Action field to REJECT the record.
<action "valid_action_field"="" ==""></action>	Required. ACCEPT or REJECT. This instructs Content Gateway to either accept or reject records satisfying the condition of the filter.

Related concepts

Logging format cross-reference on page 358

LogObject

The following table lists the **LogObject** specifications.

Field	Allowed Inputs
<format "valid_format_name"="" ==""></format>	Required. Valid format names include the predefined logging formats: squid, common, extended, and extended2, as well as any previously- defined custom log formats. There is no default for this tag.

Field	Allowed Inputs
<filename "file_name"="" ==""></filename>	Required. The filename to which the given log file is written. No local log file will be created if you fail to specify this tag. All filenames are relative to the default logging directory.
	If the name does not contain an extension (for example, "squid"), an extension is added: .log for ASCII logs or .blog for binary logs. (See <mode =<br="">"valid_logging_mode"/> below.) If you do not want an extension to be added, end the filename with a single dot (.).</mode>
<mode "valid_logging_mode"="" ==""></mode>	Valid logging modes include ascii, binary, and ascii_pipe. The default is ascii.
	 Use ascii to create event log files in human- readable form (plain ASCII).
	Use binary to create event log files in binary format. Binary log files generate lower system overhead and occupy less space on the disk (depending on the information being logged). You must use the logcat utility to translate binary log files to ASCII format before you can read them.
	Use ascii_pipe to write log entries to a UNIX named pipe (a buffer in memory). Other processes can then read the data using standard I/O functions. Content Gateway does not have to write to disk, freeing disk space and bandwidth for other tasks. In addition, writing to a pipe does not stop when logging space is exhausted because the pipe does not use disk space.
	Note
	If you are using a collation server, the log is written to a pipe on the collation server. A local pipe is created even before a transaction is processed so that you can see the pipe right after Content Gateway starts. However, pipes on a collation server <i>are</i> created when Content Gateway starts.
<filters "list_of_valid_filter_="" =="" names"=""></filters>	A comma-separated list of names of any previously defined log filters. If more than one filter is specified, all filters must accept a record for the record to be logged.
<protocols "list_of_valid_="" =="" protocols"=""></protocols>	A comma-separated list of the protocols this object should log. Valid protocol names include HTTP.
<serverhosts "list_of_valid_="" =="" servers"=""></serverhosts>	A comma-separated list of valid hostnames. This tag indicates that only entries from the named servers will be included in the file.

Field	Allowed Inputs
<collationhosts "list_of_valid_="" =="" hostnames"=""></collationhosts>	A comma-separated list of collation servers to which all log entries (for this object) are forwarded. Collation servers can be specified by name or IP address. Specify the collation port with a colon after the name (for example, host:port).
<header "header"="" ==""></header>	The header text you want the log files to contain. The header text appears at the beginning of the log file, just before the first record.
<rollingenabled "truth="" =="" value"=""></rollingenabled>	Enables or disables log file rolling for the LogObject. This setting overrides the value for the configuration setting Log Rolling: Enabled/ Disabled in the Content Gateway manager or proxy.config.log2.rolling_enabled in the records. config file.
	Set "truth value" to 1 or true to enable rolling; set it to 0 or false to disable rolling for this particular LogObject.
<rollingintervalsec "seconds"="" ==""></rollingintervalsec>	Specifies the seconds between log file rolling for the LogObject. This setting overrides the value for the configuration setting Log Rolling: Interval in the Content Gateway manager or proxy.config. log2.rolling_interval_sec in the records.config file. This option allows you to specify different rolling intervals for different LogObjects.
<rollingoffsethr "hour"="" ==""></rollingoffsethr>	Specifies an hour (from 0 to 23) at which rolling is guaranteed to align. Rolling may start before then, but a rolled file will be produced only at that time. The impact of this setting is only noticeable if the rolling interval is larger than one hour. This setting overrides the configuration setting Log Rolling: Offset Hour in the Content Gateway manager or proxy.config.log2.rolling_offset_hr in the records. config file.

Examples

The following is an example of a LogFormat specification collecting information using three common fields:

```
<LogFormat>
<Name = "minimal"/>
<Format = "%<chi> : %<cqu> : %<pssc>"/>
</LogFormat>
```

The following is an example of a LogFormat specification using aggregate operators:

```
<LogFormat>
<Name = "summary"/>
<Format = "%<LAST(cqts)> : %<COUNT(*)> : %<SUM(psql)>"/>
<Interval = "10"/>
</LogFormat>
```

The following is an example of a LogFilter that will cause only REFRESH_HIT entries to be logged:

```
<LogFilter>
<Name = "only_refresh_hits"/>
<Action = "ACCEPT"/>
<Condition = "%<pssc> MATCH REFRESH_HIT"/>
</LogFilter>
```



Note

When specifying the field in the filter condition, you can omit the %<>. This means that the following filter is equivalent to the example directly above:

```
<LogFilter>
<Name = "only_refresh_hits"/>
<Action = "ACCEPT"/>
<Condition = "pssc MATCH REFRESH_HIT"/>
</LogFilter>
```

The following is an example of a LogObject specification that creates a local log file for the minimal format defined earlier. The log filename will be minimal.log because this is an ASCII log file (the default).

```
<LogObject>
<Format = "minimal"/>
<Filename = "minimal"/>
</LogObject>
```

The following is an example of a LogObject specification that includes only HTTP requests served by hosts in the domain company.com or by the specific server server. somewhere.com. Log entries are sent to collation host logs.company.com on port 4000 and to collation host 209.131.52.129 on port 5000.

```
<LogObject>
<Format = "minimal"/>
<Filename = "minimal"/>
<ServerHosts = "company.com,server.somewhere.com"/>
<Protocols = "http"/>
<CollationHosts = "logs.company.com:4000,209.131.52.
129:5000"/>
</LogObject>
```

WebTrends Enhanced Log Format (WELF)

Content Gateway supports WELF so that you can analyze Content Gateway log files with WebTrends reporting tools. A predefined <LogFormat> that is compatible with WELF is provided at the end of the logs.config file (shown below). To create a WELF format log file, create a <LogObject> that uses this predefined format.

```
<LogFormat>
<Name = "welf"/>
<Format = "id=firewall time=\"%<cqtd> %<cqtt>\" fw=%<phn>
pri=6 proto=%<cqus> duration=%<ttmsf> sent=%<psql>
rcvd=%<cqhl> src=%<chi> dst=%<shi> dstname=%<shn>
user=%<caun> op=%<cqhm> arg=\"%<cqup>\" result=%<pssc>
ref=\"%<{Referer}cqh>\" agent=\"%<{user-agent}cqh>\"
cache=%<crc>"/>
</LogFormat>
```

mgmt_allow.config

The **mgmt_allow.config** file specifies the IP addresses of remote hosts allowed access or denied access to the Content Gateway manager.



Important

After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Format

Each line in the mgmt_allow.config file has the following format:

src_ip=<ipaddress> action=<ip_allow|ip_deny>

Here, <ipaddress> is the IP address or range of IP addresses allowed to access the Content Gateway manager.

Use "action" to specify either ip_allow (grant access to the Content Gateway manager) or ip_deny (block access).

By default, the **mgmt_allow.config** file contains the following line, which allows all remote hosts to access the Content Gateway manager. Comment out or delete this line before adding rules to restrict access.

src_ip=0.0.0.0-255.255.255.255 action=ip_allow

Examples

The following example configures Content Gateway to allow only one user to access the Content Gateway manager:

src_ip=123.12.3.123 action=ip_allow

The following example configures Content Gateway to allow a range of IP addresses to access the Content Gateway manager:

src_ip=123.12.3.000-123.12.3.123 action=ip_allow

The following example configures Content Gateway to deny one IP address access to the Content Gateway manager:

src_ip=123.45.67.8 action=ip_deny

parent.config

The **parent.config** file identifies the HTTP parent proxies used in an HTTP cache hierarchy. Use this file to perform the following configuration:

Set up parent cache hierarchies, with multiple parents and parent failover

Configure selected URL requests to bypass parent proxies

Rules are applied from the list top-down; the first match is applied. Bypass rules are usually placed above parent proxy designation rules.

Content Gateway uses the **parent.config** file only when the HTTP parent caching option is enabled. See *Configuring Content Gateway to use an HTTP parent cache*.



After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Related tasks

Important

Configuring Content Gateway to use an HTTP parent cache on page 80

Format

Each line in the **parent.config** file must contain a parent caching rule. Content Gateway recognizes three spacedelimited tags:

primary_destination=value secondary_specifier=value action=value

The following table lists the possible primary destinations and their allowed values.

Primary Destination	Allowed Value
dest_domain	A requested domain name
dest_host	A requested hostname
dest_ip	A requested IP address or range of IP addresses separated by a dash (-).
url_regex	A regular expression to be found in a URL. See Specifying URL regular expressions (url_regex) for information about using regular expressions.

Secondary specifiers are optional in the parent.config file. The following table lists the possible secondary specifiers and their allowed values.

Secondary Specifier	Allowed Value
time	A time range, such as 08:00-14:00, during which the parent cache is used to serve requests
prefix	A prefix in the path part of a URL
suffix	A file suffix in the URL
src_ip	A client IP address.
port	A requested URL port

Secondary Specifier	Allowed Value
scheme	A request URL protocol; one of the following:
	HTTP
	■ FTP
method	A request URL method; one of the following:
	■ get
	■ post
	• put
	• trace
user_agent	A request header User-Agent value.
	Takes a regular expression that is applied to the user- agent string. See <i>Specifying URL regular expressions</i> (<i>url_regex</i>) for information about using regular expressions.

The following table lists the possible actions and their allowed values.

Action	Allowed Value
parent	An ordered list of parent servers. If the request cannot be handled by the last parent server in the list, it will be routed to the origin server. You can specify either a hostname or an IP address. You must specify the port number.
round_robin	One of the following values:
	 true - Content Gateway goes through the parent cache list in a round-robin based on client IP address.
	 strict - Content Gateway machines serve requests strictly in turn. For example, machine proxy1 serves the first request, proxy2 serves the second request, and so on.
	false - round-robin selection does not occur.
go_direct	One of the following values:
	 true - requests bypass parent hierarchies and go directly to the origin server.
	 false - requests do not bypass parent hierarchies.

Related reference

Specifying URL regular expressions (url_regex) on page 364

Examples

The following rule configures a parent cache hierarchy consisting of Content Gateway (which is the child) and two parents, p1.x.com and p2.x.com. The proxy forwards the requests it cannot serve to the parent servers p1.x.com and p2.x.com in a round-robin fashion because round_robin=true.

dest_domain=. method=get parent="pl.x.com:8080; p2.y. com:8080" round_robin=true

The following rule configures Content Gateway to route all requests containing the regular expression **politics** and the path /viewpoint directly to the origin server (bypassing any parent hierarchies):

url_regex=politics prefix=/viewpoint go_direct=true

The following rule is a typical destination bypass rule:

dest_domain=example.com go_direct=true



Important

Every line in the parent.config file must contain either a parent= or go_direct= directive.

A bypass rule that includes parent= **and** go_direct=true, causes the specified dest_domain to be sent to the parent while all other domains are bypassed (the opposite of the usual intended action).

records.config

The records.config file is a list of configurable variables used by Content Gateway.

Most values are set using controls in the Content Gateway manager. Some options can be set only by editing variables in the records.config file.



Warning

Do not change the records.config variables unless you are certain of the effect. Many variables are coupled, meaning that they interact with other variables. Changing a single variable in isolation can cause Content Gateway to fail. Whenever possible, use the Content Gateway manager to configure Content Gateway.



Important

After you modify this file, run the following command to apply the changes:

/opt/WCG/bin/content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Format

Each variable has the following format:

CONFIG <variable_name> <DATATYPE> <variable_value>

Here, <DATATYPE> is one of the following:

- INT (an integer)
- STRING (a string)
- FLOAT (a floating point)

Examples

In the following example, the variable **proxy.config.proxy_name** is of datatype **STRING** and its value is **contentserver1**. This means that the name of the Content Gateway proxy is **contentserver1**.

CONFIG proxy.config.proxy_name STRING contentserver1

In the following example, the variable **proxy.config.winauth.enabled** is a yes/no flag. A value of 0 (zero) disables the option. A value of 1 enables the option.

CONFIG proxy.config.winauth.enabled INT 0

In the following example, the variable sets the cluster startup timeout to 10 seconds.

CONFIG proxy.config.cluster.startup_timeout INT 10

Configuration variables

The following list describe the configuration variables listed in the records.config file.

Related reference

System variables on page 389 Virtual IP manager on page 396 ARM on page 397 Authentication basic realm on page 404 **RADIUS** authentication on page 409 Integrated Windows Authentication on page 413 HTTP engine on page 415 FTP engine on page 427 SOCKS processor on page 432 Cluster subsystem on page 433 DNS on page 434 HostDB on page 436 URL remap rules on page 443 WCCP configuration on page 444 ICAP on page 453 Local manager on page 392 Alarm configuration on page 396 Load shedding configuration (ARM) on page 403 LDAP on page 406 NTLM on page 410 Transparent authentication on page 414 Parent proxy configuration on page 418 Anonymous FTP password on page 426 FTP transfer mode on page 426 Customizable user response pages on page 426 Net subsystem on page 433 Cache on page 434 DNS proxy on page 436 Logging configuration on page 437 SSL Decryption on page 445 Connectivity, analysis, and boundary conditions on page 456

System variables

Variable	Data Type	Description
proxy.config.proxy_name	STRING	Default: (none)
		The name of the Content Gateway node.
proxy.config.bin_path	STRING	Default: bin
		The location in which the Content Gateway binary files are placed by the installer.

Variable	Data Type	Description
proxy.config.proxy_binary	STRING	Default: content_gateway
		The name of the executable that runs the content_gateway process.
proxy.config.proxy_binary_ opts	STRING	Default: -M
		The command-line options for starting content_gateway.
proxy.config.manager_binary	STRING	Default: content_manager
		The name of the executable that runs the content_manager process.
proxy.config.cli_binary	STRING	Default: content_line
		The name of the executable that runs the content_line interface.
proxy.config.watch_script	STRING	Default: content_cop
		The name of the executable that runs the content_cop process.
proxy.config.env_prep	STRING	Default: example_prep.sh
		The script that is executed before the content_ manager process spawns the content_ gateway process.
proxy.config.config_dir	STRING	Default: config
		The directory, relative to bin_path (above), that contains the Content Gateway configuration files.
proxy.config.temp_dir	STRING	Default: /tmp
		The directory used for Content Gateway temporary files
proxy.config.alarm_email	STRING	Default: <install user=""></install>
		The email address to which Content Gateway sends alarm messages.
		During installation, you can specify the email address; otherwise, Content Gateway uses the Content Gateway user account name as the default value.
proxy.config.syslog_facility	STRING	Default: LOG_DAEMON
		The facility used to record system log files. See <i>Working With Log Files</i> .

Variable	Data Type	Description	
proxy.config.cop.core_ signal	INT	Default: 3 The signal sent by content_cop to its managed processes – content_manager and content_ gateway – to stop them.	
		Note Do not change the value of this variable.	
proxy.config.cop.sleep_ time	INT	Default: 45 The interval, in seconds, between heartbeat tests performed by content_cop to test the health of the content_manager and content_gateway processes. Image: Content_gateway protent proceses. Image: Content_g	
proxy.config.cop.linux_ min_swapfree_kb	INT	Default: 10240 This variable is not used.	
proxy.config.cop.linux_ min_memfree_kb	INT	Default: 10240 This variable is not used.	
proxy.config.output.logfile	STRING	Default: content_gateway_out The name and location of the file that contains warnings, status messages, and error messages produced by the Content Gateway processes. If no path is specified, Content Gateway creates the file in its logging directory.	
proxy.config.output.logfile. log_dir_usage_percent	INT	Default: 35 The percentage of space allocated by proxy. config.log2.max_space_ mb_for_logs , that can be used for logs in /opt/WCG/logs/ except for content_gateway.out . Content_gateway. out can use up to the log directory limit.	

Variable	Data Type	Description
proxy.config.snapshot_dir	STRING	Default: snapshots
		The directory in which Content Gateway stores configuration snapshots on the local system. Unless you specify an absolute path, this directory is located in the Content Gateway config directory.
proxy.config.attach_	STRING	Default: NULL
debugger_script		This variable should be used only on the direction of Technical Support.
		If set, when the content_gateway process resets, a debug script (in / opt/WCG/bin) is run.
proxy.config.healthcheck_	INT	Default: 0
force_offline		When enabled (1), forces URL health checks to report proxy down. See, <i>Health Check URLs</i> .

Related concepts

Health Check URLs on page 350

Related information

Working With Log Files on page 217

Local manager

Configuration Variable	Data Type	Descript	lion
proxy.config.lm.sem_id	INT	Default: ²	11452
		Specifies local mar	the semaphore ID for the nager.
			Note
			Do not change the value of this variable.
proxy.local.cluster.type	INT	Default: 3	3
		Sets the	clustering mode:
		 2 = m 3 = nc 	anagement-only mode o clustering

Configuration Variable	Data Type	Description
proxy.config.cluster.rsport	INT	Default: 8087
		Specifies the reliable service port. The reliable service port is used to send configuration information between the nodes in a cluster. All nodes in a cluster must use the same reliable service port.
proxy.config.cluster.mcport	INT	Default: 8088
		Specifies the multicast port. The multicast port is used for node identification. All nodes in a cluster must use the same multicast port.
proxy.config.cluster.mc_	STRING	Default: 224.0.1.37
group_addr		Specifies the multicast address for cluster communications. All nodes in a cluster must use the same multicast address.
proxy.config.cluster.mc_ttl	INT	Default: 1
		Specifies the multicast Time-To-Live for cluster communications.
proxy.config.cluster.log_	INT	Default: 1
bogus_mc_msgs		Enables (1) or disables (0) logging of invalid multicast messages.
proxy.config.admin.html_doc_root	STRING	Default: ui
		Specifies the document root for the Content Gateway manager.
proxy.config.admin.web_	INT	Default: 8081
interface_port		Specifies the Content Gateway manager port.
proxy.config.admin.autoconf_port	INT	Default: 8083
		Specifies the autoconfiguration port.
proxy.config.admin.	INT	Default: 5
aulocont_port_timeout		Specifies a timeout, in seconds, for port 8083 when a connection has been established with that port but no data is sent. Valid values are 1 - 300 seconds.

Configuration Variable	Data Type	Description
proxy.config.admin. overseer_port	INT	Default: 0
		retrieving and setting statistics and configuration variables. Java applets, which are not supported on all browsers, are required to display statistics.
		This port is disabled by default.
proxy.config.admin.admin_ user	STRING	Default: admin
		Specifies the administrator ID that controls access to the Content Gateway manager.
proxy.config.admin. heartbeat_	INT	Default: 8079
port		Specifies the port used by the content_cop process to send heartbeat requests to Content Gateway manager for a system health check.
proxy.config.admin.admin_	STRING	Default (none)
password		Specifies the encrypted administrator password that controls access to the Content Gateway manager. You cannot edit the password; however, you can specify a value of NULL to clear the password.
		See Accessing the Content Gateway manager if you forget the master administrator password.
proxy.config.admin.use_ssl	INT	Default: 1
		Enables the Content Gateway manager SSL option for secure communication between a remote host and the Content Gateway manager.
proxy.config.admin.ssl_cert_file	STRING	Default: server.pem
		Specifies the filename of the SSL certificate installed on the Content Gateway system for secure communication between a remote host and the Content Gateway manager.
		Note that the contents of this file must not be password protected.

Configuration Variable	Data Type	Description
proxy.config.admin. number_config_bak	INT	Default: 3 Specifies the maximum number of copies of rolled configuration files to keep.
proxy.config.admin.user_id	STRING	Default: root Specifies the non-privileged user account designated to Content Gateway.
proxy.config.admin.ui_ refresh_rate	INT	Default: 30 Specifies the refresh rate for the display of statistics in the Monitor pages of the Content Gateway manager.
proxy.config.admin.log_ mgmt_access	INT	Default: 0 Enables (1) or disables (0) logging of all Content Gateway manager transactions to the Im.log file.
proxy.config.admin.log_ resolve_hostname	INT	Default: 1 When enabled (1), the hostname of the client connecting to the Content Gateway manager is recorded in the Im.log file. When disabled (0), the IP address of the client connecting to the Content Gateway manager is recorded in the Im.log file.
proxy.config.admin. subscription	STRING	Default: NULL Not used.
proxy.config.admin. supported_cipher_list	STRING	Default: AES128-SHA, DHE-RSA- AES128-SHA, DHE-DSS-AES128- SHA, DES-CBC3-SHA, EDH-RSA- DES-CBC3- SHA, EDH-DSS-DES- CBC3-SHA A comma-separated list, no spaces,
		of ciphers supported by Content Gateway. No validation is performed on the
		string.
proxy.config.lm.display_ reset_alarm	INT	Default: 0 When enabled (1), email is sent to the administrator (proxy.config.alarm_email) whenever Content Gateway resets.

Configuration Variable	Data Type	Description
proxy.local.install.type	INT	Default: 1
		Indicates that Content Gateway is installed as a component of Forcepoint Web Security (1) or Forcepoint DLP without Forcepoint Web Security (2)

Related tasks

Accessing the Content Gateway manager if you forget the master administrator password on page 19

Process manager

Configuration Variable	Data Type	Description
proxy.config.process_manager. mgmt_port	INT	Default: 8084 Specifies the port used for internal communication between the content_manager process and the content_gateway process

Virtual IP manager

Configuration Variable	Data Type	Description
proxy.config.vmap.enabled	INT	Default: 0
		Enables (1) or disables (0) the virtual IP option.

Alarm configuration

Configuration Variable	Data Type	Description
proxy.config.alarm.bin	STRING	Default: example_alarm_bin.sh
		Specifies the name of the script file that can execute certain actions when an alarm is signaled. Edit the example script to suit your needs.
Configuration Variable	Data Type	Description
-----------------------------	-----------	---
proxy.config.alarm.abs_path	STRING	Default: NULL Specifies the full path to
		the script file specified by proxy.config.alarm.bin (prior entry).

ARM

Configuration Variable	Data Type	Description
proxy.config.arm.enabled	INT	Default: 1
		Specifies that the ARM is enabled or disabled.
		Warning: Do not disable the ARM. In all deployments, it must be running to support proper proxy function.
proxy.config.arm.ignore_ifp	INT	Default: 1 When redirection rules are applied, configures Content Gateway to use any available interface when sending packets back to the client, rather than the one that triggered the redirection rule.

Configuration Variable	Data Type	Description
proxy.config.arm.always_ guery_dest	INT	Default: 0
4,		When enabled (1), Content Gateway always asks the ARM for the original destination IP address of incoming requests. This is done instead of doing a DNS lookup on the hostname of the request.
		When enabled, domain names are logged, instead of IP addresses, unless proxy. config.arm.use_hostname_for_
		wisp_and_ reporting (see below) is disabled.
		When disabled, domain names are logged. See <i>Reducing DNS lookups</i> for additional information.
		It is recommended that you do not enable this variable if Content Gateway is running in <i>both</i> explicit proxy and transparent proxy modes. In explicit proxy mode, the client does not perform a DNS lookup on the hostname of the origin server, so Content Gateway must do it.
proxy.config.arm.use_	INT	Default: 1
hostname_for_wisp_and_ reporting		Enables (1) or disables (0) the ability to capture hostname (instead of IP address) when Always Query Destination is enabled for transparent proxy deployments. See preceding entry. Note: This variable must be manually added to the config file.
proxy.config.http.outgoing_ip_	INT	Default: 0
spoofing_enabled		Enables (1) or disables (0) the IP spoofing option that allows Content Gateway to establish connections to origin servers with the client IP address instead of the Content Gateway IP address. See Content Gateway IP spoofing
		cos comon caloway in spooling.

Configuration Variable	Data Type	Descri	ption
proxy.config.arm.bypass_ dynamic_enabled	INT	Default Enable adaptiv the pro origin s servers Dynam	:: 0 s (1) or disables (0) the ve bypass option to bypass xy and go directly to the server when clients or s cause problems. See hicbypass rules.
proxy.config.arm.bypass_use_ and_rules_bad_client_request	INT	Default Enable source, event c 80.	:: 0 s (1) or disables (0) dynamic /destination bypass in the of non-HTTP traffic on port
		E,	Note The variable proxy.config.arm. bypass_on_
		E,	Note bad_client_request must also be enabled for this option to work.
proxy.config.arm.bypass_use_ and_rules_400	INT	Default Enable genera bypass returns	:: 0 s (1) or disables (0) dynamic tion of source/destination rules when an origin server a 400 error.
		F	Note The variable proxy.config.arm. bypass_on_400 must also be enabled for this option to work.

Configuration Variable	Data Type	Description
proxy.config.arm.bypass_use_ and_rules_401	config.arm.bypass_use_ INT De les_401 En ge byj ret	Default: 0 Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 401 error.
		Note The variable proxy.config.arm. bypass_on_401 must also be enabled for this option to work.
proxy.config.arm.bypass_use_ and_rules_403	INT	Default: 0 Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 403 error. Note The variable proxy.config.arm. bypass_on_403 must also be enabled for this option to work.
proxy.config.arm.bypass_use_ and_rules_405	INT	Default: 0 Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 405 error. Note Note The variable proxy.config.arm. bypass_on_405 also be enabled for this option to work.

Configuration Variable	Data Type	Description
proxy.config.arm.bypass_use_ and_rules_406	INT	Default: 0Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 406 error.Image: the state of the server returns a 406 error.Image: the server returns a 406 error.
proxy.config.arm.bypass_use_ and_rules_408	INT	Default: 0 Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 408 error. Note Note The variable proxy.config.arm. bypass_on_408 must also be enabled for this option to work.
proxy.config.arm.bypass_use_ and_rules_500	INT	Default: 0 Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 500 error. Note The variable proxy.config.arm. bypass_on_500 must also be enabled for this option to work.
proxy.config.arm.bypass_on_ bad_client_request	INT	Default: 0 Enables (1) or disables (0) dynamic destination bypass in the event of non- HTTP traffic on port 80.

Configuration Variable	Data Type	Description
proxy.config.arm.bypass_on_ 400	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 400 error.
proxy.config.arm.bypass_on_ 401	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 401 error.
proxy.config.arm.bypass_on_ 403	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 403 error.
proxy.config.arm.bypass_on_ 405	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 405 error.
proxy.config.arm.bypass_on_ 406	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 406 error.
proxy.config.arm.bypass_on_ 408	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 408 error.
proxy.config.arm.bypass_on_ 500	INT	Default: 0 Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 500 error.
proxy.config.arm.forward_ unwanted_traffic	INT	Default: 0 Enables (1) or disables (0) the forwarding of traffic that is neither HTTP, HTTPS, nor FTP.

Related concepts Reducing DNS lookups on page 61 Content Gateway IP spoofing on page 63 Dynamic bypass rules on page 59

Load shedding configuration (ARM)

Configuration Variable	Data Type	Description
proxy.config.arm.loadshedding. max_connections	INT	Default: 1000000 Specifies the maximum number of client connections allowed before the proxy starts forwarding incoming requests directly to the origin server.
proxy.config.http.client. connection_control.enabled	INT	Default: 1 Disables (0) or enables (1) the ability to limit the number of connections from a single computer.
proxy.config.http.client. concurrent_connection_control. close.enabled	INT	Default: 1 Disables (0) or enables (1) closing connections on reaching the concurrent connection limit.
proxy.config.http.client. concurrent_connection_control. alert.enabled	INT	Default: 0 Disables (0) or enables (1) alerting on violation of the concurrent connection limit.
proxy.config.http.client. concurrent_connection_control. max_connections	INT	Default: 1000 Configures the maximum number of concurrent connections allowed from one client IP address.
proxy.config.http.client. connection_rate_control.close. enabled	INT	Default: 0 Disables (0) or enables (1) closing connections on reaching the connection rate limit.
proxy.config.http.client. connection_rate_control.alert. enabled		Default: 1 Disables (0) or enables (1) alerting on exceeding the connection rate limit.

Configuration Variable	Data Type	Description
proxy.config.http.client. connection_rate_control.second	INT	Default: 100 Configures the maximum connections per second allowed from one client IP.
proxy.config.http.client. connection_control.exceptions	STRING	Default: NULL Specifies a comma separated list of IP addresses for which the connection limits do not apply.

Authentication basic realm

Configuration Variable	Data Type	Description
proxy.config.proxy.authenticate. basic.realm	STRING	Default: NULL Specifies the authentication realm name. If the default of NULL is specified, Content Gateway is used.
proxy.config.auth_type	INT	 Default: 0 Specifies the type of client authentication. 0 = None 1 = LDAP 2 = RADIUS 3 = Legacy NTLM 4 = Integrated Window Authentication 5 = Rule-Based Authentication
proxy.config.multiauth.enabled	INT	Default: 0 Enables (1) or disables (0) rule- based authentication. Tells Content Gateway to use the auth_rules.config file.
proxy.config.multiauth.domain. max	INT	Default: 50 Specifies the maximum number of domains that can be added or joined on Configure > Security > Access Control > Domains

Configuration Variable	Data Type	Description
proxy.config.auth.form_ filename	STRING	Default: auth_form.html
		Specifies the file that defines the Captive Portal authentication page.
		This variable must be added manually. Changing this filename is not recommended.
proxy.config.internal.file.path	STRING	Default: /opt/WCG/config/ui_files
		Specifies the location of any css and image files used to define the Captive Portal authentication page. The full default path is /opt/WCG/ config/ui_files. Image files are located in an /images sub- directory. This variable must be added manually.
proxy.config.ssl.auth_server_ port	INT	Default: 4443
		Specifies the local port used for the HTTPS Captive Portal page.
proxy.config.ssl.use_custom_	INT	Default: 0
cert_for_captve_portal		Enables (1) or disables (0) the use of a custom certificate with Captive Portal.
proxy.config.auth.sharecookie	INT	Default:0
		Enables (1) or disables (0) authentication cookie sharing. This is automatically enabled when cookie cashing is enabled.
proxy.config.auth.reauth_for_	INT	Default: 0
null_user		Enabled (1) or disables (0) re- authentication with a NULL user is used with a valid password.
		This variable must be added manually and is only valid for LDAP authentication

Configuration Variable	Data Type	Description
proxy.config.auth_user_ip_sess_	INT	Default: 60
timeout		Specifies the number of seconds a device IP address and a user are associated.
		Content Gateway will associate an IPv4 address with the same user name for this length of time.
		This variable works with proxy.config. auth_user_ip_max_num and must be manually added.
proxy.config.auth_user_ip_	INT	Default: 1
max_num		Specifies the number of devices from which a single user can access the Internet for the number of seconds configured by proxy.config.auth_
		user_ip_sess_timeout.
		This variable must be manually added.
proxy.config.auth.ssl_auth_url	INT	Default: 1
		Disables (0) or enables (1) authentication of HTTPS requests over HTTPS, using port 8443. When disabled, authentication for HTTPS requests is done over HTTP, using port 8080.
proxy.config.auth.sync_auth_ config	INT	Default: 0
		Disables (0) or enables (1) the control of the syncing of auth_rules.config files across the cluster.
		Note that, to avoid invalid rules, the domain identifier needs to be the same across the cluster.

LDAP

Configuration Variable	Data Type	Description
proxy.config.ldap.auth.enabled	INT	Default: 0
		Enables (1) or disables (0) LDAP proxy authentication. See <i>LDAP authentication</i> .

Configuration Variable	Data Type	Description
proxy.config.ldap.cache.size	INT	Default: 5000
		The maximum number of entries allowed in the LDAP cache.
		If this value is modified, you must update the value of proxy.config.ldap.cache. storage_size proportionally. For example, if you double the cache size, also double the cache storage size.
proxy.config.ldap.cache.	INT	Default: 24582912
storage_size		The size of the LDAP cache in bytes. This is directly related to the number of entries in the cache.
		If this value is modified, you must update the value of proxy.config.ldap.cache.size proportionally. For example, if you double the storage size, also double the cache size.
		Modifying this variable without modifying proxy.config.ldap.cache.size can cause the LDAP subsystem to stop functioning.
proxy.config.ldap.auth.ttl_value	INT	Default: 3000
		The amount of time (in minutes) that entries in the cache remain valid.
proxy.config.ldap.auth.purge_	INT	Default: 1
		When enabled (1), configures Content Gateway to delete the authorization entry for the client in the LDAP cache if authorization fails.
proxy.config.ldap.proc.ldap.	STRING	Default: NULL
server.name		The LDAP server name.
proxy.config.ldap.proc.ldap.	INT	Default: 398
		The LDAP server port.
proxy.config.ldap.proc.ldap. base.dn	STRING	Default: NULL
		The LDAP Base Distinguished Name (DN). Obtain this value from your LDAP administrator.

Configuration Variable	Data Type	Description
proxy.config.ldap.proc.ldap.uid_	STRING	Default: sAMAccountName
filter		The LDAP login name/ID. Use this as a filter to search the full DN database.
		userPrincipalName is also valid for Microsoft Active Directory.
		For eDirectory or other directory services, enter uid in this field.
proxy.config.ldap.secure.bind.	INT	Default: 0
enabled		When enabled (1), configures the proxy to use secure LDAP (LDAPS) to communicate with the LDAP server.
		Secure communication is usually performed on port 636 or 3269.
proxy.config.ldap.proc.ldap.	STRING	Default: NULL
server.bind_dn		The Full Distinguished Name (fully qualified name) of a user in the LDAP- based directory service. For example:
		CN=John Smith,CN=USERS, DC=MYCOMPANY,DC=COM
		Enter a maximum of 128 characters in this
		field.
		If no value is specified for this field, the proxy attempts to bind anonymously.
proxy.config.ldap.proc.ldap.	STRING	Default: NULL
server.bind_pwd		Specifies a password for the user identified by the proxy.config.ldap.proc. Idap.server.bind_dn variable.
proxy.config.ldap.proc.encode_	INT	Default: 0
convert		Enables (1) or disables (0) the support of passwords with special characters.
		The variable proxy.config.ldap.proc. encode_name is required when this variable is enabled.
		This variable must be added manually. See this page for additional information.

Configuration Variable	Data Type	Description
proxy.config.ldap.proc.encode_ name	STRING	Default: NULL The encoding name to be used when proxy.config.ldap.proc.encode_ convert is enabled.
		This variable must be added manually.

Related concepts

LDAP authentication on page 180

RADIUS authentication

Configuration Variable	Data Type	Description
proxy.config.radius.auth.enabled	INT	Default: 0 Enables (1) or disables (0) RADIUS proxy authentication.
proxy.config.radius.proc.radius. primary_server.name	STRING	Default: NULL The hostname or IP address of the primary RADIUS authentication server.
proxy.config.radius.proc.radius. primary_server.auth_port	INT	Default: 1812 The RADIUS server port that Content Gateway uses to communicate with the RADIUS server.
proxy.config.radius.proc.radius. primary_server.shared_key	STRING	Default: NULL The key used for encoding with the first RADIUS authentication server.
proxy.config.radius.proc.radius. secondary_server.name	STRING	Default: NULL The hostname or IP address of the secondary RADIUS authentication server.
proxy.config.radius.proc.radius. secondary_server.auth_port	INT	Default: 1812 The port that the proxy uses to communicate with the secondary RADIUS authentication server.

Configuration Variable	Data Type	Description
proxy.config.radius.proc.radius. secondary_server.shared_key	STRING	Default: NULL The key used for encoding with the secondary RADIUS authentication server.
proxy.config.radius.auth.min_ timeout	INT	Default: 10 The amount of time the connection to the RADIUS server can remain idle before Content Gateway closes the connection.
proxy.config.radius.auth.max_ retries	INT	Default: 10 The maximum number of times Content Gateway tries to connect to the RADIUS server.
proxy.config.radius.cache.size	INT	Default: 1000 The number of entries allowed in the RADIUS cache. The minimum value is 256 entries.
proxy.config.radius.cache. storage_size	INT	Default: 15728640 The maximum amount of space that the RADIUS cache can occupy on disk. This value must be at least one hundred times the number of entries. It is recommended that you provide the maximum amount of disk space possible.
proxy.config.radius.auth.ttl_ value	INT	Default: 60 The number of minutes that Content Gateway stores username and password entries in the RADIUS cache.

NTLM

Configuration Variable	Data Type	Description
proxy.config.ntlm.auth.enabled	INT	Default: 0
		Enables (1) or disables (0) NTLM proxy authentication.

Configuration Variable	Data Type	Description
proxy.config.ntlm.dc.list	STRING	Default: NULL
		A comma-separated list of domain controller hostnames. The format is:
		host_name[:port] [%netbios_name]
		or
		IP_address[:port] [%netbios_name]
		If you are using Active Directory 2008,
		you must include the netbios_name or use SMB port 445.
proxy.config.ntlm.dc.load_ balance	INT	Default: 0
		Enables (1) or disables (0) load balancing. When enabled, Content Gateway balances the load when sending authentication requests to the domain controllers.
		Note
		When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short- term failover provision, until such time that the primary domain controller can accept new connections.
proxy.config.ntlm.dc.max_	INT	Default: 10
connections		Specifies the maximum number of connections Content Gateway can have open to the domain controller.

Configuration Variable	Data Type	Description
proxy.config.ntlm.cache.enabled	INT	Default: 1 Enables (1) or disables (0) the NTLM cache. Applies only when Content Gateway is an explicit proxy. When disabled, Content Gateway does not store any credentials in the NTLM cache for future use. Content Gateway always sends the credentials to the domain server to be validated.
proxy.config.ntlm.cache.ttl_ value	INT	Default: 900 Specifies the number of seconds that Content Gateway stores entries in the NTLM cache. The supported range of values is 300 to 86400 seconds.
proxy.config.ntlm.cache.size	INT	Default: 5000 Specifies the number of entries allowed in the NTLM cache.
proxy.config.ntlm.cache. storage_size	INT	Default: 15728640 Specifies the maximum amount of space that the NTLM cache can occupy on disk. This value should be proportionate to number of entries in the NTLM cache. For example, if each entry in the NTLM cache is approximately 128 bytes and the number of entries allowed in the NTLM cache is 5000, the cache storage size should be at least 64000 bytes.
proxy.config.ntlm.cache_ 0exception.list	STRING	Default: NULL Holds the list of IP addresses and IP address ranges that will not be cached. This variable gets its value from the Content Gateway manager NTLM Multi-Host IP addresses field. The exception list is a comma separated list that can contain up to: 64 IPv4 addresses 32 IPv4 addresses 24 IPv6 addresses 12 IPv6 address ranges

Configuration Variable	Data Type	Description
proxy.config.ntlm.fail_open	INT	 Default: 1 Enables (1) or disables (0) whether client requests are allowed to proceed when authentication fails due to: no response from the domain controller badly formed messages from the client invalid SMB responses Image: Password authentication failures are always failures.
proxy.config.ntlm.check_ account_passwd	INT	Default: 0 Enables (1) or disables (0) whether Content Gateway will create a log file entry when users are locked out after multiple failed password errors. Filter. config can be edited for user agents causing the lockout. Image: State of the

Integrated Windows Authentication

Configuration Variable	Data Type	Description
proxy.config.winauth.enabled	INT	Default: 0
		Enables (1) or disables (0) Integrated Windows Authentication (Kerberos).
proxy.config.winauth.basic_	INT	Default: 0
auth_disable		Disables (1) or enables (0) basic authentication.

Configuration Variable	Data Type	Description
proxy.config.winauth.realm	STRING	Default: NULL
		The name of the Windows Active Directory domain. By entering "*", all domain controllers found in the DNS SRV records will be used.
proxy.config.winauth.dc.list	STRING	Default: NULL
		A comma separated list of domain controllers.
proxy.config.winauth.log_	INT	Default: 1
denied_requests		Enables (1) or disables (0) logging of denied authentication requests.

Transparent authentication

Configuration Variable	Data Type	Description
proxy.config.http.transparent_	STRING	Default: NULL
auth_hostname		An alternate hostname for the proxy that can be resolved for all clients via DNS. This is needed if the regular hostname of the Content Gateway machine cannot be resolved for all users via DNS.
		When Cookie Sharing is enabled, this value must be the FQDN of the load balancer and must be the same value for all proxies in the cluster.
proxy.config.http.transparent_	INT	Default: 1
auth_type		 0 associates a session ID with the username after the user session is authenticated. Use this setting to uniquely identify users who share a single IP address, such as in proxy- chaining.
		 1 associates a client IP address with a username after the user session is authenticated.
		In either mode, the length of time before a client must re- authenticate is determined by the value of proxy.config.http. transparent_auth_session_time .

Configuration Variable	Data Type	Description
proxy.config.http.transparent_ auth_session_time	INT	Default: 15 The length of time (in minutes) before the browser must re- authenticate. This value is used in both IP and cookie modes.

HTTP engine

Configuration Variable	Data Type	Description
proxy.config.http.server_port	INT	Default: 8080
		The port that Content Gateway uses when acting as a proxy server for web traffic or when serving web traffic transparently.
proxy.config.http.server_port_ attr	STRING	Default: X
		Specifies the server port options. You can specify one of the following:
		C=SERVER_PORT_COMPRESSED
		X=SERVER_PORT_DEFAULT
		T=SERVER_PORT_BLIND_TUNNE
proxy.config.http.server_other_	STRING	Default: NULL
ports		Specifies the ports other than the port specified by the variable proxy.config. http.server_port to bind for incoming HTTP requests.
proxy.config.http.ssl_ports	STRING	Default: 443 563 8081 8071 9443 9444
		8443 9447
		The ports used for tunneling.This is a space-separated list that can also include ranges of ports, e.g. 1-65535.
		Content Gateway allows tunnels only to the specified ports.

Configuration Variable	Data Type	Description
proxy.config.http.insert_	INT	Default: 1
		0 = no extra information is added to the string
		 1 = all extra information is
		 added. 2 = some extra information is
		added.
proxy.config.http.insert_	INT	Default: 1
		 0 = no extra information is added to the string.
		 1 = all extra information is added
		 2 = some extra information is
		added.
proxy.config.http.enable_url_	INT	Default: 1
		Lnables (1) or disables (0) .com domain expansion, which
		configures Content Gateway to attempt to resolve ungualified
		hostnames by redirecting them to
		with www. and appended with
		a request to host , Content Gateway
		redirects the request to www. host.com.
proxy.config.http.no_dns_just_	INT	Default: 0
forward_to_parent		When enabled (1), and if HTTP
		Gateway does no DNS lookups on
nrovu config http://poochoohlo		Default: 0
requests_bypass_parent		When enabled (1). Content
		Gateway bypasses the parent proxy for a request that is not cacheable.
proxy.config.http.keep_alive_	INT	Default: 1
		Enables (1) or disables (0) the use of keep- alive connections to either origin servers or clients.

Configuration Variable	Data Type	Description
proxy.config.http.chunking_enabled	INT	Default: 1
		Specifies whether Content Gateway will generate a chunked response:
		• 0 = Never
		1 = Always
proxy.config.http.send_http11_	INT	Default: 3
requests		Configures Content Gateway to use HTTP Version 1.1 when communicating with origin servers. You can specify one of the following values:
		 0 = Never use HTTP 1.1 when communicating with origin servers.
		 1 = Always use HTTP 1.1 when communicating with origin servers.
		 2 = Use HTTP 1.1 if the origin server has previously used HTTP 1.1.
		 3 = Use HTTP 1.1 if the client request is HTTP 1.1 and the origin server has previously used HTTP 1.1.
		Note If HTTP 1.1 is used, Content Gateway can use keep-alive connections with pipelining to origin servers. If HTTP 0.9 is used, Content Gateway does not use keep- alive connections to origin servers. If HTTP 1.0 is used, a Content Gateway can use keep-alive connections without pipelining to origin servers.

Configuration Variable	Data Type	Description
proxy.config.http.send_http11_ asfirstrequest	INT	Default: 1 When enabled (1), specifies that Content Gateway send HTTP 1.1 in the first request to server. Otherwise, the default behavior is specified by proxy.config. http. send_http11_requests .
proxy.config.http.share_server_ sessions	INT	Default: 1 Enables (1) or disables (0) the re- use of server sessions. Note When IP spoofing is enabled, Content Gateway automatically disables this variable.
proxy.config.http.share_server_ sessions_max	INT	Default: 2500 The maximum number of server sessions that can be reused.
proxy.config.http.ftp_enabled	INT	Default: 1 Enables (1) or disables (0) Content Gateway from serving FTP requests sent via HTTP.
proxy.config.http.record_heartbeat	INT	Default: 0 Enables (1) or disables (0) content_cop heartbeat logging.
proxy.config.http.large_file_ support	INT	Default: 1 When enabled (1), Content Gateway supports downloading of files larger than 2 GB.

Parent proxy configuration

Configuration Variable	Data Type	Description
proxy.config.http.parent_proxy_	INT	Default: 1
routing_enable		Enables (1) or disables (0) the HTTP parent caching option.
		See Hierarchical Caching.

Configuration Variable	Data Type	Description
proxy.config.http.parent_proxy.	INT	Default: 300
retry_time		The amount of time allowed between connection retries to a parent cache that is unavailable.
proxy.config.http.parent_proxy.	INT	Default: 10
fail_threshold		The number of times the connection to the parent cache can fail before Content Gateway considers the parent unavailable.
proxy.config.http.parent_proxy. total_connect_attempts	INT	Default: 4 The total number of connection attempts allowed to a parent cache before Content Gateway bypasses the parent or fails the request (depending on the go_direct option in the bypass.config file).
proxy.config.http.parent_proxy.	INT	Default: 2
per_parent_connect_attempts		The total number of connection attempts allowed per parent if multiple parents are used.
proxy.config.http.parent_proxy.	INT	Default: 30
connect_attempts_timeout		The timeout value, in seconds, for parent cache connection attempts.
proxy.config.http.forward.	INT	Default: 0
proxy_auth_to_parent		When enabled (1), the Proxy- Authorization header is <i>not</i> stripped from requests sent to a parent proxy.
		Enable this when Content Gateway is a child proxy and the parent proxy performs authentication.
proxy.config.http.child_proxy.	INT	Default: 0
read_auth_from_header		When Content Gateway is the parent proxy, read X-Authenticated- User and X- Forwarded-For fields from incoming request headers.
		1 = enabled
		0 = disabled
proxy.local.http.parent_proxy.	INT	Default: 0
		When enabled (1), HTTPS requests bypass the parent proxy.

Configuration Variable	Data Type	Description
proxy.local.http.parent_proxy. disable_unknown_connect_ tunneling	INT	Default: 0 When enabled (1), non-HTTPS tunnel requests bypass the parent proxy.

Related information

Hierarchical Caching on page 79

HTTP connection timeouts (secs)

Configuration Variable	Data Type	Description
proxy.config.http.keep_alive_ no_activity_timeout_in	INT	Default: 60 How long Content Gateway keeps connections to clients open for a subsequent request after a transaction ends.
proxy.config.http.keep_alive_ no_activity_timeout_out	INT	Default: 60 How long Content Gateway keeps connections to origin servers open for a subsequent transfer of data after a transaction ends.
proxy.config.http.transaction_ no_activity_timeout_in	INT	Default: 120 How long Content Gateway keeps connections to clients open if a transaction stalls.
proxy.config.http.transaction_ no_activity_timeout_out	INT	Default: 120 How long Content Gateway keeps connections to origin servers open if the transaction stalls.
proxy.config.http.transaction_ active_timeout_in	INT	Default: 0 How long Content Gateway remains connected to a client. If the transfer to the client is not complete before this timeout expires, Content Gateway closes the connection. The default value of 0 specifies that there is no timeout.

Configuration Variable	Data Type	Description
proxy.config.http.transaction_	INT	Default: 0
active_timeout_out		How long Content Gateway waits for fulfillment of a connection request to an origin server. If Content Gateway does not complete the transfer to the origin server before this timeout expires, the connection request is terminated.
		The default value of 0 specifies that there is no timeout.
proxy.config.http.accept_no_	INT	Default: 120
activity_timeout		The timeout interval in seconds before Content Gateway closes a connection that has no activity.
proxy.config.http.background_	INT	Default: 60
fill_active_timeout		How long Content Gateway continues a background fill before giving up and dropping the origin server connection.
proxy.config.http.background_	FLOAT	Default: 0.50000
fill_completed_threshold		The proportion of total document size already transferred when a client aborts at which the proxy continues fetching the document from the origin server to get it into the cache (a <i>background fill</i>).

Origin server connection attempts

Configuration Variable	Data Type	Description
proxy.config.http.connect_ attempts_max_retries	INT	Default: 1 The maximum number of
		connection retries Content Gateway makes when the origin server is not responding.
proxy.config.http.connect_	INT	Default: 1
attempts_max_retries_dead_ server		The maximum number of connection retries Content Gateway makes when the origin server is unavailable.

Configuration Variable	Data Type	Description
proxy.config.http.connect_	INT	Default: 2
attempts_rr_retries		The maximum number of failed connection attempts allowed before a round-robin entry is marked as down if a server has round-robin DNS entries.
proxy.config.http.connect_	INT	Default: 60
		The timeout value in seconds for an origin server connection.
proxy.config.http.streaming_	INT	Default: 1800
connect_attempts_timeout		The timeout value in seconds for a streaming content connection.
proxy.config.http.down_server.	INT	Default: 30
cache_time		How long in seconds Content Gateway remembers that an origin server was unreachable.
proxy.config.http.down_server.	INT	Default: 10
abort_threshold		The number of seconds before Content Gateway marks an origin server as unavailable when a client abandons a request because the origin server was too slow in sending the response header.

Negative response caching

Configuration Variable	Data Type	Description
proxy.config.http.negative_	INT	Default: 0
caching_enabled		When enabled (1), Content Gateway caches negative responses, such as 404 Not Found, if a requested page does not exist. The next time a client requests the same page, Content Gateway serves the negative response from the cache.
		Content Gateway caches the following negative responses:
		204 No Content
		305 Use Proxy
		400 Bad Request

Configuration Variable	Data Type	Description
		403 Forbidden
		404 Not Found
		405 Method Not Allowed
		500 Internal Server Error
		501 Not Implemented
		502 Bad Gateway
		503 Service Unavailable
		504 Gateway Timeout
proxy.config.http.negative_	INT	Default: 1800
caching_lifetime		Specifies how long Content Gateway
		keeps the negative responses as valid in
		cache.

Proxy users variables

Configuration Variable	Data Type	Description
proxy.config.http.anonymize_ remove_from	INT	Default: 0 When enabled (1), Content Gateway removes the From header that accompanies transactions to protect the privacy of your users.
proxy.config.http.anonymize_ remove_referer	INT	Default: 0 When enabled (1), Content Gateway removes the Referer header that accompanies transactions to protect the privacy of your site and users.
proxy.config.http.anonymize_ remove_user_agent	INT	Default: 0 When enabled (1), Content Gateway removes the User- Agent header that accompanies transactions to protect the privacy of your site and users.

Configuration Variable	Data Type	Description
proxy.config.http.anonymize_ remove_cookie	INT	Default: 0 When enabled (1), Content Gateway removes the Cookie header that accompanies transactions to protect the privacy of your site and users.
proxy.config.http.anonymize_ remove_client_ip	INT	Default: 1 When enabled (1), Content Gateway removes Client-IP headers for more privacy.
proxy.config.http.anonymize_ insert_client_ip	INT	Default: 0 When enabled (1), Content Gateway inserts Client-IP headers to retain the client's IP address.
proxy.config.http.anonymize_ other_header_list	STRING	Default: NULL Specifies the headers that Content Gateway will remove from outgoing requests. Can be specified in a comma separated list.
proxy.config.http.snarf_ username_from_authorization	INT	Default: 0 When enabled (1), Content Gateway takes the username and password from the authorization header for LDAP if the authorization scheme is <i>Basic</i> .
proxy.config.http.insert_squid_ x_forwarded_for	INT	Default: 0 When enabled (1), Content Gateway adds the client IP address to the X-Forwarded- For header when the outbound request is sent to a configured parent proxy.
proxy.config.http.insert_xff_to_ external	INT	Default: 0 When enabled (1), Content Gateway adds the client IP address to the X-Forwarded- For header to outbound requests sent to the Internet. Note: This variable must be manually added to the config file.

Configuration Variable	Data Type	Description
proxy.config.http.insert_x_	INT	Default: 0
authenticateduser		When enabled (1), Content Gateway inserts the X- Authenticated-User header to advertise the proxy authenticated user. When enabled, the user name will be sent only to a configured parent proxy.
proxy.config.http.insert_xua_to_	INT	Default: 0
external		When enabled (1), Content Gateway inserts the X - Authenticated-User header to advertise the proxy authenticated user to all outbound requests.
		Note This variable must be manually added to the config file.

Security

Configuration Variable	Data Type	Description
proxy.config.http.push_method_ enabled	INT	Default: 0 When enabled (1), filter.config rules can be used to push content directly into the cache without a user request. You must add a filtering rule with the PUSH action to ensure that only known source IP addresses implement PUSH requests to the cache. This variable must be enabled before PUSH is available in the Method drop down list in the configuration file editor.

Anonymous FTP password

Configuration Variable	Data Type	Description
proxy.config.http.ftp. anonymous_passwd	STRING	Default: the administrator email address supplied during installation The anonymous password for FTP servers that require a password for access. Content Gateway uses the Content
		Gateway user account name as the default value for this variable.

FTP transfer mode

Configuration Variable	Data Type	Description
proxy.config.http.ftp.binary_ transfer_only	INT	Default: 0 When enabled (1), all FTP documents requested from HTTP clients are transferred in binary mode only. When disabled (0), FTP documents requested from HTTP clients are transferred in ASCII or binary mode_depending on the document
		type.

Customizable user response pages

Configuration Variable	Data Type	Description
proxy.config.body_factory.	INT	Default: 0
enable_customizations		Specifies whether customizable response pages are enabled or disabled and which response pages are used:
		 0 = disable customizable user response pages
		 1 = enable customizable user response pages in the default directory only
		 2 = enable language-targeted user response pages

Configuration Variable	Data Type	Description
proxy.config.body_factory.	INT	Default: 0
enable_logging		Enables (1) or disables (0) logging for customizable response pages. When enabled, Content Gateway records a message in the error log each time a customized response page is used or modified.
proxy.config.body_factory.	STRING	Default: config/body_factory
template_sets_dir		Specifies the customizable response page default directory.
proxy.config.body_factory.	INT	Default: 0
response_suppression_mode		Specifies when Content Gateway suppresses generated response pages:
		 0 = never suppress generated response pages
		 1 = always suppress generated response pages
		 2 = suppress response pages only for intercepted traffic

FTP engine

Configuration Variable	Data Type	Description
FTP over HTTP		
proxy.config.ftp.data_ connection_mode	INT	Default: 1 Specifies the FTP connection mode: 1 = PASV then PORT 2 = PORT only 3 = PASV only
proxy.config.ftp.control_ connection_timeout	INT	Default: 300 Specifies how long Content Gateway waits for a response from the FTP server.

X		
proxy.config.ftp.rc_to_switch_	STRING	Default: NULL
to_PORT		Specifies the response codes for which Content Gateway automatically fails over to the PORT command when PASV fails if the configuration variable proxy.config. ftp.data_connection_mode is set to 1. This variable is used for FTP
		requests from HTTP clients only.
FTP Proxy		
proxy.config.ftp.ftp_enabled	INT	Default: 0
		Enables (1) or disables (0) processing of FTP requests from FTP clients.
proxy.config.ftp.logging_ enabled	INT	Default: 1
		Enables (1) or disables (0) logging of FTP transactions.
proxy.config.ftp.proxy_server_ port	INT	Default: 2121
		Specifies the port used for FTP connections.
proxy.config.ftp.open_lisn_	INT	Default: 1
port_mode		Specifies how FTP opens a listening port for a data transfer:
		 1 = The operating system chooses an available port. Content Gateway sends 0 and retrieves the new port number if the listen succeeds.
		 2 = The listening port is determined by the range of ports specified by the Content Gateway variables proxy. config.ftp.min_lisn_port and proxy. config.ftp.max_lisn_port, described below.
proxy.config.ftp.min_lisn_port	INT	Default: 32768
		Specifies the lowest port in the range of listening ports used by Content Gateway for data connections when the FTP client sends a PASV or Content Gateway sends a PORT to the FTP server.

Λ		
proxy.config.ftp.max_lisn_port	INT	Default: 65535
		Specifies the highest port in the range of listening ports used by Content Gateway for data connections when the FTP client sends a PASV or Content Gateway sends a PORT to the FTP server.
proxy.config.ftp.server_data_	INT	Default: 1
default_pasv		Specifies the default method used to set up server side data connections:
		 1 = Content Gateway sends a PASV to the FTP server and lets the FTP server open a listening port.
		 0 = Content Gateway tries PORT first (sets up a listening port on the proxy side of the connection).
proxy.config.ftp.different_	INT	Default: 0
client_port_ip_allowed		When enabled (1), Content Gateway can connect to a machine other than the one on which the FTP client is running to establish a data connection.
		The FTP client uses PORT to set up a listening port on its side and allows Content Gateway to connect to that port to establish the data connection (used to transfer files). When setting up the listening port, an FTP client specifies the IP address and port number for the listening port. If this variable is set to 0 (zero), Content Gateway cannot connect to the FTP client if the IP address sent by the client is different from the IP address of the machine running the FTP client.
proxy.config.ftp.try_pasv_times	INT	Default: 1024
		Specifies the number of times Content Gateway can try to open a listening port when the FTP client sends a PASV.
proxy.config.ftp.try_port_times	INT	Default: 1024
		Specifies the maximum number of times Content Gateway can try to open a listening port when sending a PORT to the FTP server.

proxy.config.ftp.try_server_ctrl_	INT	Default: 6
		Specifies the maximum number of times Content Gateway can try to connect to the FTP server's control listening port.
proxy.config.ftp.try_server_	INT	Default: 3
data_connect_times		Specifies the maximum number of times Content Gateway can try to connect to the FTP server's data listening port when it sends a PASV to the FTP server and gets the IP/ listening port information.
proxy.config.ftp.try_client_	INT	Default: 3
data_connect_times		Specifies the maximum number of times Content Gateway can try to connect to the FTP client's data listening port when the FTP client sends a PORT with the IP/ listening port information.
proxy.config.ftp.client_ctrl_no_	INT	Default: 900
activity_timeout		Specifies the inactivity timeout, in seconds, for the FTP client control connection.
proxy.config.ftp.client_ctrl_	INT	Default: 14400
active_timeout		Specifies the active timeout, in seconds, for the FTP client control connection.
proxy.config.ftp.server_ctrl_no_	INT	Default: 120
activity_timeout		Specifies the inactivity timeout, in seconds, for the FTP server control connection.
proxy.config.ftp.server_ctrl_	INT	Default: 14400
active_timeout		Specifies the active timeout, in seconds, for the FTP server control connection.
proxy.config.ftp.client_data_no_	INT	Default: 120
activity_timeout		Specifies the maximum time, in seconds, that a client FTP data transfer connection can be idle before it is aborted.
proxy.config.ftp.client_data_	INT	Default: 14400
		Specifies the maximum time, in seconds, of an FTP data transfer connection from a client.

proxy.config.ftp.server_data_ no_activity_timeout	INT	Default: 120 Specifies the maximum time, in seconds, that a server FTP data transfer connection can be idle before it is aborted.
proxy.config.ftp.server_data_ active_timeout	INT	Default: 14400 Specifies the maximum time, in seconds, of an FTP data transfer connection from a server.
proxy.config.ftp.pasv_accept_ timeout	INT	Default: 120 Specifies the timeout value for a listening data port in Content Gateway (for PASV, the client data connection).
proxy.config.ftp.port_accept_ timeout	INT	Default: 120 Specifies the timeout value for a listening data port in Content Gateway (for PORT, the server data connection).
proxy.config.ftp.share_ftp_ server_ctrl_enabled	INT	Default: 1 Enables (1) or disables (0) sharing the server control connections among multiple anonymous FTP clients.
proxy.config.ftp.share_only_ after_session_end	INT	 Default: 1 How an FTP server control connection is shared between different FTP client sessions: 1 = the FTP server control connection can be used by another FTP client session only when the FTP client session is complete (typically, when the FTP client sends out a QUIT command). 0 = the FTP server control connection can be used by another FTP client session only if the FTP client session is not actively using the FTP server connection: for example, during an idle session.
proxy.config.ftp.server_ctrl_ keep_alive_no_activity_timeout	INT	The timeout value when the FTP server control connection is not used by any FTP clients.

proxy.config.ftp.reverse_ftp_ enabled	INT	Default: 0 Not supported.
proxy.config.ftp.data_source_ port_20_enabled	INT	Default: 0 When enabled (1), bind to source port 20 for outgoing data transfer connections to Active mode FTP clients.

SOCKS processor

Configuration Variable	Data Type	Description	
proxy.config.socks.socks_ needed	INT	Default: 0 Enables (1) or disables (0) the SOCKS option. See <i>Configuring SOCKS firewall</i>	
proxy.config.socks.socks_version	INT	Default: 4 The SOCKS version.	
proxy.config.socks.default_ servers	STRING	Default: s1.example. com:1080;socks2:4080 The names and ports of the SOCKS servers with which Content Gateway communicates.	
proxy.config.socks.accept_ enabled	INT	Default: 0 Enables (1) or disables (0) the SOCKS proxy option. As a SOCKS proxy, Content Gateway receives SOCKS traffic (usually on port 1080) and forwards all requests directly to the SOCKS server.	
proxy.config.socks.accept_port	INT	Default: 1080 The port on which Content Gateway accepts SOCKS traffic.	
proxy.config.socks.socks_ server_enabled	INT	Default: 0 Note: Configure only if Content Gateway is installed on an appliance.	
Configuration Variable	Data Type	Descri	ption
---------------------------	-----------	---------	--
proxy.config.socks.socks_	INT	Default	:: 61080
server_port		Ę	Note
			Configure only if Content Gateway is installed on an appliance.

Related concepts

Configuring SOCKS firewall integration on page 160

Net subsystem

Configuration Variable	Data Type	Description
proxy.config.net.connections_ throttle	INT	Default: 45000 The maximum number of connections that Content Gateway can handle. If Content Gateway receives additional client requests, they are queued until existing requests are served. Do not set this variable below 100.

Cluster subsystem

Configuration Variable	Data Type	Description
proxy.config.cluster.cluster_port	INT	Default: 8086 The port used for cluster communication.
proxy.config.cluster.ethernet_ interface	STRING	Default: your_interface The network interface used for cluster traffic. All nodes in a cluster must use the same network interface.

Cache

Configuration Variable	Data Type	Description
proxy.config.cache.ram_cache. size	INT	Default: -1
		The size of the RAM cache, in bytes.
proxy.config.cache.limits.http.	INT	Default: 3
max_alts		The maximum number of HTTP alternates that Content Gateway can cache.
proxy.config.cache.max_doc_ size	INT	Default: 0
		The maximum size of documents in the cache (in bytes):
		0 = there is no size limit.

DNS

Configuration Variable	Data Type	Description
proxy.config.dns.search_	INT	Default: 1
default_domains		Enables (1) or disables (0) local domain expansion so that Content Gateway can attempt to resolve unqualified hostnames by expanding to the local domain; for example, if a client makes a request to an unqualified host named host_x , and if the Content Gateway local domain is y.com , Content Gateway expands the hostname to host_x.y.com .
proxy.config.dns.splitDNS. enabled	INT	Default: 0
		Enables (1) or disables (0) DNS server selection. When enabled, Content Gateway refers to the splitdns.config file for the selection specification. See Using the Split DNS option

Configuration Variable	Data Type	Description
proxy.config.dns.splitdns.def_ domain	STRING	Default: NULL The default domain for split DNS requests. This value is appended automatically to the hostname if it does not include a domain before split DNS determines which DNS server to use.
proxy.config.dns.splitdns.fast_ match.count	INT	Default: 4 The maximum number of rules that can exist in order for the NOT logical operator (!) to be applied to any of the rules.
proxy.config.dns.url_expansions	STRING	Default: NULL A list of extensions automatically added to the hostname after a failed lookup; for example, if you want Content Gateway to add the hostname extension .org, specify org as the value for this variable (Content Gateway automatically adds the dot (.) before the extension.) Note: If proxy.config.http.enable_url_ expandomatic is set to 1 (default), you do not have to add www.and .com to this list; Content Gateway tries www.and .com automatically after trying the values you specify.
proxy.config.dns.lookup_ timeout	INT	Default: 20 The DNS lookup timeout duration in seconds. When the timeout period expires, the lookup attempt is terminated. The default value is lower than proxy. config.hostdb.lookup_timeout and, therefore, takes precedence.
proxy.config.dns.retries	INT	Default: 1 The number of times a DNS lookup is retried before giving up.
proxy.config.dns.prefer_ipv4	INT	Default: 1 When a name resolves to both IPv4 and IPv6 addresses, specifies the preferred address type.

Configuration Variable	Data Type	Description
proxy.config.ipv6.ipv6_enabled	INT	Default: 0
		Used to enable (1) or disable (0) support for IPv6.

Related tasks

Using the Split DNS option on page 163

DNS proxy

Configuration Variable Data Type	Data Type	Description
proxy.config.dns.proxy.enabled	INT	Default: 0
		Enables (1) or disables (0) the DNS proxy caching option that lets you resolve DNS requests on behalf of clients. This option off-loads remote DNS servers and reduces response time for DNS lookups. See DNS Proxy Caching.
proxy.config.dns.proxy_port	INT	Default: 5353
		The port that Content Gateway uses for DNS traffic.

Related information

DNS Proxy Caching on page 85

HostDB

Configuration Variable	Data Type	Description
proxy.config.hostdb.size	INT	Default: 200000
		The maximum number of entries allowed in the host database.

Configuration Variable	Data Type	Description
proxy.config.hostdb.ttl_mode	INT	Default: 0
		The host database time to live (ttl) mode.
		0 = obey the ttl values set by the name servers
		1 = ignore the ttl values set by name servers and use the value set by the Content Gateway configuration variable proxy.config.hostdb.timeout .
		2 = use the lower of the two values (the one set by the name server or the one set by Content Gateway)
		3 = use the higher of the two values (the one set by the name server or the one set by Content Gateway)
proxy.config.hostdb.timeout	INT	Default: 86400
		The foreground timeout, in seconds.
proxy.config.hostdb.fail.timeout	INT	Default: 60
		The time for which a failed DNS will be cached in seconds.
proxy.config.hostdb.strict_	INT	Default: 0
round_robin		When disabled (0), Content Gateway always uses the same origin server for the same client as long as the origin server is available.

Logging configuration

Configuration Variable	Data Type	Description
proxy.config.log2.logging_ enabled	INT	Default: 1
		Enables and disables event logging:
		0 = logging disabled
		1 = log errors only
		2 = log transactions only
		 3 = full logging (errors + transactions) See Working With Log Files.

Configuration Variable	Data Type	Description
proxy.config.log2.max_secs_ per_buffer	INT	Default: 5 Specifies the maximum amount of time before data in the buffer is flushed to disk.
proxy.config.log2.max_space_ mb_for_logs	INT	Default: 5120 <i>or</i> 20480 The amount of space allocated to the logging directory, in megabytes. When Content Gateway is on an appliance, the size is 5120 (5 GB) and cannot be changed. When Content Gateway is installed on a standalone server, the default size is 20480 (20 GB) and the size is configurable.
proxy.config.log2.max_space_ mb_for_orphan_logs	INT	Default: 25 The amount of space allocated to the logging directory, in megabytes, if this node is acting as a collation client.
proxy.config.log2.max_space_ mb_headroom	INT	Default: 100 The tolerance for the log space limit in bytes. If the variable proxy.config.log2. auto_delete_rolled_file is set to 1 (enabled), auto-deletion of log files is triggered when the amount of free space available in the logging directory is less than the value specified here.
proxy.config.log2.hostname	STRING	Default: localhost The hostname of the machine running Content Gateway.
proxy.config.log2.logfile_dir	STRING	Default: /opt/WCG/logs The full path to the logging directory.

Configuration Variable	Data Type	Description
proxy.config.log2.logfile_perm	STRING	Default: rw-rr
		Specifies the log file permissions. The standard UNIX file permissions are used (owner, group, other). Valid values are:
		 - = no permission
		r = read permission
		w = write permission
		 x = execute permission Permissions are subject to the umask
		settings for the Content Gateway process. This means that a umask setting of 002 will not allow write permission for others, even if specified in the configuration file.
		Permissions for existing log files are not changed when the configuration is changed.
		Linux only.
proxy.config.log2.custom_logs_	INT	Default: 0
enabled		When enabled (1), supports the definition and generation of custom log files according to the specifications in logs_xml.config .
		See logs_xml.config.
proxy.config.log2.xml_logs_ config	INT	Default: 1
		Specifies the size, in MB, which when reached causes the log files to roll. See <i>Rolling event log files</i> .
proxy.config.log2.squid_log_	INT	Default: 0
enabled		Enables (1) or disables (0) the squid log file format.
proxy.config.log2.squid_log_is_	INT	Default: 1
ascii		The squid log file type:
		■ 1 = ASCII
		 0 = binary
proxy.config.log2.squid_log_ name	STRING	Default: squid
		Specifies the squid log filename.
proxy.config.log2.squid_log_	STRING	Default: NULL
header		The squid log file header text.

Configuration Variable	Data Type	Description
proxy.config.log2.common_log_ enabled	INT	Default: 0 Enables (1) or disables (0) the Netscape common log file format.
proxy.config.log2.common_log_ is_ascii	INT	Default: 1 The Netscape common log file type: 1 = ASCII 0 = binary
proxy.config.log2.common_log_ name	STRING	Default: common The Netscape common log filename.
proxy.config.log2.common_log_ header	STRING	Default: NULL The Netscape common log file header text.
proxy.config.log2.extended_ log_enabled	INT	Default: 1 Enables (1) or disables (0) the Netscape extended log file format.
proxy.confg.log2.extended_log_ is_ascii	INT	Default: 1 The Netscape extended log file type: 1 = ASCII 0 = binary
proxy.config.log2.extended_ log_name	STRING	Default: extended Specifies the Netscape extended log filename.
proxy.config.log2.extended_ log_header	STRING	Default: NULL Specifies the Netscape extended log file header text.
proxy.config.log2.extended2_ log_enabled	INT	Default: 0 Enables (1) or disables (0) the Netscape Extended-2 log file format.
proxy.config.log2.extended2_ log_is_ascii	INT	Default: 1 The Netscape Extended-2 log file type: 1 = ASCII 0 = binary

Configuration Variable	Data Type	Description
proxy.config.log2.extended2_ log_name	STRING	Default: extended2 The Netscape Extended-2 log filename.
proxy.config.log2.extended2_ log_header	STRING	Default: NULL The Netscape Extended-2 log file header text.
proxy.config.log2.separate_ host_logs	INT	Default: 0 When enabled (1), configures Content Gateway to create a separate log file for HTTP/FTP transactions for each origin server listed in the log_hosts.config file (see <i>HTTP host log splitting</i>).
proxy.local.log2.collation_mode	INT	 Default: 0 The log collation mode: 0 = Collation disabled. 1 = This host is a log collation server. 2 = This host is a collation client and sends entries using standard formats to the collation server. For information about sending XML- based custom formats to the collation server, see logs_xml.config.
proxy.confg.log2.collation_host	STRING	Default: NULL The hostname of the log collation server.
proxy.config.log2.collation_port	INT	Default: 8085 The port used for communication between the collation server and client.
proxy.config.log2.collation_secret	STRING	Default: foobar The password used to validate logging data and prevent the exchange of unauthorized information when a collation server is being used.

Configuration Variable	Data Type	Description
proxy.config.log2.collation_ host_tagged	INT	Default: 0 When enabled (1), configures Content Gateway to include the hostname of the collation client that generated the log entry in each entry.
proxy.config.log2.collation_ retry_sec	INT	Default: 5 The number of seconds between collation server connection retries.
proxy.config.log2.rolling_enabled	INT	Default: 1 Enables (1) or disables (0) log file rolling. See <i>Rolling event log files</i> .
proxy.config.log2.rolling_ interval_sec	INT	Default: 21600 The log file rolling interval, in seconds. The minimum value is 300 (5 minutes). The maximum value is 86400 seconds (one day).
proxy.config.log2.rolling_offset_hr	INT	Default: 0 The file rolling offset hour. The hour of the day that starts the log rolling period.
proxy.config.log2.rolling_size_ mb	INT	Default: 10 The size, in megabytes, which when reached causes the current file to be closed and a new file to be created.
proxy.config.log2.auto_delete_ rolled_files	INT	Default: 1 Enables (1) or disables (0) automatic deletion of rolled files.
proxy.config.log2.sampling_ frequency	INT	 Default: 1 Configures Content Gateway to log only a sample of transactions rather than every transaction. You can specify the following values: 1 = log every transaction 2 = log every second transaction 3 = log every third transaction and so on

Related concepts		
logs_xml.config on page 377		
Rolling event log files on page 227		
HTTP host log splitting on page 230		

Related information

Working With Log Files on page 217

URL remap rules

Configuration Variable	Data Type	Description
proxy.config.url_remap.default_ to_server_pac	INT	Default: 0 Enables (1) or disables (0) requests for a PAC file on the proxy service port (8080 by default) to be redirected to the PAC port
		For this type of redirection to work, proxy. config.reverse_proxy.enabled must be set to 1.
proxy.config.url_remap.default_	INT	Default: -1
to_server_pac_port		PAC requests made to the Content Gateway proxy service port are redirected to this port.
		-1 sets the PAC port to the autoconfiguration port (default 8083).
		This variable can be used with proxy. config.url_remap.default_
		to_server_pac to get a PAC file from a different port. You must create and run a process that serves a PAC file on this port; for example, if you create a Perl script that listens on port 9000 and writes a PAC file in response to any request, you can set this variable to 9000, and browsers that request the PAC file from a proxy server on port 8080 will get the PAC file served by the Perl script.

Configuration Variable	Data Type	Description
proxy.config.url_remap.remap_ required	INT	Default: 0 Set this variable to 1 if you want Content Gateway to serve requests only from origin servers listed in the mapping rules of the remap.config file. If a request does not match, the browser receives an error.
proxy.config.url_remap. pristine_host_hdr	INT	Default: 0 Set this variable to 1 if you want to retain the client host header in a request during remapping.

SNMP configuration

Configuration Variable	Data Type	Default Value
proxy.config.snmp.master_agent_ena	INT	0
proxy.config.snmp_encap_enabled	INT	0

Plug-in configuration

Configuration Variable	Data Type	Description
proxy.config.plugin.plugin_dir	STRING	Default: config/plugins
		The directory in which plugins are located.

WCCP configuration

Configuration Variable	Data Type	Description
proxy.config.wccp.enabled	INT	Default: 0
		Enables (1) or disables (0) WCCP.

FIPS (Security Configuration)

Configuration Variable	Data Type	Description
proxy.config.fips.security_ enabled	INT	Default: 0 Warning: Do not enable FIPS mode in records.config. Use the Content Gateway manager: Configure > Security > FIPS Security . FIPS mode cannot be disabled without reinstalling Content Gateway.
proxy.config.fips.security_ enabled_ui	INT	Default: 0 Warning: Do not enable FIPS mode in records.config. Use the Content Gateway manager: Configure > Security > FIPS Security . FIPS mode cannot be disabled without reinstalling Content Gateway.

SSL Decryption

Configuration Variable	Data Type	Description
proxy.config.ssl.enabled	INT	Default: 1
		When enabled (1), Content Gateway accepts SSL connections and performs URL filtering before establishing a connection with the origin server.
		See proxy.config.ssl_decryption.use_ decryption to enable SSL decryption.
proxy.config.ssl_decryption.	INT	Default: 0
use_decryption		When enabled (1), Content Gateway accepts and decrypts SSL traffic. See <i>Working With Encrypted Data</i> .
proxy.config.ssl_decryption_ ports	INT	Default: 443
		The HTTPS ports. Content Gateway allows SSL decryption and policy lookup only to the specified ports.
proxy.config.ssl_decryption.	INT	Default: 1
tunnel_ILSv13		When enabled (1), allows for tunneling of TLSV1.3-only connections (SSL connections that offer TLSv1.3, and no other protocols, as their "client hello").
proxy.config.administrator_id	STRING	Default: NULL Do not change.
		Holds the encrypted administrator ID.

Configuration Variable	Data Type	Description	
proxy.config.ssl_decryption. tunnel_unknown_protocols	INT	Default: 0 Enables (1) or disables the tunneling of unrecognized protocols using SSL ports.	
proxy.config.ssl_decryption. tunnel_unknown_protocols_ timeout	INT	Default: 10 Specifies the time in seconds that Content Gateway waits for the "client hello" response before tunneling the request as an unknown protocol.	
proxy.config.ssl_decryption. mirror_enabled	INT	Default: 0 Enables (1) or disables SSL Decryption for Port Mirroring. Note that this feature is only available when SSL decryption is enabled and when Content Gateway is installed on an appliance. This variable should be edited only by using the appliance CLI.	
proxy.config. ssl_decryption. mirror_interface	STRING	Default: NULL The appliance interface that will be used to mirror decrypted SSL traffic. This variable should be edited only by using the appliance CLI.	
proxy.config. ssl_decryption. custom_request_header	STRING	Default: X-Proxy-HTTPS:1 The custom header name and value that Port Mirroring inserts into each HTTP request header sent to the monitor network interface. This variable should be edited only by using the appliance CLI.	
proxy.config.ssl.server.TLSv1	INT	Default: 0 When enabled (1), Content Gateway accepts TLSv1 connections from clients. (In this case, "server" refers to Content Gateway's role as server to the client.)	
proxy.config.ssl.server.TLSv11	INT	Default: 0 When enabled (1), Content Gateway accepts TLSv1.1 connections from clients. (In this case, "server" refers to Content Gateway's role as server to the client.)	
proxy.config.ssl.server.TLSv12	INT	Default: 1 When enabled (1), Content Gateway accepts TLSv1.2 connections from clients. (In this case, "server" refers to Content Gateway's role as server to the client.)	

Configuration Variable	Data Type	Description
proxy.config.ssl.server.TLSv13	INT	Default: 1
		When enabled (1), Content Gateway accepts TLSv1.3 connections from clients. (In this case, "server" refers to Content Gateway's role as server to the client.)
proxy.config.ssl.client.TLSv1	INT	Default: 0
		When enabled (1), Content Gateway accepts TLSv1 connections from origin servers. (In this case, "client" refers to Content Gateway's role as client to the origin server.)
proxy.config.ssl.client.TLSv11	INT	Default: 0
		When enabled (1), Content Gateway accepts TLSv1.1 connections from origin servers. (In this case, "client" refers to Content Gateway's role as client to the origin server.)
proxy.config.ssl.client.TLSv12	INT	Default: 1
		When enabled (1), Content Gateway accepts TLSv1.2 connections from origin servers. (In this case, "client" refers to Content Gateway's role as client to the origin server.)
proxy.config.ssl.client.TLSv13	INT	Default: 1
		When enabled (1), Content Gateway accepts TLSv1.3 connections from origin servers. (In this case, "client" refers to Content Gateway's role as client to the origin server.)
proxy.config.ssl.client.TLS_ padding	INT	Default: 1
		When enabled (1), Content Gateway will add padding to ensure a "client hello" does not hang the connection
proxy.config.ssl.server.custom_	INT	Default: 0
ciphersuites_enabled		When enabled (1), Content Gateway will consider customer cipher suites configured in proxy.config.ssl.server.cipher_suite (for TLS1.2 and below) and proxy.config.ssl.server.TLSv1_3.cipher_suites (for TLS1.3).

Configuration Variable	Data Type	Description	
proxy.config.ssl.server.cipher_suite	STRING	Default: DEFAULT	
		Specifies the client-to-proxy cipher setting. Values are:	
		DEFAULT	
		HIGH	
		MEDIUM:HIGH	
		In the user interface, when MEDIUM is selected, this value (MEDIUM:HIGH) is actually set in the config file.	
		These above options must be in uppercase.	
		In addition to the predefined options, you can apply custom cipher suites using the format recommended by OpenSSL. To enable this, set proxy.config.ssl.server.custom_ciphersuites_enabled to 1; otherwise, the configuration will be considered invalid.	
		See SSL configuration settings for inbound traffic.	
proxy.config.ssl.server.cipherlist_suffix	STRING	Default:	
		:!ADH:!RC4:!3DES:!EXP:!DES:!IDEA- CBC- SHA:@STRENGTH	
		List of ciphers not allowed for use in client-to-proxy (inbound) communication.	
		The cipher list is determined by combining the corresponding cipherlist_ option with this list.	
		Note these entries are case-sensitive and require the leading colon (:).	
proxy.config.ssl.server.	STRING	Default:	
TLSv1_3.cipher_suites		TLS_AES_256_GCM_SHA384:	
		TLS_CHACHA20_POLY1305_SHA256:	
		TLS_AES_128_GCM_SHA256:	
		TLS_AES_128_CCM_8_SHA256:	
		TLS_AES_128_CCM_SHA256	
		List of cipher suites to be used for TLS1.3 when proxy.config.ssl.server.custom_ciphersuites_enabled is set to 1. Use the OpenSSL format for specifying TLS1.3 cipher suites.	
proxy.config.ssl.client.custom_	INT	Default: 0	
ciphersuites_enabled		When enabled (1), Content Gateway will consider customer cipher suites configured in proxy.config.ssl.client.cipher_suite (for TLS1.2 and below) and proxy.config.ssl.client.TLSv1_3.cipher_suites (for TLS1.3).	

Configuration Variable	Data Type	Description	
proxy.config.ssl.client.cipher_suite	STRING	Default: DEFAULT	
		Specifies the client-to-proxy cipher setting. Values are:	
		DEFAULT	
		HIGH	
		MEDIUM:HIGH	
		In the user interface, when MEDIUM is selected, this value (MEDIUM:HIGH) is actually set in the config file.	
		These above options must be in uppercase.	
		In addition to the predefined options, you can apply custom cipher suites using the format recommended by OpenSSL. To enable this, set proxy.config.ssl.client.custom_ciphersuites_enabled to 1; otherwise, the configuration will be considered invalid. See <i>SSL configuration settings for inbound traffic</i> .	
proxy config ssl client cinherlist suffix	STRING	Default:	
		:!ADH:!RC4:!3DES:!EXP:!DES:!IDEA- CBC- SHA:@STRENGTH	
		List of ciphers not allowed for use in proxy-to-server (outbound) communication.	
		The cipher list is determined by combining the corresponding cipherlist_ option with this list.	
		Note these entries are case-sensitive and require the leading colon (:).	
proxy.config.ssl.client.	STRING	Default:	
TLSv1_3.cipher_suites		TLS_AES_256_GCM_SHA384:	
		TLS_CHACHA20_POLY1305_SHA256:	
		TLS_AES_128_GCM_SHA256:	
		TLS_AES_128_CCM_8_SHA256:	
		TLS_AES_128_CCM_SHA256	
		List of cipher suites to be used for TLS1.3 when proxy.config.ssl.client.custom_ciphersuites_enabled is set to 1. Use the OpenSSL format for specifying TLS1.3 cipher suites.	
proxy.config.ssl.client.	INT	Default: 0	
		Whether client certificates are not needed, optional, or required. certification level should be:	
		0 = no client certificates	
		1 = client certificates optional	
		2 = client certificates required	

Configuration Variable	Data Type	Description
proxy.config.ssl.client.set_sni	IINT	Default: 1
		Enables (1) or disables (0) a feature that forces the proxy to add an outbound SNI (server name indication) when requesting a server certificate be added to the Incident List.
proxy.config.ssl_skip_dns_on_ sni	INT	Default: 0
		Enables (0) or disables (1) a DNS lookup for the CONNECT hostname when X- Server-IP is present in the header
proxy.config.ssl.server.cert. filename	STRING	Default: server.crt.pem
		The server certificate filename.
proxy.config.ssl.server.private_	STRING	Default: Domainkey.pem
key.filename		The private key for the server certificate.
proxy.config.ssl.server.private_	STRING	Default: /config
key.path		The private key path for the server certificate.
proxy.config.ssl.CA.cert. filename	STRING	Default: NULL
		Te name of the file containing the list of CAs that Content Gateway will accept from a client.
		When the connection is from the client to Content Gateway and the value of proxy. config.ssl.client.certification_level is 1 or 2, Content Gateway sends the CA list to client.
proxy.config.ssl.CA.cert.path	STRING	Default: NULL
		The path to the CA list files. See the preceding entry.
proxy.config.ssl.catree_update	INT	Default: 1
		Enables (1) or disables (0) automatic updates of the Certificate Authority tree. See <i>Automatic certificate updates</i> .
proxy.config.ssl.client.cert. policy	INT	For SSL certificate incidents, specifies whether to tunnel an incident (0), or block the request and create an entry in the incident list (1).
proxy.config.ssl.client.verify. server	INT	Enables (1) or disables the Certificate Verification Engine (CVE). See <i>Validating certificates</i> .
proxy.config.ssl.cert.verify.	INT	Default: 0
denycnmismatch		Enables (1) or disables the CVE check: "Deny certificates where the common name does not match the URL"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify.add_	INT	Default: 1
cert_to_database		Enables (1) or disables the automatic adding of new certificates to the certificate database.

Configuration Variable	Data Type	Description
proxy.config.ssl.cert.verify. allowcnwild	INT	Default: 0
		Enables (1) or disables the CVE check: "Allow wildcard certificates"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify. denyexpired	INT	Default: 0
		Enables (1) or disables the CVE check: "No expired or not yet valid certificates"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify.	INT	Default: 1
denyselfsigned		Enables (1) or disables the CVE check: "Deny self-signed certificates"
		This setting applies only when the CVE is enabled
proxy.config.ssl.cert.verify.	INT	Default: 1
denysha1cert		Enables (1) or disables a feature that invalidates SHA-1 intermediate certificates for HTTPS traffic. A block page is served if a SHA-1 certificate is encountered.
proxy.config.ssl.cert.verify. certchain	INT	Default: 1
		Enables (1) or disables the CVE check: "Verify entire certificate chain"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify. checkcrl	INT	Default: 0
		Enables (1) or disables the CVE check: "Check certificate revocation by CRL"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify. checkocsp	INT	Default: 0
		Enables (1) or disables the CVE check: "Check certificate revocation by OCSP"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify.	INT	Default: 0
blockunknownocsp		Enables (1) or disables the CVE check: "Block certificates with Unknown OCSP state"
		The setting applies only when the CVE is enabled.
proxy.config.ssl.cert.verify.	INT	Default: 0
denymd5cert		Enables (1) denial of certificates that use an MD5 signiture.

Configuration Variable	Data Type	Description
proxy.config.ssl.cert.verify. revprefer	INT	Default: 1 The preferred method for the certificate revocation check. 1 = CRL
proxy config ssl cert verify blocknouri		2 = OCSP
proxy.comig.ssi.cen.veniy. blockhoun		Enables (1) or disables the CVE check: "Block certificates with no CRL URI and with no OCSP URI"
proxy.config.ssl.cert.verify. bypassfail	INT	Default: 1 Enables (1) the certificate check failure bypass option that allows users to proceed to a site after the certificate check has failed.
proxy.config.ssl.cert.verify. bypasscache	INT	Default: 1 Enables (1) the verification timeout cache.
proxy.config.ssl.cert.verify. bypasscachetimeout	INT	Default: 6 The time, in seconds, that an entry in verification bypass cache times out and is purged.
proxy.config.ssl_decryption_ bypass.tunnel_non-ssl_traffic	INT	Default: 0 Enables (1) or disables (0) tunneling of non-ssl traffic. This variable must be added manually.

Related concepts

Automatic certificate updates on page 126 Validating certificates on page 131

Related tasks

SSL configuration settings for inbound traffic on page 127

Related information

Working With Encrypted Data on page 111

ICAP

Configuration Variable	Data Type	Description
proxy.config.icap.enabled	INT	Default: 0
		Enables (1) or disables (0) ICAP support with Data Security Suite (DSS). See <i>Working With Web</i> <i>DLP</i> .
proxy.config.icap.ICAPUri	STRING	Default: NULL
		The Uniform Resource Identifier for the ICAP service.
		A backup server can be specified in a comma-separated list.
		Obtain the identifier from your Forcepoint DLP administrator. Enter the URI in the following format:
		icap://hostname:port/path
		<hostname> is the IP address or hostname of the Protector appliance.</hostname>
		<port> is 1344 by default.</port>
		<path> is the path of the ICAP service on the host machine.</path>
		For example:
		icap://ICAP_machine:1344/ opt/icap_services
		You do not need to specify the port if you
		are using the default ICAP port 1344.
proxy.config.icap.FailOpen	INT	Default: 1
		 1 allows traffic when the ICAP servers are down
		 0 sends a block page if the ICAP servers are down
proxy.config.icap.	INT	Default: 0
BlockHugeContent		 0 sends a block page if a file larger than the Forcepoint DLP size limit (default 50 MB) is sent. 1 allows traffic

Configuration Variable	Data Type	Description
proxy.config.icap. AnalyzeSecureContent	INT	Default: 1
		to its destination.
		 1 sends decrypted traffic to Forcepoint DLP for analysis.
proxy.config.icap.AnalyzeFTP	INT	Default: 1
		When enabled (1), send native FTP upload file transfers to ICAP server for analysis.
proxy.config.icap. ActiveTimeout	INT	Default: 5
		The read/response timeout in seconds. The activity is considered a failure if the timeout is exceeded.
proxy.config.icap.RetryTime	INT	Default: 5
		The recovery interval, in seconds, to test whether a down server is back up.
proxy.config.icap.LoadBalance	INT	Default: 1
		When to ICAP servers are specified:
		 1 distributes requests to all available servers
		 0 sends requests to only the primary server.
proxy.config.icap. MaxConnection	INT	Default: 5
		When the ICAP server does not specify the maximum number of connections that the ICAP client can have, this value is used.
		Valid values are 1 - 100.

Related information Working With Web DLP on page 101

Web DLP

Configuration Variable	Data Type	Description
proxy.config.dss.enabled	INT	Default: 0 Enables (1) or disables (0) support for on- box Web DLP. See <i>Working</i> <i>With Web DLP</i> .
proxy.config.dss.AnalyzeFTP	INT	Default: 1 When enabled (1), send native FTP upload file transfers to the on-box Web DLP policy engine for analysis.
proxy.config.dss. AnalyzeSecureContent	INT	 Default: 1 0 sends decrypted traffic directly to its destination. 1 sends decrypted traffic to Forcepoint DLP for analysis.
proxy.config.dss.analysis_ timeout	INT	Default: 10000 The maximum length of time, in milliseconds, that a single file analysis can take before analysis is aborted.
proxy.config.dss.UsingLoginID	INT	Default: 0 Enables (1) or disables (0) sending Login ID rather than full user name to Forcepoint DLP. This variable must be added manually.
proxy.config.dss.domain_user_ format	INT	Default: 0 When proxy.config.dss.UsingLoginID is enabled, enables (1) or disables (0) the domain/user format. This variable must be added manually.
proxy.config.dss.large_file_ threshold	INT	Default: 5 (MB) Determines how large a file should be before a longer period of time than the current default of 10 seconds is given for analysis time. Files that exceed this size are give the time set in proxy.config.dss. analysis_timeout_for_large_file.

Configuration Variable	Data Type	Description
proxy.config.dss.analysis_ timeout_for_large_file	INT	Default: 20 (seconds) Determines the period of time given for file analysis to files that exceed the size set in proxy.config.dss.large_file_threshold.
proxy.config.dss_and_icap. enabled	INT	Default: 0 Enables (1) the selection of both ICAP and Web DLP as integration options when Content Gateway is deployed with the DLP Module. Only one option can be selected when disabled (0).

Related information Working With Web DLP on page 101

Connectivity, analysis, and boundary conditions

Configuration Variable	Data Type	Description
wtg.config.subscription_key	STRING	Default: NULL
		The Forcepoint Web Security subscription key value.
wtg.config.download_server_ip	STRING	Default: download.forcepoint.com
		The hostname or IP address of the download server.
wtg.config.download_server_ port	INT	Default: 443
		The port number of the download server. Download also does a license check.
wtg.config.policy_server_ip	STRING	The IP address of the Policy Server.
wtg.config.policy_server_port	INT	Default: 55806
		The port number of the Policy Server.
wtg.config.wse_server_ip	STRING	The IP address of the Filtering Service.

Configuration Variable	Data Type	Description
wtg.config.wse_server_port	INT	Default: 15868 The port number of the Filtering Service WISP interface.
wtg.config.wse_server_timeout	INT	Default: 5000 The maximum timeout period, in milliseconds, for communication with Filtering Service.
wtg.config.ssl_bypassed_ categories	STRING	Default: NULL A list of category identifiers that will bypass SSL decryption. Do not change the value of this variable. It is included strictly as a troubleshooting aid. Use the Web Security module of the Forcepoint Security Manager to specify categories to bypass SSL decryption.
wtg.config.ssl_decryption_ bypass_ip_based	INT	Default: 0 Whether the SSL category bypass process uses only the IP address (not the hostname) when performing a category lookup. 0 = disabled 1 = enabled
wtg.config.ssl_fail_open	INT	Default: 0 Whether SSL sites are decrypted if Filtering Service becomes unreachable. 0 = all SSL sites are decrypted when Filtering Service is unreachable. 1 = no SSL sites are decrypted when Filtering Service is unreachable
wtg.config.fail_open	INT	Default: 1 Whether Content Gateway permits or blocks requests when Filtering Service is unavailable. • 0 sends a block page • 1 permits the request

Configuration Variable	Data Type	Description
wtg.config.fail_open_analytic_ scan	INT	 Default: 1 Specifies how Content Gateway behaves should analytic scanning become non- functional or exceeds the maximum scan time. Set to: 0 to block traffic 1 to perform a lookup in the URL database and apply policy Note An alarm is raised whenever analytics scanning becomes non-functional.
wtg.config.fail_open_analytic_ scan_size_exceeded	INT	 Default: 0 How Content Gateway handles files that are not fully scanned because they exceed the Scan Size Limit set in the Forcepoint Security Manager. 0 blocks access to the file 1 permits access to the file
wtg.config.archive_depth	INT	Default: 5 The maximum depth of analysis performed on archive files.
wtg.config.max_decompressions	INT	Default: 10 The maximum number of total decompressions to be performed on archive files (per transaction). The value should not exceed 25.
wtg.config.max_subsamples	INT	Default: 10000 The maximum number of discrete files within an archive file that Content Gateway may decompress and analyze to classify a given transaction.
wtg.config.zipbomb_action	INT	Default: 1 For internal use. Indicates zip bomb analysis status. Do not change the value of this variable.

Configuration Variable	Data Type	Description
wtg.config.rdnsclients	INT	Default: 0
		Enables (1) or disables (0) logging of clients' hostnames in the log records via reverse DNS.
wtg.config.ip_ranges_not_to_ scan	STRING	Default: 10.0.0.0-10.255.255.255,
		172.16.0.0-172.31.255.255,
		192.168.0.0-192.168.255.255
		Internal IP address ranges not to scan. By default, the list is the standard private non- routable IP addresses. Address ranges are hyphenated with each range separated by a comma.
		This is especially helpful in explicit proxy deployments in which a PAC file is not used and you want to exclude the standard internal IP addresses from being scanned.
wtg.config.scan_ip_ranges	INT	Default: 1
		Enables (1) or disables (0) bypass of the internal IP address ranges specified in wtg.config.
		ip_ranges_not_to_scan. See above.
wtg.config.feedback.enabled	INT	Default: 1
		Enables (1) or disables (0) analytic/ category feedback to Forcepoint.
		Set at install time.
wtg.config.scan_uncat_block	INT	Default: 1
		Enables (1) or disables (0) the scanning of blocked, uncategorized URLs.
wtg.config.filter_unknown_file	INT	Default: 0
		When enabled (1), unknown is sent as a valid file type to Filtering Service.
wtg.config.respond_with_303_	INT	Default: 0
on_redirect		When enabled (1), user requests that are not GET or POST are blocked as expected when scanning is enabled, regardless of the User-agent in use.

remap.config

The **remap.config** file contains mapping rules that Content Gateway uses to redirect HTTP requests permanently or temporarily without Content Gateway having to contact any origin server:



Important

After you modify this file, restart the proxy or run the following command from the Content Gateway **bin** directory (/opt/WCG/bin) to apply the changes:

content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Format

Each line in the **remap.config** file must contain a mapping rule. Content Gateway recognizes three spacedelimited fields: type, target, and replacement. The following table describes the format of each field.

Field	Description
type	Enter one of the following:
	 map provides the same function as redirect. Use redirect instead.
	 redirect: redirects HTTP requests permanently without having to contact the origin server. Permanent redirects notify the browser of the URL change (by returning an HTTP status code 301) so that the browser can update bookmarks.
	 redirect_temporary: redirects HTTP requests temporarily without having to contact the origin server. Temporary redirects notify the browser of the URL change for the current request only (by returning an HTTP status code 307).
	Note reverse_map is not supported.
target	Enter the origin or <i>from</i> URL. You can enter up to four components:
	<pre>scheme://host:port/path_prefix</pre>
	<scheme> can be http, https, or ftp.</scheme>

Field	Description
strict URL matching flag	Enable Match URL Exactly to force matching to be exact against the entire requested URL.
	Without this option, the URL is compared up to the end of the target (From Path Prefix). If there is a match, the redirect is applied. This can cause unwanted matching, when the redirect URL includes the base URL. See <i>Mapping and Redirection</i> .
replacement	Enter the destination or <i>to</i> URL. You can enter up to four components:
	<pre>scheme://host:port/path_prefix <scheme> can be http, https, or ftp.</scheme></pre>



The scheme type (HTTP, HTTPS, FTP) of the target and replacement must match.

Related reference

Note

Mapping and Redirection on page 298

Examples

The following rule **permanently** redirects all HTTP requests for www.company.com to www.company2.com:

redirect http://www.company.com http://www.company2.com

The following rule temporarily redirects all HTTP requests for www.company1.com to www.company2.com:

redirect_temporary http://www.company1.com http://www.company2.com

socks.config

The socks.config file specifies:

- SOCKS servers that the proxy must use to access specific origin servers, and the order in which the proxy
 goes through the SOCKS server list.
- Origin servers that Content Gateway accesses directly, without going through a SOCKS server.



Note

It is recommended that all SOCKS configuration be performed in the Content Gateway manager.



Important

After you modify this file, you must restart the proxy.

Traffic that does not match a manually configured rule is handled via a default rule. A default rule is constructed for each SOCKS server with the **default** option enabled in the **Socks Servers** table. Default rules are created automatically and displayed on the SOCKS Server page. Default rules are not written in the **socks.config** file. The destination IP address is "All."

Format

To specify SOCKS servers that the proxy must use to reach specific origin servers, add rules to the **socks.config** file in the following format:

dest_ip=<ipaddress> socksparent="<alias1>" [round_ robin=<value>]

Here:

<ipaddress> is the origin server IP address or range of IP addresses separated by - or /.

<alias1> is the alias name of the SOCKS server named in the SOCKS Servers list.

<value> is either strict if you want Content Gateway to try the SOCKS servers one by one, or false if you do not want round-robin selection to occur.

To specify origin servers that you want Content Gateway to access directly, *without* going through the SOCKS servers, enter a rule in **socks.config** in the following format:

```
no_socks <ipaddress>
```

Here, *<ipaddress>* is a comma-separated list of the IP addresses or IP address ranges associated with the origin servers that you want Content Gateway to access directly. Do not specify the all networks broadcast address: 255.255.255.255.



Note

Each rule in **socks.config** can consist of a maximum of 400 characters. The order of the rules in the **socks.config** file is not significant.

Examples

The following example configures the proxy to send requests to the origin servers associated with the range of IP addresses 123.15.17.1 - 123.14.17.4 through the SOCKS server aliases "alias1" and "alias2." Because the optional specifier **round_ robin** is set to **strict**, the proxy sends the first request to alias1, the second request to alias2, the third request to alias1, and so on.

dest_ip=123.14.15.1 - 123.14.17.4

socksparent="alias; alias2" round_robin=strict

The following example configures the proxy to access the origin server associated with the IP address 11.11.11.1 directly, without going through the SOCKS server:

no_socks 11.11.11.1

The following example configures Content Gateway to access the origin servers associated with the range of IP addresses 123.14.15.1 - 123.14.17.4 and the IP address 113.14.18.2 directly, without going through the SOCKS server:

no_socks 123.14.15.1 - 123.14.17.4, 113.14.18.2

socks_server.config

The socks_server.config file specifies the SOCKS servers available to Content Gateway.

Format

To specify SOCKS servers use the following format:

```
alias=<name> host=<IP_address|domain_name> port=<port> [username=<user_name>
password=<password>] default=true|false
```

Here:

<name> is the name of a SOCKS server.

<IP_address | domain_name> is an IP address or a domain name that can be resolved by your DNS service.

<port_number> is the port on which the SOCKS server is listening.

<username> and <password> are the username/password pair for SOCKS 5 authentication. The password is encrypted.

Set default to **true** to make the specified server a default SOCKS server. When the default server option is on, the SOCKS server is used when no SOCKS rule matches.

If no SOCKS server is designated a default server, traffic that doesn't match a rule is not routed through a SOCKS server.

Examples

This example adds the SOCKS server "default1" at 127.0.0.1 on port 61080. It is designated a default SOCKS server.

alias=default1 host=127.0.0.1 port=61080 default=true

This example adds a SOCKS server that uses authentication. Note that the password ("465751475058") is not the real password. It is encrypted.

```
alias=test1 host=socks5.example.com port=1080 username=test password=465751475058 default=false
```

If this file is modified, you must restart Content Gateway.

Note

Each rule in socks_server.config cannot exceed 400 characters.

splitdns.config

The **splitdns.config** file enables you to specify the DNS server that Content Gateway should use for resolving hosts under specific conditions.

To specify a DNS server, you must supply the following information in each active line within the file:

- A primary destination specifier in the form of a destination domain, a destination host, or a URL regular expression
- A set of server directives, listing one or more DNS servers with corresponding port numbers

You can also include the following optional information with each DNS server specification:

- A default domain for resolving hosts
- A search list specifying the domain search order when multiple domains are specified

For more information, see Using the Split DNS option.



Note

After you modify this file, restart the proxy or run the following command from the Content Gateway **bin** directory (/opt/WCG/bin) to apply the changes:

content_line -x

When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Related tasks

Using the Split DNS option on page 163

Format

Each line in the **splitdns.config** file uses one of the following formats:

dest_domain=dest_domain | dest_host | url_regex named=dns_ server

def_domain=def_domain search_list=search_list

The following table describes each field.

Field	Allowed Value
dest_domain	A valid domain name. This specifies that the DNS server selection be based on the destination domain. You can prefix the domain with an exclamation mark (!) to indicate the NOT logical operator.
dest_host	A valid hostname. This specifies that the DNS server selection be based on the destination host. You can prefix the host with an exclamation mark (!) to indicate the NOT logical operator.
url_regex	A valid URL regular expression. This specifies that the DNS server selection be based on a regular expression.
	See Specifying URL regular expressions (url_regex) for information about using regular expressions.

Field	Allowed Value
dns_server	This is a required directive. It identifies the DNS server for Content Gateway to use with the destination specifier. You can specify a port using a colon (:). If you do not specify a port, 53 is used. You can specify multiple DNS servers separated by spaces or by semicolons (;). You must specify the domains using IP addresses in dot notation.
def_domain	A valid domain name. This optional directive specifies the default domain name to use for resolving hosts. Only one entry is allowed. If you do not provide the default domain, the system determines its value from / etc/resolv.conf.
search_list	A list of domains separated by spaces or semicolons (;). This specifies the domain search order. If you do not provide the search list, the system determines the value from /etc/resolv.conf .

Related reference

Specifying URL regular expressions (url_regex) on page 364

Examples

Consider the following DNS server selection specifications:

dest_domain=internal.company.com named=255.255.255.255.255.255.255.255.254
def_domain=company.com search_list=company. com company1.com

dest_domain=!internal.company.com named=255.255.255.253

Now consider the following two requests:

http://minstar.internal.company.com

This request matches the first line and select DNS server 255.255.255.255 on port

212. All resolver requests will use **company.com** as the default domain, and **company.com** and **company1.com** as the set of domains to search first.

http://www.microsoft.com

This request will match the second line. Therefore, Content Gateway selects DNS server 255.255.255.253. No **def_domain** or **search_list** was supplied, so Content Gateway retrieves this information from **/etc/resolv.conf**.

wccp.config

The **wccp.config** file stores the WCCP configuration information and service group settings. When WCCP is enabled on the **Configure > MyProxy > Basic** page, WCCP service group settings can be configured on the

Configure > Networking > WCCP page. Service groups must be defined if WCCP is to be used for transparent redirection to Content Gateway.

For more information, see *Transparent interception with WCCP v2 devices*.

Related concepts Transparent interception with WCCP v2 devices on page 38

Appendix F Content Gateway Error Messages

Contents

- Error messages in log files on page 467
- Content Gateway alarm messages on page 469
- Content Gateway HTML messages sent to clients on page 472
- Content Gateway standard HTTP response messages on page 475

Error messages in log files

The following table lists messages that can appear in system log files. This list is not exhaustive; it describes warning messages that can occur and might require your attention. For information about warning messages not included in the list below, go to www.forcepoint.com and then navigate to Support and Knowledge Base.

Process fatal errors

Message	Description
Accept port is not between 1 and 65535. Please check configuration.	The port specified in the records.config file that accepts incoming HTTP requests is not valid.
Ftp accept port is not between 1 and 65535.	The port specified in the records.config file that accepts incoming FTP requests is not valid.
Self loop is detected in parent proxy configuration.	The name and port of the parent proxy are the same as that of Content Gateway. This creates a loop when Content Gateway attempts to send requests to the parent proxy.
Could not open the ARM device	The ARM failed to load. The most common reason for this is that the host system has an incompatible system kernel. To see if the ARM is loaded, run: /sbin/lsmod grep arm
content_manager failed to set cluster IP address	The content_manager process could not set the cluster IP address. Check the cluster IP address. Make sure that it is not already used by another device in the network.

Warnings

Message	Description
Logfile error: error_number	Generic logging error.
Bad cluster major version range	Incompatible software versions causing a problem.
version1-version2 for node	
<i>IP address</i> connect failed	
can't open config file <i>filename</i> for reading custom formats	Custom logging is enabled, but Content Gateway cannot find the logs.config file.
connect by disallowed client	The specified client is not allowed to connect to
IP address, closing connection	Content Gateway. The client IP address is not listed in the ip_allow.config file.
Could not rename log <i>filename</i> to	System error when renaming log file during roll.
rolled filename	
Did <i>this_amount</i> of backup still to do	Congestion is approaching.
remaining_amount	
Different clustering minor versions <i>version 1, version 2</i> for node <i>IP address</i> continuing	Incompatible software versions causing a problem.
log format symbol symbol_name not found	Custom log format references a field symbol that does not exist. See <i>Event Logging Formats</i> .
missing field for field marker	Error reading a log buffer.
Unable to accept cluster connections on port: <i>cluster_port_number</i>	Contact Technical Support. Go to support.forcepoint.com for Technical Support contact information
Unable to open log file <i>filename</i> , errno=error_number	Cannot open the log file.
Error accessing disk <i>disk_name</i>	Content Gateway might have a read problem. You might have to replace the disk.
Too many errors accessing disk	Content Gateway is not using the disk because it
<i>disk_name:</i> declaring disk bad	encountered too many errors. The disk might be corrupt and might have to be replaced.
Missing DC parameter	A required parameter was not specified. Please
<missing_param> on auth.profile line</missing_param>	provide a value for the missing parameter.
Bad DC parameter <bad_param> -</bad_param>	A specified Domain Controller parameter is invalid.
<dc_name></dc_name>	Please enter a valid value for the cited parameter.
[ParentSelection] <error_description> for default parent proxy</error_description>	Proxy chaining is not working due to misconfiguration of the parent proxy in the child proxy. Please check the chaining configuration of parent proxy values in the child proxy.
Message	Description
--	--
WCCP2: Cannot find Interface name. Please check that the variable proxy.local.wccp2. ethernet_interface is set correctly	No value is specified for the WCCP interface. In the Content Gateway manager, check Configure > Networking > WCCP > General , or assign a value to proxy.local.wccp2.ethernet_interface in records.config .
ARMManager: Unable to read network interface configuration	There is a format or configuration error in ipnat.conf . In the Content Gateway manager, go to Configure > Networking > ARM > General and click Edit File to view and correct ipnat.conf .

Related information

Event Logging Formats on page 355

Content Gateway alarm messages

The following table describes alarm messages that you may see in the Content Gateway manager.

Message	Description/Solution
The Content Gateway subscription has expired.	Please refer to the Knowledge base article 33797.
Content Gateway subscription download failed.	Content Gateway was unable to connect to the download server to verify the subscription information. Please check your connection to the download server.
After several attempts, Content Gateway failed to connect to the Database Download Service. Please troubleshoot the connection.	Verify that Content Gateway is able to access the Internet. Check firewall and upstream proxy server settings that might prevent Content Gateway from connecting to the download server.
After several attempts, Content Gateway failed to connect to the Policy Server. Please troubleshoot the connection.	Verify that there is network connectivity between Content Gateway and the Policy Server machine. Sometimes firewall settings block connectivity. Also confirm that Policy Server is running.
After several attempts, Content Gateway failed to connect to the Policy Broker. Please troubleshoot the connection.	Verify that there is network connectivity between Content Gateway and Policy Broker. Sometimes firewall settings block connectivity. Also confirm that Policy Broker is running.
After several attempts, Content Gateway failed to connect to Filtering Service. Please troubleshoot the connection. Please refer to the Knowledge base article 41457.	Verify that there is network connectivity between Content Gateway and the Filtering Service machine. Sometimes firewall settings block connectivity. Also confirm that Filtering Service is running.
Communication with the analytics engine has failed. Please restart Content Gateway.	Restart Content Gateway.
SSL decryption has been disabled due to an internal error, please restart Content Gateway.	There was a fatal error in SSL Support. Please restart Content Gateway.

Message	Description/Solution
[Rollback::Rollback] Config file is read-only: <i>filename</i>	Go to the Content Gateway config directory (default location is /opt/WCG/config) and check the indicated file permissions; change them if necessary.
[Rollback::Rollback] Unable to read or write config file <i>filename</i>	Go to the Content Gateway config directory and make sure the indicated file exists. Check its permissions and change them if necessary.
[Content Gateway Manager] Configuration File Update Failed <i>error_number</i>	Go to the Content Gateway config directory and check the indicated file permissions; change them if necessary.
Access logging suspended - configured space allocation exhausted.	The space allocated to the event log files is full. You must either increase the space or delete some log files to enable access logging to continue. To prevent this from happening, consider rolling log files more frequently and enabling the autodelete feature. See <i>Rolling event log files</i> .
Access logging suspended - no more space on the logging partition.	The entire partition containing the event logs is full. You must delete or move some log files to enable access logging to continue. To prevent this from happening, consider rolling log files more frequently and enabling the autodelete feature. See <i>Rolling event</i> <i>log files</i> .
Created zero length place holder for config file <i>filename</i>	Go to the Content Gateway config directory and check the indicated file. If it is indeed zero in length, use a backup copy of the configuration file.
Content Gateway can't open <i>filename</i> for reading custom formats	Make sure that the <i>proxy.config.log2.config_file</i> variable in the records.config file contains the correct path to the custom log configuration file (the default is logging/logs.config).
Content Gateway could not open logfile <i>filename</i>	Check permissions for the indicated file and the logging directory.
Content Gateway failed to parse line <i>line_number</i> of the logging config file <i>filename</i>	Check your custom log configuration file. There may be syntax errors. See <i>Custom logging fields</i> for correct custom log format fields.
vip_config binary is not setuid root, manager will be unable to enable virtual ip addresses	The content_manager process is not able to set virtual IP addresses. You must setuid root for the vip_config file in the Content Gateway bin directory.
Content Gateway cannot parse the ICAP URI. Please ensure that the URI is entered correctly in Content Gateway Manager or in the <i>proxy.config.icap.ICAPUri</i> configuration variable.	The Universal Resource Identifier (URI) is not in the correct format. Enter the URI as follows: icap://hostname:port/path See <i>Working With Web DLP</i> for additional details on the format of the URI.

Message	Description/Solution
The specified ICAP server does not have a DNS entry. Please ensure that a valid DSS hostname is entered correctly in Content Gateway Manager or in the <i>proxy.config.icap.ICAPUri</i> configuration variable.	The hostname in the records.config file does not match any entries in the DNS. Ensure that the name of a valid Forcepoint DLP server is entered correctly in the Content Gateway manager.
	See <i>Working With Web DLP</i> for information about the format of the URI.
Content Gateway is not able to communicate with the DSS server. Please try again.	Ensure that the Forcepoint management server is up and running, and accepting connections on the port specified in the <i>proxy.config.icap.ICAPUri</i> variable. Contact your Forcepoint DLP administrator if this message persists.
Domain controller <i>domain_controller_name:port</i> is down.	The named NTLM domain controller is not responding to requests and has been marked as down. Investigation the status of the domain controller.
Windows domain [domain name] unreachable or bad	This alarm can indicate any of the following:
membership status	 The Active Directory is unreachable. The AD server is either down or there is a network connectivity problem.
	2) The AD is reachable, but there is a configuration problem that prevents it from communicating with Content Gateway. For example, the alarm is generated if the AD has multiple Sites and the subnet that Content Gateway resides on has not been added to one of them.
The Scanning Data Files Update option (My Proxy > Subscription) is set to 'suspend updates'. To get the best protection, set it to 'no delay', or, on a backup system, use a time- based option.	This alarm is a reminder that downloads of the security scanning data files used by Content Gateway analysis has been suspended.
	It is recommended that you not clear this alarm until the delay time has been reset.
Port Mirroring cannot work unless SSL decryption is	(Appliance deployments only)
enabled also. Please enable SSL decryption (HTTPS) if you want to use the Port Mirroring feature.	Ensure the SSL decryption (HTTPS) is enabled before attempting to use Port Mirroring.
The mirror interface <int> cannot be connected for Port</int>	(Appliance deployments only)
Mirroring. Please check the interface configuration or edit the interface value.	The interface configured for Port Mirroring is not valid, is not active, or requires configuration.
An unexpected error has occurred with the Office365 Bypass feature. Please refer to the Knowledge base article 36961.	A processing error has occurred and requests to Office 365 products may not be bypassing authentication or the proxy. Technical Support assistance is required to correct the problem.
An error message has occurred when you exceeded the 1000 concurrent connection limit	Please refer to the Knowledge base article 41458.

Related concepts

Rolling event log files on page 227

Related reference

Custom logging fields on page 355

Related information

Working With Web DLP on page 101

Content Gateway HTML messages sent to clients

Content Gateway returns detailed error messages to browser clients when there are problems with the HTTP transactions requested by the browser. These response messages correspond to standard HTTP response codes, but provide more information. A list of the more frequently encountered HTTP response codes is provided in *Content Gateway standard HTTP response messages*. You can customize the response messages.

The following table lists the Content Gateway hard-coded HTTP messages, their corresponding HTTP response codes, and their corresponding customizable files.

HTTP Code and Title	Description	Customizable Filename
403	You are not allowed to access the	access#denied
Access Denied	document at location URL.	
400	Bad HTTP request for FTP object.	ftp#bad_request
Bad HTTP request for FTP Object		
500	Error reading from cache. Please	cache#read_error
Cache Read Error	retry request.	
504	Server has not sent any data for too	timeout#inactivity
Connection Timed Out	long a time.	
400	Could not process this request	request#no_content_length
Content Length Required	specified.	
400	Your request is prohibited because	request#cycle_detected
Cycle Detected	It would cause an HTTP proxy cycle.	

HTTP Code and Title	Description	Customizable Filename
403 Forbidden	port_number is not an allowed port for SSL connections. (You have made a request for a secure SSL connection to a forbidden port number.)	access#ssl_forbidden
401 FTP Authentication Required	You need to specify a correct user name and password to access the requested FTP document URL.	ftp#auth_required
502 FTP Connection Failed	Could not connect to the server server_name.	connect#failed_connect
502 FTP Error	The FTP server server_name returned an error. The request for document URL failed.	ftp#error
400 Host Header Required	An attempt was made to transparently proxy your request, but this attempt failed because your browser did not send an HTTP Host header.	interception#no_host
	Manually configure your browser to use the following URL as an HTTP proxy:	
	https:// proxy_name:proxy_port	
	See your browser's documentation for details.	
	Alternatively, end users can upgrade to a browser that supports the HTTP Host header field.	
400 Host Header Required	Your browser did not send a Host HTTP header field and therefore the virtual host being requested could not be determined. To access this website, upgrade to a browser that supports the HTTP Host header field.	request#no_host
505 HTTP Version Not Supported	The origin server server_name is using an unsupported version of the HTTP protocol.	response#bad_version
400 Invalid HTTP Request	Could not process this client_request HTTP method request for URL.	request#syntax_error
502 Invalid HTTP Response	The host server_name did not return the document URL correctly.	response#bad_response
502 Malformed Server Response	The host server_name did not return the document URL correctly.	response#bad_response

HTTP Code and Title	Description	Customizable Filename
502 Malformed Server Response Status	The host server_name did not return the document URL correctly.	response#bad_response
504 Maximum Transaction Time exceeded	Too much time has passed transmitting document URL.	timeout#activity
502 No Response Header From Server	The host server_name did not return the document URL correctly.	response#bad_response
504 Not Cached	This document was not available in the cache, and you (the client) accept cached copies only.	cache#not_in_cache
404 Not Found on Accelerator	The request for URL on host server_name was not found. Check the location and try again.	urlrouting#no_mapping
502 NULL	The host hostname did not return the document UR <i>L</i> correctly.	response#bad_response
407 Proxy Authentication Required	Please log in with user name and password.	access#proxy_auth_required
502 Server Hangup	The server hostname closed the connection before the transaction was completed.	connect#hangup
302 Temporarily Moved	The document you requested, URL, has moved to a new location. The new location is new_URL.	redirect#moved_temporarily
406 Transcoding Not Available	Unable to provide the document URL in the format requested by your browser.	transcoding#unsupported
502 Tunnel Connection Failed	Could not connect to the server hostname.	connect#failed_connect
502 Unknown Error	The host hostname did not return the document URL correctly.	response#bad_response
500 Unknown Host	Unable to locate the server named hostname. The server does not have a DNS entry. Perhaps there is a misspelling in the server name or the server no longer exists. Double- check the name and try again.	connect#dns_failed
400 Unsupported URL Scheme	Cannot perform your request for the document URL because the protocol scheme is unknown.	request#scheme_unsupported

Related reference

Content Gateway standard HTTP response messages on page 475

Content Gateway standard HTTP response messages

The following standard HTTP response messages are provided for your information. For a more complete list, see the *Hypertext Transfer Protocol — HTTP/1.1 Specification*.

Message	Description
200	ОК
202	Accepted
204	No Content
206	Partial Content
300	Multiple Choices
301	Moved Permanently
302	Found
303	See Other
304	Not Modified
400	Bad Request
401	Unauthorized; retry
403	Forbidden
404	Not Found
405	Method Not Allowed
406	Not acceptable
408	Request Timeout
500	Internal server error
501	Not Implemented
502	Bad Gateway
504	Gateway Timeout