



Online Help

Websense Content Gateway

v7.5

Websense Content Gateway Online Help

April 2010

R050410750

Copyright © 1996-2010 Yahoo, Inc., and Websense, Inc. All rights reserved.

This document contains proprietary and confidential information of Yahoo, Inc and Websense, Inc. The contents of this document may not be disclosed to third parties, copied, or duplicated in any form, in whole or in part, without prior written permission of Websense, Inc.

Websense, the Websense Logo, Threatseeker and the YES! Logo are registered trademarks of Websense, Inc. in the United States and/or other countries. Websense has numerous other unregistered trademarks in the United States and internationally. All other trademarks are the property of their respective owners.

Every effort has been made to ensure the accuracy of this manual. However, Websense Inc., and Yahoo, Inc. make no warranties with respect to this documentation and disclaim any implied warranties of merchantability and fitness for a particular purpose. Websense Inc. 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.

Traffic Server is a trademark or registered trademark of Yahoo! Inc. in the United States and other countries.

Red Hat is a registered trademark of Red Hat Software, Inc.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, Windows NT, and Active Directory are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation in the United States and in other countries.

UNIX is a registered trademark of AT&T.

All other trademarks are property of their respective owners.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Websense, Inc, 10240 Sorrento Valley Parkway, San Diego, CA 92121.

Portions of Websense Content Gateway include third-party technology used under license. Notices and attribution are included elsewhere in this manual.

Contents

Topic 1	Overview	1
	Deployment options	1
	As a Web proxy cache	1
	In a cache hierarchy	2
	In a cluster	2
	As an SSL server	2
	Components	3
	Cache	3
	RAM cache	3
	Adaptive Redirection Module	3
	DNS resolver	3
	Processes	4
	Administration tools	4
	Traffic analysis options	5
	Security options	5
	Online Help	6
	Technical Support	7
Topic 2	Getting Started	9
	Starting and stopping Content Gateway	9
	Starting Content Gateway Manager	10
	Entering your subscription key	11
	Providing system information	11
	Verifying that the proxy is running	11
	Using the command-line interface	12
Topic 3	Web Proxy Caching	15
	Cache requests	15
	Ensuring cached object freshness	16
	HTTP object freshness	16
	FTP object freshness	20
	Scheduling updates to local cache content	21
	Configuring the Scheduled Update option	22
	Forcing an immediate update	23
	Pinning content in the cache	23
	Setting cache pinning rules	23
	Enabling cache pinning	24
	To cache or not to cache?	24
	Caching HTTP objects	24
	Client directives	25

	Origin server directives	26
	Configuration directives	28
	Forcing object caching	29
	Caching HTTP alternates	30
	Configuring how Content Gateway caches alternates.	30
	Limiting the number of alternates for an object	31
	Caching FTP objects.	31
	Disabling FTP over HTTP caching.	32
	Disabling FTP proxy object caching.	32
	Disabling caching of full or simple directory listings.	32
Topic 4	Explicit Proxy Caching	35
	Manual browser configuration	35
	Using a PAC file.	36
	Sample PAC file	37
	Using WPAD	37
	Configuring FTP clients in an explicit proxy environment	39
Topic 5	Transparent Proxy Caching and ARM.	43
	Enabling the ARM option	44
	Interception strategies.	45
	Using a Layer 4 switch	45
	Using a WCCP-enabled router	46
	Using multicast mode	55
	Using policy-based routing	56
	Using software-based routing	56
	Interception bypass	57
	Dynamic bypass rules	58
	Static bypass rules.	59
	Viewing the current set of bypass rules	60
	Configuring ARM security.	60
	Connection load shedding	61
	Reducing DNS lookups	61
	IP spoofing	62
	IP spoofing and the flow of HTTP traffic	63
	Enabling IP spoofing:	64
Topic 6	Clusters.	65
	Management-only clustering.	66
	Full clustering	66
	Changing clustering configuration	67
	Adding nodes to a cluster.	68
	Deleting nodes from a cluster	70

	Virtual IP failover	70
	What are virtual IP addresses?	70
	Enabling and disabling virtual IP addressing	71
	Adding and editing virtual IP addresses	71
Topic 7	Hierarchical Caching	73
	HTTP cache hierarchies	73
	Parent failover	74
	Configuring Content Gateway to use an HTTP parent cache	74
	ICP cache hierarchies	75
	Configuring Content Gateway to use an ICP cache hierarchy	75
Topic 8	Configuring the Cache	77
	Changing cache capacity	77
	Querying cache size	78
	Increasing cache capacity	78
	Reducing cache capacity	79
	Partitioning the cache	80
	Making changes to partition sizes and protocols	80
	Partitioning the cache according to origin server or domain	80
	Configuring cache object size limit	81
	Clearing the cache	82
	Changing the size of the RAM cache	82
Topic 9	Configuring the System	85
	Content Gateway Manager	85
	Starting Configure mode	85
	Using Configure mode	86
	Command-line interface	89
	Configuration files	89
	Saving and restoring configurations	90
	Taking configuration snapshots	91
	Restoring configuration snapshots	92
Topic 10	Monitoring Traffic	93
	Viewing statistics	93
	Starting Monitor mode	93
	Using Monitor mode	94
	Viewing statistics from the command line	96
	Working with alarms	97
	Clearing alarms	98
	Configuring Content Gateway to email alarms	98
	Using a script file for alarms	98
	Using Performance graphs	98

	Creating reports with SSL Manager	99
	Certificate Authorities	100
	Incidents	101
Topic 11	Working With Websense Data Security	105
	Registering and configuring on-box Data Security	106
	Unregistering on-box Data Security	108
	Stopping and starting Data Security processes	108
	Configuring the ICAP client	109
Topic 12	Working With Encrypted Data	111
	Running in explicit proxy mode	112
	Enabling SSL Manager	114
	Tasks	114
	Certificates	115
	Internal Root CAs	115
	Importing your Root CA	116
	Creating your new Root CA	116
	Creating a sub-certificate authority	117
	Backing up your internal Root CA	119
	Managing certificates	119
	View a certificate	119
	Delete a certificate	120
	Change the allow/deny status of a certificate	120
	Adding new certificate authorities	120
	Backing up certificates	121
	Restoring certificates	121
	Decryption and Encryption	121
	Configuring SSL Manager for inbound traffic	122
	Configuring SSL Manager for outbound traffic	122
	Validating certificates	124
	Configuring validation	124
	Bypassing verification	126
	Keeping revocation information up to date	127
	Certificate revocation lists	127
	Online certification status protocol (OCSP)	127
	Managing Web HTTPS site access	128
	Viewing incidents	128
	Changing the status of an incident	130
	Deleting an incident	130
	Changing the text of a message	130
	Viewing incident details	131
	Adding Web sites to the incident list	131

	Client certificates	132
	When a client certificate is requested	132
	Importing client certificates	132
	When a client certificate is always required: the hostlist	133
	Deleting client certificates.	133
	Configuring logging	133
	How long should log files be kept?	134
	How big can log files grow?	134
	What fields should appear in the access log files?	135
	Customizing messages	136
	Certificate validation failed.	136
	Connection failure.	137
Topic 13	Security	139
	Controlling client access to the proxy cache	139
	Controlling host access to the proxy server	140
	Controlling access to Content Gateway Manager	141
	Setting the administrator ID and password	142
	Creating a list of user accounts	142
	Controlling host access to Content Gateway Manager	143
	Using SSL for secure administration	143
	Filtering Rules	144
	Creating filtering rules	144
	Configuring SOCKS firewall integration.	146
	Configuring the proxy to use a SOCKS firewall.	147
	Setting SOCKS proxy options	147
	Setting authentication.	148
	Setting SOCKS server bypass	148
	Example.	149
	Using the Split DNS option	149
	Configuring proxy authentication	150
	Browser limitations.	150
	Transparent proxy authentication.	151
	Using LDAP proxy authentication	152
	Using RADIUS proxy authentication	155
	Using NTLM proxy authentication.	158
	Multiple authentication realms	161
Topic 14	Working With Log Files.	167
	Event log files.	168
	Managing event log files	169
	Choosing the logging directory.	169
	Controlling logging space	169

Event log file formats	170
Using standard formats	171
Custom format	171
Choosing binary or ASCII	174
Using logcat to convert binary logs to ASCII	175
Rolling event log files	176
Rolled log filename format	177
Rolling intervals	178
Setting log file rolling options	178
Splitting event log files	179
ICP log splitting	179
NNTP log splitting	180
HTTP host log splitting	180
Setting log splitting options	181
Collating event log files	182
Configuring Content Gateway to be a collation server	183
Configuring Content Gateway to be a collation client	183
Using a stand-alone collator	184
Viewing logging statistics	185
Viewing log files	185
Example event log file entries	187
Squid format	187
Netscape examples	188
Appendix A Statistics	191
My Proxy	191
Summary	191
Node	193
Graphs	194
Alarms	194
Protocols	194
HTTP	194
FTP	196
Content Routing (ICP)	197
Security	198
ARM security	198
LDAP	198
NTLM	199
SOCKS	199
Data Security	200
Subsystems	200
Cache	201

	Clustering	202
	Logging	202
	Networking	203
	System	203
	ARM	204
	ICAP	204
	WCCP	205
	DNS Resolver	206
	Virtual IP	206
	Performance	207
	SSL	207
	SSL Key Data	208
	CRL Statistics	209
	Reports	209
Appendix B	Commands and Variables	211
	Websense Content Gateway commands	211
	Websense Content Gateway variables	212
	Statistics	212
Appendix C	Configuration Options	219
	My Proxy	219
	Basic	220
	Subscription	223
	UI Setup	223
	Snapshots	226
	Logs	227
	Protocols	228
	HTTP	229
	HTTP Responses	237
	HTTP Scheduled Update	238
	HTTPS	240
	FTP	240
	Content Routing	241
	Hierarchies	242
	Mapping and Redirection	245
	Browser Auto-Config	247
	Security	247
	Connection Control	247
	Data Security	250
	Access Control	251
	SOCKS	258
	Subsystems	260

	Cache	260
	Logging	261
	Networking	265
	Connection Management	265
	ARM	267
	WCCP	271
	DNS Resolver	273
	ICAP	275
	Virtual IP	276
	SSL	276
Appendix D	Event Logging Formats	279
	Custom logging fields	279
	Logging format cross-reference	282
	Squid logging formats	282
	Netscape Common logging formats	283
	Netscape Extended logging formats	283
	Netscape Extended-2 logging formats	284
Appendix E	Configuration Files	285
	Specifying URL regular expressions (url_regex)	285
	Examples	286
	arm_security.config	287
	Format	288
	Examples	288
	auth.config	289
	Format	289
	Examples	291
	bypass.config	291
	Format	293
	Dynamic deny bypass rules	293
	Examples	294
	cache.config	294
	Format	295
	Examples	296
	filter.config	297
	Format	298
	Examples	299
	hosting.config	300
	Format	300
	Examples	301
	icap.config	302
	Format	302

Examples	303
ip_allow.config	303
Format	303
Examples	303
ipnat.conf	304
Format	304
Examples	304
log_hosts.config	305
Format	305
Examples	305
logs_xml.config	306
Format	306
Examples	311
WELF (WebTrends Enhanced Log Format)	313
mgmt_allow.config	313
Format	313
Examples	314
parent.config	314
Format	314
Examples	316
partition.config	316
Format	317
Examples	317
records.config	318
Format	318
Examples	318
Configuration variables	319
System variables	320
Local manager	322
Process manager	325
Virtual IP manager	325
Alarm configuration	325
ARM (transparency configuration)	326
Load shedding configuration (ARM)	329
Authentication basic realm	329
LDAP	330
RADIUS authentication	331
NTLM	333
Transparent authentication	335
HTTP engine	336
Parent proxy configuration	339
HTTP connection timeouts (secs)	340

Origin server connection attempts	341
Congestion control	342
Negative response caching	343
Proxy users variables.	343
Security	345
Cache control.	345
Heuristic expiration.	347
Dynamic content and content negotiation.	348
Anonymous FTP password	348
Cached FTP document lifetime.	348
FTP transfer mode.	349
Customizable user response pages	349
FTP engine	350
SOCKS processor	355
Net subsystem	356
Cluster subsystem	356
Cache.	356
DNS.	357
HostDB	358
Logging configuration.	359
URL remap rules.	364
ICP configuration	365
Scheduled update configuration	365
WCCP configuration.	366
ARM (security configuration).	369
SSL Decryption.	369
ICAP	370
Data Security	371
Connectivity, analysis, and boundary conditions	372
remap.config	374
Format	374
Examples	375
socks.config	375
Format	375
Examples	376
splitdns.config	377
Format	377
Examples	378
storage.config	379
Format	379
update.config	379
Format	380

	Examples	381
	congestion.config	381
	Format	382
	Destination specifier tags	382
	Congestion control rule variables	383
Appendix F	Error Messages	385
	Websense Content Gateway error messages	385
	Process fatal.	385
	Warnings	385
	Alarm messages	386
	HTML messages sent to clients	388
	Standard HTTP response messages	391
Appendix G	The req_ca.cnf File	393
Appendix H	FAQs and Troubleshooting Tips	395
	Frequently Asked Questions (FAQs)	395
	How do disk I/O errors affect the cache and what does Content Gateway do when a cache disk fails?	395
	If a client disconnects during the time that Content Gateway is downloading a large object, is any of the object saved in the cache?	396
	Can Content Gateway cache Java applets, JavaScript programs, or other application files like VBScript?	396
	How do you access Content Gateway Manager if you forget the master administrator password?	396
	How do you apply changes to the logs_xml.config file to all nodes in a cluster?	397
	In Squid- and Netscape-format log files, what do the cache result codes mean?	397
	What does the cqtx field record in a custom log file?	399
	Does Content Gateway refresh entries in its host database after a certain period of time if they have not been used?	399
	Can you improve the look of your custom response pages by using images, animated gifs, and Java applets?	399
	How do you configure Content Gateway to serve only transparent requests?	400
	Troubleshooting tips	401
	When you clear the cache, the statistics in Content Gateway Manager do not reset to zero	401
	The throughput statistic is inaccurate in Content Gateway Manager	401
	You are unable to execute Content Gateway commands	402
	You observe inconsistent behavior when one node obtains an object from another node in the cluster	402
	Web browsers may display an error document with a data missing message	403

	Content Gateway does not resolve any Web sites.	403
	Maximum document size exceeded message in the system log file	404
	DrainIncomingChannel message in the system log file	404
	No cop file message in the system log file	404
	Warning in system log file when editing vadders.config (Linux) . .	405
	Non transparent requests fail after enabling always_query_destination	405
	Content Gateway is running but no log files are created	406
	Content Gateway error indicates too many network connections. .	406
	Low memory symptoms	407
	Connection timeouts with the origin server	407
	IBM Web servers do not work with Content Gateway	407
	Content Gateway does not start (or stop)	408
Appendix I	Glossary	409
Appendix J	Copyrights	415
Index	419

1

Overview

Websense® Content Gateway is a high-performance Web proxy cache that improves network efficiency and performance by caching frequently accessed information at the edge of the network. This brings content physically closer to end users for faster delivery and reduced bandwidth usage.

Content Gateway can be deployed:

- ◆ *As a Web proxy cache*
- ◆ *In a cache hierarchy*
- ◆ *In a cluster*
- ◆ *As an SSL server*

Related topics:

[Deployment options, page 1](#)

[Components, page 3](#)

[Traffic analysis options, page 5](#)

[Security options, page 5](#)

[Online Help, page 6](#)

[Technical Support, page 7](#)

Deployment options

As a Web proxy cache

When Content Gateway is deployed as a Web proxy, user requests for Web content pass through Content Gateway on the way to the destination Web server (origin server). If the Content Gateway cache contains the requested content, Content Gateway serves the content directly. If the Content Gateway cache does not have the requested content, Content Gateway acts as a proxy, fetching the content from the origin server on the user's behalf, while keeping a copy to satisfy future requests.

Content Gateway provides the following proxy caching options:

- ◆ *Explicit proxy caching*, where the user's client software must be configured to send requests directly to Content Gateway. See [Explicit Proxy Caching, page 35](#).

- ◆ *Transparent proxy caching*, where user requests are invisibly routed through the Content Gateway cache on their way to the destination server. Users request Internet content as usual, without any browser configuration, and Content Gateway serves their requests. The user's client software (typically a browser) is unaware that it is communicating with a proxy. See [Transparent Proxy Caching and ARM](#), page 43.

In a cache hierarchy

Websense 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 cache, either to other Content Gateway systems or to other caching products.

Content Gateway can be a member of an HTTP cache hierarchy and an ICP (Internet Cache Protocol) cache hierarchy. See [Hierarchical Caching](#), page 73.

In a cluster

Websense Content Gateway scales from a single node into multiple nodes that form a cluster, improving system performance and reliability. The proxy detects the addition or removal of nodes. If Content Gateway's 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 hard node failures (such as power supply or CPU failures) and reassign IP addresses of the failed node to the operational nodes.

Content Gateway has the following clustering modes:

- ◆ *Management-only mode*, where you can administer all the nodes in a cluster at the same time. Nodes automatically share configuration information.
- ◆ *Full-clustering mode*, where the node caches act as a single aggregate cache. A Content Gateway cluster distributes its cache across its nodes into a single, virtual object store, rather than replicating the cache node by node.

A fully clustered Content Gateway system provides a single system image to both users and administrators, appearing as a single virtual server. Full-clustering mode includes management-only mode.

See [Clusters](#), page 65.

As an SSL server

If SSL Manager is enabled, HTTPS data is decrypted and then re-encrypted as it travels from the client to the origin server. Content Gateway does not cache HTTPS data.

SSL Manager includes a complete set of certificate-handling capabilities. See [Working With Encrypted Data](#), page 111.

Components

Cache

The *cache* consists of a high-speed object database called the object store. The object store indexes objects according to URLs and associated headers. The object store can cache alternate versions of the same object, varying on spoken language or encoding type, and can store small and large documents, minimizing wasted space. When the cache is full, the proxy removes stale data, ensuring that frequently requested objects are fresh.

Content Gateway tolerates disk failure on any cache disk. If the disk fails completely, Content Gateway marks the disk as corrupt and continues using the remaining disks. If all cache disks fail, Content Gateway goes into proxy-only mode.

You can partition the cache to reserve disk space for storing data for specific protocols and origin servers. See [Configuring the Cache, page 77](#).

RAM cache

Content Gateway maintains a small RAM memory cache of extremely popular objects. This RAM cache serves the most popular objects quickly and reduces load on disks, especially during traffic peaks. You can configure the RAM cache size. See [Changing the size of the RAM cache, page 82](#).

Adaptive Redirection Module

The Adaptive Redirection Module (ARM) is used in transparent proxy caching to redirect user requests destined for an origin server to the proxy. Before the traffic is redirected by the ARM, it is intercepted by a Layer 4 switch or router.

To redirect user requests to the proxy, the ARM changes an incoming packet's address. The packet's destination IP address is changed to the IP address of the proxy, and the packet's destination port is changed according to the protocol used. For example, for HTTP, the packet's destination port is changed to the proxy's HTTP port (usually 8080).

The ARM supports automatic bypass of sites that do not function properly with proxy caches.

The ARM also 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, page 61](#).

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.

Processes

The **content_gateway** process accepts connections, processes protocol requests, and serves documents from the cache or origin server.

The **content_manager** process 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 queueing shields users from server restart downtime.

The **content_cop** process 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 a timeout interval or if an incorrect response is received, **content_cop** restarts **content_manager** and **content_gateway**.

Administration tools

Related topics:

[Content Gateway Manager, page 85](#)

[Command-line interface, page 89](#)

[Configuration files, page 89](#)

Websense Content Gateway offers administration alternatives:

- ◆ *Content Gateway Manager* is a Web-based interface accessible through a browser. 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. Content Gateway Manager offers password-protected, SSL-encrypted, single-point administration for an entire Content Gateway cluster.

- ◆ A *command-line interface* enables you to monitor Content Gateway performance and network traffic, and configure the proxy. You can execute individual commands or script a series of commands in a shell.
- ◆ *Configuration files* allow complete administration through a file-editing and signal-handling interface. You can change configuration options by editing configuration files instead of using Content Gateway Manager or the command-line interface. Any changes you make through Content Gateway Manager or the command-line interface are automatically made to the configuration files.

Traffic analysis options

Websense Content Gateway provides options for network traffic analysis and monitoring:

- ◆ *Manager statistics and graphs* show network traffic information. View graphs and statistics from 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, document hit rates, and so on. View *Performance* graphs in the Content Gateway Manager.
- ◆ *Manager alarms* are presented in 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.
- ◆ *Transaction logging* lets you record information in a log file about every request the proxy cache receives and every error it detects. Use the logs to determine how many people use the proxy cache, how much information each person requested, and which pages are most popular. You can see why a transaction was in error and see the state of the proxy cache at a particular time. 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, page 93](#). For logging options, see [Working With Log Files, page 167](#).

Security options

Websense Content Gateway combined with Websense Web Filter or Websense Web Security Suite provides protection and filtering along with Web proxy technology. Not only can you provide a proxy and caching to improve responsiveness, you can also provide an increased level of security for Internet usage.

In addition, Content Gateway enables you to establish secure communication between the Content Gateway system and other computers on the network. You can:

- ◆ Control client access to the proxy cache.
- ◆ Control which hosts are allowed to access the proxy cache machine (ARM security).
- ◆ Configure Content Gateway integration into your firewall and control traffic through a SOCKS server.
- ◆ Configure Content Gateway to use different DNS servers, depending on whether it 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.
- ◆ Configure the proxy to ensure that clients are authenticated before they can access content from the cache. Content Gateway supports LDAP, RADIUS, and NTLM proxy authentication.
- ◆ Configure Content Gateway support for ICAP to enable sites using Websense Data Security Suite to examine outbound material such as Web postings, and either block or allow the posting based on company policy. See [Working With Websense Data Security](#), page 105.
- ◆ Control access to Content Gateway Manager using:
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access
 - An access control list (ACL) that defines which hosts are allowed to access Content Gateway Manager
 - User accounts that define which users can access Content Gateway Manager and which activities they can perform (for example, view statistics only or view statistics and configure Content Gateway).

See [Security](#), page 139.

Online Help

Select the **Help** option within the program to display detailed information about using the product.



Important

Default Microsoft 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. (Check **Allow active content to run in files on My Computer** under Security options.)

To access a PDF version of the online help, or to access Release Notes, the Installation Guide, the Deployment Guide, and knowledge base articles:

1. Log into your [MyWebsense](#) account.
2. Click the **Support** link and select **Websense Security Gateway** from the Knowledge Base drop-down list.
3. Search or browse the “v7 Websense Web Security Gateway” category.

Technical Support

Technical information about Websense software and services is available 24 hours a day at:

www.websense.com/support/

- ◆ the latest release information
- ◆ the searchable Websense Knowledge Base
- ◆ Support Forums
- ◆ Support Webinars
- ◆ show-me tutorials
- ◆ product documents
- ◆ answers to frequently asked questions
- ◆ Top Customer Issues
- ◆ in-depth technical papers

For additional questions, click the **Contact Support** tab at the top of the page.

If your issue is urgent, please call one of the offices listed below. You will be routed to the first available technician, who will gladly assist you.

For less urgent cases, use our online **Support Request Portal** at ask.websense.com.

For faster phone response, please use your **Support Account ID**, which you can find in the Profile section at [MyWebsense](#).

Location	Contact information
North America	+1-858-458-2940
France	Contact your Websense Reseller. If you cannot locate your Reseller: +33 (0) 1 5732 3227
Germany	Contact your Websense Reseller. If you cannot locate your Reseller: +49 (0) 69 517 09347
UK	Contact your Websense Reseller. If you cannot locate your Reseller: +44 (0) 20 3024 4401
Rest of Europe	Contact your Websense Reseller. If you cannot locate your Reseller: +44 (0) 20 3024 4401

Location	Contact information
Middle East	Contact your Websense Reseller. If you cannot locate your Reseller: +44 (0) 20 3024 4401
Africa	Contact your Websense Reseller. If you cannot locate your Reseller: +44 (0) 20 3024 4401
Australia/NZ	Contact your Websense Reseller. If you cannot locate your Reseller: +61 (0) 2 9414 0033
Asia	Contact your Websense Reseller. If you cannot locate your Reseller: +86 (10) 5884 4200
Latin America and Caribbean	+1-858-458-2940

For telephone requests, please have ready:

- ◆ Websense subscription key
- ◆ Access to Web Security Manager and Content Gateway Manager
- ◆ Access to the machine running Filtering Service, the machine running reporting tools, and the database server (Microsoft SQL Server or MSDE)
- ◆ Access to the machine running Websense Content Gateway
- ◆ Permission to access the Websense Log Database
- ◆ Familiarity with your network's architecture, or access to a specialist
- ◆ Specifications of machines running Filtering Service and Web Security Manager
- ◆ A list of other applications running on the Filtering Service machine

2

Getting Started

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

Refer to the following procedures to get started:

- ◆ [Starting and stopping Content Gateway, page 9](#)
- ◆ [Starting Content Gateway Manager, page 10](#)
- ◆ [Entering your subscription key, page 11](#)
- ◆ [Verifying that the proxy is running, page 11](#)
- ◆ [Using the command-line interface, page 12](#)

Starting and stopping Content Gateway

Websense Content Gateway starts automatically when installation is completed. To stop or start Content Gateway from the command line:

1. Become root:

```
su
```

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

To start the proxy:

```
./WCGAdmin start
```

To stop the proxy

```
./WCGAdmin stop
```

To restart the proxy

```
./WCGAdmin restart
```

To see what Content Gateway services are running:

```
./WCGAdmin status
```

**Note**

Always use the `./WCGAdmin stop` command to stop Content Gateway.

After you have installed Content Gateway, open Content Gateway Manager (the management interface) to verify that the proxy is running. See [Starting Content Gateway Manager, page 10](#) and [Verifying that the proxy is running, page 11](#).

Starting Content Gateway Manager

Content Gateway Manager is the management user interface to Websense Content Gateway. Content Gateway Manager is supported on Microsoft Internet Explorer 7 and later, and on Firefox 3.0.x - 3.5.x. Use of another browser version may result in unexpected behavior.

Java and JavaScript must be enabled in your browser. See your browser documentation for information on enabling Java and JavaScript.

To access Content Gateway Manager:

1. Open your Web browser.
2. Enter the following location in your browser:

```
https://nodename:adminport
```

where *nodename* is the IP address and *adminport* is the number assigned to the Content Gateway Manager port (default: 8081). For more information on using HTTPS to start Content Gateway Manager, see [Using SSL for secure administration, page 143](#).

3. Log on to Content Gateway Manager with the administrator ID (default: admin) and password, or your user account.

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 Content Gateway Manager, page 141](#).

Content Gateway Manager opens to the **Monitor > My Proxy > Summary** page. This page provides information on the features of your subscription and details of your Content Gateway system. See [Viewing statistics, page 93](#), for additional information on the Monitor tab and [Configuring the System, page 85](#) for information on the configuration options in Content Gateway Manager.

If you receive a message indicating that Content Gateway Manager is finding an unrecognized or invalid certificate, install the **content_gateway_ca.cer** certificate provided by Websense. It is located in the **/home/Websense** directory.

Use the **Configure > SSL > Certificates > Add Root CA** page to import this certificate.

1. Browse to `/home/Websense`, and select `content_gateway_ca.cer`.
2. Click **Open**.
3. Click **Add Certificate Authority**.

Entering your subscription key

Related topic:

[Providing system information, page 11](#)

1. On the **Configure > My Proxy > Subscription > Subscription Management** tab, enter the subscription key that Websense provided to you.
2. Click **Apply**.
3. Click **Restart** on **Configure > My Proxy > Basic > General** page.

Providing system information

If your subscription includes the scanning features (Web Security Gateway or Web Security Gateway Anywhere), you must provide information to facilitate communicating with Websense filtering software.

1. Select the **Configure > My Proxy > Subscription > Scanning** tab. Notice the IP address and port for the filtering service. This is the information that you entered when you installed the filtering product.



Note

The Scanning tab appears only if you have subscribed to Web Security Gateway or Web Security Gateway Anywhere. It can take several minutes after you enter your subscription key before this tab appears in the TRITON Web Security (Web Security Manager) user interface.

2. Select the appropriate check box to permit or block traffic if communication with Policy Server or Filtering Service fails.
3. Click **Apply**.

Verifying that the proxy is running

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

1. Open Content Gateway Manager. See [Starting Content Gateway Manager](#), page 10.
2. Navigate to **Monitor > Protocol > HTTP > General** to display the General HTTP Statistics table.
3. Note the current **Total Document Bytes** statistic in the **Client** section of the table.

Check the value of this statistic.

General HTTP Statistics	
Attribute	Current Value
Client	
Total Document Bytes	0
Total Header Bytes	0
Total Connections	0
Current Connections	0
Transactions in Progress	0
Server	
Total Document Bytes	0
Total Header Bytes	0
Total Connections	0
Current Connections	0
Transactions in Progress	0

4. Set your browser to the proxy port.
5. Browse the Internet.
6. 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.

You can execute individual commands or script multiple commands in a shell. See [Websense Content Gateway commands](#), page 211.

1. Become root:
`su`
2. Change to the Content Gateway **bin** directory (default = `/opt/WCG`). Run Content Gateway commands from this directory.

Commands take the form:

```
content_line -command argument
```

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
```

3

Web Proxy Caching

Web proxy caching stores copies of frequently accessed Web objects (such as documents, images, and articles) close to users and serves this information to them. Internet users get their information faster, and Internet bandwidth is freed for other tasks.

Internet users direct their requests to Web servers all over the Internet. For a caching server to serve these requests, it must act as a Web proxy server. A Web proxy server receives user requests for Web objects and either serves the requests or forwards them to the *origin server* (the Web server that contains the original copy of the requested information).

Content Gateway supports both *transparent proxy caching*, where the user's client software (typically a browser) is unaware that it is communicating with a proxy, and *explicit proxy caching*, where the user's client software must be configured to send requests directly to the proxy.

Cache requests

Related topics:

[Ensuring cached object freshness, page 16](#)

[Scheduling updates to local cache content, page 21](#)

[Pinning content in the cache, page 23](#)

[To cache or not to cache?, page 24](#)

[Caching HTTP objects, page 24](#)

[Forcing object caching, page 29](#)

[Caching HTTP alternates, page 30](#)

[Caching FTP objects, page 31](#)

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

1. Content Gateway receives a user request for a document, image, news article, or other Web object.
2. Using the object address, the proxy tries to locate the requested object in its object database (cache).

3. If the object is in the cache, the proxy checks to see if the object is fresh enough to serve (see [Ensuring cached object freshness](#), page 16). If the object is fresh, the proxy serves it to the user as a *cache hit*.
4. If the data in the cache is stale, the proxy connects to the origin server and asks if the object is still fresh (a revalidation). If the object is still fresh, the proxy sends the cached copy to the user immediately.
5. If the object is not in the cache (a cache miss) or the server indicates that the cached copy is no longer valid, the proxy obtains the object from the origin server, simultaneously streaming it to the user and the cache. Subsequent requests for the object will be served faster because the object will come directly from the cache.

Ensuring cached object freshness

When Content Gateway receives a request for a Web object, it tries to locate the requested object in its cache. If the object is in the cache, the proxy checks to see if the object is fresh enough to serve.

The protocol determines how the proxy handles object freshness in the cache:

- ◆ HTTP objects support author-specified expiration dates. The proxy adheres to these expiration dates; otherwise, it picks an expiration date based on how frequently the object is changing and on administrator-chosen freshness guidelines. In addition, objects can be revalidated, checking with the origin server if an object is still fresh. See [HTTP object freshness](#), page 16.
- ◆ FTP objects stay in the cache for a specified time period. See [FTP object freshness](#), page 20.

HTTP object freshness

Content Gateway determines whether an HTTP object in the cache is fresh by:

- ◆ Checking the **Expires** or **max-age** header
Some HTTP objects contain **Expires** headers or **max-age** headers that define how long the object can be cached. Comparing the current time with the expiration time tells the proxy whether or not the object is fresh.
- ◆ Checking the **Last-Modified / Date** headers

If an HTTP object has no **Expires** header or **max-age** header, the proxy can calculate a freshness limit using the following formula:

$$\text{freshness_limit} = (\text{date} - \text{last_modified}) * 0.10$$

where *date* is the date in the object's server response header, and *last_modified* is the date in the **Last-Modified** header. If there is no **Last-Modified** header, the proxy uses the date that the object was written to cache. You can increase or reduce the value 0.10 (10 percent). See [Modifying the aging factor for freshness computations](#), page 17.

The computed freshness limit is bound by minimum and maximum boundaries. See [Setting an absolute freshness limit](#), page 17.

- ◆ Checking the absolute freshness limit
For HTTP objects that do not have **Expires** headers or do not have both **Last-Modified** and **Date** headers, the proxy uses a maximum and minimum freshness limit. See [Setting an absolute freshness limit](#), page 17.
- ◆ Checking revalidate rules in the **cache.config** file
Revalidate rules apply freshness limits to specific HTTP objects. You can set freshness limits for objects originating from particular domains or IP addresses, objects with URLs that contain specified regular expressions, and objects requested by particular clients, for example. See [cache.config](#), page 294.

Modifying the aging factor for freshness computations

If an object does not contain any expiration information, Content Gateway can estimate its freshness from the **Last-Modified** and **Date** headers. By default, the proxy stores an object for 10% of the time that elapsed since it last changed. You can increase or reduce the percentage.

1. Open the **records.config** file located in the Content Gateway **config** directory.
2. Edit the following variable:

Variable	Description
<code>proxy.config.http.cache.heuristic_lm_factor</code>	Specify the aging factor for freshness computations. The default value is 0.10 (10 percent).

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

```
content_line -x
```

Setting an absolute freshness limit

Some objects do not have **Expires** headers or do not have both **Last-Modified** and **Date** headers. You can control how long these objects are considered fresh in the cache by specifying an absolute freshness limit. A longer lifetime means objects are kept in the cache longer. Performance can improve if pages are taken from the cache rather than going out to the network.

1. Navigate to the **Configure > Protocols > HTTP > Cacheability** tab.
2. In the **Minimum Heuristic Lifetime** area of the **Freshness** section, specify the minimum amount of time that HTTP objects without an expiration date can remain fresh in the cache before being considered stale. The default value is 3600 seconds (1 hour).

3. In the **Maximum Heuristic Lifetime** field, specify the maximum amount of time that HTTP objects without an expiration date can remain fresh in the cache before being considered stale. The default value is 86400 seconds (1 day).
4. Click **Apply**.

Specifying header requirements

To ensure freshness of the objects in the cache, configure Content Gateway to cache only objects with specific headers.



Warning

By default, the proxy caches all objects (including objects with no headers). Websense recommends that you change the default setting only for specialized proxy situations. If you configure the proxy to cache only HTTP objects with **Expires** or **max-age** headers, the cache hit rate will be seriously reduced (very few objects have explicit expiration information).

1. Navigate to the **Configure > Protocols > HTTP > Cacheability** tab.
2. In the **Required Headers** area of the **Behavior** section, select one of the following:
 - **An Explicit Lifetime Header** to cache only HTTP objects with **Expires** or **Cache-Control** headers.
 - **A Last-Modified Header** to cache only HTTP objects with **Expires** or **Last-Modified** headers.
 - **No Required Headers** to cache all HTTP objects (no specific headers are required). This is the default.
3. Click **Apply**.

Cache-Control headers

Even though an object might be fresh in the cache, clients or servers might have constraints that prevent them from retrieving the object from the cache. For example, a client might request that a object not come from a cache, or if it does, it cannot have been cached for more than 10 minutes.

Content Gateway bases the servability of a cached object on **Cache-Control** headers. **Cache-Control** headers can appear in both client requests and server responses.

The following **Cache-Control** headers affect whether objects are served from the cache:

- ◆ The **no-cache** header, sent by clients, tells the proxy to serve *no* objects directly from the cache; always obtain the object from the origin server. You can configure the proxy to ignore client **no-cache** headers (see [Configuring the proxy to ignore client no-cache headers](#), page 25).

- ◆ The **max-age** header, sent by servers, is compared to the object age; if the age is less than **max-age**, the object is fresh and can be served.
- ◆ The **min-fresh** header, sent by clients, is an *acceptable freshness tolerance*. The client wants the object to be at least this fresh. If a cached object does not remain fresh at least this long in the future, it is revalidated.
- ◆ The **max-stale** header, sent by clients, permits the proxy to serve stale objects provided they are not too old. Some browsers might be willing to take slightly old objects in exchange for improved performance, especially during periods of poor Internet availability.

The proxy applies Cache-Control servability criteria *after* HTTP freshness criteria. For example, an object might be considered fresh, but if its age is greater than its *max-age*, it is not served.

Revalidating HTTP objects

When a client requests an HTTP object that is stale in the cache, Content Gateway revalidates the object, querying the origin server to check if the object is unchanged. Revalidation results in one of the following:

- ◆ If the object is still fresh, the proxy resets its freshness limit and serves the object.
- ◆ If a new copy of the object is available, the proxy caches the new object, replacing the stale copy, and serves the object to the user simultaneously.
- ◆ If the object no longer exists on the origin server, the proxy does not serve the cached copy.
- ◆ If the origin server does not respond to the revalidation query, the proxy does not perform any validation; it serves the stale object from the cache.

By default, the proxy revalidates a requested HTTP object in the cache if it considers the object to be stale. The proxy evaluates object freshness as described in [HTTP object freshness](#), page 16. You can configure how often you want the proxy to revalidate an HTTP object.

1. Navigate to the **Configure > Protocols > HTTP > Cacheability** tab.
2. In the **When to Revalidate** area of the **Behavior** section, select:
 - **Never Revalidate** to never verify the freshness of a requested HTTP object with the origin server.
 - **Always Revalidate** to always verify the freshness of a requested HTTP object with the origin server.
 - **Revalidate if Heuristic Expiration** to verify the freshness of a requested HTTP object with the origin server if the object contains no **Expires** or **Cache-Control** headers. Content Gateway considers all HTTP objects without **Expires** or **Cache-Control** headers to be stale.
 - **Use Cache Directive or Heuristic** to verify the freshness of a requested HTTP object with the origin server when Content Gateway considers the object in the cache to be stale. This is the default.

3. Click **Apply**.

**Note**

You can also set specific revalidation rules in the `cache.config` file. See [cache.config](#), page 294.

FTP object freshness

FTP objects carry no time stamp or date information and remain fresh in the cache for the period of time you specify (from 15 minutes to 2 weeks), after which they are considered stale.

FTP objects can be requested from either an HTTP client (such as a browser) or an FTP client (such as WS_FTP). Content Gateway caches the FTP objects requested from HTTP clients in HTTP format and the FTP objects requested from FTP clients in a proprietary format.

The procedure for specifying the freshness limit for objects requested by HTTP clients is different from the procedure for objects requested by FTP clients. Follow the procedure appropriate for your needs.

FTP objects requested by HTTP clients

You can set an absolute freshness limit for FTP objects requested by HTTP clients (FTP over HTTP objects).

**Note**

In addition to setting an absolute freshness limit for all FTP objects requested by HTTP clients, you can set freshness rules for specific FTP objects in the `cache.config` file (see [cache.config](#), page 294).

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **FTP Document Lifetime** area of the **Freshness** section, enter the amount of time that FTP objects requested by HTTP clients can remain fresh in the cache before being considered stale. The default value is 259200 seconds (3 days).
3. Click **Apply**.

FTP objects requested by FTP clients

You can set freshness limits for different types of FTP objects that are requested from FTP clients: directory listings, login messages, and FTP files.

1. Navigate to **Configure > My Proxy > Basic > General**.
2. Select **On** for FTP, if it is not already enabled. Then click **Apply** and click **Restart**.
3. Navigate to **Configure > Protocols > FTP > Cacheability**.

4. In the **Login Information** area of the **Freshness** section, enter the amount of time that FTP login messages can remain fresh in the cache before the proxy considers them stale. The default value is 2592000 seconds (30 days).
5. In the **Directory Listings** field, enter the amount of time that FTP directory listings can remain in the cache before the proxy considers them stale. The default value is 604800 seconds (7 days).
6. In the **Files** field, enter the amount of time that FTP files can remain fresh in the cache before the proxy considers them stale. The default value is 259200 seconds (3 days).
7. Click **Apply**.

Scheduling updates to local cache content

To further increase performance and to ensure that HTTP and FTP objects (requested from HTTP clients) are fresh in the cache, you can use the Scheduled Update option to configure the proxy to load specific objects into the cache at scheduled times.

**Note**

The Scheduled Update option does not update FTP objects requested from FTP clients.

To use the Scheduled Update option:

- ◆ Specify the list of URLs that contain the objects you want to schedule for update, the time the update should take place, and the recursion depth for the URL.
- ◆ Enable the Scheduled Update option and configure optional retry settings.

See [Configuring the Scheduled Update option, page 22](#) for more information.

Content Gateway uses the information you specify to determine the URLs for which it is responsible and, for each URL, derives all recursive URLs if applicable. It then generates a unique URL list. Using this list, the proxy initiates an HTTP **GET** for each unaccessed URL, ensuring that it remains within the user-defined limits for HTTP concurrency at any given time.

**Note**

The system logs the completion of all HTTP **GET** operations, enabling you to monitor the performance of this feature.

The Force Immediate Update option that enables you to update URLs without waiting for the specified update time to occur. You can use this option to test your scheduled update configuration. See [Forcing an immediate update, page 23](#).

Configuring the Scheduled Update option

1. Navigate to **Configure > Protocols > HTTP Scheduled Update > Update URLs**.
2. In the **Scheduled Object Update** area, click **Edit File** to open the configuration file editor for the **update.config** file.
3. Enter the following information:
 - In the **URL** field, enter the URL you want to schedule for update.
 - *Optional.* In the **Request Headers** field, enter the semicolon-separated list of headers passed in each **GET** request. You can define any request header that conforms to the HTTP specification.
 - In the **Offset Hour** field, enter the base hour used to derive the update periods. You can specify a value in the range 00 to 23.
 - In the **Interval** field, enter the interval (in seconds) at which updates occur, starting at the offset hour.
 - In the **Recursion Depth** field, enter the depth to which referenced URLs are recursively updated, starting at the given URL. For example, a recursion depth of 1 updates the given URL, as well as all URLs immediately referenced by links from the original URL.
4. Click **Add**, and then click **Apply**.
5. Click **Close**.
6. Click the **General** tab.
7. Enable **Scheduled Update**.
8. In the **Maximum Concurrent Updates** field, enter the maximum number of simultaneous update requests allowed at any time to prevent the scheduled update process from overburdening the host. The default is 100.
9. In the **Count** field of the **Retry on Update Error** section, enter the number of times you want to retry the scheduled update of a URL in the event of failure. The default value is 10.
10. In the **Interval** field of the **Retry on Update Error** section, enter the delay in seconds between each scheduled update retry for a URL in the event of failure. The default value is 2.
11. Click **Apply**.

Forcing an immediate update

The Force Immediate Update option lets you verify the URLs listed in the **update.config** file immediately. This option disregards the offset hour and interval set in the **update.config** file and updates the URLs listed.



Important

When you enable the Force Immediate Update option, the proxy continually updates the URLs specified in the **update.config** file until you disable the option.

1. Navigate to **Configure > Protocols > HTTP Scheduled Update > General**.
2. Ensure that **Scheduled Update** is enabled.
3. Click the **Update URLs** tab.
4. Enable **Force Immediate Update**.
5. Click **Apply**.

Pinning content in the cache

The cache pinning option configures Content Gateway to keep certain HTTP objects (and FTP objects requested from HTTP clients) in the cache for a specified time. Use this option to ensure that the most popular objects are in the cache when needed and that the proxy does not delete important objects from the cache.



Note

The proxy observes Cache-Control headers and pins an object in the cache only if it is cacheable.

To use cache pinning, perform the following tasks:

- ◆ Set cache pinning rules in the **cache.config** file. See [Setting cache pinning rules, page 23](#).
- ◆ Enable the cache pinning option. See [Enabling cache pinning, page 24](#).

Setting cache pinning rules

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. Click **Edit File** at the end of the page to display the configuration file editor for the **cache.config** file.
3. In the fields provided, supply the following information:
 - From the **Rule Type** drop-down box, select **pin-in-cache**.

- From the **Primary Destination Type** drop-down box, select **url_regex**.
- In the **Primary Destination Value** field, specify the URL you want to pin in the cache.
- In the **Time Period** field, specify the amount of time that the proxy pins the object in the cache.

In addition, you can add secondary specifiers (such as **Prefix** and **Suffix**) to the rule. All the fields are described under [HTTP](#), page 229.

4. Click **Add** to add the rule to the list, and then click **Apply**.
5. Click **Close**.

Enabling cache pinning

1. On **Configure > Subsystems > Cache > General**, enable **Allow Pinning**.
2. Click **Apply**.

To cache or not to cache?

When Content Gateway receives a request for a Web object that is not in the cache, it retrieves the object from the origin server and serves it to the client. At the same time, the proxy checks if the object is cacheable before storing it in its cache to serve future requests.

Content Gateway determines if an object is cacheable based on protocol:

- ◆ For HTTP objects, the proxy responds to caching directives from clients and origin servers. In addition, you can configure the proxy not to cache certain objects. See [Caching HTTP objects](#), page 24.
- ◆ For FTP objects, the proxy responds to caching directives you specify through configuration options and files. See [Caching FTP objects](#), page 31.

Caching HTTP objects

Content Gateway responds to caching directives from clients and origin servers, as well as directives you specify through configuration options and files.

This section discusses the following topics:

- ◆ [Client directives](#), page 25
- ◆ [Origin server directives](#), page 26
- ◆ [Configuration directives](#), page 28

Client directives

By default, Content Gateway does *not* cache objects with the following request headers:

- ◆ **Cache-Control: no-store**
- ◆ **Cache-Control: no-cache**



Note

You can configure the proxy to ignore the **Cache-Control: no-cache** header. See [Configuring the proxy to ignore client no-cache headers](#), page 25.

- ◆ **Cookie:** (for text objects)

By default, the proxy caches objects served in response to requests that contain cookies unless the object is text. You can configure the proxy to *not* cache cookie content of any type, cache all cookie content, or cache cookie content that is of image type only. See [Caching cookie objects](#), page 29.

- ◆ **Authorization:**



Note

FTP objects requested from HTTP clients can also contain **Cache-Control: no-store**, **Cache-Control: no-cache**, or **Authorization** headers. If an FTP object requested from an HTTP client contains such a header, the proxy does not cache it unless explicitly configured to do so.

Configuring the proxy to ignore client no-cache headers

By default, Content Gateway observes client **Cache Control: no-cache** directives. If a requested object contains a **no-cache** header, the proxy forwards the request to the origin server even if it has a fresh copy in the cache.

You can configure the proxy to ignore client **no-cache** directives. In this case, the proxy ignores **no-cache** headers from client requests and serves the object from its cache.



Important

The default behavior of observing **no-cache** directives is appropriate in most cases. Configure Content Gateway to ignore client **no-cache** directives only if you are knowledgeable about HTTP 1.1.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Behavior** section, enable the **Ignore “no-cache” in Client Requests** option.

3. Click **Apply**.

**Note**

Certain versions of Microsoft Internet Explorer do not request cache reloads from transparent caches when the user presses the browser **Refresh** button. This can prevent content from being loaded directly from the origin server. You can configure Content Gateway to treat Microsoft Internet Explorer requests more conservatively, providing fresher content at the cost of serving fewer documents from the cache. You can configure the proxy to add **no-cache** headers to requests from Microsoft Internet Explorer in Content Gateway Manager (in the **Behavior** section **Configure** > **Protocols** > **HTTP** > **Cacheability** tab).

Origin server directives

By default, Content Gateway does not cache objects with the following response headers:

- ◆ **Cache-Control: no-store**
- ◆ **Cache-Control: private**
- ◆ **WWW-Authenticate:**

**Note**

You can configure the proxy to ignore **WWW-Authenticate** headers. See [Configuring the proxy to ignore WWW-Authenticate headers, page 27](#).

- ◆ **Set-Cookie:**
- ◆ **Cache-Control: no-cache**

**Note**

You can configure the proxy to ignore **no-cache** headers. See [Configuring the proxy to ignore server no-cache headers, page 27](#).

- ◆ **Expires:** header with value of 0 (zero) or a past date

Configuring the proxy to ignore server no-cache headers

By default, Content Gateway observes **Cache-Control:no-cache** directives. A response from an origin server with a **no-cache** header is not stored in the cache, and any previous copy of the object in the cache is removed.



Important

If you configure the proxy to ignore **no-cache** headers, it also ignores **no-store** headers.



Important

The default behavior of observing **no-cache** directives is appropriate in most cases. Configure the proxy to ignore origin server **no-cache** headers only if you are knowledgeable about HTTP 1.1.

You can configure the proxy to ignore origin server **no-cache** headers.

1. Open the **records.config** file located in the Content Gateway **config** directory.
2. Edit the following variable:

Variable	Description
<i>proxy.config.http.cache.ignore_server_no_cache</i>	Set to 1 to ignore server directives to bypass the cache.

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

```
content_line -x
```

Configuring the proxy to ignore WWW-Authenticate headers

By default, Content Gateway does not cache objects that contain **WWW-Authenticate** response headers. The **WWW-Authenticate** header contains

authentication parameters that the client uses when preparing the authentication challenge response to an origin server.



Important

The default behavior of not caching objects with **WWW-Authenticate** headers is appropriate in most cases. Configure the proxy to ignore server **WWW-Authenticate** headers only if you are knowledgeable about HTTP 1.1.

You can configure the proxy to ignore origin server **WWW-Authenticate** headers, in which case, objects with **WWW-Authenticate** headers are stored in the cache for future requests.

1. Open the **records.config** file located in the Content Gateway **config** directory.
2. Edit the following variable:

Variable	Description
<code>proxy.config.http.cache.ignore_authentication</code>	Set to 1 to cache objects with WWW-Authenticate headers.

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

```
content_line -x
```

Configuration directives

In addition to client and origin server directives, Content Gateway responds to directives you specify through configuration options and files.

You can configure the proxy to:

- ◆ *Not* cache any HTTP objects. See [Disabling HTTP object caching, page 28](#).
- ◆ Cache dynamic content (objects with URLs that contain a question mark (?), a semicolon (;), or cgi, or that end in .asp). See [Caching dynamic content, page 29](#).
- ◆ Cache objects served in response to the **Cookie:** header. See [Caching cookied objects, page 29](#).
- ◆ Observe never-cache rules in the **cache.config** file. See [cache.config, page 294](#).

Disabling HTTP object caching

By default, Content Gateway caches all HTTP objects except those for which you have set never cache rules in the **cache.config** file. You can disable HTTP object caching so that all HTTP objects are served from the origin server and never cached.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. Disable **HTTP Caching**.
3. Click **Apply**.

Caching dynamic content

A URL is considered dynamic if it contains a question mark (?), a semicolon (;), or cgi, or if it ends in .asp. By default, Content Gateway does *not* cache dynamic content. However, you can configure the proxy to cache this content.



Warning

It is recommended that you configure the proxy to cache dynamic content for specialized proxy situations only.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Dynamic Caching** section, enable **Caching Documents with Dynamic URLs**.
3. Click **Apply**.

Caching cookieed objects

By default, Content Gateway caches objects served in response to requests that contain cookies *unless* the object is text. The proxy does not cache cookieed text content, because object headers are stored as well as the object, and personalized cookie header values could be saved with the object.

With non-text objects, personalized headers are unlikely to be delivered or used.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Caching Response to Cookies** area of the **Dynamic Caching** section, select a caching option:
 - Select **Cache All but Text** to cache all cookieed content except content that is text (this is the default setting).
 - Select **Cache Only Image Types** to cache cookieed content that is an image.
 - Select **Cache Any Content Type** to cache cookieed content of all types.
 - Select **No Cache on Cookies** to *not* cache cookieed content of any type.
3. Click **Apply**.

Forcing object caching

You can force Content Gateway to cache specific URLs (including dynamic URLs) for a specified duration regardless of **Cache-Control** response headers.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.

2. Click **Edit File** at the end of the page to display the configuration file editor for the **cache.config** file.
3. In the fields provided, supply the following information:
 - From the **Rule Type** drop-down box, select **t1-in-cache**.
 - From the **Primary Destination Type** drop-down box, select **url_regex**.
 - In the **Primary Destination Value** field, specify the URL you want to force cache.
 - In the **Time Period** field, specify the amount of time that the proxy can serve the URL from the cache.

In addition, you can add secondary specifiers (such as **Prefix** and **Suffix**) to the rule. All the fields are described in [HTTP, page 229](#).
4. Click **Add**, and then click **Apply**.
5. Click **Close**.

Caching HTTP alternates

Some origin servers answer requests to the same URL with a variety of objects. The content of these objects can vary, according to whether a server delivers content for different languages, targets different browsers with different presentation styles, or provides different document formats (HTML, PDF). Different versions of the same object are termed *alternates* and are cached by Content Gateway based on **Vary** response headers.

Configuring how Content Gateway caches alternates

You can specify additional request and response headers for specific content types that the proxy will identify as alternates for caching.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Vary Based on Content Type** section, click **Enabled** to cache alternate versions of HTTP documents that do not contain the **Vary** header.
3. Specify additional request and response headers for the proxy server to identify:
 - In the **Vary by Default on Text** field, enter the HTTP header field on which you want to vary if the request is for text (for example, an HTML document).
 - In the **Vary by Default on Images** field, enter the HTTP header field on which you want to vary if the request is for images (for example, a **.gif** file).

- In the **Vary by Default on Other Document Types** field, enter the HTTP header field on which you want to vary if the request is for anything other than text or images.

**Note**

If you specify **Cookie** as the header field on which to vary in the above fields, make sure that the appropriate option is enabled in the **Caching Response to Cookies** area of the **Dynamic Caching** section. For example, if you enable the **Cache Only Image Types** option in the **Caching Response to Cookies** area and you enable the **Vary by Default on Text** option in the **Vary Based on Content Type** section, alternates by cookie will not apply to text.

4. Click **Apply**.

Limiting the number of alternates for an object

You can limit the number of alternates Content Gateway can cache per object. The default number of alternates is 3.

**Note**

Large numbers of alternates can affect proxy performance because all alternates have the same URL. Although Content Gateway can look up the URL in the index very quickly, it must scan sequentially through available alternates in the object store.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Maximum Alternates** field, enter the maximum number of alternate versions of an object you want the proxy to cache. The default value is 3.
3. Click **Apply**.

Caching FTP objects

FTP objects can be requested from either an HTTP client (such as a browser) or an FTP client (such as WS_FTP).

For FTP objects requested from HTTP clients (FTP over HTTP), perform the following configuration to determine what the proxy stores in the cache:

- ◆ Disable FTP over HTTP caching so that the proxy does not cache any FTP objects requested from HTTP clients (see [Disabling FTP over HTTP caching, page 32](#)).
- ◆ Set never cache rules in the **cache.config** file (see [cache.config, page 294](#)).

- ◆ Configure the proxy to ignore client **Cache-Control: no-store** or **Cache-Control: no-cache** headers (see [Configuring the proxy to ignore client no-cache headers](#), page 25).

For FTP objects requested from FTP clients (FTP proxy objects), perform the following configuration to determine what Content Gateway stores in the cache:

- ◆ Disable FTP proxy object caching so that Content Gateway does not cache any FTP objects requested from FTP clients (see [Disabling FTP proxy object caching](#), page 32).
- ◆ Disable caching of simple and/or full directory listings (see [Disabling caching of full or simple directory listings](#), page 32).

**Note**

Only FTP objects that are retrieved anonymously (username = anonymous) are cached. Objects fetched by a named login are not cached.

Disabling FTP over HTTP caching

You can configure Content Gateway not to cache any FTP objects that are requested from HTTP clients by disabling the FTP over HTTP option. The proxy processes the requests by forwarding them to the FTP server but does not cache any requested objects.

1. Navigate to **Configure > Protocols > HTTP > Cacheability**.
2. In the **Caching** section, disable **FTP over HTTP Caching**.
3. Click **Apply**.

Disabling FTP proxy object caching

You can configure Content Gateway not to cache any FTP objects that are requested from FTP clients.

The proxy processes FTP requests by forwarding them to the FTP server (as long as the FTP processing option is enabled in the **Features** table on the **Configure/My Proxy/Basic** tab) but does not cache any requested objects.

1. Navigate to **Configure > Protocols > FTP > Cacheability**.
2. Disable **FTP Caching**.
3. Click **Apply**.

Disabling caching of full or simple directory listings

You can configure Content Gateway to not cache simple and/or full directory listings. A simple directory listing contains no arguments. For example:

```
dr/ls
```

A full directory listing does contain arguments. For example:

```
ls *.txt
```

1. Navigate to **Configure > Protocols > FTP > Cacheability**.
2. In the **Directory Caching** section:
 - Disable the **Simple** option if you do not want to cache simple directory listings.
 - Disable the **Full** option if you do not want to cache full directory listings.
3. Click **Apply**.

4

Explicit Proxy Caching

If Internet requests are not transparently routed to Content Gateway via a Layer 4 switch or router (see [Transparent Proxy Caching and ARM, page 43](#)), 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 configuring their browsers to send requests directly to the proxy. See [Manual browser configuration, page 35](#).
- ◆ By configuring their browsers to download proxy configuration instructions from a PAC (Proxy Auto-Config) file. See [Using a PAC file, page 36](#).
- ◆ By using WPAD (Web Proxy Auto-Discovery Protocol) to download proxy configuration instructions from a WPAD server (Microsoft Internet Explorer only). See [Using WPAD, page 37](#).

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, page 39](#).

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 host name or IP address
- ◆ 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 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 version 7.0 and later, 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 2.0 and later, proxy configuration settings are in **Tools > Options > Advanced > Network > Settings > 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.

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. If you have a **proxy.pac** file, copy it into the Content Gateway **config** directory.



Note

The PAC file can reside on any server in your network.

If you are using SSL Manager, refer to [Running in explicit proxy mode, page 112](#), for information on a PAC file to use with HTTPS traffic.

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.
2. Navigate to the **Configure > Content Routing > Browser Auto-Config > PAC** tab.
3. 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](#), page 37.
5. Click **Apply**.
 6. Click **Restart** on **Configure > My Proxy > Basic > General**.
 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 host name **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
```

The procedures for specifying the PAC file location vary among browsers. For example, for Microsoft Internet Explorer, you set the location of the PAC file in the **Use automatic configuration script** field under **Tools > Internet Options > Connections > LAN Settings**. For Mozilla Firefox, proxy configuration settings are in **Tools > Options > Advanced > Network > Settings > Connection Settings > Automatic proxy configuration URL**. 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 version 7 and later to automatically detect a server that will 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 version 7 or later browser starts up, it searches for a WPAD server that will supply it with proxy server configuration settings. It prepends the host name 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 version 7 and later are set to automatically detect WPAD servers. However, browser users can disable this setting.

You can configure Content Gateway to be a WPAD server:

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.
2. Navigate to **Configure > My Proxy > Basic > General**.
3. In the **Features** table, ensure that **ARM** is **On** (the default) in the **Networking** section.
4. Navigate to **Configure > Content Routing > Browser Auto-Config > WPAD** to display the **wpad.dat** file.
5. 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 (default location is in **/opt/WCG**), 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, page 37](#) (a **wpad.dat** file can contain the same script as a **proxy.pac** file).
6. Click **Apply**.
7. Navigate to **Configure > Networking > ARM**.
8. In the **Network Address Translation (NAT)** section, click **Edit File** to add a special remap rule to the **ipnat.conf** file.
9. Enter information in the fields provided, and then click **Add**:
 - In the **Ethernet Interface** field, enter the network interface that receives browser WPAD requests (for example **hme0** or **eth0**).
 - From the **Connection Type** drop-down list, select **tcp**.
 - In the **Source 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 followed by **/32**; for example: **123.456.7.8/32**.
 - In the **Source Port** field, enter **80**.

- In the **Destination IP** field enter the same IP address you entered in the **Source IP** field but omit /32.
 - In the **Destination Port** field, enter **8083**.
10. Use the arrow keys on the left side to move the new rule to the first line in the file.
 11. Click **Apply**, and then click **Close**.
 12. Click **Restart** on the **Configure > My Proxy > Basic > General**.

Configuring FTP clients in an explicit proxy environment

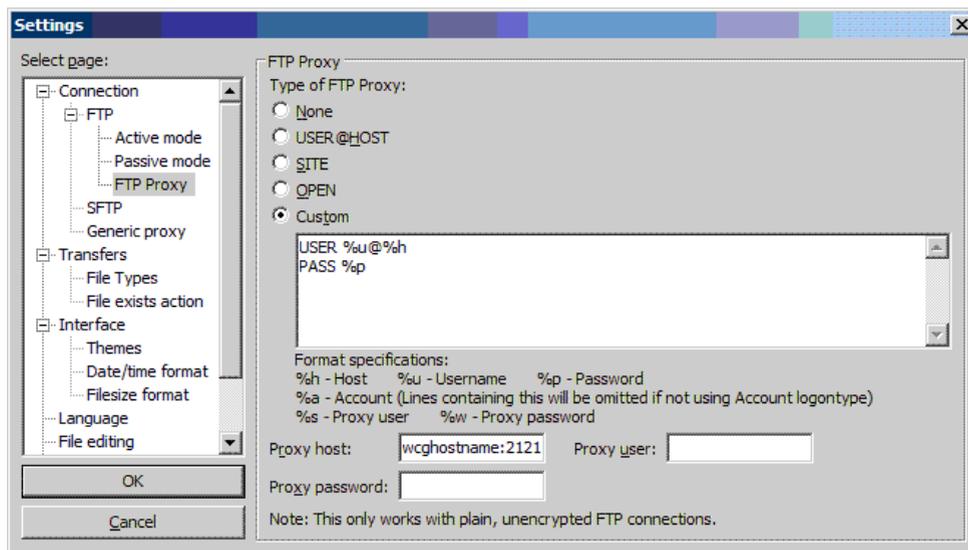
When Content Gateway is configured to proxy FTP traffic (see [FTP, page 240](#)), FTP client applications, such as FileZilla or WS_FTP, should be configured to send FTP requests to the proxy. When so configured, the user works with the FTP client application as if no proxy were present.

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

From:	To:
FTP server hostname	FTP <i>proxy</i> hostname
FTP server port number	FTP <i>proxy</i> 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.

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



In the **FTP Proxy** area:

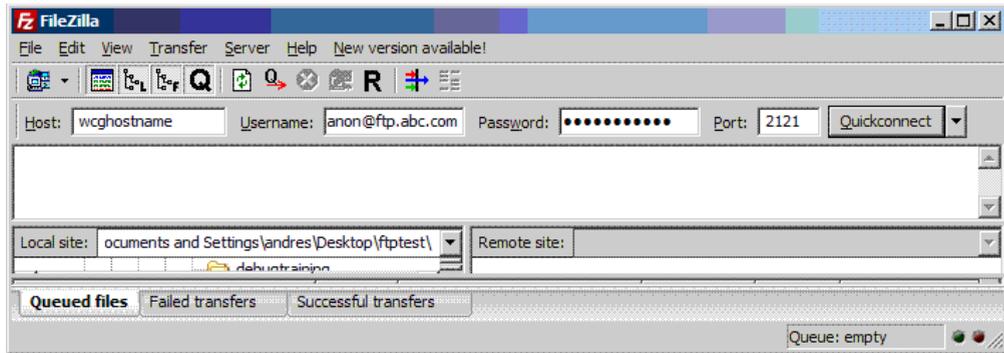
1. Set **FTP Proxy** to **Custom** and define **USER** and **PASS** as shown.
2. Set **Proxy host** to the Content Gateway FTP proxy hostname and port number.
3. Accept the settings by clicking **OK**.

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

```
Host:          ftp.abc.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.abc.com
Password:     123acb
Port:         2121
```



5

Transparent Proxy Caching and ARM

Related topics:

[Connection load shedding, page 61](#)

[Reducing DNS lookups, page 61](#)

[IP spoofing, page 62](#)

The transparent proxy caching option enables Content Gateway to respond to Internet requests without requiring users to reconfigure their browser settings. It does this by redirecting the traffic flow into the proxy cache after the traffic has been intercepted, typically by a Layer 4 (L4) switch or router. (This is sometimes referred to as a *transparent proxy deployment*.)

In transparent interception:

1. The proxy intercepts client requests to origin servers by a switch or router. See [Interception strategies, page 45](#).
2. The Adaptive Redirection Module (ARM) changes the destination IP address of an incoming packet to the proxy's IP address and the destination port to the proxy port (if different). See [Enabling the ARM option, page 44](#). The ARM is enabled by default.
3. The proxy receives and begins processing the intercepted client requests. If a request is a cache hit, the proxy serves the requested object. If a request is a miss, the proxy retrieves the object from the origin server and serves it to the client.
4. On the way back to the client, the ARM changes the source IP address to the origin server IP address and the source port to the origin server port.



Note

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 these clients and servers, passing their traffic to the origin server. You can create static rules to exempt clients and servers from caching. See [Interception bypass, page 57](#).

To prevent unauthorized access to machines running Content Gateway, you can configure the ARM to use an access control list employing administrator-specified rules to allow or deny other computers from communicating with the machine. This enables you to effectively create a firewall in front of Content Gateway, denying potentially malicious packets from reaching the TCP/IP stack on the machine. See [Controlling host access to the proxy server](#), page 140.

Enabling the ARM option

The Content Gateway ARM inspects incoming packets before the IP layer sees them and readdresses the packets to the proxy server so that they can be served from the cache.

The ARM can make two changes to an incoming packet's address: its destination IP address and its destination port. For example, the destination IP address of an HTTP packet is readdressed to the IP address of the proxy and the destination HTTP port is readdressed to the Content Gateway HTTP proxy port (usually port 8080).

On the way back to the client, the ARM changes the source IP address to the origin server IP address and the source port to the origin server port.

The ARM component consists of several files and a kernel module, which are installed during product installation. The installation program also creates redirection rules to readdress packets using the IP address of the proxy machine and default port assignments. The ARM is enabled by default when Content Gateway software is installed.

For the proxy to serve HTTP or FTP requests transparently, you must check the redirection rules in the **ipnat.conf** file and edit them if necessary.

1. Check that the ARM has been enabled by looking at the **Configure > My Proxy > Basic > General** tab. Click **ARM On** if it is not already selected. (ARM is enabled by default during installation.)
2. Navigate to the **Configure > Networking > ARM > General** tab.

The **Network Address Translation (NAT)** section displays the redirection rules in the **ipnat.conf** file. Check the redirection rules and make changes if necessary.

- a. To change a redirection rule, click **Edit File** to open the configuration file editor for the **ipnat.conf** file.
- b. Select the rule you want to edit and then modify the appropriate fields. Click **Set** and then click **Apply** to apply your changes. Click **Close** to exit the configuration file editor.

All fields are described in [ARM](#), page 267.

3. Click **Restart** on **Configure > My Proxy > Basic > General**.

Interception strategies

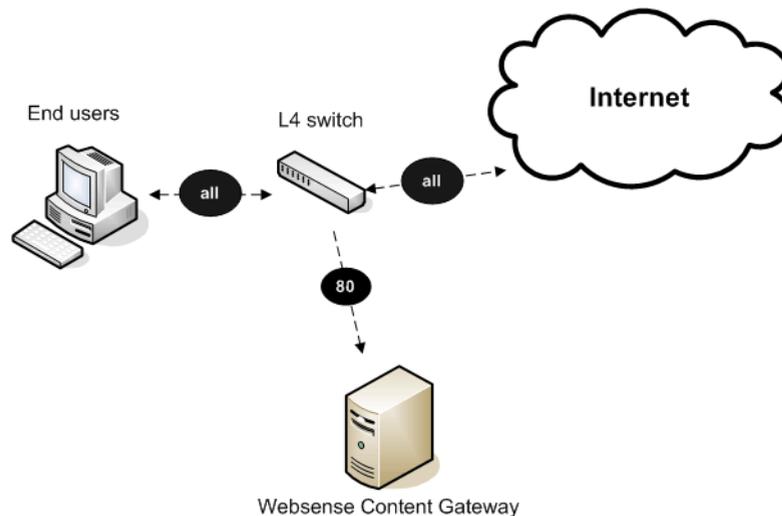
Websense Content Gateway supports the following transparency routing solutions:

- ◆ A Layer 4 switch. See [Using a Layer 4 switch, page 45](#).
- ◆ A Cisco IOS-based router using the Web Cache Control Protocol (WCCP). See [Using a WCCP-enabled router, page 46](#).
- ◆ Policy-based routing. See [Using multicast mode, page 55](#).
- ◆ Software routing. See [Using software-based routing, page 56](#).

How client requests reach Content Gateway depends on network topology. In a complex network, you must decide which clients are to be served transparently and make sure that the proxy is 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.

Using 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, so that the node cache need not be repopulated; this feature is called *cache affinity*.



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. Also, if the switch provides a load-balancing option that has cache affinity, it is recommended that you use management-only clustering mode.

Using a WCCP-enabled router

Related topics:

[ARM bypass and WCCP, page 47](#)

[WCCP load balancing, page 48](#)

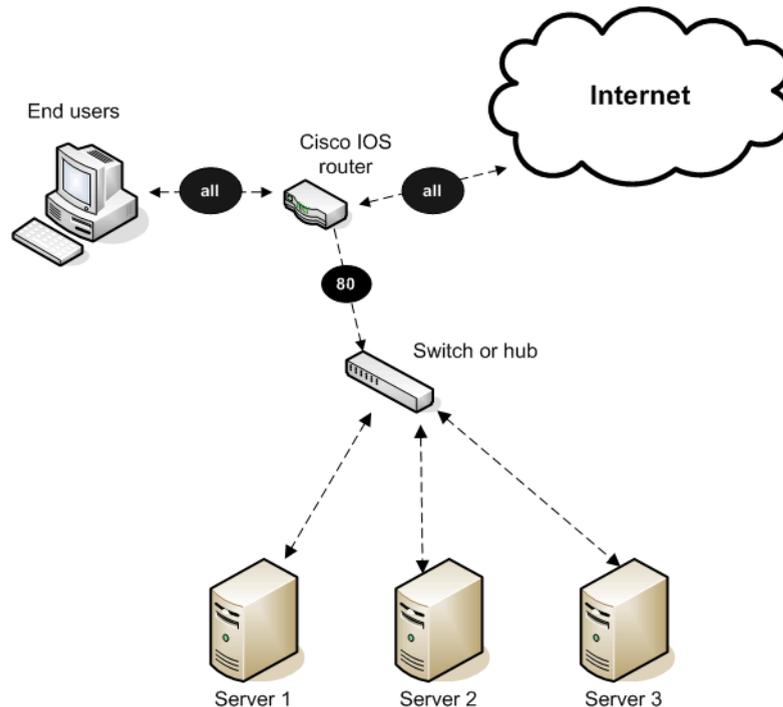
[Configuring WCCP routers, page 49](#)

[Enabling WCCP in Content Gateway Manager, page 52](#)

Content Gateway supports WCCP v1 and WCCP v2.

A WCCP v1-enabled router can send all port 80 (HTTP) traffic to the proxy server, as shown below. The ARM readdresses port 80 to the Content Gateway proxy port (by default, port 8080). The proxy processes the request as usual, retrieving the requested document from the cache if it is a hit and sending the response back to the client. Along the way, the ARM readdresses the proxy port in the response header to port 80 (undoing the readdressing it did on the way to the proxy). The user then sees the response as if it had been sent directly from the origin server.

A WCCP v2-enabled router works in the same way. In addition to port 80 (HTTP) traffic, WCCP v2 supports additional protocols, including FTP and HTTPS.



WCCP provides the following routing benefits:

- ◆ The WCCP-enabled router and the proxy exchange heartbeat messages, letting each other know they are running. The WCCP router reroutes traffic if the proxy server becomes unavailable.
- ◆ If several Content Gateway servers receive traffic from a WCCP router, WCCP balances the load among the servers. The group of servers is called a *WCCP cache farm*. See [WCCP load balancing, page 48](#).
- ◆ Content Gateway handles node failure in WCCP cache farms. If one node becomes unavailable, its load is redistributed among the remaining nodes.
- ◆ In WCCP v2, you can use multiple routers. Traffic flowing through multiple routers can share the same pool of caches.

ARM bypass and WCCP

If Content Gateway has an ARM bypass rule (discussed in [Interception bypass, page 57](#)), Content Gateway forwards particular client requests directly to the origin server, bypassing the cache.

Bypassed requests are unchanged by the ARM; they retain their client source IP addresses. In WCCP v1, ARM bypass rules cannot work if the WCCP router is also the Content Gateway default gateway router. The WCCP router sends port 80 traffic to the proxy server *and* it serves as the servers' default gateway or *next hop* to the Internet.

Bypassed requests go to the WCCP router, which sends them back to Content Gateway.

In WCCP v2, you can exclude certain router interfaces from redirection. Content Gateway bypass rules can work if you exclude the router interface on which Content Gateway is connected from using WCCP. You can do this by setting the router configuration command **ip wccp redirect exclude** in (see *Content Gateway Installation Guide* for information about router configuration).

WCCP load balancing

If a WCCP router serves several nodes, the router balances load among the proxy caching servers. The router sends each node requests aimed at a particular range of IP addresses, so that each node is responsible for caching content residing at particular IP addresses. If a node becomes unavailable, its traffic is redistributed.

Content Gateway also supports cache affinity. If a node becomes unavailable and then recovers, Content Gateway returns the node to its former load distribution. The node's cache does not need to be repopulated.

The WCCP cache farm acts as a simple form of distributed cache, which is sufficient for many applications. A WCCP-enabled network device distributes traffic to individual proxy servers based on the IP address of the origin server. Each node caches objects requested from a set of origin servers that belong to that node's assigned range of destination IP addresses.

Busy origin servers are often mapped to several IP addresses (using a DNS round-robin mechanism). Using WCCP-based load balancing alone, each of these different IP addresses could be allotted to different Content Gateway nodes. This can result in a slightly lower hit rate and wasted cache space, since the same content is being replicated across nodes. The Content Gateway full-clustering mode ensures that all requests to a specific page on that origin server (no matter which IP address is used) are cached on the same node.

With full clustering, objects are distributed among nodes according to their URLs; WCCP distributes objects according to destination IP address. If a particular IP address is receiving many requests, WCCP load balancing might lead to a hot spot, where all of that site's traffic is cached on one node, instead of being distributed among the nodes. Full-clustering mode distributes different pages from the busy site to different Content Gateway nodes.

In general, if load-handling capacity and latency are most important, management-only clustering is recommended in WCCP environments.

If you are running clustered Content Gateway proxy caching servers, it is recommended that you do not enable virtual IP failover in WCCP environments. The Content Gateway WCCP failover mechanism handles node failures and restarts. See [Virtual IP failover](#), page 70, for details.

Configuring WCCP routers

If the proxy must support HTTPS or FTP, you must deploy with a WCCP v2 router. Otherwise, you can deploy either a WCCP v1 or WCCP v2 router.

To use a WCCP v1 router, you need only enable WCCP processing in Content Gateway Manager (see [Enabling WCCP v1](#), page 53). There is just one protocol (HTTP) and port (80), and it is already configured in the router.

To prepare a WCCP v2 router for use with the proxy:

1. Configure a service group for each protocol you intend to use (see [Configuring service groups](#), page 49).
2. Configure the router to enable WCCP processing for these service groups (see [Enabling WCCP processing for a service group](#), page 50).
3. Optionally, enable router security (see [Enabling WCCP v2 security on the router](#), page 51).



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/univercd/cc/td/doc/product/core/> and search for your IOS and device version, for example, IOS 12.2(4) 3750.

4. When you are done configuring the router, you must also enable WCCP processing in Content Gateway Manager. (See [Enabling WCCP v2](#), page 53.)

Configuring service groups

When configuring a WCCP v2 router, you must specify *service group ID numbers* when executing certain commands. (This is not required for WCCP v1.)

A service group ID number represents a group of Content Gateway servers and routers that support a specific protocol. The service group ID also represents the router port and Content Gateway port(s) associated with that particular protocol. For example, service group ID 0 is used for port 80 HTTP traffic, and service group ID 1 is used for port 119.

Following is a table of default service group IDs for WCCP v2. If you are configuring for IP spoofing, see the table in [IP spoofing](#), page 62 for the default reverse service groups IDs.

Service ID	Port	Traffic Type
0	80	HTTP
5	21	FTP
70	443	HTTPS (requires SSL Manager)

You must configure the service group IDs on your router so it can send traffic to Content Gateway servers. Follow the instructions in your router documentation for specifics, but in general:

1. Specify the WCCP version you want to use, for example:

```
ip wccp version 2
```

2. If you used another proxy cache with your router before, disable the service ID that was previously used. For example, if you have a Cisco router, disable the service ID **webcache** by issuing this command:

```
no ip wccp webcache
```

To see what has been configured on the router for WCCP, type:

```
show running-config | include wccp
```

3. 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.

All Content Gateway servers (and routers) assigned to a particular service group must support the protocol associated with that service group. If you want to use a service ID, other than the default, for destination port HTTP traffic, or source port HTTP traffic, you can assign a different service ID.

1. To change the HTTP destination port service group ID, add the variable **proxy.config.wccp2.HTTP_svc_id** to the **records.config** file (default location /opt/WCG/config) and assign the appropriate service ID. For example:

```
CONFIG proxy.config.wccp2.HTTP_svc_id STRING 24
```

2. To change the HTTP source port service group ID, for example to support IP spoofing, add the variable **proxy.config.wccp2.rev_HTTP_svc_id** to the **records.config** file and assign the appropriate service ID. For example:

```
CONFIG proxy.config.wccp2.rev_HTTP_svc_id STRING 20
```

3. Save the file and from the Content Gateway **bin** directory, run:

```
content_line -x
```

to apply the changes.

Note that on a given router, service group IDs are unique; there is only one service ID 0, and only one service ID 1, and so on.

Enabling WCCP processing for a service group

For each service group that you configure (WCCP v2 only), you must enable WCCP processing. WCCP v2 routers contain multiple network interfaces, including:

- ◆ an interface dedicated to outbound traffic, that is, aimed at the Internet
- ◆ one or more interfaces connected to Content Gateway
- ◆ one or more interfaces that receive inbound client traffic

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

1. Turn on the WCCP feature:

```
ip wccp <service group ID> password [0-7] <passwd>
```
2. On the router/switch interface, enable redirection for incoming packets or outgoing packets.

For *incoming* packets:

```
ip wccp <service group ID> redirect in
```

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

```
ip wccp 0 redirect in
```

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

```
ip wccp 20 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
```

Run the command above for each protocol you want to support, but *only for the interface dedicated to inbound traffic*.

For *outgoing* packets:

```
ip wccp <service group ID> redirect out
```

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

```
ip wccp 0 redirect out
```

Enter the following to turn on redirection for HTTPS:

```
ip wccp 70 redirect out
```

To turn on redirection for FTP enter:

```
ip wccp 5 redirect out
```

Run the command above for each protocol you want to support, but *only for the interface dedicated to outbound traffic*.

3. Turn off redirection for packets received on the interface:

```
ip wccp redirect exclude in
```

Run the command above when dynamic or static bypass are enabled. With bypass enabled, the proxy redirects bypassed traffic to the Internet. This command prevents looping back of packets by stopping the router from redirecting this same bypassed traffic back to Content Gateway. Run the command above for each interface connected to the proxy.

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 on your routers, you can enable security on the Content Gateway node so that the proxy caching server and your routers can authenticate each

other. You must enable security for each service group that the router supports individually. You cannot configure a router globally as you would Content Gateway.

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

The authentication password you specify must match the authentication password configured on the router for each service group being intercepted (HTTP, FTP, SOCKS). The following procedure provides an example of how to set an authentication password for different service groups on a WCCP v2-enabled router.

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 processing 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 password
```

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

**Important**

Content Gateway supports only one password. Therefore, you must use the same password for all service groups. Passwords longer than 8 characters are truncated.

4. Exit and save the router configuration.

Enabling WCCP in Content Gateway Manager

Related topics:

[Enabling WCCP v1, page 53](#)

[Enabling WCCP v2, page 53](#)

[Configuring WCCP routers, page 49](#)

[Enabling WCCP v2 security on the router, page 51](#)

Once you have configured your WCCP router to send traffic to the proxy server, you must also enable WCCP on Content Gateway. You do this through the Content Gateway Manager.



Important

Before you enable WCCP, make sure that your configuration meets the following requirements:

- ◆ The WCCP router is running the correct Cisco IOS release.
- ◆ If you are using several Content Gateway nodes, determine whether you want the nodes to have management-only clustering or full clustering. See [WCCP load balancing, page 48](#) and [Clusters, page 65](#).
- ◆ Each proxy server must have the ARM enabled. Ensure that ARM is **On** on **Configure > My Proxy > Basic > General**.

Enabling WCCP v1

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the **Features** table, click **WCCP On** in the **Networking** section.
3. Click **Apply**.
4. Navigate to **Configure > Networking > WCCP > General**.
5. Enable the **WCCP v1.0** option, and then click **Apply**.
6. Select the **WCCP v1.0** tab.
7. In the **WCCP Network Interface** field, enter the interface on the Content Gateway system that receives traffic from the router.
8. In the **WCCP Router IP Address** field, enter the IP address of the router that sends traffic to Content Gateway.
9. Click **Apply**.
10. Click **Restart** on **Configure > My Proxy > Basic > General**.

Enabling WCCP v2

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the **Features** table, click **WCCP On** in the **Networking** section.
3. Click **Apply**.
4. Navigate to the **Configure > Networking > WCCP > General** tab.
5. Enable the **WCCP v2.0** option, and then click **Apply**.
6. Select the **WCCP v2.0** tab.
7. In the **WCCP Network Interface** field, enter the interface on the Content Gateway system that receives traffic from the router.
8. (*Optional.*) In the **Security** area, enable the **Security** option if you want Content Gateway and your routers to authenticate each other.

In the **Password** field, enter the password used for service group authentication with the router. The password must match the authentication password for the service group specified on the router and must be a maximum of eight characters long. See [Enabling WCCP processing for a service group](#), page 50, step 1. See also [Enabling WCCP v2 security on the router](#), page 51.

9. To run in multicast mode, enable the **Multicast** option in the **Multicast** area, and then enter the multicast IP address for WCCP in the **Multicast IP Address** field. The multicast IP address must match the multicast IP address specified on the router. See [Using multicast mode](#), page 55.
10. To run in unicast mode, specify the number of WCCP routers that direct traffic to Content Gateway in the **Number of WCCP Routers** field of the **Routers** area. Then specify the IP addresses of each router in the fields provided. For more information on unicast mode, see the *Content Gateway Installation Guide*.
11. In the **Services Enable/Disable** section, make sure that the appropriate services are enabled so that the router is allowed to send HTTP, FTP, SOCKS, or DNS traffic to the proxy.
12. In the **Mode Negotiation** section, select the preferred Packet Forward Method, Assignment Method, and Packet Return Method.

The Packet Forward Method determines how intercepted traffic is transmitted from the WCCP server to the proxy.

The Packet Return Method specifies the method used to return intercepted traffic to the WCCP server.

The Assignment Method is the method used by the router to distribute intercepted traffic across multiple proxies.

The proxy always adjusts to use a method that the router supports. If the router supports more than one method, the preferred method specified in this area is used.

Cisco provides guidance on the best methods for different classes of devices. For example, Cisco recommends the following methods with WCCP on software-based platforms (includes ISR and 7200 series platforms):

- GRE
- Hash

Cisco recommends these methods with WCCP on hardware-based platforms (includes Catalyst switches and 7600 series routers):

- L2
- Mask

13. Click **Apply**.

- Click **Restart** on **Configure > My Proxy > Basic > General**.

**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.

**Note**

It may take more than 30 seconds for the router to report that a new proxy caching server has joined a service group.

Using multicast mode

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

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.

- Telnet to the router and switch to Enable mode.
- At the prompt, enter the following command to configure the router from the terminal:

```
configure terminal
```

- 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
```

where *hostname* is the host name 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
```

where *interface_name* is the network interface on the router that is being intercepted and redirected.

- At the prompt, enter the following command for each service group that the router intercepts:

```
hostname(config-if)# ip wccp service_group group-listen
```

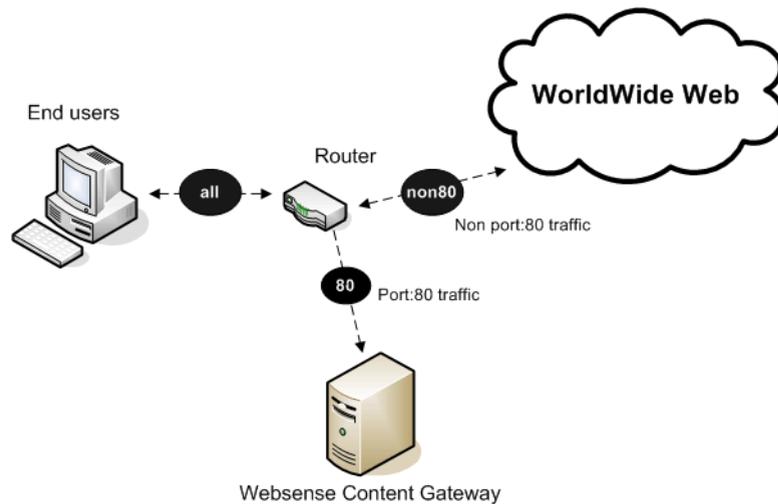
- Exit and save the router configuration.

Using 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 translates intercepted requests into Content Gateway requests.
- ◆ Translated requests are sent to the proxy.
- ◆ Web objects to be served transparently are readdressed 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, page 70](#).



Using 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; it routes port 80 HTTP requests to the proxy cache.
- ◆ The ARM translates intercepted requests into proxy requests.
- ◆ Translated requests are sent to the proxy.
- ◆ Web objects to be served transparently are readdressed 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, page 70](#).) The Content Gateway cluster failover mechanism is similar to the Hot Standby Router Protocol (HSRP).

Interception bypass

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

- ◆ Client software bugs (customized, non-commercial browsers).
- ◆ Server software bugs.
- ◆ Applications that send non-HTTP traffic over HTTP ports as a way of defeating security restrictions.
- ◆ Server IP 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. However, Content Gateway contains an adaptive learning module that recognizes interoperability problems caused by transparent proxy caching 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 the cache when it detects non-HTTP traffic

on port 80 or when it encounters certain HTTP errors. See [Dynamic bypass rules](#), page 58.

- ◆ *Static* bypass rules must be manually configured in the **bypass.config** file. See [Static bypass rules](#), page 59.



Note

Do not confuse bypass rules with client access control lists. Bypass rules are generated in response to interoperability problems. Client access control is simply restriction of the client IP addresses that can access the proxy cache, as described in [Controlling client access to the proxy cache](#), page 139.

Dynamic bypass rules

Related topics:

[Setting dynamic bypass rules](#), page 59

[Viewing dynamic bypass statistics](#), page 59

When configured to do so, the proxy watches for protocol interoperability errors. As it detects errors, it configures the ARM to bypass the proxy for those clients and/or servers causing the errors.

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 caching server so that they can continue to function (but without caching).

You can configure the proxy to dynamically bypass the cache for any of the following 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

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.

In another example, 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.

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](#), page 60.

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](#), page 291.

Setting dynamic bypass rules

By default, Content Gateway is not configured to bypass the cache when it encounters HTTP errors or non-HTTP traffic on port 80. You must enable dynamic bypass rules by setting the appropriate options.

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the **Features** table, make sure that **ARM** is **On** in the **Networking** section.
3. Navigate to **Configure > Networking > ARM > Dynamic Bypass**.
4. Enable the **Dynamic Bypass** option.
5. In the **Behavior** section, select the dynamic bypass rules you want to use.
6. Click **Apply**.
7. Click **Restart** on the **Configure > My Proxy > Basic > General** tab.

Viewing dynamic bypass statistics

Content Gateway tallies bypassed requests for each type of dynamic bypass trigger. For example, Content Gateway counts all requests that are bypassed in response to a 401 error.

- ▶ Navigate to **Monitor > Networking > ARM**.

The statistics are displayed in the **HTTP Bypass Statistics** section of the table.

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.

For example, you might want client IP addresses that did not pay for a caching service to be steered around the cache, while paying clients can obtain the benefits of caching.

Or you might want to remove some servers from caching lists because they do not want to have their pages cached.

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. For example, you can use this solution to bypass clients who want to opt out of a caching solution.
- ◆ 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. Destination bypass rules prevent Content Gateway from caching an entire site. You will experience hit rate impacts if the site you bypass is popular.
- ◆ 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 when cached.

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](#), page 291).

Viewing the current set of bypass rules

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

To view all current dynamic and static bypass rules:

1. Log on to a Content Gateway node as the Content Gateway administrator, and then change to the Content Gateway **bin** directory (default location is in **/opt/WCG**).
2. Enter the following command at the prompt and press **Return**:

```
./print_bypass
```

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

Configuring ARM security

To prevent unauthorized access to machines running Content Gateway, you can configure the ARM to use an access control list employing administrator-specified rules to allow or deny other computers from communicating with the machine. This enables you to effectively create a firewall in front of Content Gateway, denying

potentially malicious packets from reaching the TCP/IP stack on the machine. See [Controlling host access to the proxy server](#), page 140.

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.

1. Navigate to **Configure > Networking > Connection Management > Load Shedding**.
2. 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. Click **Restart** on **Configure > My Proxy > Basic > General**.

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 driver. Content Gateway then uses that IP address to determine the origin server instead of doing a DNS lookup on the host name of the request. Because the client already performed a DNS lookup, Content Gateway does not have to.



Note

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. See [How do you configure Content Gateway to serve only transparent requests?](#), page 400, for information about running Content Gateway in transparent proxy mode only. In explicit proxy mode, the client does not perform a DNS lookup on the host name of the origin server, so the proxy must perform a DNS lookup.

In addition, do not enable the Always Query Destination option if you want domain names, rather than IP addresses, to be captured in the log server.

To enable the Always Query Destination option:

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is in **/opt/WCG**).
2. Edit the following variable:

Variable	Description
<code>proxy.config.arm.always_query_dest</code>	Set to 0 to disable the Always Query Destination option. Domain names are captured. Set to 1 to enable the Always Query Destination option. IP addresses are captured; domain names are not.

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

```
content_line -x
```

IP spoofing

The IP spoofing option configures the proxy to always use the IP address of the client when communicating with the origin server, instead of the proxy's own IP address. As a result, HTTP requests appear to come from the client rather than the proxy.

IP spoofing is applied to HTTP traffic only.

IP spoofing affects both explicit and transparent traffic.

Use of the ARM is required with IP spoofing. Before IP spoofing can be enabled, ARM must be enabled. (ARM translates the IP addresses of outgoing requests and incoming responses.)



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.

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

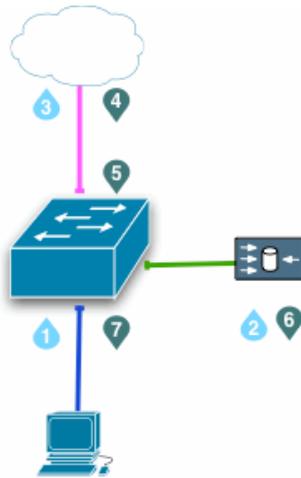


Important

For a discussion of how the proxy kernel routing table impacts transparent proxy deployment, see Knowledge Base article 4747, "Web sites in the Static or Dynamic bypass list fail to connect".

IP spoofing and the flow of HTTP traffic

The following describes the flow of HTTP traffic when IP spoofing is used with WCCP. Policy-based routing can be implemented to achieve the same results. The numbers in the diagram correspond to the actions described in the numbered list.



1. A client request arrives at a routed port or Switched Virtual Interface (SVI) looking for traffic with a destination port of HTTP (80).
2. The switch redirects the client request to Content Gateway (the proxy), applying a traditional transparent proxy deployment. The proxy internally routes the traffic to port 8080 (for HTTP) of its own IP address.
If needed, the proxy creates a connection to the Web origin server using the original client IP address.
3. The request is sent to the Web origin server through the switch, NAT and/or firewall.
4. When the origin server response is returned, the IP packet has the client IP address as the destination.
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). See the note below.
6. The switch redirects the origin server response to the proxy, completing the proxy-to-Web server TCP connection.
7. A proxy response to the client is generated and returned to 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.

The default WCCP service group IDs are:

Service ID	Port	Traffic Type
0	destination port 80	HTTP
5	21	FTP
20	source port 80	HTTP
70	443	HTTPS (requires SSL Manager)

Enabling IP spoofing:

1. Navigate to **Configure > Networking > ARM > General**.
(If ARM is not a choice, go to **Configure > My Proxy > Basic > General** and ensure that ARM is **On**.)
2. Select **IP Spoofing**.
3. Click **Apply**.
4. Click **Restart** on **Configure > My Proxy > Basic > General**.

For information about configuring WCCP routers, see [Configuring WCCP routers](#), page 49.

6

Clusters

Related topics:

[Changing clustering configuration, page 67](#)

[Adding nodes to a cluster, page 68](#)

[Deleting nodes from a cluster, page 70](#)

[Virtual IP failover, page 70](#)

Websense Content Gateway scales from a single node to multiple nodes that form a cluster, allowing you to improve system performance and reliability. The nodes in a cluster share configuration information and can form a single logical cache or Content Gateway node.

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.

When the *virtual IP failover* feature (see [Virtual IP failover, page 70](#)) is enabled, the live nodes in a cluster can assume a failed node's responsibilities.

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



Important

To establish communication between cluster nodes, network routing tables must be configured correctly. When clustering is configured during the installation process (see the *Content Gateway Installation Guide*), Content Gateway adds the appropriate route to the server's routing table. However, if clustering is enabled after installation, or clustering settings must change, you must manually update your server's routing table. See [Changing clustering configuration, page 67](#).

If Content Gateway is running on a V-Series Appliance, please see the *V-Series Appliance Manager Help* for instructions on how to modify routes.

Content Gateway has two clustering modes:

- ◆ Management-only mode
- ◆ Full-clustering mode

Both modes require a dedicated network interface for cluster communication.



Important

In a proxy hierarchy, the nodes in the cluster cannot be ICP siblings or be a mixture of HTTP parents and children; you must configure each node in a Content Gateway cluster as a single node in the hierarchy, because they share a common configuration.

Management-only clustering

In management-only clustering mode, Content Gateway cluster nodes share configuration information. You can administer all the nodes at the same time.

Content Gateway uses a multicast management protocol to provide a single system image of your Content Gateway cluster. Information about cluster membership, configuration, and exceptions is shared across all nodes, and the **content_manager** process automatically propagates configuration changes to all nodes.



Important

A second network interface is required for cluster communication.

Full clustering

In full-clustering mode, in addition to sharing configuration information, a Content Gateway cluster distributes its cache across its nodes into a single, virtual object store, rather than replicating the cache node by node. Content Gateway can provide an aggregate cache size and can maximize the cache hit rate by storing objects only once across the entire cluster.

A fully clustered Content Gateway maps objects to specific nodes in the cluster. When a node receives a request, it checks to see if the request is a hit somewhere in the cluster. If the request is a hit on a different node, the node handling the request fetches the object from the hit node and serves it to the client. Content Gateway uses a proprietary communication protocol to fetch an object from other cluster nodes.

If a node fails or is shut down and removed, Content Gateway removes references to the missing node on all nodes in the cluster. If virtual IP failover is enabled (see

[Virtual IP failover, page 70](#)), requests destined for the missing node are handled by another node.



Important

A second network interface is required for cluster communication.

Changing clustering configuration

Clustering is usually configured when you install the proxy. You can, however, configure clustering through Content Gateway Manager as well.

1. At the command line, enter the following:

```
route add multicast.group address/32 dev interface_name
```

where *interface_name* is the name of the interface used for cluster communication. For example:

```
route add 224.0.1.37/32 dev eth1
```

You need the multicast group address and interface name to configure through Content Gateway Manager.



Important

If Content Gateway is running on a V-Series Appliance, please see the *V-Series Appliance Manager Help* for instructions on how to modify routes.

2. In Content Gateway Manager, go to the **Configure > My Proxy > Basic > Clustering** tab.
3. In the **Cluster Type** area, select the clustering mode:
 - Select **Single Node** if this node is not part of a cluster.
 - Select **Management Clustering** if cluster nodes share configuration information only.
 - Select **Full Cache Clustering** to share configuration information among cluster nodes and to distribute the cache across all the cluster nodes into a single, virtual object store.
4. For management or full cache clustering, enter the communication interface and the multicast group address. These are the same interface and address values that you entered in [Step 1](#).



Important

Content Gateway does not force clustering communication on this interface; it follows the network routing table for your server or V-Series Appliance.

5. Click **Apply**.
6. Click **Restart** on **Configure > My Proxy > Basic > General**.

**Note**

Content Gateway does not apply the clustering mode change to all 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 run the same version of the same operating system, and Content Gateway must be installed in the same directory (including path) on each node. For example, the default installation directory for Content Gateway is **/opt/WCG**. If the proxy is installed elsewhere on one node in the cluster, either that node should have the proxy reinstalled at the default location, or all nodes must have the proxy installed in the alternate location.

1. Install the appropriate hardware and connect it to your network. (Consult the hardware documentation for hardware installation instructions.)
2. Install the Content Gateway software using the appropriate procedure for installing a cluster node. See the *Content Gateway Installation Guide*. 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 cluster.
 - The port assignments for the new node are the same as the port assignments used by the other nodes in the cluster.
 - You have added multicast addresses and multicast route settings.
3. Restart Content Gateway. See [Starting and stopping Content Gateway](#), page 9.

If you have an existing Content Gateway installation and you want to add that server to the cluster, you do not have to reinstall the Content Gateway on the node. Instead, you can edit configuration variables on the existing Content Gateway node.

1. Open the **records.config** file located in the **config** directory (default location is **/opt/WCG/config**) on the node you want to add to the cluster.
2. Edit the following variables:

Variable	Description
<code>proxy.config.cluster.type</code>	Specify the clustering mode used by the other nodes in the cluster: 1 for full-clustering mode 2 for management-only mode
<code>proxy.config.proxy_name</code>	Specify the name of the Content Gateway cluster. All nodes in a cluster must use the same name.
<code>proxy.config.cluster.mc_group_addr</code>	Specify the multicast address for cluster communications. All nodes in a cluster must use the same multicast address.
<code>proxy.config.cluster.rsport</code>	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.
<code>proxy.config.cluster.mcport</code>	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 number is 8088.
<code>proxy.config.cluster.ethernet_interface</code>	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**).
5. To change from Full mode to Management mode or vice versa:
 - a. Access Content Gateway Manager.
 - b. Navigate to **Configure > My Proxy > Basic > Clustering**.
 - c. In the **Cluster Type** area, select the appropriate type (**Full** or **Management**).
 - d. Click **Apply**.
 - e. Click **Restart** on **Configure > My Proxy > Basic > General**.

Deleting nodes from a cluster

On the node you want to remove from the cluster:

1. Navigate to **Configure > My Proxy > Basic > Clustering**.
2. In the **Cluster Type** area, select **Single Node**.
3. Click **Apply**.
4. Click **Restart** on **Configure > My Proxy > Basic > General**.

Virtual IP failover

Through the virtual IP failover feature, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes in the cluster as necessary. 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 reassignment is done through a process called *ARP rebinding*.

What are virtual IP addresses?

Related topics:

[Enabling and disabling virtual IP addressing, page 71](#)

[Adding and editing virtual IP addresses, page 71](#)

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.

Enabling and disabling virtual IP addressing

1. Navigate to **Configure > My Proxy > Basic > General**.
2. Under the Networking section in the Features table, select **On** or **Off** for **Virtual IP** to enable or disable Virtual IP addressing.
3. Click **Apply**.
4. Click **Restart** on **Configure > My Proxy > Basic > General** to restart Content Gateway on all the nodes in the cluster.

Adding and 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.

1. Navigate to **Configure > Networking > Virtual IP**.

The **Virtual IP Addresses** area displays the virtual IP addresses managed by Content Gateway.



Note

The Virtual IP button is displayed only if you have enabled the Virtual IP option in the Features table on **Configure > My Proxy > Basic > General**.

2. Click **Edit File** to add new or edit existing virtual IP addresses.

3. 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**.

4. Click **Apply**, and then click **Close**.
5. Click **Restart** on **Configure > My Proxy > Basic > General**.

7

Hierarchical Caching

Websense Content Gateway can participate in cache hierarchies, where 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.

Content Gateway can be a member of the following cache hierarchies:

- ◆ An HTTP cache hierarchy (described in [HTTP cache hierarchies, page 73](#))
- ◆ An ICP (Internet Cache Protocol) cache hierarchy (described in [ICP cache hierarchies, page 75](#))

HTTP cache hierarchies

In an HTTP cache hierarchy, if a Content Gateway node cannot find a requested object in its cache, it can search a parent cache—which itself can search other caches—before resorting to retrieving the object from the origin server.

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 in [Parent failover, page 74](#).



Note

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, page 314](#)).



Note

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 cached and served to the client.

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, page 314](#). By default, the parent caches are queried in the order in which they are listed in the configuration file.

Configuring Content Gateway to use an HTTP parent cache

1. On the **Configure > Content Routing > Hierarchies > Parenting** page, enable **Parent Proxy**.
2. Click **Edit File** to open the configuration file editor for the [parent.config](#) file.
3. Enter information in the fields provided, and then click **Add**. All the fields are described in Appendix C: [Configuration Options](#).
4. Click **Apply**, and then click **Close**.
5. 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.

ICP cache hierarchies

The Internet Cache Protocol (ICP) is a protocol used by proxy caches to exchange information about their content. ICP query messages ask other caches if they are storing a particular URL. ICP response messages reply with a hit (the URL is stored) or miss (the URL is not stored) answer.

A cache exchanges ICP messages only with specific ICP peers, which are neighboring caches that can receive ICP messages. An ICP peer can be a sibling cache, which is at the same level in the hierarchy, or a parent cache, which is one level up in the hierarchy.

If Content Gateway has ICP caching enabled, it sends out ICP queries to its ICP peers in the event of a cache miss on an HTTP request. If there are no hits and parents exist, a parent is selected using a round-robin policy. If no ICP parents exist, the proxy forwards the request to its HTTP parents. If there are no HTTP parent caches, the proxy forwards the request to the origin server.



Note

If Content Gateway receives a hit message from an ICP peer, the proxy sends the HTTP request to that peer.

Under some circumstances, what appears to be a cache hit is really a miss; this happens when the original HTTP request contains header information that is not communicated by the ICP query. For example, the hit might not be the requested alternate.

If an ICP hit turns out to be a miss, the proxy forwards the request to either its HTTP parent caches or to the origin server.

Configuring Content Gateway to use an ICP cache hierarchy

1. Navigate to the **Configure > Content Routing > Hierarchies > ICP Peering** tab.
2. In the ICP mode area, select:
 - **Only Receive Queries** to configure Content Gateway to only receive ICP queries from other ICP peers.
 - **Send/Receive Queries** to configure Content Gateway to both send and receive ICP queries.
 - **Disabled** to configure Content Gateway to neither send nor receive ICP queries.
3. Enter the name of the interface to use for ICP messages. This field is populated automatically, but if the machine is running with more than one interface, you can specify a different one here.

4. Enter the port to use for ICP messages. The default is 3130.



Note

If clustering is **not** enabled, each peer must be configured to use the same port. If clustering is enabled, then ports are automatically kept in sync.

5. Enable the **ICP Multicast** option to send ICP messages through multicast if your proxy server has a multicast channel connection to its ICP peers.
6. Enter the timeout for ICP queries. The default is 2 seconds.
7. Click **Edit File** to open the configuration file editor for the **icp.config** file.
8. Enter information in the fields provided, and then click **Add**. All the fields are described in Appendix B: *Configuration Options*.
9. Click **Apply**, and then click **Close**.
10. On the **ICP Peering** tab, click **Apply** to save your configuration.

8

Configuring the Cache

The cache consists of a high-speed object database called the **object store**. The object store indexes objects according to URLs and associated headers, enabling Websense Content Gateway to store, retrieve, and serve Web pages, and also parts of Web pages, providing optimum bandwidth savings. Using object management, the object store can cache alternate versions of the same object, varying on spoken language or encoding type, and can store small and large documents, minimizing wasted space. When the cache is full, Content Gateway removes stale data.

Content Gateway can tolerate disk failures on any cache disk. If the disk fails, Content Gateway marks the disk as corrupt and continues using the remaining disks. An alarm is sent to Content Gateway Manager, indicating which disk failed. If all cache disks fail, Content Gateway goes into proxy-only mode.

You can perform the following cache configuration tasks:

- ◆ Change the total amount of disk space allocated to the cache. See [Changing cache capacity, page 77](#).
- ◆ Partition the cache by reserving cache disk space for specific protocols and origin servers and domains. See [Partitioning the cache, page 80](#).
- ◆ Specify a size limit for objects allowed in the cache. See [Configuring cache object size limit, page 81](#).
- ◆ Delete all data in the cache. See [Clearing the cache, page 82](#).
- ◆ Change the size of the RAM cache. See [Changing the size of the RAM cache, page 82](#).

RAM cache

Content Gateway maintains a small RAM cache of popular objects. This RAM cache serves the most popular objects as fast as possible and reduces load on disks, especially during temporary traffic peaks. You can configure the RAM cache size. See [Changing the size of the RAM cache, page 82](#).

Changing cache capacity

The maximum aggregate disk cache size is limited to 147 GB. This size makes best use of system resources, while also providing an excellent end-user experience.

Because the rapidly growing majority of Internet sites are composed of dynamic, uncacheable content, a large disk cache does not improve the performance of an end-user's Web browsing experience.

The minimum disk cache size is 2 GB.

Related topics:

[Querying cache size, page 78](#)

[Increasing cache capacity, page 78](#)

[Reducing cache capacity, page 79](#)

Querying cache size

To view the configured aggregate cache size, open the Content Manager and go to **Monitor > Subsystems > Cache**. The cache size is displayed, in bytes, in the **Current Value** column of the **Cache Size** field.

Alternatively, display the cache size with the following command, executed from the Content Gateway **bin** directory (default location: **/opt/WCG/bin**).

```
content_line -r proxy.process.cache.bytes_total
```

Increasing cache capacity

To increase the total disk space allocated to the cache on existing disks, or to add new disks to a Content Gateway node:

1. Stop Content Gateway. See [Starting and stopping Content Gateway, page 9](#).
2. Add hardware, if necessary.

- a. Set up the raw device and modify the permissions. For example:

```
mknod /etc/udev/devices/raw c 162 0
chmod 600 /etc/udev/devices/raw
```

- b. Identify the physical device name and note the size in bytes (used later). For example:

```
fdisk -l | grep "^Disk"
```

```
Disk /dev/cciss/c0d1: 146.7 GB, 146778685440 bytes
```

- c. For each real disk, create a node, change the owner of the node, and map that raw node to a physical disk. Note that the final argument increments by 1 for each disk added.

To create a node:

```
mknod /etc/udev/devices/raw_c0d1 c 162 1
```

You can change the device name to the name that is returned from the **fdisk -l** command in step b.

To change the owner:

```
chown Websense /etc/udev/devices/raw_c0d1
```

The owner is the installation user (default is Websense). Use the device name used in the mknod statement.

To map the raw node to a physical disk:

```
/usr/bin/raw /etc/udev/devices/raw_c0d1 /dev/cciss/c0d1
```

Use the device name used in the mknod statement.

- d. Add the same **/usr/bin/raw** commands to the **/etc/init.d/content_gateway** file to make the changes effective on reboot. For example, at line 6 add:

```
...
case "$1" in
'start')
/usr/bin/raw /etc/udev/devices/raw_c0d1 /dev/cciss/c0d1
```

3. Edit the **storage.config** file located in the Content Gateway **config** directory (default location is **/opt/WCG/config**) to increase the amount of disk space allocated to the cache on existing disks or add the new disk devices. See [storage.config](#), page 379.
4. Restart Content Gateway.

Reducing cache capacity

You can reduce the total amount of disk space allocated to the cache on an existing disk or remove disks from a Content Gateway node.

1. Stop Content Gateway.
2. Remove hardware, if necessary.
3. Edit the **storage.config** file to reduce the amount of disk space allocated to the cache on existing disks or to delete the reference to the hardware you are removing. See [storage.config](#), page 379.
4. If you remove a disk, you must edit the **/etc/rc.d/init.d/content_gateway** file to remove the raw disk binding for the disk.
5. Restart Content Gateway.



Important

In the **storage.config** file, a formatted or raw disk must be at least 2 GB.

Partitioning the cache

You can manage your cache space more efficiently and restrict disk usage by creating cache partitions of different sizes for specific protocols. You can further configure these partitions to store data from specific origin servers and domains.

**Important**

HTTP is the only protocol supported at this time.

**Important**

The partition configuration must be the same on all nodes in a cluster.

Cache partitioning is specified in the [partition.config](#) file.

Making changes to partition sizes and protocols

After you have configured your cache partitions based on protocol, you can make changes to the configuration at any time. Before making changes, note the following:

- ◆ You must stop Content Gateway before you change the cache partition size and protocol assignment.
- ◆ When you increase the size of a partition, the contents of the partition are *not* deleted. However, when you reduce the size of a partition, the contents of the partition *are* deleted.
- ◆ When you change the partition number, the partition is deleted and then re-created, even if the size and protocol type remain the same.
- ◆ When you add new disks to your Content Gateway node, the partition sizes specified in percentages increase proportionately.
- ◆ A lot of changes to the partition sizes might result in disk fragmentation, which affects performance and hit rate. It is recommended that you clear the cache (see [Clearing the cache, page 82](#)) before making many changes to cache partition sizes.

Partitioning the cache according to origin server or domain

After you have partitioned the cache according to size and protocol, you can assign the partitions you created to specific origin servers and domains.

You can assign a partition to a single origin server or multiple origin servers. However, if a partition is assigned to multiple origin servers, there is no guarantee on the space available in the partition for each origin server. Content is stored in the partition according to popularity.

In addition to assigning partitions to specific origin servers and domains, you must assign a generic partition to store content from all origin servers and domains that are not listed. This generic partition is also used if the partitions for a particular origin server or domain become corrupt.

**Important**

If you do not assign a generic partition, Content Gateway runs in proxy-only mode.

**Note**

You do *not* need to stop Content Gateway before you assign partitions to particular hosts or domains. However, this type of configuration can cause a spike in memory usage and is time consuming. It is recommended that you configure partition assignment during periods of low traffic.

You can partition the cache according to host name and domain in Content Gateway Manager.

In Content Gateway Manager:

1. Configure the cache partitions according to size and protocol, as described in [partition.config](#), page 316.

You should create a separate partition based on protocol (HTTP only) for each host and domain, and an additional generic partition to use for content that does not belong to these origin servers or domains. For example, if you want to separate content from two different origin servers, you must have at least three separate partitions: one HTTP-based partition for each origin server and a generic partition for all other origin servers not listed (the partitions do not have to be the same size).

2. On the **Configure** tab, click **Subsystems**, and then click **Cache**.
3. Click the **Hosting** tab and in the **Cache Hosting** area, click **Edit File** to open the configuration file editor for the **hosting.config** file.
4. Enter information in the fields provided, and then click **Add**. All the fields are described in [Cache](#), page 260.
5. Click **Apply**, and then click **Close**.

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.

1. Select **Configure > Subsystems > Cache > General**.

2. 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**.

Clearing the cache

When you clear the cache, you remove all data from the entire cache, which includes the data in the host database. Clear the cache before performing certain cache configuration tasks, such as partitioning.



Note

You cannot clear the cache when Content Gateway is running.

1. Stop Content Gateway. See [Starting and stopping Content Gateway, page 9](#).
2. Enter the following command to clear the cache:

```
content_gateway -Cclear
```



Warning

The **clear** command deletes all data in the object store and the host database. Content Gateway does *not* prompt you to confirm the deletion.

3. Restart Content Gateway.

Changing the size of the RAM cache

Websense Content Gateway provides a dedicated RAM cache for fast retrieval of popular small objects. The default RAM cache size is calculated based on the number

and size of the cache partitions you have configured. 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.

**Note**

If you have partitioned your cache according to protocol or hosts, the size of the RAM cache for each partition is proportional to the size of that partition.

1. Select **Configure > Subsystems > Cache > General**.
2. In the **Ram Cache Size** field, enter the amount of space (in megabytes) you want to allocate to the RAM cache. Although the user interface will accept larger values, **do not exceed 512 MB**.

**Note**

The default value of -1 means that the RAM cache is automatically sized at approximately 1 MB per GB of disk, with a cap of 512 MB.

3. Click **Apply**.
4. Click **Restart** on **Configure > My Proxy > Basic > General**.

9

Configuring the System

Websense Content Gateway provides several options for configuring the system:

- ◆ [Content Gateway Manager](#), page 85
- ◆ [Command-line interface](#), page 89
- ◆ [Configuration files](#), page 89
- ◆ [Saving and restoring configurations](#), page 90

We recommend that you restart Content Gateway any time you make configuration changes.

Content Gateway Manager

Use Configure mode to view and change your Content Gateway configuration.



Note

Certain options can be changed only by editing configuration variables either in the **records.config** file or from the command-line interface. See [Command-line interface](#), page 89 and [Configuration files](#), page 89.

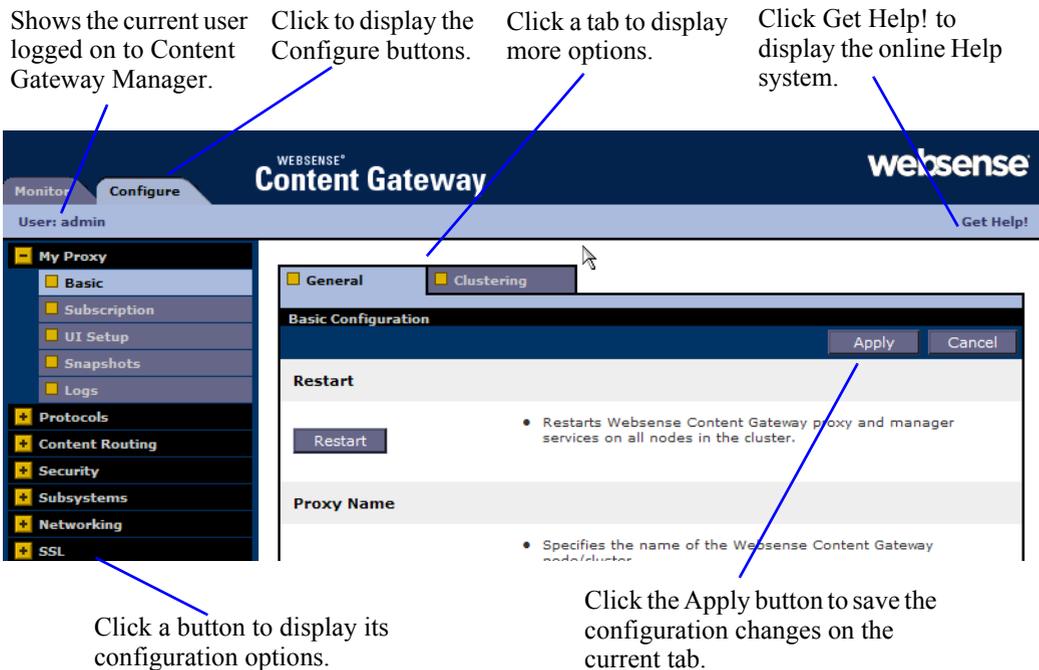
Starting Configure mode

1. Open your Web browser.
Content Gateway Manager requires Java and JavaScript; be sure to enable Java and JavaScript in your browser.
2. Enter the following location in your browser:
`https://nodename:adminport`
where **nodename** is the name of the node and **adminport** is the number assigned to the Content Gateway Manager port (8081 by default).

3. If necessary, log on to Content Gateway Manager with the administrator ID and password, or use your user account. The administrator ID and password are set during product installation. You can change the ID and password, as well as create and modify user accounts. For more information, refer to [Controlling access to Content Gateway Manager](#), page 141.

Content Gateway Manager starts by default in Monitor mode.

4. Click the **Configure** tab to display the Configure mode buttons.



Using Configure mode

In Configure mode, Content Gateway Manager displays a series of buttons. Each button represents a group of configuration options.

All the configuration options available in Configure mode are described in [Configuration Options](#).

My Proxy

- ◆ Click **Basic** to restart the proxy and manager services (you need to restart after changing certain configuration options), identify the name of the Content Gateway node, set alarm email, and enable or disable features (such as FTP processing, proxy authentication, ARM, WCCP, cluster options, and so on).
- ◆ Click **Subscription** to enter your subscription key and provide Websense Web Filter information if your subscription includes the scanning options. See the Web Security Manager Help system for more information on the scanning options.
- ◆ Click **UI Setup** to identify and change the port on which browsers connect to Content Gateway Manager, enable SSL connections to Content Gateway

Manager, specify how often Content Gateway Manager refreshes the statistics on the Monitor tab, and configure access control lists, administrator accounts, and user accounts to secure Content Gateway Manager access.

- ◆ Click **Snapshots** to take and restore configuration snapshots.
- ◆ Click **Logs** to display, delete, or copy a selected log file to the local filesystem.

Protocols

- ◆ Click **HTTP** to configure HTTP caching and tune HTTP timeouts.
- ◆ Click **HTTP Responses** to specify which HTTP responses are sent to clients when the proxy detects an HTTP problem with a client transaction (such as unavailable origin servers, authentication requirements, and protocol errors).
- ◆ Click **HTTP Scheduled Update** to configure the proxy to load specific objects into the cache at scheduled times.
- ◆ Click **FTP** to configure FTP caching and to tune FTP timeouts.

The FTP options affect requests that originate from FTP clients only. You can configure options that affect FTP requests originating from HTTP clients from the HTTP group. The FTP button appears only if you have enabled FTP processing in the **Features** table on **Configure > My Proxy > Basic > General**.
- ◆ Click **HTTPS** to specify port information for inbound and outbound HTTPS traffic.

Content Routing

- ◆ Click **Hierarchies** to configure parent caching and ICP peer options.
- ◆ Click **Mapping and Redirection** to set URL remapping rules and FTP remapping rules.
- ◆ Click **Browser Auto-Config** to identify the port used to download browser auto-configuration files, and to set PAC and WPAD options.

Security

- ◆ Click **Connection Control** to specify which hosts are allowed to communicate with the Content Gateway server (ARM security) and which clients are allowed to access the proxy cache.
- ◆ Click **Access Control** to set filtering rules and set proxy authentication options (LDAP, RADIUS, and NTLM).
- ◆ Click **Data Security** to register with the Data Security Management Server and enable the local Data Security policy engine.
- ◆ Click **SOCKS** to configure Content Gateway to use a SOCKS firewall. The SOCKS button appears only if you have enabled SOCKS in the Features table on **Configure > My Proxy > Basic > General**.



Note

To use SOCKS there must be a separate SOCKS server.

Subsystems

- ◆ Click **Cache** to enable or disable cache pinning, configure the RAM cache size, specify the maximum size of objects allowed in the cache, and partition your cache according to protocol and origin servers.
- ◆ Click **Logging** to enable or disable event logging and set logging configuration options.

Networking

- ◆ Click **Connection Management** to specify the maximum number of connections the proxy can accept. For transparent proxy caching, you can specify the maximum number of client connections allowed before the proxy starts forwarding incoming requests directly to the origin server.

Ensure that the ARM is **On** on **Configure > My Proxy > Basic > General** before setting redirection rules.

- ◆ Click **ARM** to set redirection rules that specify how incoming packets are readdressed in transparent mode. You can also set dynamic and static bypass rules. The ARM button appears only if ARM is enabled in the Features table on **Configure > My Proxy > Basic > General**.
- ◆ Click **WCCP** to set WCCP configuration settings. The WCCP button appears only if WCCP is enabled in the Features table under the **Configure > My Proxy > Basic > General** tab.
- ◆ Click **DNS Resolver** to enable or disable local domain expansion and tune host database timeouts, and configure Split DNS options.
- ◆ Click **Virtual IP** to enable or disable virtual IP failover and specify the virtual IP addresses managed by the Content Gateway node. The Virtual IP button appears only if you have enabled Virtual IP in the Features table on **Configure > My Proxy > Basic > General**.

SSL

- ◆ Click **Certificates** to view the certificate authority tree. Click any entry to view the details of that certificate.
- ◆ Click **Decryption/Encryption** to configure how SSL Manager handles inbound and outbound traffic. Inbound traffic travels from the browser to SSL Manager, where the content is decrypted and inspected. Outbound traffic travels from SSL Manager to the destination Web server. SSL Manager checks the revocation status of the site certificate before forwarding re-encrypted data to the site.
- ◆ Click **Validation** to configure certificate validation, specify what to do in the case that a certificate is invalid, set up verification bypass, and configure the handling of certificate revocation lists.
- ◆ Click **Incidents** to view a report of occurrences in which clients received an access denial message, and to identify URLs that you want to allow, blacklist, or tunnel.
- ◆ Click **Client Certificates** to configure how SSL Manager handles client certificate requests.

- ◆ Click **Logging** to select the SSL logging level, logging detail, log file names, and log file handling.
- ◆ Click **Customization** to customize the certificate validation failure message.
- ◆ Click **Internal Root CA** to import, create, or backup the internal Root Certificate Authority.

Command-line interface

As an alternative to Content Gateway Manager, you can use the command-line interface to view and change your Websense Content Gateway configuration.

1. Log on to a Content Gateway node as the Content Gateway administrator, and then make the Content Gateway **bin** directory (default location is **/opt/WCG/bin**) your working directory.
2. To view a configuration setting, enter the following command:

```
content_line -r var
```

where *var* is the variable associated with the configuration option (for a list of the variables, refer to [Configuration variables](#), page 319).

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

```
content_line -s var -v value
```

where *var* is the variable associated with the configuration option and *value* is the value you want to use.

For example, to change the FTP inactivity timeout option to 200 seconds, enter the following command at the prompt and press Return:

```
content_line -s  
proxy.config.ftp.control_connection_timeout -v 200
```



Note

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

For example:

```
./content_line -r variable
```

Configuration files

You can change Websense Content Gateway configuration options by editing specific variables in the **records.config** file, located in the Content Gateway **config** directory

(default location is `/opt/WCG/config`). Open the file in a text editor (such as `vi` or `emacs`) and change the value of the variable.



Note

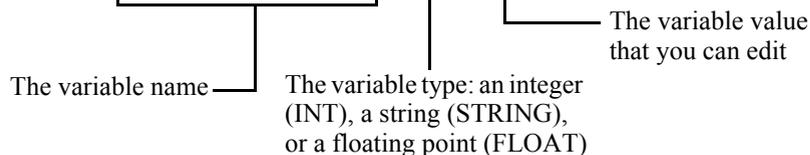
After you modify the `records.config` file, Content Gateway must reread the configuration files; from the Content Gateway `bin` directory (default location is `/opt/WCG/bin`), enter the command:

```
content_line -x
```

In some cases, you have to restart the proxy to apply the changes.

The figure below shows a sample portion of the `records.config` file:

```
##Id: records.config.v 1.617.2,27 2008/09/16 22:06:35 brilee Exp #
#
# Process Records Config File
#
# <RECORD-TYPE> <NAME> <TYPE> <VALUE (till end of line)>
#
# RECORD-TYPE: CONFIG, LOCAL
# NAME: name of variable
# TYPE: INT, STRING, FLOAT
# VALUE: Initial value for record
#
#####
#
# System Variables
#
#####
CONFIG proxy.config.proxy_name STRING ibid
CONFIG proxy.config.bin_path STRING bin
CONFIG proxy.config.proxy_binary STRING traffic_server
CONFIG proxy.config.proxy_binary_opts STRING -M
CONFIG proxy.config.manager_binary STRING traffic_manager
CONFIG proxy.config.cli_binary STRING traffic_line
CONFIG proxy.config.watch_script STRING traffic_cop
CONFIG proxy.config.env_prep STRING example_prep.sh
CONFIG proxy.config.config_dir STRING config
CONFIG proxy.config.temp_dir STRING /tmp
CONFIG proxy.config.alarm_email STRING inktomi
```



Content Gateway provides other configuration files that are used to configure specific features. All the configuration files are described in [Configuration Files, page 285](#).

Saving and restoring configurations

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 and it can save you hours of correcting configuration mistakes.

This section describes how to perform the following tasks:

- ◆ Take a snapshot of the current configuration. See [Taking configuration snapshots](#), page 91.
- ◆ Restore previously taken configuration snapshots. See [Restoring configuration snapshots](#), page 92.
- ◆ Delete configuration snapshots stored on the Content Gateway node. See [Deleting configuration snapshots](#), page 92.

Taking configuration snapshots

You can save all the current configuration settings on your Content Gateway system through Content Gateway Manager.

To take a configuration snapshot and save it on the local system

1. Navigate to **Configure > Snapshots > File System**.
2. The **Change Snapshot Directory** field displays the name of the directory where Content Gateway saves configuration snapshots. The default location is the Content Gateway **config/snapshots** directory. To change the directory, enter the full path in the **Change Snapshot Directory** field. If you enter a relative path, Content Gateway assumes that the directory is located in its **config** directory (for example, **/opt/WCG/config**).
3. In the **Save Snapshot** field, type the name you want to use for the current configuration.
4. Click **Apply**.

To take a configuration snapshot and save it on an FTP server

1. Navigate to **Configure > Snapshots > FTP Server**.
2. In the fields provided, enter the FTP server name, the login and password, and the remote directory where the FTP server stores configuration snapshots.
3. 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 the name of the configuration snapshot you want to take.
5. Click **Apply**.

Restoring configuration snapshots

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

To restore a configuration snapshot stored on the local node

1. Navigate to the **Configure > Snapshots > File System** tab.
2. From the **Restore > Delete Snapshot** drop-down list, select the configuration snapshot that you want to restore.
3. Click the **Restore Snapshot from “*directory_name*” Directory** box.
4. Click **Apply**.

The Content Gateway system or cluster uses the restored configuration.

To restore a configuration snapshot from an FTP server

1. Navigate to **Configure > Snapshots > FTP Server**.
2. In the fields provided, enter the FTP server name, the login and password, and the remote directory in which the FTP server stores configuration snapshots.
3. Click **Apply**.
After you have successfully logged on to the FTP server, the **FTP Server** tab displays additional fields.
4. In the **Restore Snapshot** drop-down list, select the configuration snapshot that you want to restore.
5. Click **Apply**.

The Content Gateway system or cluster uses the restored configuration.

Deleting configuration snapshots

1. Navigate to **Configure > Snapshots > File System**.
2. From the **Restore > Delete a Snapshot** drop-down list, select the configuration snapshot you want to delete.
3. Click the **Delete Snapshot from “*directory_name*” directory** box.
4. Click **Apply**.

The configuration snapshot is deleted.

10

Monitoring Traffic

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

- ◆ Content Gateway Manager provides statistics that show Content Gateway performance and network traffic information. See [Viewing statistics, page 93](#). The command-line interface provides an alternative method of viewing this information. See [Viewing statistics from the command line, page 96](#).
- ◆ Content Gateway Manager presents alarms that signal any detected failure conditions. See [Working with alarms, page 97](#).
- ◆ A performance graphing tool produces a variety of graphs that show historical Content Gateway performance and network traffic information. See [Using Performance graphs, page 98](#).
- ◆ Create reports through SSL Manager to see the status of certificate authorities and incidents. See [Creating reports with SSL Manager, page 99](#).

Viewing statistics

Use Content Gateway Manager to collect and interpret statistics about Content Gateway performance and Web traffic. View statistics using Monitor mode.

Starting Monitor mode

1. Open your Web browser.
Content Gateway Manager requires Java and JavaScript; be sure to enable Java and JavaScript in your browser.
2. Enter the following location in your browser:
`https://nodename:adminport`
where *nodename* is the name of the Content Gateway node and *adminport* is the number assigned to the Content Gateway Manager port.
3. If necessary, log on to Content Gateway Manager with the administrator ID and password, or use your user account. The administrator ID and password are set during installation. You can change the ID and password, as well as create and modify user accounts. For more information, see [Controlling access to Content Gateway Manager, page 141](#).

Using Monitor mode

In Monitor mode, Content Gateway Manager displays a series of buttons on the left of the display. Click a button to view its statistics.

All the statistics displayed in Monitor mode are described in detail in [Statistics](#), page 191.

My Proxy

Click **My Proxy** to see statistics about your proxy.

- ◆ Click **Summary** to see a concise view of your Content Gateway system. The top portion of the page displays information about the features in your Content Gateway or Websense Web Security Gateway subscription, including the expiration date. The middle portion of the page displays information about the scanning engines in use and their associated data files. The bottom portion of the page contains statistics on proxy nodes, displaying all cluster nodes by name and tracking essential statistics for each node. If you want to display detailed information about a particular node in a cluster, click the node's name in the Summary table, and then click one of the other buttons on the **Monitor** tab.
- ◆ Click **Node** to see information about the selected node. You can see if the node is active or inactive, the date and time that the **content_gateway** process was started, cache performance information (document hit rate, bandwidth savings, and what percentage of the cache is currently free), the number of client and server connections currently open, and the number of transfers currently in progress. You can also see name resolution information, such as the host database hit rate and the number of DNS lookups per second.



Note

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.

- ◆ Click **Graphs** to view the same statistics displayed on the **Node** page (cache performance, current connections and transfers, network, and name resolution) in graphical format. You can display multiple statistics in one graph. To display a particular statistic in graphical format, click the box next to the name of the graph, and then click **Graph**. To display multiple statistics in one graph, click the box next to the name of each graph you want to display, and then click **Graph**.
- ◆ Click **Alarms** to view the alarms that Content Gateway has signaled. See [Working with alarms](#), page 97.

Protocols

The Protocols button provides information about HTTP and FTP transactions.

- ◆ Click **HTTP** to see information about HTTP transactions and speeds (such as cache misses, cache hits, connection errors, aborted transactions) and client and server connection information. Also see information about FTP requests from HTTP clients, such as the number of open FTP server connections, the number of successful and unsuccessful PASV and PORT connections, and the number of cache lookups, hits, and misses.
- ◆ Click **FTP** to see information about FTP requests from FTP clients. For example, see the number of open FTP client and FTP server connections, the number of file hits and misses, change directory hits and misses, and list directory hits and misses.



Note

The **FTP** button appears only if you have enabled FTP processing in the **Features** table under the **Configure > My Proxy > Basic** tab.

Content Routing

Click **Content Routing** to see ICP statistics that include information about queries originating from the Content Gateway node and from ICP peers (parents and siblings).

Security

The Security button provides information about ARM, proxy authentication, and SOCKS server connections:

- ◆ Click **ARM Security** to see the number of dropped TCP and UDP connections.
- ◆ Click **LDAP** to see the number of LDAP cache hits and misses, and the number of LDAP authentication server errors and unsuccessful authentication attempts. The LDAP button appears only if you have enabled the LDAP option in the Features table on the **Configure > My Proxy > Basic > General** tab.
- ◆ Click **NTLM** to see the number of NTLM cache hits and misses, and the number of NTLM authentication server errors and unsuccessful authentication attempts. The NTLM button appears only if you have enabled the NTLM option in the Features table on the **Configure > My Proxy > Basic > General** tab.
- ◆ Click **SOCKS** to see the number of successful and unsuccessful connections to the SOCKS server and the number of connections currently in progress. The SOCKS button appears only if you have enabled the SOCKS option in the Features table on the **Configure > My Proxy > Basic > General** tab.

Subsystems

The Subsystems button provides information about the proxy cache, clusters, and event logging:

- ◆ Click **Cache** to see information about the proxy cache. See how much space in the cache is currently being used, the total size of the cache in gigabytes, the total size of the RAM cache in bytes, the number of RAM cache hits and misses, and the number of cache lookups, object reads, writes, updates, and removes.

- ◆ Click **Clustering** to see the number of nodes in the cluster, the total number of cluster operations, the number of bytes read and written to all the nodes in the cluster, and the current number of open connections in the cluster.
- ◆ Click **Logging** to see the number of log files currently open, the amount of space currently being used for log files, the number of access events and error events logged, and the number of access events skipped.

Networking

The Networking button provides information about system network configuration, the ARM, WCCP routers, domain name resolution, and virtual IP addressing.

- ◆ Click **System** to see system network configuration, including the host name assigned to the proxy machine and the default gateway, search domain, and DNS servers that the proxy machine uses.
- ◆ Click **ARM** to see information about Network Address Translation and dynamic bypass. The ARM button appears only if you have enabled ARM in the Features table under the **Configure > My Proxy > Basic** tab.
- ◆ Click **WCCP** to see WCCP version v1 or v2 statistics that include information about the routers being used, the number of active nodes, the leader's IP address, and whether WCCP is currently enabled on the Content Gateway node. The WCCP button appears only if you have enabled WCCP in the Features table on the **Configure > My Proxy > Basic > General** tab.
- ◆ Click **DNS Resolver** to see the total number of lookups and hits in the host database, and the average lookup time, the total number of lookups, and the number of successful lookups in the DNS server.
- ◆ Click **Virtual IP Address** to see the current virtual IP address mappings. The Virtual IP Address button appears only if you have enabled the Virtual IP option in the Features table on the **Configure > My Proxy > Basic > General** tab.

Performance

The Performance button displays historical performance graphs. See [Using Performance graphs, page 98](#).

Viewing statistics from the command line

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

You can also configure, stop, and restart the Content Gateway system from the command line. See [Command-line interface, page 89](#), and [Websense Content Gateway commands, page 211](#).

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

1. Become root:

```
su
```

2. Log on to a node as the Content Gateway administrator.
3. From the Content Gateway **bin** directory (default location is `/opt/WCG/bin`), enter the following command:

```
content_line -r variable
```

where *variable* is the variable that represents the information you want. For a list of the variables you can specify, see [Websense Content Gateway variables](#), page 212.

For example, the following command displays the document hit rate for the node:

```
content_line -r proxy.node.http.cache_hit_ratio
```



Note

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

For example:

```
./content_line -r variable
```

Working with alarms

Websense 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).

Not all alarms are critical. Some alarms report transient conditions. For example, a **license download failed:4** alarm can be generated by a temporary disruption in internet connectivity.

Navigate to **Monitor > My Proxy > Alarms** to see a listing of current alarms, as shown below.

The Alarm! (pending) bar appears at the top of the display when alarms exist.

The screenshot shows the Websense Content Gateway interface. The top navigation bar includes 'Monitor' and 'Configure' tabs, and the 'Monitor' tab is selected. The user is logged in as 'admin'. The left sidebar shows a tree view with 'My Proxy' expanded, containing 'Summary', 'Node', 'Graphs', and 'Alarms'. The 'Alarms' section is selected, displaying a red bar at the top with the text 'Websense Content Gateway Alarms' and a 'Clear' button. Below the bar, the current time is 'Tue Apr 7 10:17:26 2009'. A table with columns 'Node' and 'Alarm' is shown, with the text 'No currently active alarms.' below it. A 'Clear' button is also present at the bottom right of the table area.

Clearing alarms

After you have read an alarm message, you can click **Clear** in the alarm message window to dismiss the alarm. [Alarm messages, page 386](#), provides a description of some of the alarm messages that Content Gateway provides.



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 alarms

1. 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@company.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 the **bin** directory (default location is `/opt/WCG/bin`). You can modify this file.

Using Performance graphs

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

Access the Performance graphs from the Monitor tab in Content Gateway Manager.



Important

To run Multi Router Traffic Grapher (the Performance graphing tool), you must have Perl version 5.005 or later installed on your Content Gateway system.

1. If your Content Gateway node is in a cluster, select the node whose statistics you want to view from the **Monitor > My Proxy > Summary** display.
2. On the **Monitor** tab, click **Performance**.
3. 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 5-minute sample intervals for the tool to initialize statistics.

If Multi Router Traffic Grapher (MRTG) has not been configured, the system displays a message indicating that it is not available. To configure the tool:

1. Make sure Perl 5.005 is installed on your system.
2. At the command prompt, type

```
perl ./pathfix.pl 'which perl'
```

to ensure that the perl binary is in your PATH.
3. Change to the Content Gateway **bin** directory.
4. Modify the MRTG update interval by typing the following at the command prompt:

```
./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.
5. Start the MRTG cron updates:

```
./mrtgcron start
```
6. Wait about 15 minutes before accessing the performance graphs from the Content Gateway Manager.

**Note**

To stop MRTG cron updates, type the command

```
./mrtgcron stop.
```

Creating reports with SSL Manager

You can request a report detailing the status of certificate authorities (see [Certificate Authorities, page 100](#)) or listing incidents (see [Incidents, page 101](#)). Reports can be either in HTML or comma-separated format. The comma-separated reports appear as Excel spreadsheets in SSL Manager.

Certificate Authorities

1. Go to the **Monitor > SSL > Reports > Certificate Authorities** tab.
2. Select the format of the report.
 - a. HTML
 - b. Comma-separated values (CSV)

If you select CVS, the report is created as an Excel spreadsheet.
3. Specify the time period the report will cover.
 - a. A number of days
 - b. A starting date spanning to the present
 - c. All records in the log
4. Indicate the sort order for the report.
 - a. List authorities by date
 - b. List OCSP good responses first
 - c. List OCSP bad responses first

See [Keeping revocation information up to date](#), page 127.
5. Click **Generate Report** to generate the report.

HTML output looks like this:

Certificate Authorities		Incidents			
Validation Reports					
HTML Report of EVA - Certificate Authorities					
Profile: default_default					
Certificate Authority	Count good	Percentage	Count bad	Percentage	Last Access Date
Class 3 Public Primary Certification Authority	167	13.47 %	0	0.00 %	2008-02-12 12:07:17
www.verisign.com/CPS Incorpor. by Ref. LIABILITY LTD.(c)97 VeriSign	88	7.10 %	0	0.00 %	2008-02-12 12:07:17
VeriSign Class 3 Secure Server CA	75	6.05 %	0	0.00 %	2008-02-12 12:07:17
Equifax Secure Certificate Authority	535	43.15 %	0	0.00 %	2008-02-12 10:30:06
Microsoft Internet Authority	112	9.03 %	0	0.00 %	2008-02-11 19:41:58

The same report in comma-separated format appears as follows:

Certificate Authorities		Incidents				
Validation Reports						
A1		CSV Report of EVA - Certificate Authorities				
A	B	C	D	E	F	G
1	CSV Report of EVA - Certificate Authorities					
2						
3	Profile: default_default					
4						
5	Certificate	Count good	Percentage	Count bad	Percentage	Last Access Date
6	Class 3 Pu	167	13.47%	0	0.00%	#####
7	www.verisi	88	7.10%	0	0.00%	#####
8	VeriSign C	75	6.05%	0	0.00%	#####
9	Equifax Se	535	43.15%	0	0.00%	#####
10	Microsoft I	112	9.03%	0	0.00%	#####



Note

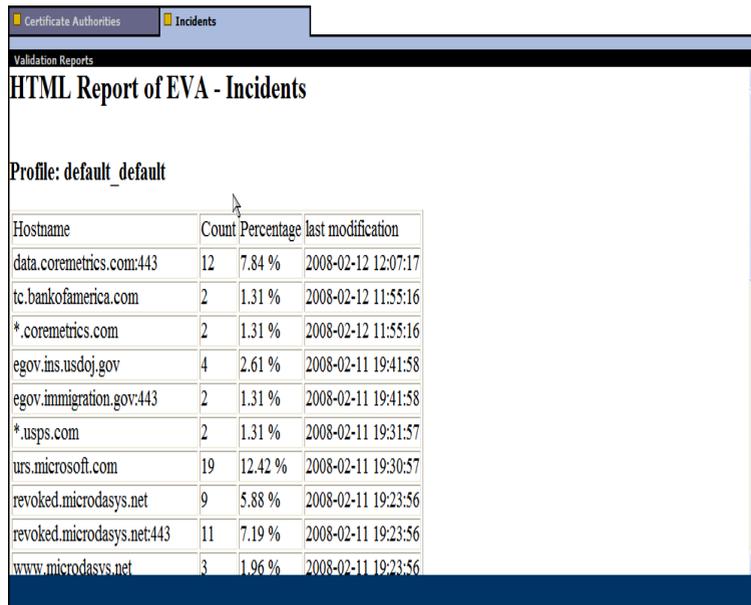
To delete the collected SSL log data, click **Reset all collected data**.

Incidents

1. Navigate to the **Monitor > SSL > Reports > Incidents** tab.
2. Select HTML or comma-separated (CSV) format. If you select comma-separated, the report is created in an Excel spreadsheet.
3. Specify the time period the report should cover. You can specify
 - a. a number of days
 - b. a date range
 - c. the period since SSL Manager was deployed
4. Indicate the sort order for the report.
 - a. Listing incidents by date
 - b. Listing incidents by URL
 - c. Listing the number of times each incident occurred

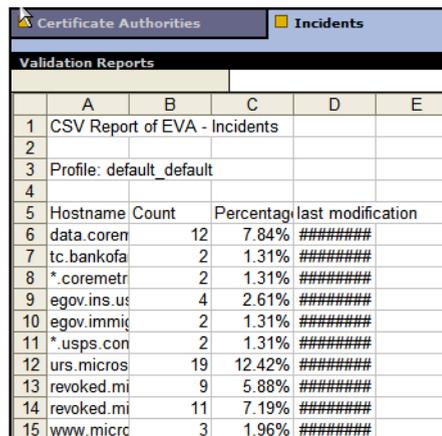
See [Managing Web HTTPS site access](#), page 128.
5. Click **Generate Report** to generate the report.

TML output looks like this:



Hostname	Count	Percentage	last modification
data.coremetrics.com:443	12	7.84 %	2008-02-12 12:07:17
tc.bankofamerica.com	2	1.31 %	2008-02-12 11:55:16
*.coremetrics.com	2	1.31 %	2008-02-12 11:55:16
egov.ins.usdoj.gov	4	2.61 %	2008-02-11 19:41:58
egov.immigration.gov:443	2	1.31 %	2008-02-11 19:41:58
*.usps.com	2	1.31 %	2008-02-11 19:31:57
urs.microsoft.com	19	12.42 %	2008-02-11 19:30:57
revoked.microdasys.net	9	5.88 %	2008-02-11 19:23:56
revoked.microdasys.net:443	11	7.19 %	2008-02-11 19:23:56
www.microdasys.net	3	1.96 %	2008-02-11 19:23:56

The same report in comma-separated format appears as follows:



	A	B	C	D	E
1	CSV Report of EVA - Incidents				
2					
3	Profile: default_default				
4					
5	Hostname	Count	Percentage	last modification	
6	data.coreme	12	7.84%	#####	
7	tc.bankofa	2	1.31%	#####	
8	*.coremetr	2	1.31%	#####	
9	egov.ins.us	4	2.61%	#####	
10	egov.immig	2	1.31%	#####	
11	*.usps.con	2	1.31%	#####	
12	urs.micros	19	12.42%	#####	
13	revoked.mi	9	5.88%	#####	
14	revoked.mi	11	7.19%	#####	
15	www.micr	3	1.96%	#####	



Note

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

11

Working With Websense Data Security

Related topics:

[Registering and configuring on-box Data Security](#), page 106

[Unregistering on-box Data Security](#), page 108

[Stopping and starting Data Security processes](#), page 108

[Configuring the ICAP client](#), page 109

Websense Content Gateway can be deployed together with Websense Data Security to provide data loss prevention (DLP) over Web channels such as HTTP, HTTPS, FTP, and FTP over HTTP. A full data security deployment can extend DLP to include channels such as mobile devices, removable media, and printers. For a complete description of Websense Data Security, please visit the Data Security product page at www.websense.com.

Web data loss prevention, as well as other data security configurations, require separate installation of the Data Security Manager and other Data Security components. Before configuring Content Gateway to work with Websense data security, please read the *Websense Web Security Gateway Anywhere Getting Started Guide*.

Content Gateway supports 2 methods of working with Data Security:

- ◆ Using the Data Security policy engine located on-box with Content Gateway (requires version 7.5 Data Security components)
- ◆ Over ICAP using a Data Security policy engine located on a separate computer (intended for use with Data Security Suite versions 7.1 and earlier)

Only 1 method can be used at a time. When 1 is configured the other is unavailable.

How it works:

1. The proxy intercepts outbound content and provides that content to Data Security.
2. Data Security analyzes the content to determine if the Web posting or FTP upload is allowed or blocked.
 - The determination is based on the Data Security policy.
 - The disposition is communicated to the proxy.
 - Data Security logs the transaction.

3. The proxy acts on the Data Security determination.
 - a. If the content is blocked, it is not transmitted to the remote host, and Data Security returns a block page to the sender.
 - b. If the content is allowed, it is forwarded to its destination.

Transactions over HTTP, HTTPS, FTP, and FTP over HTTP can be examined.

Transaction details are logged by Data Security per its configuration.

Data Security policy engine on-box with Content Gateway

When Content Gateway is installed, a Data Security policy engine is also installed, although it is disabled until registered with the Data Security Management Server (see [Registering and configuring on-box Data Security, page 106](#)).

After Data Security policies have been created and Content Gateway has registered with Data Security, Content Gateway sends content, such as postings and uploads, to the Data Security for analysis and policy enforcement.

When the on-box policy engine is used, Content Gateway collects and displays Data Security transaction statistics, such as the total number of posts, the total number of posts analyzed, the number of FTP uploads analyzed, the number of blocked requests, and more. These statistics can be viewed in the Content Gateway Manager by navigating to **Monitor > Security > Data Security**. For a complete list of statistics, see [Data Security, page 200](#).

Data Security over ICAP

When the Data Security policy engine is located on a different computer, Content Gateway can communicate with Data Security over ICAP v1.0. ICAP is the conduit for moving content from the proxy to Data Security servers for analysis. For configuration details, see [Configuring the ICAP client, page 109](#).

Registering and configuring on-box Data Security

Related topics:

[Unregistering on-box Data Security, page 108](#)

[Stopping and starting Data Security processes, page 108](#)

For an introduction to Data Security, see [Working With Websense Data Security, page 105](#).

To complete Data Security registration, you need to know the IP address of the Data Security Management Server. The Data Security Management Server must be

version 7.5. Registering Content Gateway with Data Security adds the Content Gateway module to the Data Security Manager.



Important

If Content Gateway is **not** located on a V-series appliance, registration **requires** that the Content Gateway host system have an IPv4 address assigned to the eth0 network interface. After registration, the IP address may move to another network interface on the system; however, that IP address is used for data security configuration deployment and must be available as long as the two modules are registered.

To enable the on-box Data Security policy engine:

1. Ensure that the Content Gateway and Data Security Management Server systems are running and accessible, and that their system clocks are synchronized within a few minutes.
2. Log on to Content Gateway Manager and select **Configure > Basic > General**.
3. In the list of **Features**, under **Networking** locate **Data Security**, select **On**, and then select **Integrated on-box**. A registration status link displays.
4. Click the **Not registered** link. This opens the **Configure > Security > Data Security** registration screen.
5. Enter the IP address of the **Data Security Management Server**.
6. Enter a user name and password for logging onto the Data Security Manager. This is the management interface in which Data Security policy is configured. The user must be a Data Security administrator with Deploy Settings privileges.
7. 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.

8. Before restarting Content Gateway, set the following configuration options:
 - a. **Analyze FTP Uploads**: Select this option to send FTP uploads to Data Security for analysis and policy enforcement.
 - b. **Analyze HTTPS Content**: Select this option to send decrypted HTTPS posts to Data Security for analysis and policy enforcement. SSL Manager must be enabled on Content Gateway. See [Working With Encrypted Data](#), page 111.



Note

In order for these options to have any effect, Content Gateway must be configured to proxy FTP and HTTPS traffic. Go to **Configuration > MyProxy > Basic**.

These options can be accessed whenever Data Security is registered by going to the **Configure > Security > Data Security > General** page.

9. Click **Apply** to save your settings and then go to **Configure > Basic** and click **Restart** to restart Content Gateway.
10. Log on to the Data Security Manager to configure Content Gateway module. See the section titled “Deploying the Content Gateway module” in the *Websense Web Security Gateway Anywhere Getting Started* guide.

Data Security and Content Gateway communicate over ports 5820, 8888, 8889, 8892, 9080, and 9443. If IPTables are configured on the Content Gateway host system, these ports must be open in IPTables. See the *Content Gateway Installation Guide*, or Knowledge Base article 4106, “Configuring IPTables for Websense Content Gateway”.



Note

A Content Gateway Manager alarm is generated if:

- ◆ On-box Data Security is enabled but not registered
 - ◆ On-box Data Security is enabled and registered, but not configured in the Data Security Manager
-

Unregistering on-box Data Security

To disable the integration with the on-box Data Security policy engine:

1. Log on to Content Gateway Manager and navigate to **Configure > Security > Data Security > General**. This page should indicate that the **Registration status** is **Registered**.
2. Click the **Unregister** button and confirm the action by clicking **OK** in the confirmation dialog box. When the action is complete, the **Registration status** will read **Unregistered**.
3. Go to **Configure > Basic** and click **Restart** to restart Content Gateway.

Stopping and starting Data Security processes

When Content Gateway is registered with Data Security Management Server and the on-box policy engine is running, 3 daemon processes are active on the Content Gateway machine:

- ◆ **PolicyEngine** handles transaction and data analysis.
- ◆ **PAFPREP** manages the Data Security 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 must be used for inter-operation with Data Security Suite versions 7.1 and earlier.

To configure the integration with Data Security over ICAP, log on to Content Gateway Manager and go to **Configure > My Proxy > Basic > General** page.

1. In the **Networking** section of the Features table, select Data Security **On**.
2. Click **Apply**, and then click **Restart**.
3. Navigate to **Configure > Networking > ICAP > General**.
4. In the **ICAP Service URI** field, enter the Uniform Resource Identifier (URI) for the ICAP service. A URI is similar to a URL, but the URI ends with a directory, rather than a page. Obtain the identifier from your Websense Data Security Suite administrator. Enter the URI in the following format:

```
icap://hostname:port/path
```

For *hostname*, enter the IP address or hostname of the Websense Data Security Suite Protector appliance.

The default ICAP port is 1344.

Path is the path of the ICAP service on the host machine.

For example:

```
icap://ICAP_machine:1344/REQMOD
```

You do not need to specify the port if you are using the default ICAP port 1344. For example, the above URI can also be entered without the default port:

```
icap://ICAP_machine/REQMOD
```

5. Under **Analyze HTTPS Content**, indicate if decrypted traffic should be sent to Websense Data Security Suite for analysis or sent directly to the destination. You must be running SSL Manager to send traffic to Websense Data Security Suite. See [Working With Encrypted Data, page 111](#).
6. Under **Analyze FTP Uploads**, select whether to send FTP upload requests to Websense Data Security Suite for analysis. The FTP proxy feature must be enabled to send FTP traffic to Websense Data Security Suite. See [FTP, page 240](#).
7. Under **Action for Communication Errors**, select whether to permit traffic or send a block page if Content Gateway encounters an error while communicating with Websense Data Security Suite.

8. 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 Websense Data Security Suite is sent. The default size limit for Data Security Suite version 7.0 and later is 12 MB.
9. Click **Apply**.



Note

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

12

Working With Encrypted Data

Related topics:

[Running in explicit proxy mode](#), page 112
[Tasks](#), page 114
[Enabling SSL Manager](#), page 114
[Certificates](#), page 115
[Internal Root CAs](#), page 115
[Managing certificates](#), page 119
[Configuring SSL Manager for inbound traffic](#), page 122
[Configuring SSL Manager for outbound traffic](#), page 122
[Validating certificates](#), page 124
[Managing Web HTTPS site access](#), page 128
[Client certificates](#), page 132
[Configuring logging](#), page 133
[Customizing messages](#), page 136

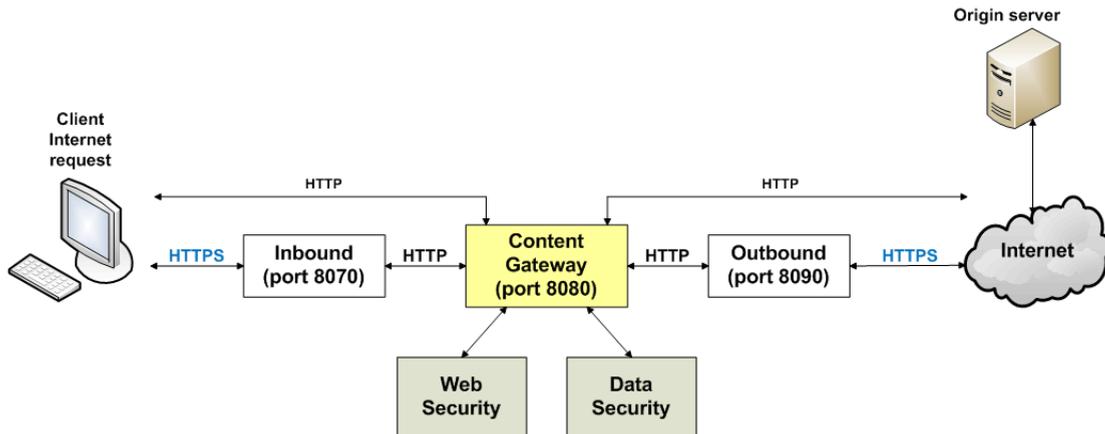
SSL (Secure Sockets Layer) is the industry standard for transmitting secure data over the Internet. It is based on a system of trusted certificates issued by certificate authorities and recognized by servers.

When SSL Manager is enabled, SSL-encrypted traffic is decrypted, inspected, and then re-encrypted before it is sent to its destination. To preserve performance, only HTTPS traffic is diverted in this manner; HTTP traffic goes through the proxy to the requested site.

Each SSL-based request consists of two separate sessions:

- ◆ From the client browser to SSL Manager. This is considered *inbound* SSL traffic.
- ◆ From SSL Manager to the Web server that will receive the secure data. This is considered *outbound* SSL traffic.

Different certificates are required for these sessions.



Content Gateway does not cache HTTPS data.

For additional information on certificates and SSL, consult any of the commercially available books on SSL for a complete description of certificates. For information on preparing your system, see the *Content Gateway Installation Guide*.

Running in explicit proxy mode

If you have an existing PAC file, replace the **proxy.pac** file located in the Content Gateway **config** directory (default location is `/opt/WCG/config`) with the existing file. If you do not have a PAC file already, see [Step 4](#) below for a script you can copy.

1. On the **Configure > My Proxy > Basic > General** tab, ensure that both ARM and HTTPS are enabled.
For ARM, check the Networking section; for HTTPS check the Protocols section. If they are disabled, set them to **On**. Then click **Apply**, followed by **Restart**.
2. Navigate to **Configure > Content Routing > Browser Auto-Config > PAC**.
3. In the **Auto-Configuration Port** field, specify the port that the proxy 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 **proxy.pac** file is empty. Copy and paste the following script for your PAC settings. You must provide the proxy domain name or IP address.

```
function FindProxyForURL(url, host)
{
    url = url.toLowerCase();
    host = host.toLowerCase();
    if(url.substring(0, 5) == "http:"){
```

```

        return "PROXY WCG_DOMAIN_NAME_or_IP_Address:8080";
    }
    else if(url.substring(0, 4) == "ftp:"){
        return "PROXY WCG_DOMAIN_NAME_or_IP_Address:2121";
    }
    else if(url.substring(0, 6) == "https:"){
        return "PROXY WCG_DOMAIN_NAME_or_IP_Address:8080";
    }
    else{
        return "DIRECT";
    }
}
}

```

5. Click **Apply**.
6. Click **Restart** on **Configure > My Proxy > Basic > General**.

Once the new PAC information is in place, you must inform your users to set their browsers to point to the PAC file. For example, if the PAC file is located on the proxy server with the host name **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
```

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

For Microsoft Internet Explorer version 7.0 and later:

1. Navigate to **Tools > Internet Options > Connections > LAN Settings**.
2. Select **Use automatic configuration script** field, and enter

```
http://WCG_Domain_Name_or_IP_Address:8083/proxy.pac
```

in the **Address** field.
3. Click **OK**.

For Mozilla Firefox 2.0 and later:

1. Navigate to **Tools > Options > Advanced > Network > Connection > Settings**.
2. Select **Automatic proxy configuration URL** field, and enter

```
http://WCG_Domain_Name_or_IP_Address:8083/proxy.pac
```
3. Click **Reload**, and then click **OK**.

See the documentation for your browser for details.

Enabling SSL Manager

1. On **Configure > My Proxy > Basic > General**, click **HTTPS On**.

**Note**

If you are running with other Websense products, such as Websense Data Security Suite, that inspect HTTPS traffic, you must select enable HTTPS here.

2. Click **Apply** and then click **Restart**.
3. On **Configure > My Proxy > UI Setup > General**, specify the port for the SSL Manager interface. The default is 8071. This must be a different port than the Content Gateway Manager interface (default 8081).
4. Enter the name of the SSL certificate file. See [Creating a sub-certificate authority, page 117](#).

Use the **Configure > Protocols > HTTPS** page to provide port information.

1. In the **HTTPS Proxy Server Port** field, enter the port for inbound (client to SSL Manager) HTTPS traffic. The default is 8070, however for explicit traffic, the value should be set to 8080.
2. In the **SSL Outbound Port** field, enter the port SSL Manager will use for outbound HTTPS traffic from SSL Manager to the destination server. The default is 8090.

Tasks

For inbound (client to SSL Manager) traffic, perform these steps in preparation to deploying SSL Manager:

1. Create an internal root CA (certificate authority). In order to sign SSL traffic, SSL Manager requires an internal SSL Certificate Authority that has the capability to sign SSL certificates. This is for traffic between the browser and SSL Manager. See [Internal Root CAs, page 115](#).
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 Web sites. Any Web site signed by a certificate authority in the certificate tree with the “allow” status is allowed through SSL Manager. See [Managing certificates, page 119](#)
3. Customize pages that browser users will see. See [Customizing messages, page 136](#). Among the pages that can be customized are a connect failure and certificate verification failure page.

Certificates

Security revolves around certificates. One role SSL Manager plays is to ensure that certificates are valid. A certificate must meet three criteria:

- ◆ It must be current (has not expired or been revoked). See [Validating certificates](#), page 124.
- ◆ It must be issued by a trusted CA (certificate authority). See [Managing certificates](#), page 119
- ◆ The URL and the certificate owner match. See [Configuring validation](#), page 124.

Traffic from the client browser to SSL Manager requires a certificate issued by an internal root certificate authority. See [Internal Root CAs](#), page 115.

Traffic from SSL Manager to the destination server requires a certificate issued by one of the authorities listed on the Certificate Authority Tree on the **Configure > SSL > Certificates > Certificate Authorities** tab. See [Managing certificates](#), page 119.

Internal Root CAs

The internal Root CA dynamically generates all certificates used between the client browser and SSL Manager. You must have an internal Root CA to pass inbound traffic to SSL Manager. You can either import or create this CA. Internal Root CAs are stored in `/opt/wcg/sxsuite/db/server/default.pem`.



Important

Back up any existing Internal Root CAs before importing or creating new ones. This enables you to return to an earlier version of the certificate, if necessary. See [Backing up your internal Root CA](#), page 119 for details.

Only one internal Root CA can be active at any time.



Important

The default Internal Root CA included with SSL Manager is not unique and should not be used in a production environment.

Replace the default Root CA with your organization's existing Root CA or create a new one. See the sections that follow.

There are three options for creating an internal Root CA:

- ◆ Leverage an existing corporate CA and import it into SSL Manager. See [Importing your Root CA](#), page 116.
- ◆ Create a new CA for proxies and make that CA available to browsers. See [Creating your new Root CA](#), page 116.
- ◆ Create a sub-certificate CA. This leverages a corporate CA, but can also be revoked by the corporate CA. See [Creating a sub-certificate authority](#), page 117.

Importing your Root CA

If your organization already has a root certificate authority, you can import it. This certificate must be trusted by all browsers in your organization. Be sure to back up any new internal Root CAs that you import. See [Backing up your internal Root CA](#), page 119 for details.

1. Navigate to **Configure > SSL > Internal Root CA > Import Root CA**.
2. Browse to select the certificate. The certificate must be in X.509 format and base64-encoded.
3. Browse to select the private key. It must correspond to the certificate you selected in Step 2.
4. Enter, and then confirm, the passphrase.
5. Click **Import Root CA**. The imported CA is stored in `/opt/WCG/sxsuite/db/`.

Creating your new Root CA

Related topic:

[Creating a sub-certificate authority](#), page 117

If you do not already have a Root CA, fill in the fields on this tab to create one. Be sure to back up any new internal Root CAs that you create. See [Backing up your internal Root CA](#), page 119 for details.

An asterisk (*) on this page indicates a required field.

1. Select **Configure > SSL > Internal Root CA**, and then select **Create Root CA**.
2. Provide requested information in the fields, particularly noting the following:
 - The fields **Organization**, **Organizational Unit**, (this field is optional) and **Common Name** comprise a *distinguished name*.
 - For **Organization**, enter the name of your company.
 - 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.

Creating a sub-certificate authority

Creating a sub-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 the certificate services in Microsoft Windows 2003.

Preparation

- Install the OpenSSL toolkit from www.openssl.org.

Creating a request

1. Create a request for the sub-CA with OpenSSL.
 - a. Create a folder to store the certificate files.
 - b. Create a **req_ca.cnf** file and put it in the folder you created in step a. You can copy the file in *The req_ca.cnf File*, page 393.
 - c. Answer all the questions, or put the defaults of your CA into the **req_ca.cnf** file. The password you enter here is needed later when you import the sub-CA into SSL Manager. This password protects the private key of your new sub-CA.
2. Create the request with the private key:

```
openssl req -newkey rsa:1024 -keyout subca.key -new -days
3650 -config ./req_ca.cnf -out subca.req
```

You now have a subca.key and a **subca.req** file. The subca.key is the private key of the new sub-CA.

Signing the request

You must sign the request with Microsoft Certificate Services.

1. Select **Start > Administrative Tools > Certification Authority**.
2. Right-click on the name of the server and select **All Tasks > Submit new request**.
3. From the Browse dialog box, select the **subca.req** file.
4. Select **Pending Requests** to see your request.
5. Right-click your request and select **All Tasks > Issue**.
6. Select **Issued Certificates** to see that your new certificate is listed there.

Exporting the public key

1. Add a column to the Issued Certificates window.
 - a. Right-click **Issued Certificates** and select **View > Add/Remove Columns**.
 - b. Click **Binary Public Key** and then click **Add** to add a heading to the new column.
2. Export the binary data.
 - a. In the Issued Certificates window, select your new certificate and then select **All Tasks > Export Binary Data**.
 - b. Select **Binary Certificate** from the drop-down list.
 - c. Mark **Save binary data to a file**.
 - d. Click **OK**.
3. Name your certificate file **subca.cer**. This file is the public key of the sub-CA in binary format.
4. Move the **subca.cer** file to the directory where the subca.key resides. See [Creating a request, page 117](#).
5. Convert the binary file into pem format using this OpenSSL command:

```
openssl x509 -in subca.cer -inform der -out subca.pem -outform pem
```
6. Through Internet Explorer, export the **rootca.pem** file. This is the public key of the root CA in base64 format.

Import the sub-CA into SSL Manager

1. Create a file called **PCAcertchain.pem**. This file contains the **rootca.pem** and the **subca.pem** in one file.

```
cat rootca.pem > PCAcertchain.pem
cat subca.pem >> PCAcertchain.pem
```
2. Import the sub-CA.
 - a. Navigate to **Configure > SSL > Internal Root CA > Import Root CA**.
 - b. Browse to select the certificate. The certificate must be in X.509 format and base64-encoded.
 - c. Browse to select the private key. It must correspond to the certificate you selected in step 2.
 - d. Enter, and then confirm, the passphrase.
 - e. Click **Import Root CA**.
3. Place the previously created **PCAcertchain.pem** into the **/opt/WCG/sxsuite/conf/CA_default/PCA** folder. See step 1 in this section.
4. Restart Content Gateway.

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. Navigate to **Configure > SSL > Internal Root CA > Backup Root CA**.
2. Click **Save Public CA Key** to view or save the public CA key. This public key must be trusted by the users' Web browsers. Consult your network administrator if you do not have the key.
3. Click **Save Private CA Key** to view or save the private CA key. Consult your network administrator if you do not have the key.

Managing certificates

Related topics:

[Adding new certificate authorities, page 120](#)

[Backing up certificates, page 121](#)

[Restoring certificates, page 121](#)

All certificate authorities trusted by Internet Explorer 7 are listed on the tab **Configure > SSL > Certificates > Certificate Authorities**; destination servers (traffic outbound from SSL Manager) can trust the Web servers with these certificates. Note that a small “i” appears before the name of some certificates validated via CRL (certificate revocation lists) or OCSP (online certification status protocol). These certificates provide URLs where you can verify their revocation status. See [Keeping revocation information up to date, page 127](#) for information on checking the revocation status of a certificate. SSL Manager checks the revocation status of a certificate for both inbound and outbound traffic.

Click on the name of a certificate authority to:

- ◆ [View a certificate, page 119](#)
- ◆ [Delete a certificate, page 120](#)
- ◆ [Change the allow/deny status of a certificate, page 120](#)

View a certificate

1. Navigate to **Configure > SSL > Certificates > Certificate Authorities**.
2. Select the name of the authority whose status you want view.
3. In the pop-up window, select **Click to view certificate**.

4. Follow the directions in the Opening window to open or save the file.

Delete a certificate

1. Navigate to **Configure > SSL > Certificates > Certificate Authorities**.
2. Select the name of the certificate authority you want to delete.
3. In the pop-up window, select **Click to delete certificate**.
4. Confirm or deny that you want to delete the certificate.
5. If you confirm that you want to delete the certificate, check that the certificate is no longer listed on **Configure > SSL > Certificates > Certificate Authorities**.

Change the allow/deny status of a certificate

1. Navigate to **Configure > SSL > Certificates > Certificate Authorities**.
2. Select the name of the authority whose status you want to change.
3. In the pop-up window, select **Click to change status to**. 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.

Adding new certificate authorities

Related topics:

[Backing up certificates, page 121](#)

[Restoring certificates, page 121](#)

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



Important

It is recommended that you back up your current certificates before making any changes, such as adding or deleting certificates. See [Backing up certificates, page 121](#). If you wish to back up your entire Content Gateway configuration, see [Saving and restoring configurations, page 90](#).

1. Click **Browse** to navigate through the directory structure to find certificates. Look for files that have a “.cer” extension. The certificate must be in X.509 format and base64-encoded.
2. Click **Add Certificate Authority**.

3. 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. See [Change the allow/deny status of a certificate](#), page 120 for additional information.

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.

Backing up certificates also backs up your SSL Manager settings.

Use the page **Configure > SSL > Certificates > Backup Certificates** to back up certificates and your SSL Manager settings.

- Click **Back Up Configuration to Database**.

To back up not only certificates, but your entire Content Gateway configuration, see [Saving and restoring configurations](#), page 90.

Restoring certificates

Restoring certificates also restores the configuration database. However, because revocation lists are updated on a regular basis, they are not restored as part of this process. See [Keeping revocation information up to date](#), page 127 for information on updating certificate revocation lists.

Use the page **Configure > SSL > Certificates > Restore Certificates** to restore the configuration database, which includes certificates and your SSL Manager settings.

1. Click **Browse** to navigate to the location of the backup certificate database.
2. Click **Restore**. You receive a message telling you that the restore was successful and indicating where the previous certificate database was backed up.

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

Decryption and Encryption

[Configuring SSL Manager for inbound traffic](#), page 122

[Configuring SSL Manager for outbound traffic](#), page 122

Configuring SSL Manager for inbound traffic

Related topics:

[Configuring SSL Manager for outbound traffic, page 122](#)

Use the page **Configure > SSL > Decryption / Encryption > Inbound** to configure how SSL Manager handles inbound traffic. Inbound traffic travels from the browser to SSL Manager, where the content is decrypted and inspected.

1. Select **IP Address** to forward authentication credentials to the next proxy.
2. Select **Send VIA-Header** to add a special header to the HTTP header to describe the proxy chain traffic passed through. This can be helpful in troubleshooting. If you do not want to include a VIA-Header, do not select this box.
3. Under **Protocol Settings**, indicate which protocols you want SSL Manager to support. Supported protocols are SSLv2 and v3, and TLS v1. Select the protocol that your enterprise browser supports; you must select at least one protocol. The default is SSLv2. These settings override the settings for these protocols in the users' browsers.

You can select different protocols for outbound traffic.

4. The cipher list describes available algorithms and level of encryption between the client and SSL Manager. The default settings indicate to use all available ciphers except the eNULL and the ADH Suite. 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. Setting encryption to a high level for inbound traffic can help ensure the integrity and security of your system.

Additional cipher settings are:

- **High** encryption cipher suites: those with key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
- **Medium** encryption cipher suites: those using 128 bit encryption.
- **Low** encryption cipher suites: those using 64- or 56-bit encryption algorithms but excluding export cipher suites.

For inbound requests (requests from a client browser in your organization to SSL Manager), consider using Low encryption to improve performance.

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

5. Click **Apply**.
6. Click **Restart** on **Configure > My Proxy > Basic > General**.

Configuring SSL Manager for outbound traffic

Use the page **Configure > SSL > Decryption / Encryption > Outbound** to configure how SSL Manager handles outbound traffic. Outbound traffic travels from SSL Manager to the destination Web server. SSL Manager checks the revocation status of this site's certificate before forwarding re-encrypted data to it.

1. Select **IP Address** to forward authentication credentials from one proxy to the next if there are multiple proxies between SSL Manager and the destination host.
2. Select **Send VIA-Header** to add a special header to the HTTP header to describe the proxy chain traffic passed through. This can be helpful in troubleshooting. If you do not want to include a VIA-Header, do not select this box.
3. Under **Protocol Settings**, indicate which protocols you want SSL Manager to support. Supported protocols are SSLv2 and v3, and TLS v1. Select the protocol that your enterprise browser supports; you must select at least one protocol. The default is SSLv2. These settings override the settings for these protocols in the users' browsers.

You can select different protocols for inbound traffic.

4. Select **Session Cache** if you want to cache keys until the time specified in Session Cache Timeout elapses. If keys are not cached, each request is negotiated again.
5. Indicate, in seconds, how long keys should be kept in the cache. The default is 300 seconds (5 minutes).
6. The cipher list describes available algorithms and level of encryption between the client and SSL Manager. The default settings indicate to use all available ciphers except the eNULL and the ADH Suite. 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. Setting encryption to a high level for outbound traffic can help ensure the integrity and security of your system.

Additional cipher settings are:

- **High** encryption cipher suites: those with key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
- **Medium** encryption cipher suites: those using 128 bit encryption.
- **Low** encryption cipher suites: those using 64- or 56-bit encryption algorithms but excluding export cipher suites.

For outbound requests (requests from SSL Manager to the destination server that is receiving the encrypted data), consider using one of the higher encryption levels to improve security.

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

7. Click **Apply**.
8. Click **Restart** on **Configure > My Proxy > Basic > General**.

Validating certificates

As part of forwarding outbound traffic, SSL Manager checks the certificates of the destination server. Use these tabs to configure certificate validation and to specify what to do in the case of invalid certificates.

Related topics:

[Bypassing verification, page 126](#)

[Keeping revocation information up to date, page 127](#)

Configuring validation

1. Navigate to the page **Configure > SSL > Validation > General**.
2. Select or clear **Enable the Certificate Verification Engine** to enable or disable verifying certificates and checking for certificate revocation. If this option is not selected, checking does not occur. It is recommended that, at initial installation, you accept the default and keep the certificate verification engine disabled, and monitor traffic through your network. After a period, you can then enable this feature and specify how SSL Manager should process frequently accessed sites by creating incidents. See [Managing Web HTTPS site access, page 128](#), for additional information.



Important

If you disable the certificate verification engine, you need to provide information only on the following pages:

- Configure > SSL > Decryption / Encryption > Inbound
 - Configure > SSL > Decryption / Encryption > Outbound
 - Configure > SSL > Logging pages
 - Configure > SSL > Customization > Connection Error
-

3. Select **Deny Certificates where the common name does not match the URL** to define how the proxy handles those cases where the common name of the certificate does not match the URL of the Web server. For example, a certificate from `www.company.com` does not match the URL `www.company.de`.
4. Select **Allow wildcard certificates** if you want to accept a single certificate for an entire domain. This means that individual servers within that domain are not verified; they are all included because of the wildcard.

5. Select **No expired or not yet valid certificates** to deny access to sites whose certificates fit that criteria. If this box is not selected, access to those sites is permitted.

**Note**

Self-signed certificates (certificates without an official certificate authority) are considered invalid and belong in this category.

6. Select **Verify entire certificate chain** to verify all certificates between a certificate and its root certificate authority.
7. Select **Check certification revocation by CRL** to use certificate revocation lists (CRLs) to check a certificate's revocation state.
8. Select **Check certification revocation by OCSP** to use the Online Certificate Status Protocol to check a certificate's revocation state. Currently, OCSP is not used as widely as CRLs, so it is recommended that you indicate CRL in this field and use OCSP as a backup to CRLs.

**Note**

Certification revocation lists are used far more widely. It is recommended that you use OCSP in addition to, rather than instead of, CRLs. See [Keeping revocation information up to date](#), page 127 for more information on CRLs and OCSP.

9. If you enable checking by both CRLs and OCSP, indicate which method SSL Manager should use first for revocation checking.
10. Indicate whether access should be permitted or denied to sites whose certificate revocation status cannot be determined. If this option is selected, access is also denied to sites whose certificates do not contain CRL or OCSP information. You can see this information when you select a certificate authority and choose view certificate. See [View a certificate](#), page 119 for details. This can result in a highly restrictive security policy, with many access denials. You can allow for exceptions by using the incident list to manage access to Web sites. See [Managing Web HTTPS site access](#), page 128.
11. For troubleshooting purposes, you can run an external program on incidents. An incident is logged whenever a client receives an access denied message. See [Managing Web HTTPS site access](#), page 128 for more information on incidents. Enter the path to the script in this field.

The minimum permissions for running this script should be as follows:

```
chmod 700 /opt/WCG/sxsuite/bin/script.sh
chown Websense /opt/WCG/sxsuite/bin/script.sh
chgrp Websense /opt/WCG/sxsuite/bin/script.sh
```

It is recommended that you copy and paste the following script for help in troubleshooting. It captures the following pieces of information and writes them to a file.

- The account that created the incident
- The client IP or the IP address of the previous proxy if the client IP address is not forwarded
- The ID of the incident as shown in the incident list
- A detailed message on what caused the incident
- The profile within the account that caused the incident
- The host section of the URL that precipitated the incident

```
#!/bin/sh
```

```
OUTFILE=/home/Websense/incidents.log
```

```
date >> $OUTFILE
```

```
echo "Account: $SCIP_INCIDENT_ACCOUNT" >> $OUTFILE
```

```
echo "Client-IP: $SCIP_INCIDENT_CLIENTIP" >> $OUTFILE
```

```
echo "Incident-ID: $SCIP_INCIDENT_ID" >> $OUTFILE
```

```
echo "Detailed Message: $SCIP_INCIDENT_MESSAGE" >> $OUTFILE
```

```
echo "Profile: $SCIP_INCIDENT_PROFILE" >> $OUTFILE
```

```
echo "Destination Host URL: $SCIP_INCIDENT_REMOTEHOST" >>
```

```
$OUTFILE
```

```
echo "User: $SCIP_INCIDENT_USER" >> $OUTFILE
```

```
echo >> $OUTFILE
```



Important

It is recommended that you do not enter any of the other commands in the `/opt/WCG/sxsuite/bin/` directory in this field, and that you exercise caution if you enter a script other than the one provided above.

Bypassing verification

Use the **Configure > SSL > Validation > Verification Bypass** page to enable users to visit a site even if the certificate is invalid.

1. Select **Permit users to visit sites with certificate failure after confirmation** to enable users to proceed to a site after they have been informed that the site has an invalid certificate. If this check box is not selected, users do not have the option to browse to the site.
2. 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.

3. For **Timeout**, specify the period of inactivity that elapses between notifications to users who bypassed this site that the site has an invalid certificate. The default is 6 minutes (360 seconds).

It is recommended that you deploy initially with bypass verification enabled. Then, as the incident rate changes, you can use the incident list to enforce policy. See [Managing Web HTTPS site access, page 128](#).

Keeping revocation information up to date

It is recommended that before your site accepts certificates, it checks the status of the certificate to ensure that it has not been revoked. There are two methods of doing this: through CRLs (see [Certificate revocation lists, page 127](#)) and through OCSP (see [Online certification status protocol \(OCSP\), page 127](#)).

Certificate revocation lists

Use the **Configure > SSL > Validation > Revocation Settings** page to configure how SSL Manager keeps revocation information current. By default, SSL Manager downloads CRLs on a daily basis.

1. For daily downloads of the CRLs, select **Download the CRL at**, and select the time when the CRL download occurs.
2. Click **Apply**.

Use this page as well if you need an immediate CRL update.

1. Click **Update CRL Now** to download the CRLs at a time other than that specified. For example, if your subscription includes SSL Manager, download the CRLs after you install the program.



Note

The CRL files can contain thousands of certifications, so downloading CRLs can take some time and consume CPU resources. It is recommended that you download CRLs at a time when Internet traffic on your system is light.

2. Click **View CRL Update Progress** to see the status of the update.

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

Online certification status protocol (OCSP)

OCSP is a protocol that operates on a request/response basis. That is, 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.

OCSP, because it is dealing with requests, rather than downloading CRLs, can provide improved performance. However, not all CAs provide responses, so CRLs can provide information about the status of more certificates.

SSL Manager 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** page to configure how SSL Manager keeps revocation information current.

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 on OCSP, see RFC 2560.

Managing Web HTTPS site access

Related topics:

[Viewing incidents, page 128](#)

[Changing the status of an incident, page 130](#)

[Deleting an incident, page 130](#)

[Changing the text of a message, page 130](#)

[Viewing incident details, page 131](#)

[Adding Web sites to the incident list, page 131](#)

These tabs can help you manage access to Web sites and can aid the Help Desk in troubleshooting access issues. Entries and changes made to this page are saved in the SSL Manager database.

When a client receives an access denial message because the Web site does not comply with security policies, SSL Manager generates an *incident*. See [Viewing incidents, page 128](#).

If you want to specify how SSL Manager treats a particular site, you can add that to the incident list as well. See [Adding Web sites to the incident list, page 131](#).

Viewing incidents

Use the **Configure > SSL > Incidents > Incident List** page to see a report of those times when clients received an access denial message. You can use the fields in this report to determine how SSL Manager treats requested access to a site in the future.

- ▶ To view incidents:
 - To view a specific incident, enter the ID number and click **Search**.
 - Click **Show All** to see the complete list of incidents.

The incident report

You can sort on any column by clicking on the small triangle next to the column heading.

The incident report contains these fields:

Field	Description
ID	Assigned by the system, this is the incident ID number, also called the Ticket ID. The Help Desk can ask the user for the Ticket ID in the error message and quickly retrieve it from the URL Incident List. The end user sees the Ticket ID and a denial message.
Status	Determines how SSL Manager will treat this Web site in the future. Four conditions are possible: <ul style="list-style-type: none"> • 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 Web site is blocked, unless Verification Bypass is configured, in which case the block page includes a “Visit site anyway” button. See Bypassing verification, page 126. • Tunnel The site is white-listed (tunneled). Traffic is not decrypted and SSL Manager does not check the certificate. Tunneling can be used to bypass inspection of trusted sites and improve performance. You can change the status of a Web site via the drop-down box in the Action column.
Type	Indicates whether the site was added based on its URL or its certificate. It is recommended that you add sites to the incident list by certificate. See Adding Web sites to the incident list, page 131 .
URL	The URL of a site whose certificate could not be validated.

Field	Description
Message	Enables you to edit the error message. See Changing the text of a message, page 130 for information on customizing error messages. The pencil and the magnifying glass each represent links. See Viewing incident details, page 131 for details on these links.
Action	Enables you to change the status of an incident. Also allows you to delete the incident. See Deleting an incident, page 130 .

Changing the status of an incident

When you change the status of an incident, you are changing how SSL Manager will treat the listed URL in the future.

1. Navigate to **Configure > SSL > Incidents > Incident List**.
2. Select one of the following from the drop-down list in the Actions column. See [The incident report, page 129](#) for an explanation of these options.
 - Tunnel
 - Block
 - Blacklist
 - Allow
3. Click **Go**. The icon in the Status column changes to reflect the new status.

Deleting an incident

1. Navigate to **Configure > SSL > Incidents > Incident List**.
2. Select the incident to delete. If the incident is not visible, you can search by ID. See [Viewing incidents, page 128](#).
3. In the Action column, select **Delete** from the Action drop-down list, and then click **Go**.

Changing the text of a message

1. Navigate to **Configure > SSL > Incidents > Incident List**.
2. Locate the incident you want to examine more closely. See [Viewing incidents, page 128](#).
3. Click the pencil to open a window where you can change the text of this error message. For example, the Help Desk can add more detail to an error message.
4. Click **Submit** when the new text is complete, or click **Close Window** if you are not making any changes.

Viewing incident details

1. Navigate to **Configure > SSL > Incidents > Incident List**.
2. Locate the incident you want to examine more closely. See [Viewing incidents, page 128](#).
3. Click the magnifying glass to see additional details about the incident, such as the:
 - Description (this is the message that appears in the incident listing)
 - Time the incident was created
 - Time the incident was modified
 - Incident count (how many times users have tried to access this site)

Adding Web sites to the incident list

Use **Configure > SSL > Incidents > Add Website** page 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, page 128](#).

1. Enter the URL of the site you are adding to the Incident List.
2. Select either **By Certificate** or **By URL**.
 - **By Certificate** provides greater security. If you add a Web site by certificate, clients cannot bypass the policy by using the IP address rather than the URL. If you select “By Certificate,” SSL Manager retrieves the server certificate and adds the site to the incident list. Every new site is whitelisted by default. To blacklist the site, find its entry in the URL Incident List and select the desired action. See [Viewing incidents, page 128](#).
If sites are blocked by certificates, wildcard certificates are not accepted, even if the common name is recognized.
 - Select **By URL** only if you want to allow the site.
3. In the Action drop-down list, specify if the site should be added with Tunnel, Allow, or Blacklist status. See [The incident report, page 129](#) for details.
4. Click **Add URL**.

It is recommended that you manually add sites to the incident list after you have monitored your network traffic for a period of time, with the certificate verification engine disabled. (See [Configuring validation, page 124](#).) This enables you to improve performance by tunneling trusted sites and blocking those you know should not be accessed. See [The incident report, page 129](#) for information about assigning a status, such as tunneling, to a site and incident.

Client certificates

For security, the destination server may request a client certificate.

Related topics:

[Importing client certificates](#), page 132

[When a client certificate is always required: the hostlist](#), page 133

[Deleting client certificates](#), page 133

When a client certificate is requested

1. Navigate to **Configure > SSL > Client Certificates > General**.
2. Select **Tunnel** or **Create incident** to specify how SSL Manager should handle that certificate and site. You must choose Create incident if you want any disposition other than tunnel (white listing). White listing will always provide the certificate to the server. See [The incident report](#), page 129 for a listing of possible dispositions.
3. Click **Apply**.

Importing client certificates

Use the **Configure > SSL > Client Certificates > Import** page to import certificates from the organization represented by the client.



Important

Remember to use only X.509-formatted, base64-encoded certificates.

1. Enter the name of the client certificate.
2. Enter the public key for the certificate. You may need to check with your network administrator for the key.
3. Enter the private key for the certificate. You may need to check with your network administrator for the key.
4. Enter, and then confirm, the passphrase. It is recommended that you use a strong passphrase, with a combination of numbers, characters, and upper- and lower-case letters. You may need to check with your network administrator for the passphrase.
5. Click **Import**.

When a client certificate is always required: the hostlist

Use the **Configure > SSL > Client Certificates > Hostlist** page to list those destination servers that always require a client certificate. Be sure to import the certificate before adding it to the hostlist. See [Importing client certificates](#), page 132.

1. Enter the URL of the destination server that requires the client certificate.
2. 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**.

Deleting client certificates

Use the **Configure > SSL > Client Certificates > Manage Certificates** page to delete imported client certificates.

1. Select the certificate you want to delete.
2. Click **Delete**.

Configuring logging

Related topics:

[How long should log files be kept?](#), page 134

[How big can log files grow?](#), page 134

[What fields should appear in the access log files?](#), page 135

SSL Manager creates 2 types of log files.

- ◆ Activity logs. These logs monitor SSL Manager activity and include messages at a level specified in the user interface
- ◆ Access logs

You can log activity for both inbound (client to SSL Manager) and outbound (SSL Manager to server) traffic. You have the option of logging data to the system log (syslog) or to a file.

Use the **Configure > SSL > Logging > General** page to specify the name and location of log files.

1. For *inbound* traffic, select the type of log files you want to keep. For activity logs, you are specifying the level of detail in the log.

- Enter a number from 1 to 7 to indicate the level of detail you want logged. Note that each level provides more information; level 7 is the most verbose. The levels of logging and granularity are:

1 (alert)	Log conditions that should be corrected immediately, such as a corrupted system file
2 (critical)	Log conditions such as device failures
3 (normal)	Log errors
4 (warning)	Log warnings
5 (notice)	Log conditions that are not error conditions, but may still require attention
6 (information)	Log informational messages
7 (debugging)	Log debugging information. Level 7 includes the most log output.

- Indicate if log data should go to the syslog or to a file.
- Repeat [Step 2](#) and [Step 3](#) for the access log file.
- For *outbound* traffic, repeat [Step 2](#) through [Step 4](#).
- Click **Apply**.

Logs are written to `/opt/WCG/sxsuite/log`.

How long should log files be kept?

A new set of log files is created every 24 hours. By default, this occurs at midnight. This rotation happens regardless of the size of the log file. In addition, the log file is rotated if it reaches its maximum size before the scheduled rotation. In that case, the scheduled rotation still takes place at midnight. See [How big can log files grow?](#), page 134 for information on specifying the maximum log size.

Use the **Configure > SSL > Logging > Options** page to specify how long to keep log files.

- Specify, in days, how long log files should be kept. The default is 3.
- Set any additional options on this page and then click **Apply**.

How big can log files grow?

Log files are rotated every night at midnight. However, a new log file is started when the file reaches its specified maximum size, even if this is before the scheduled daily rotation. Because the size of log files is checked every minute, it is possible that a log file may be larger than its maximum size for a brief period.

When a log file reaches its maximum size, it is saved with an extension of “.x” (where *x* is 1, 2, or 3, etc.), and a new file is started. If this should happen multiple times in a

24-hour period, you must indicate how many files (generations) should be kept. See [How long should log files be kept?](#), page 134 for information on log rotation.

Use the **Configure > SSL > Logging > Options** page to specify how large log files can grow.

1. Indicate, in KB, the maximum size for log files. The default is 50,000 KB.
2. For generations, indicate how many log files should be kept if the file reaches its maximum size multiple times before daily rotation. Once this number is reached and new log files are created, the oldest log file is deleted. The default is 3 generations.
3. Set any additional options on this page and then click **Apply**.

What fields should appear in the access log files?

Use the **Configure > SSL > Logging > Options** page to add or delete fields to the log file.

1. Delete or add fields in the Access log file customization box. The fields are:

time_stamp	The timestamp in the following format:[YYYY.MM.DD HH:MM:SS]
time_of_day	The timestamp in raw format: Sec.mSec starting from 1st Jan 1970 UTC
src_ip	The client's IP Address
auth_user	The user who has been authenticated
account	The account the user belongs to
profile	The user's profile
req_line	The Request in the following format: "method path protocol/version.subversion". For example: GET / HTTP/1.1
status_code	The HTTP status response code sent by the Web server
user_agent	The name of the client browser
referer	The host section of the URL
content_type	Content, such as HTML, text, image, etc.
content_length	In bytes
server_host	IP address of the Web server
bytes_from_client	Bytes transferred from the client to SSL Manager
bytes_to_client	Bytes transferred from SSL Manager to the client
bytes_from_server	Bytes transferred from the Web server to SSL Manager
bytes_to_server	Bytes transferred from SSL Manager to the Webserver

2. Click **Apply**.

Customizing messages

You can customize the message users receive when:

- ◆ They are trying to connect to a site that has an invalid certificate. See [Certificate validation failed, page 136](#).
- ◆ There is a connection failure. See [Connection failure, page 137](#).

The following variables are available within the message templates.

%P	Protocol (HTTP or HTTPS)
%h	The IP address and port of the host of the proxy that generated the message
%o	The IP address of the host of the proxy that generated the message
%H	Remote hostname of the request
%t	Time
%s	Name of the SSL Manager server
%u	Complete URL
\$\$DETAILS	Detailed error message
\$\$TICKET_ID	The ID number of the incident.

Certificate validation failed

Use the **Configure > SSL > Customization > Certificate Failure** page to customize the message users receive when certificate validation fails.



Note

You may find it helpful to click **Preview** to see how the default message appears.

1. Edit the HTML code in the window to reflect your message. See [Customizing messages, page 136](#) for a listing of variables you can use in the message.
2. Click **Preview** to see your changes.
3. Repeat steps 1 and 2 until the message appears appropriately.
4. Click **Apply** to confirm your edits or **Cancel** to return to the original message.

Connection failure

Use the **Configure > SSL > Customization > Connect Error** page to customize the message users receive when SSL Manager is unable to connect to the destination Web server.

**Note**

You may find it helpful to click **Preview** to see how the default message appears.

1. Edit the text in the window to reflect your message. See [Customizing messages, page 136](#) for a listing of variables you can use in the message.
2. Click **Preview** to see your changes.
3. Repeat steps 1 and 2 until the message appears appropriately.
4. Click **Apply** to confirm your edits or **Cancel** to return to the original message.

13

Security

Websense Content Gateway enables you to establish secure communication between the Content Gateway system and other computers on the network. You can:

- ◆ Control which clients are allowed to access the proxy cache. See [Controlling client access to the proxy cache](#), page 139.
- ◆ Control which hosts are allowed to access the proxy server. See [Controlling host access to the proxy server](#), page 140.
- ◆ Control and secure access to Content Gateway Manager using:
 - Administrator accounts (see [Setting the administrator ID and password](#), page 142 and [Creating a list of user accounts](#), page 142).
 - An access control list that defines which hosts are allowed to access Content Gateway Manager (see [Controlling host access to Content Gateway Manager](#), page 143).
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access (see [Using SSL for secure administration](#), page 143).
- ◆ Create filtering rules to control access to the Internet, specify special authentication requirements, and control other traffic transiting the proxy. See [Filtering Rules](#), page 144.
- ◆ Configure Content Gateway integration into your firewall and control traffic through the SOCKS server. See [Configuring SOCKS firewall integration](#), page 146.
- ◆ Configure the proxy to use multiple DNS servers to match your site's security configuration. See [Using the Split DNS option](#), page 149.
- ◆ Configure Content Gateway to use proxy authentication. The proxy supports LDAP, RADIUS, and NTLM proxy authentication. See [Configuring proxy authentication](#), page 150.

Controlling client access to the proxy cache

1. Navigate to the **Configure > Security > Connection Control > Proxy Access** page.
2. Click **Edit File** to open the configuration file editor for the **ip_allow.config** file.

3. Enter information in the fields provided, and then click **Add**. All the fields are described in *Configuration Options*.
4. 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 host access to the proxy server

You might want to restrict the type of communication possible with machines running Content Gateway. Using the ARM security option, you can create an access control list to either allow or deny other hosts from communicating with the Content Gateway machine on specific ports. This *firewall* prevents potentially malicious packets from disrupting the operation of the machine.

**Important**

Content Gateway requires that several ports be open to support the full set of Websense Web Security Gateway features. For a complete list of required ports, please see the *Content Gateway Installation Guide*.

When the ARM security option is enabled, the ARM examines UDP and TCP packets as they arrive at the Content Gateway machine and matches them against the access control list that you specify in a configuration file. The ARM checks all UDP packets (since UDP communication is connectionless) and looks at the first TCP packet initiating the session against the configuration file access control list. Acceptable packets using either protocol are then passed up the network stack. Only incoming UDP and TCP packets are affected. This means that it is always possible to initiate TCP and UDP connections from the proxy cache, regardless of the access control list configured.

To use the ARM security feature, you must perform the following procedures in the order listed:

1. Edit the **arm_security.config** file (default location is in **/opt/WCG/config**) to open specific ports and define the hosts that are allowed to communicate with the Content Gateway machine.

2. Enable the ARM security option.

**Important**

By default, when you enable the ARM security option, all ports on the Content Gateway machine are closed, except for the Content Gateway ports, DNS service ports, and ssh port 22. Before you enable the ARM security option, ensure that you have either console access to the Content Gateway machine or that you have added the appropriate rules to the **arm_security.config** file to allow **telnet** or **ssh** access for yourself.

When you enable the ARM security option, the ports you specify in the access control list remain closed even when Content Gateway is not running.

**Important**

You must define ports and hosts in the **arm_security.config** file before you enable the ARM Security option so that you do not lock yourself out of the Content Gateway machine.

1. Navigate to **Configure > Security > Connection Control > ARM Security**.
2. Click **Edit File** to open the configuration file editor for the **arm_security.config** file.
3. Enter information in the fields provided, and then click **Add**. All fields are described in *Configuration Options*, page 219.
4. Click **Apply**, and then click **Close**.
5. On the **ARM Security** tab, enable the **ARM Security** option.
6. Click **Apply**.
7. Click **Restart** on **Configure > My Proxy > Basic > General**.

Controlling access to Content Gateway Manager

You can restrict access to Content Gateway Manager to ensure that only authenticated users can change configuration options and view performance and network traffic statistics.

You can:

- ◆ Set the master administrator ID and password. A user who logs on to Content Gateway Manager with the administrator ID has access to all Content Gateway Manager activities. See *Setting the administrator ID and password*, page 142.

- ◆ Create and maintain a list of user accounts that determines who can log on to Content Gateway Manager and which activities they can perform. See [Creating a list of user accounts](#), page 142.
- ◆ Create an access control list of IP addresses that defines which machines can access Content Gateway Manager. See [Controlling host access to Content Gateway Manager](#), page 143.
- ◆ Use SSL for secure administration. See [Using SSL for secure administration](#), page 143.

Setting the administrator ID and password

During installation, you assign a password that controls administrative access to Content Gateway Manager. A user who logs on to 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.

You can change the administrator ID and password at any time.

1. Navigate to the **Configure > My Proxy > UI Setup > Login tab**.
2. Make sure that **Basic Authentication** is enabled.
When Basic Authentication is disabled, any user can access Content Gateway Manager unless you have set up a list of IP addresses that are denied access (see [Controlling host access to Content Gateway Manager](#), page 143).
3. To change the current administrator ID, type a new ID in the **Login** field of the **Administrator** section.
4. 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.
If you have forgotten the current administrator password, see [How do you access Content Gateway Manager if you forget the master administrator password?](#), page 396.
5. Click **Apply**.

Creating a list of user accounts

If a single administrator ID and password for Content Gateway Manager is not sufficient security for your needs, you can create a list of user accounts that define who has access to the Content Gateway Manager and which activities they can perform.

1. Navigate to **Configure > My Proxy > UI Setup > Login**.
2. Enter the name of the user allowed to access Content Gateway Manager.
3. Enter the password for the user, and then enter the password again in the New Password (Retype) field.
4. Click **Apply**.

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**.
7. Repeat [Step 2](#) through [Step 6](#) for each user allowed to access Content Gateway Manager.
8. Make sure that **Basic Authentication** is enabled.

Content Gateway checks user names and passwords only if this option is enabled.

Controlling host access to Content Gateway Manager

In addition to using an administrator ID and user accounts, you can control which hosts have access to Content Gateway Manager.

1. Navigate to **Configure > My Proxy > UI Setup Access**.
2. In the Access Control area, click **Edit File** to open the configuration file editor for the **mgmt_allow.config** file.
3. Enter information in the fields provided, and then click **Add**. All the fields are described in [UI Setup, page 223](#).
4. Click **Apply**, and then click **Close**.

Using SSL for secure administration

Websense supports the Secure Sockets Layer protocol (SSL) to provide protection for remote administrative monitoring and configuration using Content Gateway Manager. SSL security provides authentication of both ends of a network connection using certificates and provides privacy using encryption.

To use SSL, you must:

- ◆ Obtain an SSL certificate
- ◆ Enable the Content Gateway Manager SSL option

Obtaining an SSL certificate

You can obtain an SSL certificate from a recognized certificate authority (for example, VeriSign). Install the certificate in the Content Gateway **config** directory (default location is **/opt/WCG/bin**). You must either rename the certificate to the default

filename **private_key.pem**, or specify the name of the certificate using Content Gateway Manager (follow the procedure in [Enabling SSL, page 144](#)).

Enabling SSL

After you have obtained an SSL certificate, you can enable SSL.

1. Navigate to the **Configure > My Proxy > UI Setup > General** tab.
2. Enable the **HTTPS** option.
3. In the Certificate File field, specify the filename of the SSL certificate.
You have to change the filename only if the certificate file does not use the default name **private_key.pem**.
4. Click **Apply**.

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
- ◆ Keep or strip header information from client requests
- ◆ Allow specified applications or requests to specified Web sites to bypass authentication
- ◆ Prevent specified applications from transiting the proxy



Note

To create rules for NTLM and LDAP authentication, see [Multiple authentication realms, page 161](#). To get started with Content Gateway authentication options, see [Configuring proxy authentication, page 150](#).

Filtering rules are created and modified on the **Configure > Security > Access Control > Filtering** tab. 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. However, you cannot repeat a secondary specifier.

See [filter.config](#) for information about the structure of stored rules.

Creating filtering rules

1. Go to the **Configure > Security > Access Control > Filtering** tab and click **Edit File** to open [filter.config](#) in the file editor.

2. 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 caches and 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.

**Note**

The “radius” rule type is **not** supported.

3. 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.

4. 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**

5. Provide values for any 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)

**Note**

Although **rtsp** and **mms** are included in the Scheme drop-down list, they are not supported.

User-Agent — Specifies a request header User-Agent value.

You can use the User-Agent field to create application filtering rules that:

- Allow particular applications that don't properly handle authentication challenges to bypass authentication
 - Block particular client-based applications from accessing the Internet
6. When you have finished defining the rule, click **Add** to add the rule and then **Apply** to save the rule.
 7. When you are done adding rules, click **Apply** to save all the changes and then click **Close** to close the edit window.
 8. To put the new rules into effect, select the Content Gateway Manager window and restart Content Gateway.

Editing a rule:

1. Go to **Configure > Security > Access Control > Filtering** and click **Edit File** to open *filter.config* in the file editor.
2. In the list, select the rule to be modified and change the values as desired.
3. Click **Set** to update the rule and click **Apply** to save the rule.
4. Click **Close** to close the edit window.
5. To put the changes into effect, select the Content Gateway Manager window and restart Content Gateway.

Configuring SOCKS firewall integration

SOCKS is commonly used as a network firewall that allows hosts behind a SOCKS server to gain full access to the Internet and prevents unauthorized access from the Internet to hosts inside the firewall.

When the proxy receives a request for content that is not in the cache or is stale, it must request the content from the origin server. 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.

The proxy caches the content and sends it to the client.

Content Gateway can act as a SOCKS client, where it receives and serves HTTP or FTP requests as usual. It can also act as a SOCKS proxy, relaying requests to and from

the SOCKS server (usually on port 1080). Content Gateway cannot and does not 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 the proxy to use a SOCKS firewall

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the Features table, click **SOCKS On** in the Security section.
3. Click **Apply**.
4. Navigate to **Configure > Security > SOCKS > General**.
5. Specify the SOCKS version running on your SOCKS servers.
6. Click **Apply**.
7. Click the **Server** tab.
8. In the Default Servers field of the SOCKS Server section, enter the host names of your default SOCKS servers and the ports through which the proxy communicates with the SOCKS servers. Separate the host name and the port with a colon (:), and separate each entry with a semicolon (;). For example:

```
socks1:1080;socks2:4080
```
9. Click **Apply**.
10. In the SOCKS Server Rules area, click **Edit File** to perform additional SOCKS server configuration, such as SOCKS server bypass and authentication.
The configuration file editor for the **socks.config** file opens.
11. Enter information in the fields provided, and then click **Add**. All the fields are described in [Configuration Options](#).
12. Click **Apply**, and then click **Close**.
13. Click **Restart** on **Configure > My Proxy > Basic > General**.

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 forwards all 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 the proxy to use a SOCKS firewall](#), page 147.

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**.

Setting authentication

For SOCKS version 5, you can configure the proxy to use username and password authentication with the SOCKS server. You specify the username and password the proxy must use in the **socks.config** file.

1. Navigate to **Configure > Security > SOCKS > Server** to display the **socks.config** file.
2. At the end of the file, add a line using the following format:

```
auth u username password
```

username is the username and *password* is the password Content Gateway must use for authentication with the SOCKS version 5 server.
3. Click **Apply**.
4. Click **Restart** on **Configure > My Proxy > Basic > General**.

Setting SOCKS server bypass

You can configure Content Gateway to bypass the SOCKS server and access certain origin servers directly (SOCKS Server bypass).

1. Navigate to **Configure > Security > SOCKS > Server** to display the **socks.config** file.
2. Enter a line in the file specifying the IP addresses or IP address range of the origin servers that you want Content Gateway to access directly. Use the following format:

```
no_socks IPaddresses or IPaddress range
```

IPaddresses or *IPaddress range* is a comma-separated list of the IP addresses or IP address ranges associated with the origin servers you want Content Gateway to access directly.

Never specify the all networks broadcast address: 255.255.255.255.

3. Click **Apply**
4. Click **Restart** on **Configure > My Proxy > Basic > General**.

Example

To configure 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
```

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 Content Gateway Manager:

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**.

Configuring proxy authentication

Websense Content Gateway supports LDAP, RADIUS, and NTLM proxy authentication in both explicit and transparent proxy modes.



Important

Enable only one authentication option.

See the following sections for information:

- ◆ *Browser limitations*, page 150
- ◆ *Transparent proxy authentication*, page 151
- ◆ *Using LDAP proxy authentication*, page 152
- ◆ *Using RADIUS proxy authentication*, page 155
- ◆ *Using NTLM proxy authentication*, page 158
- ◆ *Multiple authentication realms*, page 161

Content Gateway does not support proxy authentication in clustered environments using WCCP load balancing.

When sending traffic through a Web Cache Control Protocol (WCCP)-enabled router, WCCP decides which proxy to use based on the destination host. This maximizes the potential that the cache will have the content it requires. One user can leverage many proxies.

Authentication on the proxy requires users to identify themselves. When WCCP is combined with load balancing, end users are prompted multiple times, once for each proxy used.

To support authentication, the load balancer needs to ensure flow affinity by the client (source) as opposed to the destination. This requires special hardware and is not supported on basic WCCP routers and switches.

For more information, see the Mask-Assignment page on the Cisco Web site:
http://www.cisco.com/en/US/docs/ios/ipapp/configuration/guide/ipapp_wccp.html

Browser limitations

Not all Web browsers fully support the 2 common authentication methods:

- Transparent NTLM (no-prompt, Integrated Windows Authentication)
- Explicit authentication (prompt the user for credentials)

Internet Explorer 7 and 8

- ◆ Full support

Mozilla Firefox 2 and 3

- ◆ Full support

Google Chrome

- ◆ Explicit authentication supported
- ◆ NTLM transparent authentication not supported
 - Always challenges users for credentials

Opera 10

- ◆ Explicit authentication partially supported
 - HTTP supported
 - HTTPS not supported
- ◆ NTLM transparent authentication not supported
 - Always challenges users for credentials

Windows Safari 4

- ◆ Explicit authentication partially supported
 - HTTP supported
 - HTTPS:
 - Supported for NTLM
 - LDAP users prompted twice
- ◆ NTLM transparent authentication not supported
 - Always challenges users for credentials



Note

When prompted for credentials, if the user does not enter a domain name, a “session timeout” error can result, or the user may be re-prompted.

Transparent proxy authentication

Whenever the NTLM, LDAP, or Radius authentication method is enabled, Content Gateway simultaneously performs authentication for users that are filtered in explicit proxy mode (for example, via a PAC file) and users that are filtered in transparent proxy mode (for example, via a WCCP redirect).

To properly perform authentication for users in transparent proxy mode, Content Gateway provides several extended configuration options. In Content Gateway Manager, go to the **Configuration > Security > Access Control > Transparent Proxy Authentication** tab.

- ◆ **Redirect Hostname** (optional) — specifies an alternate hostname for the proxy.

By default, authenticating clients are redirected to the hostname of the Content Gateway machine. If clients aren't able to resolve that hostname via DNS, or if an alternate DNS name for the proxy is defined, that hostname can be specified in the **Redirect Hostname** field.

**Note**

In order to ensure that the authentication for transparent proxy users occurs transparently (i.e., 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 will allow authentication to occur transparently, without prompting the user. Consult your browser documentation.

- ◆ **Authentication Mode** — specifies the authentication mode. Content Gateway must be set to one of the following authentication modes:
 - In **IP mode** (the default), the client IP address is associated with a username when the session is authenticated. Requests made from that IP address are not authenticated again until the **Session TTL** expires (default = 15 minutes). Requests made from that IP address within the time-to-live are considered to be made by the user associated with that IP address.
 - **Cookie mode** is used to uniquely identify users who share a single IP address, such as, for example, environments in which proxy-chaining is used or where network address translation (NAT) occurs.
- ◆ **Session TTL** — Once the user's session is authenticated, the session is valid for the duration of the time specified in **Session TTL** (time-to-live; default = 15 minutes). The supported range of values is 5-65535 minutes.

Whenever changes are made to any of these fields, click **Apply** to save your changes and then restart the proxy to put the changes into effect.

Using LDAP proxy authentication

Websense Content Gateway provides the LDAP option to ensure that users are authenticated with an LDAP server before accessing content through the proxy cache. In the proxy, LDAP authentication supports both simple and anonymous bind.

When you enable the LDAP option, the proxy acts as an LDAP client and directly challenges users who request content for a username and password. After receiving the username and password from a client, the proxy contacts the LDAP server to check that they are correct. If the LDAP server accepts the username and password, the proxy serves the client with the requested content and stores the username and password entry in the Content Gateway LDAP cache; all future authentication requests for that user are served from the LDAP cache until the entry expires. 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.

Configuring Content Gateway to be an LDAP client

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the Features table, click **LDAP On** in the Authentication section.
3. Click **Apply**.
4. Navigate to **Configure > Security > Access Control > LDAP**.
5. Enable **Purge Cache on Authentication Failure** to configure the proxy to delete the authorization entry for the client from LDAP cache if authorization fails.
6. Enter the host name of the LDAP server.
7. 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.

8. Enable Secure LDAP if you want the proxy to use secure communication with the LDAP server. Secure communication is usually performed on port 636 or 3269. (Change the port value in the previous field, if necessary.)
9. Select the type of your directory service to set the filter for searching. The default is **sAMAccountName** for Active Directory. Select **uid** for eDirectory or other directory services.
10. Enter 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.
11. Enter a password for the user specified in the previous step.
12. Enter the Base Distinguished Name (DN). Obtain this value from your LDAP administrator.
13. Click **Apply**.
14. Click **Restart** on **Configure > My Proxy > Basic > General**.

As optional steps, you can:

- ◆ Change LDAP cache options; See [Setting LDAP cache options, page 154](#).

- ◆ Configure Content Gateway to allow certain clients to access specific sites on the Internet without being authenticated by the LDAP server; See [Access Control](#), page 251).
- ◆ Configure an alternate Content Gateway hostname for authentication, set the Authentication Mode (IP Mode or Cookie Mode), and set the session time-to-live period; See [Transparent proxy authentication](#), page 151.
- ◆ Configure multiple authentication realms in which certain IP addresses use specific LDAP servers; See [Multiple authentication realms](#), page 161.

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.

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is in **/opt/WCG/config**).
2. Edit the following variables:

Variable	Description
<code>proxy.config.ldap.cache.size</code>	Specify the number of entries allowed in the LDAP cache. The default value is 5000. The minimum value is 256.
<code>proxy.config.ldap.auth.ttl_value</code>	Specify the number of minutes that Content Gateway can store username and password entries in the LDAP cache.
<code>proxy.config.ldap.cache.storage_size</code>	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 (default location is in **/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:

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is 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. Please refer to the documentation provided by your directory provider for instructions.

Using RADIUS proxy authentication

Websense Content Gateway provides the RADIUS option to ensure that users are authenticated with a RADIUS server before accessing content the proxy cache.

When you enable the RADIUS option, 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 from a client, Content Gateway contacts the RADIUS server to check that they 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 host name 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](#), page 156.

Configuring Content Gateway to be a RADIUS client

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the Features table, click Radius **On** in the Authentication section.
3. Click **Apply**.
4. Navigate to **Configure > Security > Access Control > Radius**.
5. Enter the host name of your primary RADIUS server.
6. Enter the port number through which Content Gateway communicates with the primary RADIUS server.
7. Enter the key used for encoding.
8. If you are using a secondary RADIUS server, enter the host name, port, and shared key in the appropriate fields of the **Secondary Radius Server (Optional)** area.
9. Click **Apply**.
10. 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.

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is in **/opt/WCG/config**).
2. Edit the following variables:

Variable	Description
<code>proxy.config.radius.auth.min_timeout</code>	Specify the amount of time in seconds that the Content Gateway connection to the RADIUS server remains idle before Content Gateway closes the connection.
<code>proxy.config.radius.auth.max_retries</code>	Specify the maximum number of times Content Gateway tries to connect to the RADIUS server.
<code>proxy.config.radius.cache.size</code>	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.
<code>proxy.config.radius.auth.ttl_value</code>	Specify the number of minutes that Content Gateway can store username and password entries in the RADIUS cache.
<code>proxy.config.radius.cache.storage_size</code>	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 (default location is in **/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.

Using NTLM proxy authentication

Websense 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.

When the 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 to the user.

Content Gateway supports both transparent (Single Sign-On) and explicit authentication. Transparent authentication is supported with Microsoft Internet Explorer 7 and 8, and Mozilla Firefox 2 and 3. Single Sign-On allows users to sign on only once, so that they can seamlessly 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 authentication and the user is not prompted again for a username and password. Explicit (basic) authentication is supported for other browsers. With explicit authentication, users are prompted for a username and password before they can access the protected content.

Content Gateway supports the use of backup domain controllers for failover. If the primary domain controller does not respond to proxy requests, Content Gateway contacts the next domain controller in the list (the backup domain controller). For the next request, the proxy tries to contact the primary domain controller again and then contacts the backup domain controller if the connection fails. Content Gateway does this 5 times before considering the server unavailable. After marking the primary domain controller unavailable, the proxy waits 5 minutes before trying to contact it again.

Content Gateway supports access to Windows NT domain controllers and Windows 2000, 2003, and 2008 Active Directory.

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.
3. **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](#), page 150.

6. NTLM credential caching is performed when authentication is successful in explicit mode. Transparent proxy authentication caching is handled separately and is configured on the **Configuration > Security > Access Control > Transparent Proxy Authentication** tab.

Configuring NTLM proxy authentication

1. Navigate to **Configure > My Proxy > Basic > General**.
2. In the Features table, click NTLM **On** in the Authentication section.
3. Click **Apply**.
4. Navigate to **Configure > Security > Access Control > NTLM**.
5. In the **Domain Controller Hostnames** field, enter the host name of the primary domain controller, followed, optionally, by a comma separated list of backup domain controllers. The format of the host name 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.



Note

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. **Fail Open** is enabled by default. Fail Open allows requests to proceed when authentication fails due to:
 - No response from the domain controller
 - Malformed messages from the client
 - Invalid SMB responsesWith Fail Open, when Web filtering is used with the proxy and an XID agent is configured, if NTLM authentication fails the requester can still be identified by the XID agent and appropriate policy applied.
Disable Fail Open if you want to stop requests from proceeding to the Internet when the above listed authentication failure conditions occur.
8. **Credential Caching** is enabled by default. To disable credential caching, select **Disable**.
9. **Caching TTL** sets the time-to-live from entries in the credential cache. The default TTL is 900 seconds (15 minutes). To change the TTL, enter a new value in the entry field. The range of supported values is 300 to 86400 seconds.
10. If some users use terminal servers to access the Internet through the proxy (e.g., Citrix servers), you must create a list of those servers in the **Multi-user Hostnames** field. Credentials for such users are not cached. Enter a comma-separated list of IP addresses or IP address ranges. Do not enter hostnames.
11. Click **Apply**.
12. Click **Restart** on **Configure > My Proxy > Basic > General**.

Optionally, you can also:

- ◆ Configure Content Gateway to allow certain clients to access specific sites on the Internet without being authenticated by the NTLM server; See [Access Control, page 251](#).
- ◆ Configure an alternate Content Gateway hostname for authentication, set the Authentication Mode (IP Mode or Cookie Mode), and set the session time-to-live period; See [Transparent proxy authentication, page 151](#).
- ◆ Configure multiple authentication realms in which certain IP addresses use specific domain controllers; See [Multiple authentication realms, page 161](#).

Setting NTLM cache options in records.config

On the **Configure > Security > Access Control > NTLM** page you can enable and disable NTLM credential caching, set the time-to-live (TTL) value, and specify terminal server IP addresses. You can also change these values in **records.config**, along with a few other NTLM caching parameters.

By default, the NTLM cache is configured to store 5000 entries. By default, each entry is considered fresh for 15 minutes (900 seconds).

To modify these and other values in **records.config**:

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is in **/opt/WCG/config**).
2. Edit the following variables:

Variable	Description
<code>proxy.config.ntlm.cache.enabled</code>	Set to 0 to disable the NTLM cache. When disabled, Content Gateway does not store any credentials in the NTLM cache.
<code>proxy.config.ntlm.cache.ttl_value</code>	Specify the amount of time (in seconds) that Content Gateway stores entries in the NTLM cache. The supported range of values is 300 to 86400 seconds.
<code>proxy.config.ntlm.cache.size</code>	Specify the number of entries allowed in the NTLM cache. The default is 5000.
<code>proxy.config.ntlm.cache.storage_size</code>	Specify 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.

3. Save and close the file.
4. From the Content Gateway **bin** directory (default location is **/opt/WCG/bin**), run **content_line -L** to restart Content Gateway on the local node or **content_line -M** to restart Content Gateway on all the nodes in a cluster.

Multiple authentication realms

Related topics:

[Creating an NTLM authentication realm rule, page 162](#)

[Creating an LDAP authentication realm rule, page 164](#)

[Editing an authentication realm rule, page 165](#)

[Reordering the list of rules, page 165](#)

[Deleting a rule, page 166](#)

In networks with multiple authentication realms, rules can be defined to direct sets of IP addresses to distinct authentication servers (aka domain controllers). These rules

are defined on the **Configure > Security > Access Control > Authentication Realms** tab. Rules are stored in the *auth.config* file.

- ◆ Multiple authentication realms are supported for NTLM and LDAP authentication.
- ◆ Separate rules must be created for NTLM and LDAP.
- ◆ The specifiers used in LDAP and NTLM rules differ.
- ◆ Only one authentication method can be active at a time, LDAP or NTLM, and only rules for that method are applied. Other rules remain in *auth.config* but are not used.
- ◆ Rules are applied from the list top down; only the first match is applied.

**Note**

If all of the clients in your network are authenticated by authentication servers that share trust relationships, you don't need to create rules for multiple authentication realms.

- ◆ If no rule in the **auth.config** table applies to a client IP address, the default configuration is applied. The settings on the **Access Control > LDAP** and **NTLM** configuration tabs define the default authentication realm.

**Important**

A default authentication realm is required.

What is an authentication realm?

An authentication realm is a set of client IP addresses serviced by one or more authentication servers (domain controllers). Depending on the topology of the network, a single authentication realm may service all clients on the network. However, when a network is segmented such that groups of clients must authenticate with separate authentication servers that don't share trust relationships, then the network has multiple authentication realms, one for each set of client IP addresses that must authenticate with distinct authentication servers.

Creating an NTLM authentication realm rule

Before you create a rule for an NTLM authentication realm, you need to know:

- ◆ The set of source IP addresses to send to the domain controller(s). These can be a mix of individual IP addresses and IP address ranges.
- ◆ The name or IP address, and port number of the primary domain controller and any secondary domain controllers to be used for load balancing or failover.

- ◆ If the domain controller is Active Directory 2008, the name of the NetBIOS host.

**Note**

After entering all specifiers, you must click **Add** before you click **Apply**. If **Apply** is clicked first, or the edit window is closed, all of the entry fields are cleared.

1. In Content Gateway Manager, go to **Configure > Security > Access Control > NTLM** and review (or specify) the default NTLM authentication realm settings.
2. Go to **Configure > Security > Access Control > Authentication Realms**. A list of all existing authentication realm rules is displayed at the top of the page.
3. Click **Edit file** to open the rule editor.
4. Select **NTLM** from the **Rule Type** drop down list.
5. Give the rule a unique **Profile Name**.
6. In the **Source IP** field, enter a comma-separated list of individual IP addresses and IP address ranges. For example:
10.4.1.1, 10.12.1.1-10.12.254.254
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.
7. Select **Enable** if you want the rule to be active after you **Add** the rule and click **Apply**.
8. In the **DC List** 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.
9. Select **DC Load Balance** to enable load balancing between 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.

10. Click **Add** to add the rule.

11. Click **Apply** and then restart Content Gateway to put the rule into effect.

**Warning**

If a rule has values that conflict with another rule, a warning message displays that identifies the invalid rule.

Creating an LDAP authentication realm rule

Before you create an LDAP authentication realm rule, you need to know:

- ◆ The set of source IP addresses to send to each LDAP server. These can be a mix of individual IP addresses and IP address ranges.
- ◆ The name and port number of the LDAP server.
- ◆ The LDAP base distinguished name.
- ◆ Optionally, the LDAP UID filter.
- ◆ The LDAP bind distinguished name and password.
- ◆ Optionally, an LDAP attribute name and value.

**Note**

After entering all specifiers, you must click **Add** before you click **Apply**. If **Apply** is clicked first, or the edit window is closed, all entry fields are cleared.

1. In Content Gateway Manager, go to **Configure > Security > Access Control > LDAP** and review or specify the default LDAP authentication realm settings.
2. Go to **Configure > Security > Access Control > Authentication Realms**. A list of all existing rules is displayed at the top of the page.
3. Click **Edit file** to open the rule editor.
4. Select **LDAP** from the **Rule Type** drop down list.
5. Give the rule a unique **Profile Name**.
6. In the **Source IP** field, enter a comma-separated list of individual IP addresses and IP address ranges. For example:
10.4.1.1, 10.12.1.1-10.12.254.254
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.
7. Check **Enabled** if you want the rule to be active after you **Add** the rule and click **Apply**.
8. In the **LDAP Server Name** field, enter the fully qualified domain name and port number, or IP address of the LDAP server.
9. If the LDAP server port is other than the default (389), in the **LDAP Server Port** field, enter the LDAP server port.

10. Enter the **LDAP Base Distinguished Name**. Obtain this value from your LDAP administrator.
11. Optionally, enter the LDAP UID filter. Use this field to specify the server type when it differs from the **Server Type** value specified on the **LDAP** tab (the default value). Enter **sAMAccountName** for Active Directory, or **uid** for any other service.
12. In the **Bind DN** 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
13. In the **Bind Password** field, enter the password for the name given in the **Bind DN** field.
14. Check **Secure LDAP** if you want Content Gateway to use secure communication with the LDAP server.
15. Optionally, enter an LDAP attribute name.
16. Optionally, enter an LDAP attribute value.
17. Click **Add** to add the rule.
18. Click **Apply** and then restart Content Gateway to put the rule into effect.

**Warning**

If a rule has values that conflict with another rule, a warning message displays that identifies the invalid rule.

Editing an authentication realm rule

1. On the **Configure > Security > Access Control > Authentication Realms** tab, 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 **Apply**.
4. Click **Close** to return to the **Authentication Realms** tab.
5. Restart Content Gateway to put the changes into effect.

Reordering the list of rules

Authentication realm rules are applied top-down in the list.

1. On the **Configure > Security > Access Control > Authentication Realms** tab, click **Edit File**.
2. In the table of rules, click on the rule that you want to reposition in the list, 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 Realms** tab.
5. Restart Content Gateway to put the changes into effect.

Deleting a rule

1. On the **Configure > Security > Access Control > Authentication Realms** tab, 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 Realms** tab.
5. Restart Content Gateway to put the changes into effect.

14

Working With Log Files

Related topics:

[Event log files](#), page 168
[Managing event log files](#), page 169
[Event log file formats](#), page 170
[Rolling event log files](#), page 176
[Splitting event log files](#), page 179
[Collating event log files](#), page 182
[Viewing logging statistics](#), page 185
[Viewing log files](#), page 185
[Example event log file entries](#), page 187

Websense 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 signifying error conditions on Content Gateway Manager; see [Working with alarms](#), page 97, 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 host name of the proxy server that reported the error, and a description of the error or warning.

See [Websense Content Gateway error messages](#), page 385, 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**, or **Disabled**.

Event log files

Event log files record information about every request that Websense Content Gateway processes. By analyzing the log files, you can determine how many people use the proxy cache, 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](#), page 169.
- ◆ 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](#), page 170.
- ◆ 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](#), page 176.
- ◆ Separate log files according to protocols and hosts
You can configure the proxy to create separate log files for ICP and HTTP/FTP transactions. You can also configure the proxy to generate separate log files for different protocols based on the host. See [Splitting event log files](#), page 179.
- ◆ 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 well-defined locations. See [Collating event log files, page 182](#).

- ◆ View statistics about the logging system
Content Gateway provides statistics about the logging system. You can access the statistics through Content Gateway Manager or through the command line interface. See [Viewing logging statistics, page 185](#).
- ◆ 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, page 187](#).

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 **logs** directory, which is located in the directory where you installed Content Gateway. To use a different directory, see [Setting log file management options, page 170](#).

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, page 170](#)), Content Gateway enters a low space state and takes the following actions:

- ◆ If the autodelete option (discussed in [Rolling event log files, page 176](#)) 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.

Setting log file management options

1. Navigate to **Configure > Subsystems > Logging**.
2. In the **Log Directory** field, enter the path to the directory in which you want to store event log files. This can be an absolute path or a path relative to the directory in which Content Gateway is installed. The default directory is **logs**, located in the Content Gateway installation directory.



Note

The directory you specify must already exist.

The Websense user must have read/write permissions for the directory storing the log files.

3. In the **Limit** field of the **Log Space** area, enter the maximum amount of space you want to allocate to the logging directory. The default value is 20480 MB.



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, page 176](#).

5. Click **Apply**.

Event log file formats

Websense Content Gateway supports the following log file formats:

- ◆ *Standard formats*, such as Squid or Netscape (see [Using standard formats, page 171](#))
- ◆ the Content Gateway *custom format* (see [Custom format, page 171](#))

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](#), page 174.



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.

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](#), page 171.

By default, Content Gateway is configured to use the Netscape Extended log file format only.

Setting standard log file format options

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 by through Content Gateway Manager or by editing a configuration file.

1. On **Configure > Subsystems > Logging > Custom**, enable the **Custom Logging** option.
2. 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](#), page 306.

3. Click **Apply**.

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:

1. Navigate to **Configure > Subsystems > Logging > Custom** 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"/>
</Format>
```

where:

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](#), page 306.

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"/>
</Format>
```



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>"/>
```

3. Define a **LogObject** that uses this format.
4. Click **Apply**.

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 Content Gateway Manager. See [Setting](#)

[standard log file format options](#), page 171. For the custom log format, you specify ASCII or Binary mode in the **LogObject**. Refer to [Custom format](#), page 171.

**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](#), page 306, 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.

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.

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 following table describes the command-line options.

Option	Description
<code>-o output_file</code>	Specifies where the command output is directed.
<code>-a</code>	Automatically generates the output filename based on the input filename. If the input is from stdin , this option is ignored. For example: <pre>logcat -a squid-1.blog squid-2.blog squid-3.blog</pre> generates: <pre> squid-1.log, squid-2.log, squid-3.log</pre>
<code>-S</code>	Attempts to transform the input to Squid format, if possible.
<code>-C</code>	Attempts to transform the input to Netscape Common format, if possible.
<code>-E</code>	Attempts to transform the input to Netscape Extended format, if possible.
<code>-2</code>	Attempt to transform the input to Netscape Extended-2 format, if possible.

**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

Websense 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

Websense 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 host name 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
%H	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

1. Navigate to **Configure > Subsystems > Logging > General**.
2. In the **Log Rolling** section, ensure the **Log Rolling** option is enabled (the default).
3. 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.
4. 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 Content Gateway Manager or the **records.config** file described above.

Splitting event log files

By default, Websense Content Gateway uses standard log formats and generates log files that contain HTTP, FTP, NNTP, and ICP transactions in the same file. However, you can enable log splitting if you prefer to log transactions for different protocols in separate log files.

**Note**

If you are using a standard log file format (such as Squid or Netscape), Content Gateway always records HTTP and FTP transactions in the same log file. You cannot generate separate log files for transactions using these two protocols.

ICP log splitting

When ICP log splitting is enabled, Content Gateway records ICP transactions in a separate log file with a name that contains **icp**. For example, if you enable the Squid format, all ICP transactions are recorded in the **squid-icp.log** file.

When you disable ICP log splitting, Content Gateway records all ICP transactions in the same log file as HTTP and FTP transactions.

NNTP log splitting

When NNTP log splitting is enabled, Content Gateway records all commands it sends to the Network News Transfer Protocol (NNTP) server in a separate log file. This allows you to log transactions from incoming news feeds. The log filename contains “nntp”. For example, if you enable the Squid format, all NNTP transactions are recorded in the `squid-nntp.log` file.

When you disable NNTP log splitting, Content Gateway records all NNTP transactions in the same log file as HTTP and FTP transactions.

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, page 181](#)).

When ICP, NNTP, and HTTP host log splitting are all enabled, Content Gateway generates separate log files for HTTP/FTP transactions, based on the origin server, and places all ICP and NNTP transactions in their own respective log files.

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
<code>squid-uni.edu.log</code>	All HTTP and FTP transactions for uni.edu
<code>squid-company.com.log</code>	All HTTP and FTP transactions for company.com
<code>squid-icp.log</code>	All ICP transactions for all hosts
<code>squid.log</code>	All HTTP and FTP transactions for other hosts

If you disable ICP and NNTP log splitting, these transactions are placed in the same log file as HTTP and FTP transactions. Using the previous example hosts and assuming the Squid log format is used, Content Gateway generates these log files:

Log Filename	Description
<code>squid-uni.edu.log</code>	All entries for uni.edu
<code>squid-company.com.log</code>	All entries for company.com
<code>squid.log</code>	All other entries

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, page 171](#).

Setting log splitting options

1. Navigate to **Configure > Subsystems > Logging > Splitting**.
2. Enable the **Split ICP Logs** option to record all ICP transactions in a separate log file. Disable the **Split ICP Logs** option to record all ICP transactions in the same log file as HTTP/FTP transactions.
3. Enable the **Split NNTP Logs** option to record all NNTP transactions in a separate log file. Disable the **Split NNTP Logs** option to record all NNTP transactions in the same log file as HTTP/FTP transactions.
4. 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.
5. Click **Apply**.

Editing the log_hosts.config file

The default **log_hosts.config** file is located in the Content Gateway **config** directory (default location is **/opt/WCG/config**). To record HTTP and FTP transactions for different origin servers in separate log files, you must specify each origin server's host name 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-sport.log** (if the Squid format is enabled).



Note

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.

1. Open the **log_hosts.config** file located in the Content Gateway **config** directory (default location is **/opt/WCG/config**).
2. Enter the host name of each origin server on a separate line in the file. For example:

```
webserver1  
webserver2  
webserver3
```

3. Save and close the file.
4. To apply the changes, run the following command from the Content Gateway **bin** directory (default location is **/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.



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.



Note

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

1. Navigate to **Configure > Subsystems > Logging > Collation**.
2. In the **Collation Mode** section, enable the **Be A Collation Server** option.
3. 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

1. Navigate to **Configure > Subsystems > Logging > Collation**.
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, page 171](#).

3. In the **To Collation Server** field, enter the host name of the collation server. This could be the Content Gateway collation server or a stand-alone collation server.
4. 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.

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 stand-alone 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.

1. Configure your Content Gateway nodes as log collation clients. See [Configuring Content Gateway to be a collation client](#), page 183.
2. Copy the **sac** binary from the Content Gateway **bin** directory (default location is **/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 (default location is in **/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.

- Open the **records.config** file on the stand-alone collator and edit the following variable:

Variable	Description
<code>proxy.config.log2.logfile_dir</code>	Specify the directory where you want to store the 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. Note: The directory must already exist on the machine serving as the stand-alone collator.

- Save and close the file.
- Enter the following command:

```
sac -c config
```

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 Content Gateway Manager or retrieve them through the command-line interface. See [Monitoring Traffic, page 93](#).

Viewing log files

Related topics:

[Squid format, page 187](#)

[Netscape examples, page 188](#)

You can view the system, event, and error log files that Content Gateway creates from 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.

**Note**

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 Content Gateway Manager.

1. Navigate to **Configure > My Proxy > Logs > System**.
2. 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:

- ◆ [Squid format, page 187](#)
- ◆ [Netscape examples, page 188](#)
- ◆ [Netscape Extended format, page 188](#)
- ◆ [Netscape Extended-2 format, page 189](#)

Squid format

The following figure shows a sample log entry in a **squid.log** file. The table below describes each field.

```

1      2      3      4      5      6      7
987548934.123 19 209.131.54.138 TCP_HIT/200 4771 GET http://europe.cnn.com/
EUROPE/potd/2001/04/17/tz.pullitzer.ap.jpg | NONE/- image/jpeg
7 cont'd | 8 9 10

```

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>In Squid- and Netscape-format log files, what do the cache result codes mean?</i> , page 397. The proxy response status code (the HTTP response status code from Content Gateway to client).

Field	Description
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.

Netscape examples

Netscape Common format

The following figure shows a sample log entry in a **common.log** file. The table below describes each field.

```

1      2 3      4      5
209.131.54.138 - - [17/Apr/2001:16:20:28 -0700] "GET http://europe.cnn.com/
EUROPE/potd/2001/04/17/tz.pullitzer.ap.jpg HTTP/1.0" 200 4473
                    5 cont'd                          6 7

```

Netscape Extended format

The following figure shows a sample log entry in an **extended.log** file. The table below describes each field.

```

1      2 3      4      5
209.131.54.138 - - [17/Apr/2001:16:20:28 -0700] "GET http://europe.cnn.com/EUROPE/potd/2001/
04/17/tz.pullitzer.ap.jpg HTTP/1.0" 200 4473 000 0 0 0 458 297 0 0 0
                    5 cont'd                          6 7 8 9 10 11 12 13 14 15 16

```

Netscape Extended-2 format

The following figure shows a sample log entry in an **extended2.log** file. The table below describes each field.

```

1      2 3      4      5
209.131.54.138 - - [17/Apr/2001:16:20:28 -0700] "GET http://europe.cnn.com/EUROPE/potd/2001/04/
17/tz.pullitzer.ap.jpg HTTP/1.0" 200 4473 000 0 0 0 458 297 0 0 0 NONE FIN FIN TCP_MEM_HIT
5 cont'd      6      7      8      9 10 11 12 13 14 15 16 17 18 19      20

```

Field	Description
	Netscape Common
1	The IP address of the client's host machine.
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

Field	Description
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.
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>In Squid- and Netscape-format log files, what do the cache result codes mean?</i> , page 397.

A

Statistics

This appendix describes the following statistics on the Content Gateway Manager Monitor tab:

- ◆ [My Proxy](#), page 191
- ◆ [Protocols](#), page 194
- ◆ [Content Routing \(ICP\)](#), page 197
- ◆ [Security](#), page 198
- ◆ [Subsystems](#), page 200
- ◆ [Networking](#), page 203
- ◆ [Performance](#), page 207
- ◆ [SSL Key Data](#), page 208

My Proxy

My Proxy statistics are divided into the following categories:

- ◆ [Summary](#), page 191
- ◆ [Node](#), page 193
- ◆ [Graphs](#), page 194
- ◆ [Alarms](#), page 194

Summary

Statistic/Field	Description
	Subscription Details
Feature	Lists features purchased, such as SSL Manager, and scanning options. See Working With Encrypted Data, page 111 , for information on SSL Manager, and also Analyze Content with the Scanning Options in the Websense Web Security online help.
Purchased Status	Indicates if a feature has been purchased or not.

Statistic/Field	Description
Expiration Date	If a feature has been purchased, displays the expiration date of the subscription key.
	More Details
Subscription key	Displays the subscription key entered on the Subscription Management tab. See Entering your subscription key , page 11.
Last successful subscription download time	Displays the time of the last successful validation of the subscription key. The check is made once a day.
	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 the data file was downloaded.
Last check	Displays the time and date when Content Gateway last checked for data file updates.
	Node Details
Node	Name of the Content Gateway node or cluster.
On/Off	Indicates if the proxy is running (the proxy and manager services are running).
Objects Served	Total number of objects served by the Content Gateway node.
Ops/Sec	Number of operations per second processed by the Content Gateway node.
Hit Rate	Percentage of HTTP requests served from the cache, averaged over the past 10 seconds.
Throughput (Mbit/sec)	Number of megabits per second passing through the Content Gateway node (and cluster).
HTTP Hit (ms)	Amount of time it takes for an HTTP object that is fresh in the cache to be served to the client.
HTTP Miss (ms)	Amount of time it takes for an HTTP object that is not in the cache or is stale to be served to the client.
	More Detail
cache hit rate	Percentage of HTTP requests served from the cache, averaged over the past 10 seconds. This value is refreshed every 10 seconds.
errors	Percentage of requests that end in early hangsups.
aborts	Percentage of aborted requests.
active clients	Current number of open client connections.

Statistic/Field	Description
active servers	Current number of open origin server connections.
node IP address	IP address assigned to the node. If virtual IP addressing is enabled, several virtual IP addresses could be assigned to this node.
cache free space	Amount of free space in the cache.
HostDB hit rate	Ratio of host database hits to total host database lookups, averaged over a 10-second period.

Node

Statistic	Description
	Node Summary
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.
	Cache
Document Hit Rate	Ratio of cache hits to total cache requests, averaged over 10 seconds. This value is refreshed every 10 seconds.
Bandwidth Savings	Ratio of bytes served from the cache to total requested bytes, averaged over 10 seconds. This value is refreshed every 10 seconds.
Cache Percent Free	Ratio of cache free space to total cache space.
	In Progress
Open Server Connections	Number of currently open origin server connections.
Open Client Connections	Number of currently open client connections.
Cache Transfers in Progress	Number of cache transfers (cache reads and writes) in progress.
	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

The Graphs page displays the same statistics listed on the [Node](#) page (cache performance, 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, page 93](#).

Alarms

Websense 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, page 97](#).

Protocols

Protocol statistics are divided into the following categories:

- ◆ [HTTP, page 194](#)
- ◆ [FTP, page 196](#)

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.

Statistic	Description
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	
Hits	
Fresh	Percentage of hits that are fresh and their average transaction times.
Stale Revalidated	Percentage of hits that are stale and revalidated and turn out to be still fresh and served, and their average transaction times.
Misses	
Now Cached	Percentage of requests for documents that were not in the cache (but are now) and their average transaction times.
Server No Cache	Percentage of requests for HTTP objects that were not in the cache, but have server no-cache headers (cannot be cached); and their average transaction times.
Stale Reloaded	Percentage of misses that are revalidated and turn out to be changed, reloaded, and served, and their average transaction times.
Client No Cache	Percentage of misses with client no-cache headers and their average transaction times.
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.

Statistic	Description
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	
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.
Cache Statistics	
Hits	Number of HTTP requests for FTP objects served from the cache.
Misses	Number of HTTP requests for FTP objects forwarded directly to the origin server because the object is not in the cache or is stale.
Lookups	Number of times Content Gateway looked up an HTTP request for an FTP object in the cache.

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.
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 to the cache since installation.

Statistic	Description
	Operations
File Hit	Number of FTP files served from the cache since installation.
File Misses	Number of FTP file misses since installation.
Change Directory Hits	Number of change directory hits since installation.
Change Directory Misses	Number of change directory misses since installation.
List Directory Hits	The number of list directory hits since installation.
List Directory Misses	The number of list directory misses since installation.

Content Routing (ICP)

Statistic	Description
	Queries Originating from this Node
Query Requests	Number of HTTP requests that generate ICP query messages.
Query Messages Sent	Total number of ICP query messages sent to ICP peers. This number is larger than the number of Query Requests if there are multiple ICP peers.
Peer Hit Messages Received	Number of ICP peer hit messages received in response to ICP queries from this node.
Peer Miss Messages Received	Number of ICP peer miss messages received in response to ICP queries from this node.
Total Responses Received	Number of response messages received from ICP peers (siblings and parents).
Average ICP Message Response Time (ms)	Average time for an ICP peer to respond to an ICP query message from this node. This is a cumulative average value.
Average ICP Request Time	Average time for an HTTP request (that is sent to ICP) to receive an ICP response. This is a cumulative average value.
	Queries Originating from ICP Peers
Query Messages Received	Number of ICP query messages received from remote ICP peers (siblings and parents).
Remote Query Hits	Number of successful cache lookups in response to queries from ICP peers.

Statistic	Description
Remote Query Misses	Number of unsuccessful cache lookups in response to queries from ICP peers.
Successful Response Messages Sent to Peers	Number of successful ICP messages written in response to ICP queries from remote ICP peers.

Security

Security statistics are divided into the following categories:

- ◆ [ARM security, page 198](#)
- ◆ [LDAP, page 198](#)
- ◆ [NTLM, page 199](#)
- ◆ [SOCKS, page 199](#)
- ◆ [Data Security, page 200](#)

ARM security

Statistic	Description
	Security Statistics
TCP Dropped	Number of dropped TCP connections.
UDP Dropped	Number of dropped UDP connections.

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.
	Unsuccessful Authentications
Authorization Denied	Number of times the LDAP Server denied authorization.

Statistic	Description
Authorization Timeouts	Number of times authorization 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.

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.
	Unsuccessful Authentications
Authorization Denied	Number of times the NTLM server denied authorization.
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
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.

Data Security

Statistic	Description
Total Posts	Total number of posts sent to Data Security.
Total Analyzed	Total number of posts analyzed by Data Security.
FTP Analyzed	Total number of FTP requests analyzed by Data Security.
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 Data Security 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 Data Security.

Subsystems

Subsystems statistics are divided into the following categories:

- ◆ [Cache](#), page 201
- ◆ [Clustering](#), page 202
- ◆ [Logging](#), page 202

Cache



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
	General
Bytes Used	Number of bytes currently used by the cache.
Cache Size	Number of bytes allocated to the cache.
	Ram Cache
Bytes	Total size of the RAM cache, in bytes.
Hits	Number of document hits from the RAM cache.
Misses	Number of document misses from the RAM cache. The documents may be hits from the cache disk.
	Reads
In Progress	Number of cache reads in progress (HTTP and FTP).
Hits	Number of cache reads completed since Content Gateway was started (HTTP and FTP).
Misses	Number of cache read misses since Content Gateway was started (HTTP and FTP).
	Writes
In Progress	Number of cache writes in progress (HTTP and FTP).
Successes	Number of successful cache writes since Content Gateway was started (HTTP and FTP).
Failures	Number of failed cache writes since Content Gateway was started (HTTP and FTP).
	Updates
In Progress	Number of HTTP document updates in progress. An update occurs when the Content Gateway revalidates an object, finds it to be fresh, and updates the object header.
Successes	Number of successful cache HTTP updates completed since Content Gateway was started.
Failures	Number of cache HTTP update failures since Content Gateway was started.

Statistic	Description
	Removes
In Progress	Number of document removes in progress. A remove occurs when the Content Gateway revalidates a document, finds it to be deleted on the origin server, and deletes it from the cache (includes HTTP and FTP removes).
Successes	Number of successful cache removes completed since Content Gateway was started (includes HTTP and FTP removes).
Failures	Number of cache remove failures since Content Gateway was started (includes HTTP and FTP removes).

Clustering

Statistic	Description
Bytes Read	Number of bytes read by this node from other nodes in the cluster since Content Gateway was started.
Bytes Written	Number of bytes this node has written to other cluster nodes since Content Gateway was started.
Connections Open	Total number of intra-cluster connections opened since Content Gateway was started.
Total Operations	Total number of cluster transactions since Content Gateway was started.
Network Backups	Number of times this node encountered intra-cluster network congestion and reverted to proxy-only mode since Content Gateway was started.
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.

Statistic	Description
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:

- ◆ [System](#), page 203
- ◆ [ARM](#), page 204
- ◆ [ICAP](#), page 204
- ◆ [WCCP](#), page 205
- ◆ [DNS Resolver](#), page 206
- ◆ [Virtual IP](#), page 206

System

Statistic/Field	Description
	General
Hostname	The hostname assigned to this Content Gateway machine.
Default Gateway	IP address of the default gateway used to forward packets from this Content Gateway machine to other networks or subnets.
Search Domain	Search domain that this Content Gateway machine uses.
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>
Status	Indicates whether the NIC is up or down.
Start on Boot	Indicates whether the NIC is configured to start on boot.
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 Connections Bypassed	Total number of connections bypassed by the ARM.
DNS Packets Bypassed	Number of DNS packets bypassed by the ARM.
	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.

ICAP

Statistic	Description
Total Posts	Total number of posts sent to Data Security.
Total Analyzed	Total number of posts analyzed by Data Security.

Statistic	Description
FTP Analyzed	Total number of FTP requests analyzed by Data Security.
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 Data Security 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 Data Security.

WCCP

WCCP v1 statistics are displayed only if WCCP version 1.0 is enabled. WCCP v2 statistics are displayed only if WCCP version 2.0 is enabled.

Statistic/Field	Description
WCCP v1.0 Statistics	
Router Information	
Router IP address	IP address of the WCCP router sending traffic to Content Gateway.
Router Status	Status of the router: <i>up</i> (if the proxy is able to communicate with the router) or <i>down</i> (if the proxy is unable to communicate with the router).
Node Information	
My IP address	IP address of this Content Gateway node.
Percentage of Traffic Directed to This Node	Percentage of traffic directed to this Content Gateway node.
Number of Heartbeats Received	Number of heartbeats received by this Content Gateway node.
Protocol Information	
Leader's IP Address	IP address of the leader in the WCCP cache farm.
Number of Active Nodes	Number of active Content Gateway nodes in the WCCP cache farm.
WCCP v2.0 Statistics	
WCCP v2.0 Fragmentation	
Total Fragments	Total number of WCCP fragments.
Fragmentation Table Entries	Number of entries in the fragmentation table.

Statistic/Field	Description
Out of Order Fragments	Number of fragments out of order.
Matches	Number of fragments that match a fragment in the fragmentation table.
HTTP/FTP/SOCKS/DNS Service	
Service ID	Service ID for the protocol being serviced.
My 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 to Me	Number of buckets assigned to this Content Gateway node.
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.
Router ID Received	The number of times that Content Gateway has received WCCP protocol messages from the router(s).

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.
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.

Performance

Performance graphs allow you to monitor Websense Content Gateway performance and analyze network traffic. Performance graphs also provide information about virtual memory usage, client connections, document hit rates, hit and miss rates, 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.
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.

SSL

The following tabs are supported by SSL Manager:

[SSL Key Data](#), page 208

[CRL Statistics](#), page 209

[Reports](#), page 209

SSL Key Data

These fields provide information about the status of the SSL connection and activity between the client and SSL Manager and SSL Manager and the destination server.

Statistic/Field	Description
	SSL Inbound Key Data
Is alive	Online indicates that SSL Manager is enabled
Current SSL connections	Number of active inbound (browser to SSL Manager) SSL requests
Total SSL server connections	Number of browser requests
Total finished SSL server connections	Number of browser requests where data went to SSL Manager for decryption
Total SSL server renegotiation requests	Number of browser requests renegotiated due to handshake failures or invalid certificates between the browser and SSL Manager
	SSL Outbound Key Data
Is alive	Online indicates that SSL Manager is enabled
Current SSL connections	Number of active outbound (SSL Manager to designating server) SSL requests
Total SSL client connections	Number of browser requests
Total finished SSL client connections	Number of requests where data went from SSL Manager to the destination server
Total SSL client renegotiation requests	Number of requests were renegotiated due to handshake failures or invalid certificates between SSL Manager and the destination server
Total SSL session cache hits	Number of times that a request was validated by a key in the session cache
Total SSL session cache misses	Number of times that a request could not be validated by a key in the session cache
Total SSL session cache timeouts	Number of time 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
CRL list count	The number of certificates on the Certificate Revocation List. This list is downloaded every night. See Keeping revocation information up to date , page 127.
	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 (CRL & OCSP)

Reports

See [Creating reports with SSL Manager](#), page 99 for information on creating reports on certificate authorities or incidents.

B

Commands and Variables

Websense Content Gateway commands

Use the command line to execute individual commands and when scripting multiple commands in a shell.

To run commands, become root:

```
su
```

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
<code>WCGAdmin start</code>	Starts the Content Gateway service
<code>WCGAdmin stop</code>	Stops the Content Gateway service
<code>WCGAdmin restart</code>	Stops the Content Gateway service and then starts it again
<code>WCGAdmin status</code>	Displays the status (running or not running) of the Content Gateway services: Content Gateway, Content Gateway Manager, and content_cop .
<code>WCGAdmin help</code>	Displays a list of the WCGAdmin commands
<code>content_line -p socket_path</code>	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

Command	Description
<code>content_line -r variable</code>	Displays specific performance statistics or a current configuration setting. For a list of the variables you can specify, see Websense Content Gateway variables, page 212 .
<code>content_line -s variable -v value</code>	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 records.config, page 318 , for a list of the configuration variables you can specify.
<code>content_line -h</code>	Displays the list of Content Gateway commands.
<code>content_line -x</code>	Initiates a Content Gateway configuration file reread. Executing this command is similar to clicking Apply in Content Gateway Manager.
<code>content_line -M</code>	Restarts the content_manager process and the content_gateway process on all the nodes in a cluster.
<code>content_line -L</code>	Restarts the content_manager process and the content_gateway process on the local node.
<code>content_line -S</code>	Shuts down Content Gateway on the local node.
<code>content_line -U</code>	Starts Content Gateway on the local node.
<code>content_line -B</code>	Bounces Content Gateway cluster-wide. Bouncing Content Gateway shuts down and immediately restarts the proxy cache node-by-node.
<code>content_line -b</code>	Bounces Content Gateway on the local node. Bouncing Content Gateway shuts down and immediately restarts the proxy cache on the local node.

Websense 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, page 318](#).

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, page 96](#), and [Command-line interface, page 89](#).

Statistics

The following table lists the variables you can specify on the command line to view individual statistics. See [Statistics, page 191](#) for additional information.

To view a statistic, at the prompt enter:

content_line -r variable

Statistic	Variable
	Summary
Node name	<i>proxy.node.hostname</i>
Objects served	<i>proxy.node.user_agents_total_documents_served</i>
Transactions per second	<i>proxy.node.user_agent_xacts_per_second</i>
	Node
Document hit rate	<i>proxy.node.cache_hit_ratio_avg_10s</i> <i>proxy.cluster.cache_hit_ratio_avg_10s</i>
Bandwidth savings	<i>proxy.node.bandwidth_hit_ratio_avg_10s</i> <i>proxy.cluster.bandwidth_hit_ratio_avg_10s</i>
Cache percent free	<i>proxy.node.cache.percent_free</i> <i>proxy.cluster.cache.percent_free</i>
Open origin server connections	<i>proxy.node.current_server_connections</i> <i>proxy.cluster.current_server_connections</i>
Open client connections	<i>proxy.node.current_client_connections</i> <i>proxy.cluster.current_client_connections</i>
Cache transfers in progress	<i>proxy.node.current_cache_connections</i> <i>proxy.cluster.current_cache_connections</i>
Client throughput (Mbits/sec)	<i>proxy.node.client_throughput_out</i> <i>proxy.cluster.client_throughput_out</i>
Transactions per second	<i>proxy.node.http.user_agent_xacts_per_second</i> <i>proxy.cluster.http.user_agent_xacts_per_second</i>
DNS lookups per second	<i>proxy.node.dns.lookups_per_second</i> <i>proxy.cluster.dns.lookups_per_second</i>
Host database hit rate	<i>proxy.node.hostdb.hit_ratio_avg_10s</i> <i>proxy.cluster.hostdb.hit_ratio_avg_10s</i>
	HTTP
Total document bytes from client	<i>proxy.process.http.user_agent_response_document_total_size</i>
Total header bytes from client	<i>proxy.process.http.user_agent_response_header_total_size</i>
Total connections to client	<i>proxy.process.http.current_client_connections</i>
Client transactions in progress	<i>proxy.process.http.current_client_transactions</i>
Total document bytes from origin server	<i>proxy.process.http.origin_server_response_document_total_size</i>

Statistic	Variable
Total header bytes from origin server	<i>proxy.process.http.origin_server_response_header_total_size</i>
Total connections to origin server	<i>proxy.process.http.current_server_connections</i>
Origin server transactions in progress	<i>proxy.process.http.current_server_transactions</i>
	FTP
Currently open FTP connections	<i>proxy.process.ftp.connections_currently_open</i>
Successful PASV connections	<i>proxy.process.ftp.connections_successful_pasv</i>
Unsuccessful PASV connections	<i>proxy.process.ftp.connections_failed_pasv</i>
Successful PORT connections	<i>proxy.process.ftp.connections_successful_port</i>
Unsuccessful PORT connections	<i>proxy.process.ftp.connections_failed_port</i>
	ICP
ICP query requests originating from this node	<i>proxy.process.icp.icp_query_requests</i>
ICP query messages sent from this node	<i>proxy.process.icp.total_udp_send_queries</i>
ICP peer hit messages received from this node	<i>proxy.process.icp.icp_query_hits</i>
ICP peer miss messages received from this node	<i>proxy.process.icp.icp_query_misses</i>
Total ICP responses received from this node	<i>proxy.process.icp.icp_remote_responses</i>
Average ICP message response time (ms) from this node	<i>proxy.process.icp.total_icp_response_time</i>
Average ICP request time (ms) from this node	<i>proxy.process.icp.total_icp_request_time</i>
Query messages received from ICP peers	<i>proxy.process.icp.icp_remote_query_requests</i>

Statistic	Variable
Remote query hits from ICP peers	<i>proxy.process.icp.cache_lookup_success</i>
Remote query misses from ICP peers	<i>proxy.process.icp.cache_lookup_fail</i>
Successful response messages sent to peers	<i>proxy.process.icp.query_response_write</i>
WCCP	
WCCP router's IP address	<i>proxy.node.wccp.router_ip</i>
WCCP router status	<i>proxy.node.wccp.router_status</i>
WCCP node IP address	<i>proxy.node.wccp.my_ip</i>
Percentage of WCCP traffic received	<i>proxy.node.wccp.my_share</i>
Number of WCCP heartbeats	<i>proxy.node.wccp.hbeats_received</i>
Enabled	<i>proxy.node.wccp.enabled</i>
WCCP leader's IP address	<i>proxy.node.wccp.leader_ip</i>
Number of active WCCP nodes	<i>proxy.node.wccp.number_of_caches_up</i>
Cache	
Bytes used	<i>proxy.process.cache.bytes_used</i>
Cache size	<i>proxy.process.cache.bytes_total</i>
Lookups in progress	<i>proxy.process.cache.lookup.active</i>
Lookups completed	<i>proxy.process.cache.lookup.success</i>
Lookup misses	<i>proxy.process.cache.lookup.failure</i>
Reads in progress	<i>proxy.process.cache.read.active</i>
Reads completed	<i>proxy.process.cache.read.success</i>
Read misses	<i>proxy.process.cache.read.failure</i>
Writes in progress	<i>proxy.process.cache.write.active</i>
Writes completed	<i>proxy.process.cache.write.success</i>
Write failures	<i>proxy.process.cache.write.failure</i>
Updates in progress	<i>proxy.process.cache.update.active</i>
Updates completed	<i>proxy.process.cache.update.success</i>
Update failures	<i>proxy.process.cache.update.failure</i>

Statistic	Variable
Removes in progress	<i>proxy.process.cache.remove.active</i>
Remove successes	<i>proxy.process.cache.remove.success</i>
Remove failures	<i>proxy.process.cache.remove.failure</i>
	Host DB
Total lookups	<i>proxy.process.hostdb.total_lookups</i>
Total hits	<i>proxy.process.hostdb.total_hits</i>
Time TTL (min)	<i>proxy.process.hostdb.ttl</i>
	DNS
DNS total lookups	<i>proxy.process.dns.total_dns_lookups</i>
Average lookup time (ms)	<i>proxy.process.dns.lookup_avg_time</i>
DNS successes	<i>proxy.process.dns.lookup_successes</i>
	Cluster
Bytes read	<i>proxy.process.cluster.read_bytes</i>
Bytes written	<i>proxy.process.cluster.write_bytes</i>
Connections open	<i>proxy.process.cluster.connections_open</i>
Total operations	<i>proxy.process.cluster.connections_opened</i>
Network backups	<i>proxy.process.cluster.net_backup</i>
Clustering nodes	<i>proxy.process.cluster.nodes</i>
	SOCKS
Unsuccessful connections	<i>proxy.process.socks.connections_unsuccessful</i>
Successful connections	<i>proxy.process.socks.connections_successful</i>
Connections in progress	<i>proxy.process.socks.connections_currently_open</i>
	Logging
Currently open log files	<i>proxy.process.log2.log_files_open</i>
Space used for log files	<i>proxy.process.log2.log_files_space_used</i>
Number of access events logged	<i>proxy.process.log2.event_log_access</i>
Number of access events skipped	<i>proxy.process.log2.event_log_access_skip</i>
Number of error events logged	<i>proxy.process.log2.event_log_error</i>

C

Configuration Options

The following configuration options are available on the Content Gateway Manager Configure pane:

[My Proxy](#), page 219

[Protocols](#), page 228

[Content Routing](#), page 241

[Security](#), page 247

[Subsystems](#), page 260

[Networking](#), page 265

My Proxy

The My Proxy configuration options are divided into the following categories:

[Basic](#), page 220

[Subscription](#), page 223

[UI Setup](#), page 223

[Snapshots](#), page 226

[Logs](#), page 227

Basic

When you set these options, ensure that each option is running on a unique port.

Option	Description
	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. In a cluster configuration, the Restart button restarts the proxy and manager services on all the 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). If this node is part of a cluster, this option specifies the name of the Content Gateway cluster (in a Content Gateway cluster, all nodes must share the same name).
Alarm email	Specifies the email address to which Content Gateway sends alarm notifications.
	Features
Protocols: FTP	Enables or disables processing of FTP requests from FTP clients. When this option is enabled, Content Gateway accepts FTP requests from FTP clients. When this option is disabled, Content Gateway does not accept FTP requests from FTP clients. If you change this option, you must restart Content Gateway.
Protocols: HTTPS	Enable or disable processing of HTTPS requests (encrypted data) using SSL Manager. After selecting HTTPS On , you must provide additional information on the Configure > Protocols > HTTPS page and on the Configure > SSL pages. See Working With Encrypted Data , page 111.
Networking: ARM	Enables or disables the ARM. If you change this option, you must restart Content Gateway. See Enabling the ARM option , page 44.
Networking: WCCP	Enables or disables WCCP. Enable this option if you are using a WCCP-enabled router for transparent proxy caching. See Using a WCCP-enabled router , page 46. If you change this option, you must restart Content Gateway.
Networking: Virtual IP	Enables or disables the virtual IP failover option. 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 Virtual IP failover , page 70.

Option	Description
Networking: Data Security	<p>Enables a connection to Websense Data Security. There are 2 options:</p> <ul style="list-style-type: none"> • Registration with a Data Security Management Server for use with a co-located Data Security policy engine (must be version 7.5) • ICAP communication to a remote Data Security Suite deployment (may be version 7.1, or earlier) <p>See Working With Websense Data Security, page 105.</p> <p>If you change this option, you must restart Content Gateway.</p>
Networking: Data Security: Integrated on-box	<p>Enables registration with a Data Security Management Server for use with a co-located Data Security policy engine (must be version 7.5). See Registering and configuring on-box Data Security, page 106.</p>
Networking: Data Security: ICAP	<p>Enables ICAP for use with Data Security Suite. See Configuring the ICAP client, page 109.</p>
Security: SOCKS	<p>Enables or disables the SOCKS option. When SOCKS is enabled, Content Gateway can talk to your SOCKS servers. See Configuring SOCKS firewall integration, page 146.</p> <p>If you change this option, you must restart Content Gateway.</p>
Authentication: None	<p>The proxy will not perform user authentication. This is the default setting.</p>
Authentication: LDAP	<p>Enables or disables LDAP proxy authentication. When LDAP is enabled, you can ensure that users are authenticated by an LDAP server before accessing content from the cache. See Using LDAP proxy authentication, page 152.</p> <p>If you change this option, you must restart Content Gateway.</p>
Authentication: Radius	<p>Enables or disables RADIUS proxy authentication. When RADIUS is enabled, you can ensure that users are authenticated by a RADIUS server before accessing content from the cache. See Using RADIUS proxy authentication, page 155.</p> <p>If you change this option, you must restart Content Gateway.</p>
Authentication: NTLM	<p>Enables or disables NTLM proxy authentication. When NTLM is enabled, you can ensure that users in a Windows network are authenticated by a Domain Controller before accessing content from cache. See Using NTLM proxy authentication, page 158.</p> <p>If you change this option, you must restart Content Gateway.</p>

Option	Description
Authentication: Read authentication from child proxy	<p>Enables or disables the reading of X-Authenticated-User and X-Forwarded-For header values in incoming requests. This option is disabled by default.</p> <p>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.</p>
	Clustering
Cluster: Type	<p>Specifies the clustering mode:</p> <p>Select Single Node to run this Content Gateway server as a single node. This node will not be part of a cluster.</p> <p>Select Management Clustering to run management-only clustering mode. The nodes in the cluster share configuration information and you can administer all the nodes at the same time.</p> <p>Select Full Cache Clustering to run full-clustering mode. In full-clustering mode, as well as sharing configuration information, a Content Gateway cluster distributes its cache across its nodes into a single, virtual object store, rather than replicating the cache node by node.</p> <p>Note: To establish communication between cluster nodes, network routing tables must be configured correctly. When clustering is configured during the installation process, Content Gateway adds the appropriate route to the server's routing table. If clustering is enabled after installation, or clustering settings must change, you must manually update your server's routing table. See Changing clustering configuration, page 67.</p> <p>For complete information about clustering, see Clusters, page 65.</p> <p>If you change this option, you must restart Content Gateway.</p>
Cluster: Interface	<p>Specifies the interface on which Content Gateway communicates with other nodes in the cluster. For example, eth0.</p> <p>It is recommended that you use a dedicated secondary interface for cluster communication.</p> <p><i>Before changing this option, see Changing clustering configuration, page 67.</i></p> <p>If you change this option, you must restart Content Gateway.</p>
Cluster: Multicast Group Address	<p>Specifies the multicast group address on which Content Gateway communicates with its cluster peers.</p> <p><i>Before changing this option, see Changing clustering configuration, page 67.</i></p>

Subscription

Option	Description
	Subscription Management
Subscription Key	Enter the subscription key you received from Websense. This key reflects the products you have purchased.
	Scanning
	Policy Server
IP address	Specify the IP address of the Websense Web filtering Policy Server.
Port	Specify the port used by the Websense Web filtering Policy Server.
	Filtering Service
IP address	Specify the IP address of the Websense Web filtering Filtering Service.
Port	Specify the port used by the Websense Web filtering Filtering Service.
	Action for Communication Errors
Permit traffic	Permits all pages if communication with Policy Server or Filtering Service fails.
Block traffic	Blocks all pages if communication with Policy Server or Filtering Service fails.

UI Setup

Option	Description
	General
UI Port	Specifies the port on which browsers can connect to Content Gateway Manager. The port must be on the Content Gateway system and it must be dedicated to Content Gateway use. The default port is 8081. If you change this setting, you must restart Content Gateway.

Option	Description
SSL UI Port	<p>Specifies for port for the SSL Manager user interface. Through this interface you can specify data decryption and certificate management. The default port is 8071. See Working With Encrypted Data, page 111.</p> <p>The Content Gateway Manager interface and the SSL Manager interface must be on different ports.</p> <p>If you change this setting, you must restart Content Gateway.</p>
HTTPS: Enable/Disable	<p>Enables or disables support for SSL connections to Content Gateway Manager. 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 Using SSL for secure administration, page 143.</p>
HTTPS: Certificate File	<p>Specifies the name of the SSL certificate file used to authenticate users who want to access Content Gateway Manager.</p>
Monitor Refresh Rate	<p>Specifies how often Content Gateway Manager refreshes the statistics on the Monitor pane. The default value is 30 seconds.</p>
Login	
Basic Authentication	<p>Enables or disables basic authentication. When this option is enabled, Content Gateway checks the administrator login and password or the user name and password (if user accounts have been configured) whenever a user tries to access Content Gateway Manager.</p>
Administrator: Login	<p>Specifies the administrator login. The administrator login is the master login that has access to both Configure and Monitor mode in Content Gateway Manager.</p> <p>Note: Content Gateway checks the administrator login only if the Basic Authentication option is enabled.</p>

Option	Description
Administrator: Password	<p>Lets you change the administrator password that controls access to Content Gateway Manager.</p> <p>To change the password, enter the current password in the Old Password field, and then enter the new password in the New Password field. Retype the new password in the New Password (Retype) field, and then click Apply.</p> <p>Note: Content Gateway checks the administrator login and password only if the Basic Authentication option is enabled.</p> <p>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 Content Gateway Manager, Content Gateway updates the records.config file. If you forget the administrator password and cannot access the Content Gateway Manager, see How do you access Content Gateway Manager if you forget the master administrator password?, page 396.</p>
Additional Users	<p>Lists the current user accounts and lets you add new user accounts. User accounts determine who has access 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.</p> <p>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. 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 Creating a list of user accounts, page 142.</p> <p>Note: Content Gateway checks the user login and password only if the Basic Authentication option is enabled.</p>
	<p>Access</p>
Access Control	<p>Displays a table listing the rules in the mgmt_allow.config file that specify the remote hosts allowed to access Content Gateway Manager. The entries in this file ensure that only authenticated users can change configuration options and view performance and network traffic statistics.</p> <p>Note: By default, all remote hosts are allowed to access the Content Gateway Manager.</p>
Refresh	<p>Updates the table to display the most up-to-date rules in the mgmt_allow.config file.</p>

Option	Description
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 Content Gateway Manager. An ip_deny rule denies the remote hosts specified in the Source IP field access to Content Gateway Manager.
Source IP	Specifies the IP addresses that are allowed or denied access to Content Gateway Manager. You can enter a single IP address (111.111.11.1) or a range of IP addresses (0.0.0.0-255.255.255.255).
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Snapshots

Option	Description
	File System
Change Snapshot Directory	Specifies the directory in which snapshots are stored on this Content Gateway node.
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.

Option	Description
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.
	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

Option	Description
	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.

Option	Description
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.
	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:

[HTTP](#), page 229

[HTTP Responses](#), page 237

[HTTP Scheduled Update](#), page 238

[HTTPS](#), page 240

[FTP](#), page 240

HTTP

Option	Description
	General
HTTP Proxy Server Port	<p>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.</p> <p>If you change this option, you must restart Content Gateway.</p>
URL Expandomatic	<p>Enables or disables .com domain expansion. When this option is enabled, Content Gateway attempts to resolve unqualified hostnames by redirecting them to the expanded address, prepended with <code>www.</code> and appended with .com. For example, if a client makes a request to <code>company</code>, Content Gateway redirects the request to www.company.com.</p> <p>If local domain expansion is enabled (see DNS Resolver, page 273), Content Gateway attempts local domain expansion before .com domain expansion; Content Gateway tries .com domain expansion only if local domain expansion fails.</p>
Send HTTP 1.1 by Default	<p>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.</p>
Reverse DNS	<p>Enables reverse DNS lookup when the URL has an IP address (instead of a hostname) and there are rules in filter.config, cache.config, or parent.config. This is necessary when rules are based on destination hostname and domain name.</p>
Tunnel Ports	<p>Specifies the ports to which Content Gateway allows tunneling. 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.</p>

Option	Description
HTTPS Ports	Specifies the ports on which traffic is decrypted and policy is applied when SSL is enabled. When SSL is disabled, traffic to these ports is not decrypted, and filtering policy is applied based on the client hostname.
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	<p>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.</p> <p>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.</p> <p>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.</p> <p>Select PORT only for the FTP server to initiate the data connection and for Content Gateway to accept the connection.</p> <p>The default value is PASV then PORT.</p>
Cacheability	
Caching: HTTP Caching	Enables or disables HTTP caching. When this option is enabled, Content Gateway serves HTTP requests from the cache. When this option is disabled, Content Gateway acts as a proxy server and forwards all HTTP requests directly to the origin server.
Caching: FTP over HTTP Caching	Enables or disables FTP over HTTP caching. When this option is enabled, Content Gateway serves FTP requests from HTTP clients from the cache. When this option is disabled, Content Gateway acts as a proxy server and forwards all FTP requests from HTTP clients directly to the FTP server.

Option	Description
Behavior: Required Headers	<p>Specifies the minimum header information required for an HTTP object to be cacheable.</p> <p>Select An Explicit Lifetime Header to cache only HTTP objects with Expires or max-age headers.</p> <p>Select A Last-Modified Header to cache only HTTP objects with lastmodified headers.</p> <p>Select No Required Headers to cache HTTP objects that do not have <code>Expires</code>, <code>max-age</code>, or <code>last-modified</code> headers. This is the default option.</p> <p>Caution: By default, Content Gateway caches all objects (including objects with no headers). It is recommended that you change the default setting only for specialized proxy situations. If you configure Content Gateway to cache only HTTP objects with Expires or <code>max-age</code> headers, the cache hit rate is reduced (very few objects have explicit expiration information).</p>
Behavior: When to Revalidate	<p>Specifies how Content Gateway evaluates HTTP object freshness in the cache:</p> <p>Select Never Revalidate to never revalidate HTTP objects in the cache with the origin server (Content Gateway considers all HTTP objects in the cache to be fresh).</p> <p>Select Always Revalidate to always revalidate HTTP objects in the cache with the origin server (Content Gateway considers all HTTP objects in the cache to be stale).</p> <p>Select Revalidate if Heuristic Expiration to verify the freshness of an HTTP object with the origin server if the object contains no <code>Expires</code> or <code>Cache-Control</code> headers; Content Gateway considers all HTTP objects without <code>Expires</code> or <code>Cache-Control</code> headers to be stale.</p> <p>Select Use Cache Directive or Heuristic to verify the freshness of an HTTP object with the origin server when Content Gateway considers the object in the cache to be stale according to object headers, absolute freshness limit, and/or rules in the <code>cache.config</code> file. This is the default option.</p> <p>For more information about revalidation, see Revalidating HTTP objects, page 19.</p>

Option	Description
Behavior: Add “no-cache” to MSIE Requests	<p>Specifies when Content Gateway adds <code>no-cache</code> headers to requests from Microsoft Internet Explorer. Certain versions of Microsoft Internet Explorer do not request cache reloads from transparent caches when the user presses the browser Refresh button. This can prevent content from being loaded directly from the origin servers. You can configure Content Gateway to treat Microsoft Internet Explorer requests more conservatively, providing fresher content at the cost of serving fewer documents from cache.</p> <p>Select To All MSIE Requests to always add <code>no-cache</code> headers to all requests from Microsoft Internet Explorer.</p> <p>Select To IMS MSIE Requests to add <code>no-cache</code> headers to IMS (If Modified Since) Microsoft Internet Explorer requests.</p> <p>Select Not to Any MSIE Requests to never add <code>no-cache</code> headers to requests from Microsoft Internet Explorer.</p>
Behavior: Ignore “no-cache” in Client Requests	<p>When this option is enabled, Content Gateway ignores <code>no-cache</code> headers in client requests and serves the requests from the cache.</p> <p>When this option is disabled, Content Gateway does not serve requests with <code>no-cache</code> headers from the cache but forwards them to the origin server.</p>
Freshness: Minimum Heuristic Lifetime	Specifies the minimum amount of time that an HTTP object can be considered fresh in the cache.
Freshness: Maximum Heuristic Lifetime	Specifies the maximum amount of time that an HTTP object can be considered fresh in the cache.
Freshness: FTP Document Lifetime	Specifies the maximum amount of time that an FTP file can stay in the cache. This option affects FTP requests from HTTP clients only.
Maximum Alternates	<p>Specifies the maximum number of alternate versions of HTTP objects Content Gateway can cache.</p> <p>Caution: If you enter 0 (zero), there is no limit to the number of alternates cached. If a popular URL has thousands of alternates, you might observe increased cache hit latencies (transaction times) as Content Gateway searches over the thousands of alternates for each request. In particular, some URLs can have large numbers of alternates due to cookies. If Content Gateway is set to vary on cookies, you might encounter this problem.</p>
Vary Based on Content Type: Enable/ Disable	Enables or disables caching of alternate versions of HTTP documents that do not contain the Vary header. If no Vary header is present, Content Gateway varies on the headers specified below, depending on the document’s content type.
Vary by Default on Text	Specifies the header field on which Content Gateway varies for text documents.

Option	Description
Vary by Default on Images	Specifies the header field on which Content Gateway varies for images.
Vary by Default on Other Document Types	Specifies the header field on which Content Gateway varies for anything other than text and images.
Dynamic Caching: Caching Documents with Dynamic URLs	<p>When this option is enabled, Content Gateway attempts to cache dynamic content. Content is considered dynamic if it contains a question mark (?), a semicolon (;), cgi, or if it ends in .asp.</p> <p>Caution: It is recommended that you configure Content Gateway to cache dynamic content for specialized proxy situations only.</p>
Dynamic Caching: Caching Response to Cookies	<p>Specifies how responses to requests that contain cookies are cached:</p> <p>Select Cache All but Text to cache cookies that contain any type of content except text. This is the default.</p> <p>Select Cache Only Image Types to cache cookies that contain images only.</p> <p>Select Cache Any Content-Type to cache cookies that contain any type of content.</p> <p>Select No Cache on Cookies to not cache cookies at all.</p>
Caching Policy/Forcing Document Caching	Displays a table listing the rules in the cache.config file that specify how a particular group of URLs should be cached. This file also lets you force caching of certain URLs for a specific amount of time.
Refresh	Updates the table to display the most up-to-date rules in the cache.config file. Click Refresh 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 cache.config file.
cache.config Configuration File Editor	
Rule display box	Lists the cache.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 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.

Option	Description
Rule Type	<p>Lists the type of rules you can add to the cache.config file:</p> <p>A never-cache rule configures Content Gateway to never cache specified objects.</p> <p>An ignore-no-cache rule configures Content Gateway to ignore all Cache-Control: no-cache headers.</p> <p>An ignore-client-no-cache rule configures Content Gateway to ignore Cache-Control: no-cache headers from client requests.</p> <p>An ignore-server-no-cache rule configures Content Gateway to ignore Cache-Control: no-cache headers from origin server responses.</p> <p>A pin-in-cache rule configures Content Gateway to keep objects in the cache for a specified time.</p> <p>A revalidate rule configures Content Gateway to consider objects fresh in the cache for a specified time.</p> <p>A ttl-in-cache rule configures Content Gateway to serve certain HTTP objects from the cache for the amount of time specified in the Time Period field regardless of certain caching directives in the HTTP request and response headers.</p>
Primary Destination Type	<p>Lists the primary destination types:</p> <p>dest_domain is a requested domain name.</p> <p>dest_host is a requested hostname.</p> <p>dest_ip is a requested IP address.</p> <p>url_regex is a regular expression to be found in a URL.</p>
Primary Destination Value	<p>Specifies the value of the primary destination type. For example, if the Primary Destination Type is dest_ip, the value for this field can be 123.456.78.9.</p>
Additional Specifier: Time Period	<p>Specifies the amount of time that applies to the revalidate, pin-in-cache, and ttl-in-cache rule types. The following time formats are allowed:</p> <p>d for days (for example 2d)</p> <p>h for hours (for example, 10h)</p> <p>m for minutes (for example, 5m)</p> <p>s for seconds (for example, 20s)</p> <p>mixed units (for example, 1h15m20s)</p>
Secondary Specifiers: Time	<p>Specifies a time range, such as 08:00-14:00.</p>
Secondary Specifiers: Prefix	<p>Specifies a prefix in the path part of a URL.</p>
Secondary Specifiers: Suffix	<p>Specifies a file suffix in the URL.</p>
Secondary Specifiers: Source IP	<p>Specifies the IP address of the client.</p>

Option	Description
Secondary Specifiers: Port	Specifies the port in a requested URL.
Secondary Specifiers: Method	Specifies a request URL method.
Secondary Specifiers: Scheme	Specifies the protocol of a requested URL.
Secondary Specifiers: User-Agent	Specifies a request header User-Agent value.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost.
	Privacy
Insert Headers: Client-IP	When enabled, Content Gateway inserts the Client-IP header into outgoing requests to retain the client's IP address.
Insert Headers: Via	When enabled, Content Gateway inserts a Via header into the outgoing request.
Insert Headers: X-Forwarded-For	When enabled, Content Gateway inserts an X-Forwarded-For header into the outgoing request.
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.
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.

Option	Description
	Timeouts
Keep-Alive Timeouts: Client	<p>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.</p> <p>The client can close the connection at any time.</p>
Keep-Alive Timeouts: Origin Server	<p>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.</p> <p>The origin server can close the connection at any time.</p>
Inactivity Timeouts: Client	<p>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.</p> <p>The client can close the connection at any time.</p>
Inactivity Timeouts: Origin Server	<p>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.</p> <p>The origin server can close the connection at any time.</p>
Active Timeouts: Client	<p>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.</p> <p>The default value of 0 (zero) specifies that there is no timeout.</p> <p>The client can close the connection at any time.</p>

Option	Description
Active Timeouts: Origin Server Request	<p>Specifies how long Content Gateway waits for fulfillment of a connection request to an origin server.</p> <p>If Content Gateway does not establish connection to an origin server before the timeout expires, Content Gateway terminates the connection request.</p> <p>The default value of 0 (zero) specifies that there is no timeout.</p> <p>The origin server can close the connection at any time.</p>
Active Timeouts: Origin Server Response	<p>Specifies how long Content Gateway waits for a response from the origin server.</p>
FTP Control Connection Timeout	<p>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.</p> <p>The default value is 300.</p>

HTTP Responses

Option	Description
	General
Response Suppression Mode	<p>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.</p> <p>Select Always Suppressed if you do not want to send HTTP responses to clients.</p> <p>Select Intercepted Traffic Only if you want to send HTTP responses to nontransparent traffic only. (This option is useful when Content Gateway is running transparently and you do not want to indicate the presence of a cache.)</p> <p>Select Never Suppressed if you want to send HTTP responses to all clients.</p> <p>If you change this option, you must restart Content Gateway.</p>

Option	Description
	Custom
Custom Responses	<p>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.</p> <p>Select Enabled Language-Targeted Response to send your custom responses to clients in the language specified in the <code>Accept-Language</code> header.</p> <p>Select Enabled in “default” Directory Only to send the custom responses located in the default directory to clients.</p> <p>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.</p> <p>If you change this option, you must restart Content Gateway.</p>
Custom Response Logging	<p>When enabled, Content Gateway sends a message to the error log each time custom responses are used or modified.</p> <p>If you change this option, you must restart Content Gateway.</p>
Custom Response Template Directory	<p>Specifies the directory where the custom responses are located. The default location is the Content Gateway config/body_factory directory.</p> <p>If you change this option, you must restart Content Gateway.</p>

HTTP Scheduled Update

Option	Description
	General
Scheduled Update	<p>Enables or disables the scheduled update option. When this option is enabled, Content Gateway can automatically update certain objects in the local cache at a specified time.</p>
Maximum Concurrent Updates	<p>Specifies the maximum number of simultaneous update requests allowed at any point. This option enables you to prevent the scheduled update process from overburdening the host. The default value is 100.</p>
Retry on Update Error: Count	<p>Specifies the number of times Content Gateway retries the scheduled update of a URL in the event of failure. The default value is 10 times.</p>
Retry on Update Error: Interval	<p>Specifies the delay in seconds between each scheduled update retry for a URL in the event of failure. The default value is 2 seconds.</p>

Option	Description
	Update URLs
Force Immediate Update	When enabled, Content Gateway overrides the scheduling expiration time for all scheduled update entries and initiates updates every 25 seconds.
Scheduled Object Update	Displays a table listing the rules in the <i>update.config</i> file that control how Content Gateway performs a scheduled update of specific local cache content.
Refresh	Updates the table to display the most up-to-date rules in the update.config file.
Edit File	Opens the configuration file editor so that you can edit and add rules to the update.config file.
	update.config Configuration File Editor
rule display box	Lists the update.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 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.
URL	Specifies the URL to be updated.
Request Headers (Optional)	Specifies the list of headers (separated by semi-colons) passed in each <code>GET</code> request. You can define any request header that conforms to the HTTP specification. The default is no request header.
Offset Hour	Specifies the base hour used to derive the update periods. The range is 00-23 hours.
Interval	The interval, in seconds, at which updates should occur, starting at Offset Hour.
Recursion Depth	The depth to which referenced URLs are recursively updated, starting at the given URL. For example, a recursion depth of 1 will update the given URL, as well as all URLs immediately referenced by links from the original URL.

HTTPS

Option	Description
	General
HTTPS Proxy Server Port	Specifies the port that Content Gateway uses when acting as a Web proxy server for HTTPS traffic. This is also known as the SSL Inbound Port. When the proxy is serving transparent HTTPS traffic, the value can be 8080 or 8070 (the default). When the proxy is serving explicit HTTPS traffic, set the value should be 8080.
SSL Outbound Port	Specifies the port HTTPS traffic goes to for re-encryption before being sent to its destination. The default is 8090.

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.

Option	Description
	General
FTP Proxy Server Port	Specifies the port that Content Gateway uses to accept FTP requests. The default port is 21.
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.

Option	Description
Shared Server Connections	When enabled, server control connections can be shared between multiple anonymous FTP clients.
	Cacheability
FTP Caching	Enables or disables FTP caching. When this option is enabled, the proxy serves FTP requests from FTP clients from the cache. Note: Only FTP objects retrieved anonymously (user = anonymous) are cached. When this option is disabled, Content Gateway acts as a proxy server and forwards all FTP requests from FTP clients directly to the FTP server.
Directory Caching: Simple	When this option is enabled, the proxy caches directory listings without arguments (for example, dir/ls).
Directory Caching: Full	When this option is enabled, the proxy caches directory listings with arguments (for example, ls -al, ls *.txt).
Freshness: Login Information	Specifies how long (in seconds) the 220/230 responses (login messages) can stay fresh in the cache. The default value is 604800 seconds (7 days).
Freshness: Directory Listings	Specifies how long (in seconds) directory listings can stay fresh in the cache. The default value is 86400 seconds (1 day).
Freshness: Files	Specifies how long (in seconds) FTP files can stay fresh in the cache. The default value is 259200 seconds (3 days).
	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:

[Hierarchies](#), page 242

[Mapping and Redirection, page 245](#)

[Browser Auto-Config, page 247](#)

Hierarchies

Option	Description
	Parenting
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 (either another Content Gateway server or a different caching product) to form a cache hierarchy where a child cache relies upon a parent cache in fulfilling client requests.) See HTTP cache hierarchies, page 73 .
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 does a NDNS resolution as required only when the parent is unavailable.
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 parent.config file that identify the HTTP parent proxies used in an HTTP cache hierarchy and configure selected URL requests to bypass parent proxies.
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 parent.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.

Option	Description
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.

Option	Description
Secondary Specifiers: Port	Specifies the port in a requested URL.
Secondary Specifiers: Method	Specifies a request URL method. For example: <ul style="list-style-type: none"> • <code>get</code> • <code>post</code> • <code>put</code> • <code>trace</code>
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.
ICP Peering	
ICP Mode	Specifies the ICP mode for ICP cache hierarchies: <p>Select Only Receive Queries to configure the proxy to receive ICP queries only.</p> <p>Select Send/Receive Queries to configure the proxy to both send and receive ICP queries.</p> <p>Select Disabled to turn off ICP hierarchical caching.</p> <p>See ICP cache hierarchies, page 75, for more information.</p>
ICP Interface	Specifies the interface used for ICP messages. This field is populated automatically, but if the machine is running with more than one interface, you can specify a different one here.
ICP Port	Specifies the port Content Gateway uses for ICP messages. The default port is 3130.
ICP Multicast	Enables or disables ICP multicast. Select Enabled to send ICP messages through multicast if your proxy server has a multicast channel connection to its peers.
ICP Query Timeout	Specifies the timeout for ICP queries. The default is 2 seconds.
ICP Peers	Displays a table listing the rules in the <i>icp.config</i> file that specify the ICP peers (parent and sibling caches).
Refresh	Updates the table to display the most up-to-date rules in the <i>icp.config</i> 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 <i>icp.config</i> file.
<i>icp.config</i> Configuration File Editor	
rule display box	Lists the <i>icp.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.

Option	Description
Set	Updates the rule display box at the top of the configuration file editor page.
Peer Hostname	Specifies the hostname of the ICP Peer. This field is optional if the IP address of the ICP peer is specified in the Peer IP field below.
Peer IP	Specifies the IP address of the ICP Peer. This field is optional if the hostname of the ICP peer is specified in the Peer Hostname field above.
Peer Type	Specifies the type of ICP peer: parent or sibling.
Proxy Port	Specifies the port number of the TCP port used by the ICP peer for proxy communication.
ICP Port	Specifies the port number of the UDP port used by the ICP peer for ICP communication.
Multicast	Enables or disables multicast mode.
Multicast IP	Specifies the multicast IP address if multicast is enabled.
Multicast TTL	Specifies the multicast time to live. Select single subnet if IP multicast datagrams will not be forwarded beyond a single subnetwork. Select multiple subnets to allow delivery of IP multicast datagrams to more than one subnet (if there are one or more multicast routers attached to the first hop subnet)

Mapping and Redirection

Option	Description
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 <code>Host:</code> 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 <code>Host:</code> headers to a particular server.

Option	Description
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.
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 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	Lists the type of rules you can add to the <code>remap.config</code> file: 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).
From Scheme	Specifies the protocol of the mapping rule. Note: Mapping a URL of one protocol to a different protocol is not supported.
From Host	Specifies the hostname of the URL to map from.
From Port (Optional)	Specifies the port number in the URL to map from.
From Path Prefix (Optional)	Specifies the path prefix of the URL to map from.
To Host	Specifies 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.
{undefined}	Specifies the media protocol type of the mapping rule.

Browser Auto-Config

Option	Description
	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 Using a PAC file, page 36 .
	WPAD
WPAD Settings	Lets you edit the wpad.dat file. See Using WPAD, page 37 .

Security

The Security configuration options are divided into the following categories:

[Connection Control, page 247](#)

[Access Control, page 251](#)

[SOCKS, page 258](#)

Connection Control

Option	Description
	Proxy Access
Access Control	Displays the rules in the ip_allow.config file that control which clients can access the proxy cache. By default, all remote hosts are allowed to access the proxy cache.
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 ip_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.

Option	Description
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 cache. An ip_deny rule denies the clients listed in the Source IP field access to the proxy cache.
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.
ARM Security	
ARM Security	Enables or disables the ARM security option that restricts the type of communication possible with machines running Content Gateway. For information about the ARM security option, see Controlling host access to the proxy server , page 140. If you change this option, you must restart Content Gateway.
Access Control List	Displays a table listing the rules in the <i>arm_security.config</i> file that restrict the type of communication possible with machines running Content Gateway. For more details about the ARM security feature, see Controlling host access to the proxy server , page 140.
Refresh	Updates the table to display the most up-to-date rules in the arm_security.config file.
Edit File	Opens the configuration file editor for the arm_security.config file.
arm_security.config Configuration File Editor	
rule display box	Lists the <i>arm_security.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.

Option	Description
Rule Type	<p>Specifies the rule type. See the Note below for application and precedence rules.</p> <p>Rule Types:</p> <ul style="list-style-type: none"> • open: An open rule specifies a port or set of ports to be open by default, for either TCP or UDP. Specify the ports you want to open in the Open Port field. • allow: An allow rule specifies the hosts that are allowed access to specific destination ports, for either TCP or UDP. • deny: A deny rule specifies the hosts that are denied access to specific destination ports, for either TCP or UDP. <p>Note: The governance rules for ARM rules are as follows:</p> <ol style="list-style-type: none"> 1. When ARM security is enabled, access to ports not listed as open or allow is denied. 2. For a deny rule to have meaning, the specified port must already have been opened with an open rule. Deny rules that do not have corresponding, preceding open rules are ignored. 3. Conflicts between allow and deny rules (where there is an allow/open/deny or open/allow/deny or open/deny/allow referencing the same IP/port combination) are won by the deny rule regardless of ordering in the file.
Connection Type	Specifies the type of connection used: TCP or UDP .
Source IP	Specifies the IP address or range of IP addresses of the source of the communication.
Source Port	Specifies the source port or range of source ports from which TCP traffic is allowed.
Destination IP	Specifies the IP address or range of IP addresses of the destination of the communication.
Destination Port	<p>Specifies the destination port or range of destination ports from which TCP traffic is allowed or denied.</p> <p>Note: Traffic to destination port 80 cannot be denied by an ARM security rule.</p>
Open Port	Specifies the port or series of ports that are open by default.
Apply	Applies the configuration changes.
Close	<p>Exits the configuration file editor.</p> <p>Click Apply before you click Close; otherwise, all configuration changes will be lost.</p>

Data Security



Note

The Data Security configuration options appear on the Configure pane only if you have enabled Data Security and selected **Integrated on-box** in the **Features** table on the **Configure > My Proxy > Basic > General** tab.

Option	Description
Data Security IP address	Specifies the IP address of the Data Security Management Server. This is where the Websense Data Security policy configuration and management is performed.
Analyze HTTPS Content	Select whether decrypted traffic should be sent to Websense Data Security for analysis or sent directly to the destination.
Analyze FTP Uploads	Select whether to send FTP upload requests to Websense Data Security for analysis. The FTP proxy feature must be enabled. See FTP , page 240.

Registration screen fields:

Option	Description
Data Security IP address	Specifies the IP address of the Data Security Management Server. This is where data security policy configuration and management is performed.
Data Security Manager user name	Specifies the account name of a Websense Data Security administrator. The administrator must have Deploy Settings privileges.
Data Security Manager user name	Specifies the password of the Websense Data Security administrator.
Register button	Initiate the registration action. This button is enabled only after data is entered in all fields.

Access Control

Option	Description
	Filtering
Filtering	Displays a table listing the rules in the <i>filter.config</i> file that deny or allow particular URL requests, and that keep or strip header information from client requests. 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. NTLM and LDAP authentication rules are defined on the Authentication Realms tab and stored in the <i>auth.config</i> file (see its entry later in this table).
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; the proxy caches and serves the requested content. 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. Note: 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 hostname. dest_ip is a requested IP address. url_regex is a regular expression to be found in a URL.

Option	Description
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 can 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.
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: <ul style="list-style-type: none"> – get – post – put – trace
Secondary Specifiers: Scheme	Specifies the protocol of a requested URL. Options are: <ul style="list-style-type: none"> – HTTP – HTTPS – FTP (for FTP over HTTP only) <p>Note: Although rtsp and mms are offered in the list, they are not supported.</p>
Secondary Specifiers: User-Agent	Specifies the request header User-Agent value. Use this field to create application filtering rules that: <ul style="list-style-type: none"> • 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.

Option	Description
	Transparent Proxy Authentication
Redirect Hostname	Specifies an alternate hostname for the proxy that can be resolved by DNS for all clients on the network. For additional information, see Transparent proxy authentication, page 151 .
Authentication Mode	When transparent proxy authentication is configured, Content Gateway must be set to an authentication mode: <ul style="list-style-type: none"> • IP mode (the default) causes the client IP address to be associated with a username when a session is authenticated. Requests made from that IP address are not authenticated again until the Session TTL expires. The default is 15 minutes. • Cookie mode is used to uniquely identify users who share a single IP address, such as, for example, in environments where proxy-chaining is used or where network address translation (NAT) occurs.
Session TTL	Specifies the length of time, in minutes, before the client must re-authenticate. This is required for both IP and Cookie modes. The default is 15 minutes. The supported range of values is 5-65535 minutes.
	LDAP

Note: 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 on configuring LDAP see [Using LDAP proxy authentication, page 152](#).

Purge Cache on Authentication Failure	When this option is enabled, Content Gateway deletes the authorization entry for the client in the LDAP cache if authorization fails.
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. 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 Active Directory or other directory services.

Option	Description
LDAP Server: Bind Distinguished Name	<p>Specifies the Full Distinguished Name (fully qualified name) of a user in the LDAP-based directory service. For example:</p> <pre>CN=John Smith,CN=USERS,DC=MYCOMPANY,DC=COM</pre> <p>Enter a maximum of 128 characters in this field.</p> <p>If you do not specify a value for this field, the proxy attempts to bind anonymously.</p>
LDAP Server: Password	<p>Specifies a password for the user identified in the Bind_DN field.</p>
LDAP Server: Base Distinguished Name	<p>Specifies the base Distinguished Name (DN). You can obtain this value from your LDAP administrator.</p> <p>You must specify a correct base DN; otherwise LDAP authentication will fail to operate.</p> <p>If you change this option, you must restart Content Gateway.</p>
Radius	
<p>Note: 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.</p> <p>For more information on configuring Radius, see Using RADIUS proxy authentication, page 155.</p>	
Primary Radius Server: Hostname	<p>Specifies the hostname or IP address of the primary RADIUS authentication server.</p> <p>If you change this option, you must restart Content Gateway.</p>
Primary Radius Server: Port	<p>Specifies the port that Content Gateway uses to communicate with the primary RADIUS authentication server. The default port is 1812.</p> <p>If you change this option, you must restart Content Gateway.</p>
Primary Radius Server: Shared Key	<p>Specifies the key to use for encoding.</p> <p>If you change this option, you must restart Content Gateway.</p>
Secondary Radius Server (Optional): Hostname	<p>Specifies the hostname or IP address of the secondary RADIUS authentication server.</p> <p>If you change this option, you must restart Content Gateway.</p>
Secondary Radius Server (Optional): Port	<p>Specifies the port that Content Gateway uses to communicate with the secondary RADIUS authentication server. The default port is 1812.</p> <p>If you change this option, you must restart Content Gateway.</p>
Secondary Radius Server (Optional): Shared Key	<p>Specifies the key to use for encoding.</p> <p>If you change this option, you must restart Content Gateway.</p>

Option	Description
	NTLM
	<p>Note: 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.</p> <p>For more information on configuring NTLM, see Using NTLM proxy authentication, page 158.</p>
<p>Domain Controller Hostnames</p>	<p>Specifies the hostnames of the domain controllers in a comma separated list. The format is:</p> <pre>host_name[:port][%netbios_name]</pre> <p>or</p> <pre>IP_address[:port][%netbios_name]</pre> <p>If you are using Active Directory 2008, you must include the <code>netbios_name</code> or use SMB port 445.</p> <p>If you change this option, you must restart Content Gateway.</p>
<p>Load Balancing</p>	<p>Enables or disables load balancing. When enabled, Content Gateway balances the load when sending authentication requests to the domain controllers.</p> <p>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.</p> <p>If you change this option, you must restart Content Gateway.</p>
<p>Fail Open</p>	<p>When enabled (default), allows client requests to proceed when authentication fails due to:</p> <ul style="list-style-type: none"> • no response from the domain controller • malformed messages from the client • invalid SMB responses <p>Note: Password authentication failures are always failures.</p>
<p>IP Credentials: Credential caching</p>	<p>Enables or disables NTLM credential caching.</p>
<p>IP Credentials: Caching TTL</p>	<p>Specifies the time-to-live, in seconds, for NTLM cached Credentials. The default is 900 seconds (15 minutes). The range of supported values is 300 to 86400 seconds.</p>
<p>IP Credentials: Multi-user IP Exclusions</p>	<p>Specifies a comma separated list of multi-user IP addresses and IP address ranges for terminal servers, NAT firewalls, etc.</p> <p>Credentials for these users are not cached.</p>
	Authentication Realms

Option	Description
	<p>In networks with multiple authentication realms, rules can be defined to direct sets of IP addresses to distinct authentication servers (domain controllers).</p> <p>A default authentication realm must be defined on the Configure > Security > Access Control > LDAP or NTLM page.</p> <p>For more information, see Multiple authentication realms, page 161.</p>
Authentication	Displays a table listing the rules in the <i>auth.config</i> file that direct specified IP addresses to particular domain controllers for authentication. Multiple LDAP or NTLM authentication servers can be configured.
Refresh	Updates the table to display the current rules in the auth.config file.
Edit File	Opens the configuration file editor for the auth.config file.
	auth.config Configuration File Editor
rule display box	Lists the <i>auth.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: Select NTLM to specify rules for NTLM servers. Select LDAP to specify rules for LDAP servers.
Profile Name	Specifies a descriptive name for the profile (must be unique).
Source IP	Specifies sets of single IP addresses or IP ranges to match for this profile (should 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
Enabled	Specifies if this rule is active. To leave the rule in the list, but not active, leave the Enabled check box empty.
LDAP Specifiers: LDAP Server Name	Specifies the LDAP server name. This option applies to <code>ldap</code> rule types only.
LDAP Specifiers: LDAP Server Port	Specifies the LDAP Server Port (Optional - Default 389)
LDAP Specifiers: LDAP Base Distinguished Name	Specifies the LDAP Base Distinguished Name. This option applies to <code>ldap</code> rule types only.
LDAP Specifiers: Server Type	Sets the search filter to "sAMAccountName" for Active Directory, or to "uid" for other directory services.
LDAP Specifiers: Bind DN	Specifies the LDAP bind account distinguished name.

Option	Description
LDAP Specifiers: Bind Password	Specifies the LDAP bind account password.
LDAP Specifiers: Secure LDAP	Specifies whether Content Gateway will use secure communication with the LDAP server. If enabled, you must set LDAP port to one of the secure ports: 636 or 3269.
LDAP Specifiers: LDAP Attribute Name (Optional)	Specifies the LDAP attribute name.
LDAP Specifiers: LDAP Attribute Value (Optional)	Specifies the LDAP attribute pair.
NTLM Specifiers: DC List	Specifies 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 Specifiers: DC Load Balance	Specifies whether load balancing is used: <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled <p>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.</p>
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost.

SOCKS



Note

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.

Option	Description
	General
SOCKS Version	<p>Specifies the version of SOCKS used on your SOCKS server. Content Gateway supports SOCKS version 4 and version 5.</p> <p>If you change this option, you must restart Content Gateway.</p>
	Proxy
SOCKS Proxy	<p>Enables or disables the SOCKS Proxy option. As a SOCKS proxy, WCG can receive SOCKS packets (usually on port 1080) from the client and forwards all requests directly to the SOCKS server.</p> <p>For more information about the SOCKS Proxy option, see Configuring SOCKS firewall integration, page 146.</p> <p>If you change this option, you must restart Content Gateway.</p>
SOCKS Proxy Port	<p>Specifies the port on which Content Gateway accepts SOCKS traffic. This is usually port 1080.</p> <p>If you change this option, you must restart Content Gateway.</p>
	Server
SOCKS Server: Default Servers	<p>Specifies the names and the ports of the default SOCKS servers with which Content Gateway communicates. Each entry must be separated by a semicolon (;); for example: socks1:1080;socks2:4080</p> <p>If you change this option, you must restart Content Gateway.</p> <p>You can perform additional SOCKS server configuration in the socks.config file.</p>
Socks Server Rules	<p>Displays a table listing the rules in the <code>socks.config</code> file that control the SOCKS servers 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 the SOCKS server, and the user name and password used by Content Gateway to connect to a SOCKS version 5 server.</p>

Option	Description
Refresh	Updates the table to display the most up-to-date 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	<p>Specifies the rule type:</p> <p>Select no_socks to specify the origin servers you want the proxy to access directly without going through the SOCKS server. Enter the IP addresses of the origin servers in the Destination IP field.</p> <p>Select auth to specify the user name and password Content Gateway uses for authentication with a SOCKS version 5 server. Enter the username in the Username field and the password in the Password field.</p> <p>Select multiple_socks to specify the SOCKS servers Content Gateway must go through to reach specific origin servers. Enter the hostnames or the IP addresses of the SOCKS servers in the SOCKS Servers field and the IP addresses of the origin servers in the Destination IP field. Select how strictly Content Gateway should follow round robin from the Round Robin drop-down list.</p>
Username	<p>Specifies the username Content Gateway must use to authenticate with a SOCKS version 5 server.</p> <p>This field applies to an <code>auth</code> rule type only.</p>
Password	<p>Specifies the password Content Gateway must use to authenticate with a SOCKS version 5 server.</p> <p>This field applies to an <code>auth</code> rule type only.</p>
Destination IP	<p>For a <code>multiple_socks</code> rule, specify either a single IP address <i>or</i> a range of IP addresses of the origin servers with which Content Gateway must use the SOCKS servers specified in the SOCKS Servers field below.</p> <p>For a <code>no_socks</code> rule, 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: <code>255.255.255.255</code>.</p>
SOCKS Servers	<p>Specify the hostnames or the IP addresses of the SOCKS servers you want to go through to service requests for the origin servers listed in the Destination IP field above. Separate each entry with a semicolon.</p> <p>This field applies to a <code>multiple_socks</code> rule type only.</p>

Option	Description
Round Robin	Specifies how strictly Content Gateway should follow round robin. You can select strict , or false . This field applies to a <code>multiple_socks</code> rule type only.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Subsystems

The Subsystems configuration options are divided into the following categories:

[Cache](#), page 260

[Logging](#), page 261

Cache

Option	Description
	General
Allow Pinning	Enables or disables the cache pinning option, which lets you keep objects in the cache for a specified time. Set cache pinning rules in the <code>cache.config</code> file.
Ram Cache Size	Specifies the size of the RAM cache, in bytes. A value of -1, the default, directs Content Gateway to automatically size the RAM cache to be approximately 1 MB per 1 GB of disk cache. 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.
	Hosting
Cache Hosting	Displays a table listing the rules in the <code>hosting.config</code> file that controls which cache partitions are assigned to specific origin servers and domains.
Refresh	Updates the table to display the most up-to-date rules in the <code>hosting.config</code> file.
Edit File	Opens the configuration file editor for the <code>hosting.config</code> file. The configuration file editor page is described below.

Option	Description
	hosting.config Configuration File Editor
rule display box	Lists the hosting.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 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	Specifies the primary destination rule type: Select domain if you want to partition the cache according to domain. Select hostname if you want to partition the cache according to hostname
Primary Destination Value	Specifies the domain or origin server's hostname whose content you want to store on a particular partition.
Partitions	Specifies a comma-separated list of the partitions on which you want to store the content that belongs to the origin server or domain specified.
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost.

Logging

Option	Description
	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

Option	Description
Log Space: Limit	<p>Specifies the maximum amount of space (in megabytes) allocated to the logging directory for the log files.</p> <p>Note: Transaction logs can consume a lot of space. You should set this limit high enough to accommodate at least a single day's worth of uncompressed transaction logs. Also, make sure that this limit is smaller than the actual space available on the partition that contains the logging directory.</p>
Log Space: Headroom	<p>Specifies the tolerance for the log space limit. If the Auto-Delete Rolled Files option is enabled, autodeletion is triggered when the amount of free space available in the logging directory is less than the headroom.</p>
Log Rolling: Enable/Disable	<p>Enables or disables log file rolling. To keep log files down to manageable sizes, you can roll them at regular intervals. See Rolling event log files, page 176.</p>
Log Rolling: Offset Hour	<p>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.</p>
Log Rolling: Interval	<p>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 86400 seconds (1 day).</p>
Log Rolling: Auto-Delete Rolled Files	<p>Enables autodeletion of rolled log files when available space in the log directory is low. Autodeletion is triggered when the amount of free space available in the log directory is less than the Log Space Headroom.</p>
Formats	
Squid Format: Enable/Disable	<p>Enables or disables the Squid log format.</p>
Squid Format: ASCII/Binary	<p>Select ASCII or Binary as the type of log files to be created.</p>
Squid Format: Filename	<p>Specifies the name used for Squid log files. The default filename is squid.log.</p>
Squid Format: Header	<p>Specifies the text header you want Squid log files to contain.</p>
Netscape Common Format: Enable/Disable	<p>Enables or disables the Netscape Common log format.</p>
Netscape Common Format: ASCII/Binary	<p>Select ASCII or Binary as the type of log file to be created.</p>
Netscape Common Format: Filename	<p>Specifies the name used for Netscape Common log files. The default filename is common.log.</p>
Netscape Common Format: Header	<p>Specifies the text header you want Netscape Common log files to contain.</p>

Option	Description
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.
	Splitting
Split ICP Logs	When this option is enabled, Content Gateway records ICP transactions in a separate log file. When this option is disabled, Content Gateway records ICP transactions in the same log file with HTTP and FTP entries.
Split NNTP Logs	When this option is enabled, Content Gateway records all commands it sends to the Network News Transfer Protocol (NNTP) server in a separate log file. This allows you to log transactions from incoming news feeds. When this option is disabled, Content Gateway records NNTP transactions in the same log file with HTTP and FTP entries.
Split Host Logs	When this option is enabled, Content Gateway creates a separate log file for each of the hosts listed in the log_hosts.config file. When this option is disabled, Content Gateway records transactions for all hosts in the same log file.

Option	Description
	Collation
Collation Mode	<p>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 Collating event log files, page 182.</p> <p>Select Collation Disabled to disable log collation on this Content Gateway node.</p> <p>Select Be a Collation Server to configure this Content Gateway node to be the collation server.</p> <p>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.</p> <p>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).</p> <p>Log collation consumes cluster bandwidth in sending all log entries to a single node. It can therefore affect the performance of the cluster.</p> <p>If you want Content Gateway as a collation client to send custom (XML-based) log files, you must specify a <code>LogObject</code> in the logs_xml.config file.</p>
Log Collation Server	Specifies the hostname of the log collation server to which you want to send log files.
Log Collation Port	<p>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.</p> <p>Note: Do not change the port number unless there is a conflict with another service already using the port.</p>
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.

Option	Description
	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.

Networking

The Networking configuration options are divided into the following categories:

[Connection Management](#), page 265

[ARM](#), page 267

[WCCP](#), page 271

[DNS Resolver](#), page 273

[ICAP](#), page 275

[Virtual IP](#), page 276

Connection Management

Option	Description
	Throttling
Throttling Net Connections	Specifies the maximum number of network connections that Content Gateway accepts. 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.
	Load Shedding
Maximum Connections	Specifies the maximum number of client connections allowed before the ARM (transparent mode) 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.

Option	Description
Enabled/Disabled	<p data-bbox="685 247 916 279">Congestion Control</p> <p data-bbox="685 296 1235 380">The congestion control feature tracks which hosts are congested so Content Gateway will not forward requests to those hosts.</p> <p data-bbox="685 390 1235 474">When this option is enabled, rather than forwarding requests to a congested destination host, the proxy instructs the requestor to try again at a later time.</p> <p data-bbox="685 485 1235 569">When this option is disabled, requests are forwarded to the destination host regardless of host congestion.</p> <p data-bbox="685 579 1235 636">Congestion is determined by rules you set up in the congestion.config file on this screen.</p>
Congestion Rules	<p data-bbox="685 657 1235 772">Displays the congestion.config file so that you can define specific congestion configuration rules. Rules are defined by primary and secondary rule pairs.</p> <p data-bbox="685 783 1235 840">By default, the following Congestion Control Rule tags are defined:</p> <pre data-bbox="685 850 1016 1228"># max_connection_failures=5 # fail_window=120 # proxy_retry_interval=10 # client_wait_interval=300 # wait_interval_alpha=30 # live_os_conn_timeout=60 # live_os_conn_retries=2 # dead_os_conn_timeout=15 # dead_os_conn_retries=1 # max_connection=-1 # error_page="congestion#retryAfter" # congestion_scheme="per_ip" # snmp=on</pre> <p data-bbox="685 1239 1235 1354">Every line must have a primary destination specifier. Secondary specifiers are optional. Each line can have one or more rule pairs. Following is an example of a line in the congestion.config file:</p> <pre data-bbox="685 1365 1170 1396">dest_host=gumby port=80 fail_window=100</pre> <p data-bbox="685 1407 1235 1423">Press Submit when you are done editing the file.</p> <p data-bbox="685 1434 1235 1549">For a description of the Destination Specifiers, see Destination specifier tags, page 382. For a description of the Congestion Control Rule tags, see Congestion control rule variables, page 383.</p>

ARM



Note

The ARM configuration options appear on the Configure pane only if you have enabled ARM in the Features table on the **Configure > My Proxy > Basic > General** tab.

Option	Description
	General
IP spoofing	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 IP spoofing, page 62 .
Network Address Translation (NAT)	Displays the redirection rules in the <i>ipnat.conf</i> file that specify how incoming packets are readdressed when the proxy is serving traffic transparently. Content Gateway creates redirection rules during installation. You can modify these rules.
Refresh	Updates the table to display the most up-to-date rules in the <i>ipnat.conf</i> file.
Edit File	Opens the configuration file editor for the <i>ipnat.conf</i> 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, <code>eth0</code> on Linux.
Connection Type	Specifies the connection type that applies for the rule: TCP or UDP.
Source IP	Specifies the IP address from which traffic is sent.
Source CIDR	Specifies the IP address in CIDR (Classless Inter-Domain Routing) format, such as <code>1.1.1.0/24</code> . Entering a value in this field is optional.
Source Port	Specifies the traffic destination port: for example, <code>80</code> for HTTP traffic.
Destination IP	Specifies the IP address of your Content Gateway server.

Option	Description
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 will be lost.
	Static Bypass
Static Bypass	Displays a table listing the rules in the <i>bypass.config</i> file that specify static transparency bypass rules. When transparency is enabled, the proxy uses these rules to determine whether to bypass incoming client requests or attempt to serve them transparently.
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

Option	Description
Destination IP	<p>Specifies the destination IP address in incoming requests that the proxy must bypass or deny bypass. The IP address can be one of the following:</p> <ul style="list-style-type: none"> 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	<p>Exits the configuration file editor.</p> <p>Click Apply before you click Close; otherwise, all configuration changes will be lost.</p>
	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	<p>Select Enabled to enable dynamic bypass when Content Gateway encounters non-HTTP traffic on port 80.</p> <p>Select Disabled to disable dynamic bypass when Content Gateway encounters non-HTTP traffic on port 80.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when Content Gateway encounters non-HTTP traffic on port 80.</p> <p>Select Destination Only to enable dynamic destination bypass when Content Gateway encounters non-HTTP traffic on port 80.</p>
Behavior: HTTP 400	<p>Select Enabled to enable dynamic bypass when an origin server returns a 400 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 400 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 400 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 400 error.</p>
Behavior: HTTP 401	<p>Select Enabled to enable dynamic bypass when an origin server returns a 401 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 401 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 401 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 401 error.</p>

Option	Description
Behavior: HTTP 403	<p>Select Enabled to enable dynamic bypass when an origin server returns a 403 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 403 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 403 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 403 error.</p>
Behavior: HTTP 405	<p>Select Enabled to enable dynamic bypass when an origin server returns a 405 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 405 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 405 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 405 error.</p>
Behavior: HTTP 406	<p>Select Enabled to enable dynamic bypass when an origin server returns a 406 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 406 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 406 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 406 error.</p>
Behavior: HTTP 408	<p>Select Enabled to enable dynamic bypass when an origin server returns a 408 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 408 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 408 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 408 error.</p>
Behavior: HTTP 500	<p>Select Enabled to enable dynamic bypass when an origin server returns a 500 error.</p> <p>Select Disabled to disable dynamic bypass when an origin server returns a 500 error.</p> <p>Select Source-Destination to enable dynamic source/destination bypass when an origin server returns a 500 error.</p> <p>Select Destination Only to enable dynamic destination bypass when an origin server returns a 500 error.</p>

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.

Option	Description
	General
WCCP Version	Specifies the version of WCCP running on the router: WCCP v1 or WCCP v2. If you change this option, you must restart Content Gateway.
	WCCP v1.0
WCCP Network Interface	Specifies the Ethernet interface Content Gateway uses to talk to the WCCP v1-enabled router. If you change this option, you must restart Content Gateway.
WCCP Router IP Address	Specifies the IP address of the WCCP-enabled router sending traffic to Content Gateway. If you change this option, you must restart Content Gateway.
	WCCP v2.0
WCCP Network Interface	Specifies the Ethernet interface Content Gateway uses to talk to the WCCP v2-enabled router. If you change this option, you must restart Content Gateway.
Security: Enable/Disable	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 Cisco router documentation.) If you change this option, you must restart Content Gateway.
Security: Password	Specifies the password used for authentication. This password must be the same password configured on the router and can be a maximum of eight characters long. If you change this option, you must restart Content Gateway.
Multicast: Enable/Disable	Enables or disables WCCP multicast mode. If you change this option, you must restart Content Gateway.

Option	Description
Multicast: IP Address	<p>Specifies the multicast IP address.</p> <p>If you change this option, you must restart Content Gateway.</p>
Routers: Number of WCCP Routers	<p>If multicast is not enabled, the routers on your network are not automatically discovered. You must specify the number of routers that direct traffic to Content Gateway. WCCP v2 supports a maximum of 32 routers.</p> <p>If you change this option, you must restart Content Gateway.</p>
Routers: IP Address of 1st Router	<p>Specifies the IP address of the first WCCP v2-enabled router.</p> <p>If you change this option, you must restart Content Gateway.</p>
Routers: IP Address of 2nd Router	<p>Specifies the IP address of the second WCCP v2-enabled router.</p> <p>If you change this option, you must restart Content Gateway.</p>
Routers: IP Address of 3rd Router	<p>Specifies the IP address of the third WCCP v2-enabled router.</p> <p>If you change this option, you must restart Content Gateway.</p>
Services: HTTP	<p>When enabled, a WCCP v2-enabled router can send HTTP traffic to Content Gateway.</p> <p>If you change this option, you must restart Content Gateway.</p>
Services: HTTPS	<p>When enabled, a WCCP v2-enabled router can send secure HTTP traffic to Content Gateway. If you change this option, you must restart Content Gateway.</p>
Services: NNTP	<p>When enabled, a WCCP v2-enabled router can send Network News Transfer Protocol (NNTP) traffic to Content Gateway. If you change this option, you must restart Content Gateway.</p>
Services: FTP	<p>When enabled, a WCCP v2-enabled router can send FTP traffic to Content Gateway.</p> <p>If you change this option, you must restart Content Gateway.</p>
Services: SOCKS	<p>When enabled, a WCCP v2-enabled router can send SOCKS traffic to Content Gateway.</p> <p>If you change this option, you must restart Content Gateway.</p>
Services: DNS	<p>When enabled, a WCCP v2-enabled router can send DNS traffic to Content Gateway.</p> <p>If you change this option, you must restart Content Gateway.</p>

Option	Description
Mode Negotiation: Packet Forward Method	Specifies the encapsulation method used by the WCCP server to transmit intercepted traffic to the proxy. Note: The proxy always adjusts to use a method that the router supports. If the router supports more than one method, the method specified in this area is used.
Mode Negotiation: Assignment Method	The method used by the router to distribute intercepted traffic across multiple proxies.
Mode Negotiation: Packet Return Method	The encapsulation method used to return intercepted traffic to the WCCP server.

DNS Resolver

Option	Description
	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 <code>hostx</code> , and if the WCG local domain is <code>y.com</code> , Content Gateway expands the hostname to <code>hostx.y.com</code> .
	Host Database
DNS Lookup Timeout	Specifies the maximum number of seconds the proxy can wait for a lookup response from the Domain Name Server.
Foreground Timeout	Specifies how long DNS entries can remain in the database before they are flagged as stale. For example, if this timeout is 24 hours and a client requests an entry that has been in the database for 24 hours or longer, the proxy refreshes the entry before serving it. Caution: Setting the foreground timeout too low might slow response time. Setting it too high risks accumulation of incorrect information.
	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 Using the Split DNS option, page 149 .

Option	Description
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 splitdns.config 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 splitdns.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 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 (<code>dest_domain</code>), destination host (<code>dest_host</code>), or on a regular expression (<code>url_regex</code>).
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.
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 .

Option	Description
Apply	Applies the configuration changes.
Close	Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes are lost.

ICAP



Note

The ICAP configuration options appear on the Configure pane only if you have enabled **ICAP** in the **Features** table on the **Configure > My Proxy > Basic > General** tab.

Option	Description
ICAP Service URI	Specifies the Uniform Resource Identifier for the ICAP service. The format is: <code>icap://hostname:port/path</code> . For example: <code>icap://ICAP_machine:1344/REQMOD</code> The default ICAP port is 1344. If you are using the default port, you need not specify it in the URI.
Analyze HTTPS Content	Select whether decrypted traffic should be sent to Data Security Suite for analysis or sent directly to the destination.
Analyze FTP Uploads	Select whether to send FTP upload requests to Websense Data Security Suite for analysis. The FTP proxy feature must be enabled. See FTP , page 240.
Action for Communication Errors	Select whether to allow traffic or send a block page if Content Gateway receives an error while communication with Websense Data Security Suite.
Action for Large files	Select whether to allow traffic or send a block page if a file larger than the size limit specified in DSS is sent. The default size limit in DSS is 12 MB.

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.

Option	Description
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.

SSL

The SSL configuration options are divided into the following categories:

- ◆ Certificates (see *Managing certificates*, page 119)

- ◆ Decryption/Encryption (see [Configuring SSL Manager for inbound traffic](#), page 122 and [Configuring SSL Manager for outbound traffic](#), page 122)
- ◆ Validation (see [Validating certificates](#), page 124)
- ◆ Incidents (see [Managing Web HTTPS site access](#), page 128)
- ◆ Client certificates (see [Client certificates](#), page 132)
- ◆ Logging (see [Configuring logging](#), page 133)
- ◆ Customization (see [Customizing messages](#), page 136)
- ◆ Internal Root CA (see [Internal Root CAs](#), page 115)

D

Event Logging Formats

Custom logging fields

Related topic:

[Logging format cross-reference, page 282](#)

%<field symbol>	Description
<i>{HTTP header field name}cqh</i>	Logs the information in the requested field of the client request HTTP header; for example, %<{Accept-Language}cqh> logs the Accept-Language: field in client request headers.
<i>{HTTP header field name}pqh</i>	Logs the information in the requested field of the proxy request HTTP header; for example, %<{Authorization}pqh> logs the Authorization: field in proxy request headers.
<i>{HTTP header field name}psh</i>	Logs the information in the requested field of the proxy response HTTP header; for example, %<{Retry-After}psh> logs the Retry-After: field in proxy response headers.
<i>{HTTP header field name}ssh</i>	Logs the information in the requested field of the server response HTTP header; for example, %<{Age}ssh> logs the Age: field in server response headers.
caun	The client authenticated user name; result of the RFC931/ident lookup of the client user name.
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.

%<field symbol>	Description
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 <i>yyyy-mm-dd</i> , where <i>yyyy</i> is the 4-digit year, <i>mm</i> is the 2-digit month, and <i>dd</i> 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 <i>hh:mm:ss</i> , where <i>hh</i> is the 2-digit hour in 24-hour format, <i>mm</i> is the 2-digit minutes, and <i>ss</i> 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).
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.

%<field symbol>	Description
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, <code>img/gif</code>) 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 transfer length; the body 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).
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.

%<field symbol>	Description
<code>ttmsf</code>	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 <code>ttms</code>), 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 <code>ttms</code> field displays 1500 and the <code>ttt</code> field displays 1.
<code>ttt</code>	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.

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 Field Symbols
<code>time</code>	<code>cqts</code>
<code>elapsed</code>	<code>ttms</code>
<code>client</code>	<code>chi</code>
<code>action/code</code>	<code>crp/ssc</code>
<code>size</code>	<code>psql</code>
<code>method</code>	<code>cqhm</code>
<code>url</code>	<code>cquc</code>
<code>ident</code>	<code>caun</code>
<code>hierarchy/from</code>	<code>phr/pqsn</code>
<code>content</code>	<code>psct</code>

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, page 171](#), for more information about defining custom log files.

Netscape Common logging formats

Netscape Common	Content Gateway Field Symbols
host	chi
usr	caun
[time]	[cqtn]
“req”	“cctx”
s1	pssc
c1	pscl

Netscape Extended logging formats

Netscape Extended	Content Gateway Field Symbols
host	chi
usr	caun
[time]	[cqtn]
“req”	“cctx”
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 Field Symbols
host	chi
usr	caun
[time]	[cqtn]
“req”	“cctx”
s1	pscc
c1	pscl
s2	sscc
c2	sscl
b1	cqbl
b2	pqbl
h1	cqhl
h2	pshl
h3	pqhl
h4	sshl
xt	tts
route	phr
pfs	cfsc
ss	pfsc
crc	crc

E

Configuration Files

Websense Content Gateway contains the following configuration files that you can edit to customize the proxy.

- ◆ [arm_security.config](#), page 287
- ◆ [auth.config](#), page 289
- ◆ [bypass.config](#), page 291
- ◆ [cache.config](#), page 294
- ◆ [filter.config](#), page 297
- ◆ [hosting.config](#), page 300
- ◆ [icp.config](#), page 302
- ◆ [ip_allow.config](#), page 303
- ◆ [ipnat.conf](#), page 304
- ◆ [log_hosts.config](#), page 305
- ◆ [logs_xml.config](#), page 306
- ◆ [mgmt_allow.config](#), page 313
- ◆ [parent.config](#), page 314
- ◆ [partition.config](#), page 316
- ◆ [records.config](#), page 318
- ◆ [remap.config](#), page 374
- ◆ [socks.config](#), page 375
- ◆ [splitdns.config](#), page 377
- ◆ [storage.config](#), page 379
- ◆ [update.config](#), page 379
- ◆ [congestion.config](#), page 381

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
<code>x</code>	Matches the character <code>x</code> .
<code>.</code>	Match any character.
<code>^</code>	Specifies beginning of line.
<code>\$</code>	Specifies end of line.
<code>[xyz]</code>	A <i>character class</i> . In this case, the pattern matches either <code>x</code> , <code>y</code> , or <code>z</code> .
<code>[abj-oZ]</code>	A <i>character class</i> with a range. This pattern matches <code>a</code> , <code>b</code> , any letter from <code>j</code> through <code>o</code> , or <code>Z</code> .
<code>[^A-Z]</code>	A <i>negated character class</i> . For example, this pattern matches any character except those in the class.
<code>r*</code>	Zero or more <code>r</code> 's, where <code>r</code> is any regular expression.
<code>r+</code>	One or more <code>r</code> 's, where <code>r</code> is any regular expression.
<code>r?</code>	Zero or one <code>r</code> , where <code>r</code> is any regular expression.
<code>r{2,5}</code>	From two to five <code>r</code> 's, where <code>r</code> is any regular expression.
<code>r{2,}</code>	Two or more <code>r</code> 's, where <code>r</code> is any regular expression.
<code>r{4}</code>	Exactly 4 <code>r</code> 's, where <code>r</code> is any regular expression.
<code>"[xyz]"images"</code>	The literal string <code>[xyz]"images"</code>
<code>\X</code>	If <code>X</code> is <code>a</code> , <code>b</code> , <code>f</code> , <code>n</code> , <code>r</code> , <code>t</code> , or <code>v</code> , then the ANSI-C interpretation of <code>\x</code> ; Otherwise, a literal <code>X</code> . This is used to escape operators such as <code>*</code> .
<code>\0</code>	A NULL character.
<code>\123</code>	The character with octal value 123.
<code>\x2a</code>	The character with hexadecimal value 2a.
<code>(r)</code>	Matches an <code>r</code> ; where <code>r</code> is any regular expression. You can use parentheses to override precedence.
<code>rs</code>	The regular expression <code>r</code> , followed by the regular expression <code>s</code> .
<code>r s</code>	Either an <code>r</code> or an <code>s</code> .
<code>#<n>#</code>	Inserts an <i>end</i> node causing regular expression matching to stop when reached. The value <code>n</code> is returned.

Examples

You can specify `dest_domain=mydomain.com` to match any host in `mydomain.com`. Likewise, you can specify `dest_domain=.` to match any request.

arm_security.config

The **arm_security.config** file defines which hosts are allowed to communicate with the Content Gateway machine using TCP and UDP through defined ports. Content Gateway uses this configuration file when the ARM security option is enabled. For information about enabling the ARM security option, see [Controlling host access to the proxy server](#), page 140.



Important

Content Gateway requires that several ports be open to support the full set of Websense Web Security Gateway features. These ports are allowed in the default `arm_security.config` file. For a complete list of required ports, please see the *Content Gateway Installation Guide*.

You can specify the following information:

- ◆ The ports that are open, for either TCP or UDP



Note

Traffic to destination port 80 cannot be denied by an ARM security rule.

- ◆ The hosts that are denied access to specific destination ports, for either TCP or UDP
- ◆ The hosts that are allowed access to specific destination ports, for either TCP or UDP



Warning

Before you enable the ARM security option, ensure that you have either console access to the Content Gateway machine or that you have added the appropriate rules to the configuration file to allow telnet or ssh access for yourself.

The ports you specify in the access control list remain closed even when Content Gateway is not running.

ARM security rules are applied as follows:

- ◆ When ARM security is enabled, ports that are not explicitly opened or allowed with an open or allow rule are denied.
- ◆ A deny rule is applied only if the specified port has already been opened with an open rule; otherwise the deny rule is discarded.

- ◆ Conflicts between allow and deny rules (where an allow/open/deny or open/allow/deny or open/deny/allow references the same IP/port combination) are always won by the deny rule, regardless of rule ordering in the file.



Important

After you modify this file, you must restart the proxy.

Format

Each line in the `arm_security.config` file uses one of the following formats:

```
open tcp | udp ports o_ports
deny tcp | udp dport d_ports src src_ipaddress
allow tcp | udp src src_ipaddress dst dst_ipaddress
dport d_ports sport s_ports
```

Field	Allowed Inputs
<i>o_ports</i>	The port or series of ports that are open by default.
<i>d_ports</i>	The destination port, or range of destination ports, through which TCP traffic is either allowed or denied. Note: Port 80 is always open.
<i>s_ports</i>	The source port, or range of source ports, from which TCP traffic is allowed.
<i>src_ipaddress</i>	The IP address, or range of IP addresses, specifying the source of the communication.
<i>dst_ipaddress</i>	The IP address, or range of IP addresses, specifying the destination of the communication.



Note

The `arm_security.config` file does not support spaces before or after the hyphen in an IP address range. For example, you must specify `12.34.56.1-12.34.56.7`.

Examples

The following example defines ports 119, 23, and 554 as open for TCP communication. All other ports are closed.

```
open tcp ports 119 23 554
```

This example specifies that host 209.1.2.2 is denied access to destination port 90 using UDP (the example assumes that port 90 has already been opened).

```
deny udp dport 90 src 209.1.2.2
```

In the following example, the first line specifies that host 111.11.11.1 using source port 20 is allowed to communicate with host 123.45.67.8 on destination ports 127-130 using TCP. The second line specifies that all hosts are allowed to communicate with host 123.12.3.4 using UDP.

```
allow tcp src 111.11.11.1 dst 123.45.67.8 dport 127-130
sport 20
allow udp dst 123.12.3.4
```

auth.config

The **auth.config** file stores rules that direct specified IP addresses and IP address ranges to distinct authentication servers (aka domain controllers). These rules are defined on the **Configure > Security > Access Control > Authentication Realms** tab. See [Multiple authentication realms](#), page 161.

- ◆ Multiple authentication realms are supported for NTLM and LDAP authentication only.
- ◆ Separate rules must be created for NTLM and LDAP.
- ◆ The specifiers used in LDAP and NTLM rules differ.
- ◆ Only one authentication method can be active at a time, LDAP or NTLM, and only rules for that method are applied. Other rules remain in auth.config but are not used.
- ◆ Rules are applied from the list top-down; only the first match is applied.



Note

If all of the clients in your network are authenticated by authentication servers that share trust relationships, you don't need to create rules for multiple authentication realms.

- ◆ If no rule in the **auth.config** table applies to a client IP address, the default realm is applied. The default authentication realm is defined by the settings on the **Access Control > LDAP and NTLM** configuration tabs.



Important

A default authentication realm is required.

Format

Each line in **auth.config** contains an authentication rule that consists of a set of tags, each followed by its value. Authentication rules have the format:

```
type=<auth_type> name=<profile_name> src_ip=<IP addresses> <additional tags>
```

The following table lists the tags used in an NTLM rule.

NTLM Tag	Allowed Value
type	NTLM
name	Takes a simple, unique name used in logging.
src_ip	Takes a comma separated list of IP addresses and IP address ranges.
enabled (optional)	Specifies whether the rule will be active: <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled
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.
dc_load_balance (optional)	Specifies whether load balancing is used: <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled <p>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.</p>

The following table lists the tags used in an LDAP rule.

LDAP Tag	Allowed Value
type	LDAP
name	Takes a simple, unique name used in logging.
src_ip	Takes a comma separated list of IP addresses and IP address ranges.
enabled	Specifies whether the rule will be applied: <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled
server_name	Specifies the fully qualified domain name of the LDAP server.
server_port (optional)	Specifies the LDAP server port. The default is 389. 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.

LDAP Tag	Allowed Value
uid_filter (optional)	Specifies the type of service, if different from that configured on the LDAP tab. Enter sAMAccountName 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. <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled If enabled, set the LDAP port to 636 or 3269 (secure LDAP ports).
attr	Specifies an LDAP attribute name.
attr_value	Specifies an LDAP attribute value.

Examples

NTLM:

```
type=ntlm name=CorpHQ src_ip=10.1.1.1,12.13.0.0-12.13.0.128
dc_list=HQdc1.BigCorp.com,HQdc2.BigCorp.com
```

LDAP:

```
type=ldap name=CorpHQ src_ip=10.1.1.1,12.13.0.0-12.13.0.128
server_name=HQldap1.BigCorp.com server_port=389
```



Note

Rules are applied by first match in the order listed.

bypass.config

The **bypass.config** file contains *static* bypass rules that Content Gateway uses in transparent proxy caching 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, page 293](#).

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. For example, you can bypass clients that do not want to use caching.
- ◆ *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.



Important

Destination bypass rules prevent the proxy from caching an entire site. You will experience hit rate impacts if the site you bypass is popular.

- ◆ *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 or out-of-band HTTP traffic problems when cached. Source/destination bypass rules can be preferable to destination rules because they block a destination server only for users that experience problems.

Format

Bypass rules have the following format:

```
bypass src ipaddress | dst ipaddress | src ipaddress AND dst
ipaddress
```

Option	Description
<code>src <i>ipaddress</i></code>	<p>Specifies the source (client) IP address in incoming requests that the proxy must bypass.</p> <p><i>ipaddress</i> can be one of the following:</p> <p>A simple IP address, such as 123.45.67.8</p> <ul style="list-style-type: none"> • 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
<code>dst <i>ipaddress</i></code>	<p>Specifies the destination (origin server) IP address in incoming requests that the proxy must bypass.</p> <p><i>ipaddress</i> can be one of the following:</p> <p>A simple IP address, such as 123.45.67.8</p> <ul style="list-style-type: none"> • 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
<code>src <i>ipaddress</i> AND dst <i>ipaddress</i></code>	<p>Specifies the source and destination IP address pair that the proxy must bypass.</p> <p><i>ipaddress</i> must be a single IP address, such as 123.45.67.8</p>

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, page 293](#).

**Note**

For the dynamic deny bypass rules to work, you must enable the **Dynamic Bypass** option in Content Gateway Manager or set the variable `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.

Examples

The following example shows source, destination, and source/destination *bypass* rules:

```
bypass src 1.1.1.0/24, 25.25.25.25, 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 111.11.1.1
```

cache.config

The **cache.config** file defines how the proxy caches Web objects. You can add caching rules to specify the following configuration:

- ◆ Not to cache objects from specific IP addresses
- ◆ How long to pin particular objects in the cache
- ◆ How long to consider cached objects as fresh

- ◆ Whether to ignore no-cache directives from the server



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 **cache.config** file contains a caching rule. Content Gateway recognizes three space-delimited 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
url_regex	A regular expression to be found in a URL

Secondary specifiers are optional in the **cache.config** file. The following table lists the possible secondary specifiers and their allowed values.



Note

You can use more than one secondary specifier in a rule. However, you cannot repeat a secondary specifier.

Secondary Specifier	Allowed Value
port	A requested URL port
scheme	A request URL protocol; one of the following: <ul style="list-style-type: none"> • HTTP • FTP
prefix	A prefix in the path part of a URL
suffix	A file suffix in the URL

Secondary Specifier	Allowed Value
method	A request URL method; one of the following: <ul style="list-style-type: none"> • get • put • trace
time	A time range, such as 08:00-14:00
src_ip	A client IP address.
user_agent	A request header User-Agent value.

The following table lists the possible actions and their allowed values.

Action	Value
action	One of the following values: <ul style="list-style-type: none"> • <code>never-cache</code> configures the proxy to never cache specified objects. • <code>ignore-no-cache</code> configures the proxy to ignore all <code>Cache-Control: no-cache</code> headers. • <code>ignore-client-no-cache</code> configures the proxy to ignore <code>Cache-Control: no-cache</code> headers from client requests. • <code>ignore-server-no-cache</code> configures the proxy to ignore <code>Cache-Control: no-cache</code> headers from origin server responses.
pin-in-cache	The amount of time you want to keep the object(s) in the cache. The following time formats are allowed: <ul style="list-style-type: none"> • <i>d</i> for days (for example 2d) • <i>h</i> for hours (for example, 10h) • <i>m</i> for minutes (for example, 5m) • <i>s</i> for seconds (for example, 20s) • mixed units (for example, 1h15m20s)
revalidate	The amount of time you want to consider the object(s) fresh. Use the same time formats as <code>pin-in-cache</code> .
ttl-in-cache	The amount of time you want to keep objects in the cache regardless of <code>Cache-Control</code> response headers. Use the same time formats as <code>pin-in-cache</code> and <code>revalidate</code> .

Examples

The following example configures the proxy to never cache FTP documents requested from the IP address 112.12.12.12:

```
dest_ip=112.12.12.12 scheme=ftp action=never-cache
```

The following example configures the proxy to keep documents with URLs that contain the regular expression `politics` and the path **prefix/viewpoint** in the cache for 12 hours:

```
url_regex=politics prefix=/viewpoint pin-in-cache=12h
```

The following example configures the proxy to revalidate `gif` and `jpeg` objects in the domain `mydomain.com` every 6 hours and all other objects in `mydomain.com` every hour:

```
dest_domain=mydomain.com suffix=gif revalidate=6h
dest_domain=mydomain.com suffix=jpeg revalidate=6h
dest_domain=mydomain.com revalidate=1h
```

**Note**

The rules are applied in the order listed.

filter.config

Filtering rules stored in **filter.config** allow you to:

- ◆ Deny or allow URL requests
- ◆ Keep or strip header information from client requests
- ◆ Allow specified applications or requests to specified web sites to bypass authentication
- ◆ Prevent specified applications from transiting the proxy

Filtering rules are most conveniently defined in Content Gateway Manager on the **Configure > Security > Access Control > Filtering** tab. See [Creating filtering rules](#), page 144.

**Note**

Filtering rules for NTLM and LDAP are defined on the **Access Control > Authentication Realms** tab and stored in the [auth.config](#) file.

**Important**

After you modify the file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 **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 destinations.

Primary Destination	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

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: <ul style="list-style-type: none"> • get • post • put • trace

Secondary Specifier	Allowed Value
scheme	A request URL protocol. You can specify one of the following: <ul style="list-style-type: none"> • HTTP • HTTPS • FTP (for FTP over HTTP only)
user_agent	A request header User-Agent value.

The following table lists the possible actions and their allowed values.

Action	Allowed Value
action	Specify one of the following: <ul style="list-style-type: none"> • allow - to allow particular URL requests to bypass authentication. The proxy caches and 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: <ul style="list-style-type: none"> • date • host • cookie • client_ip
strip_hdr	The client request header information that you want to strip. You can specify the same options as with keep_hdr .

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
```

hosting.config

The **hosting.config** file lets you assign cache partitions to specific origin servers and domains so that you can manage your cache space more efficiently and restrict disk usage.

For step-by-step instructions on partitioning the cache according to origin servers and domains, see [Partitioning the cache according to origin server or domain](#), page 80.



Note

Before you can assign cache partitions to specific origin servers and domains, you must partition your cache according to size and protocol in the **partition.config** file. For more about cache partitioning, see [Partitioning the cache](#), page 80. For a description of the **partition.config** file, see [partition.config](#), page 316.

After you modify the **hosting.config** file, run **content_line -x** from the Content Gateway **bin** directory to apply the changes. When you apply the changes to a node in a cluster, Content Gateway automatically applies the changes to all nodes in the cluster.



Important

The partition configuration must be the same on all nodes in a cluster.

Format

Each line in the **hosting.config** file must have one of the following formats:

```
hostname=hostname partition=partition_numbers
domain=domain_name partition=partition_numbers
```

where:

hostname is the fully qualified hostname of the origin server whose content you want to store on a particular partition (for example, `www.myhost.com`).

domain_name is the domain whose content you want to store on a particular partition (for example, `mydomain.com`).

partition_numbers is a comma-separated list of the partitions on which you want to store the content that belongs to the origin server or domain listed. The partition numbers must be valid numbers listed in the **partition.config** file (see [partition.config](#), page 316).

**Note**

If you want to allocate more than one partition to an origin server or domain, enter the partitions in a comma-separated list on one line. The **hosting.config** file cannot contain multiple entries for the same origin server or domain.

Generic Partition

When configuring the **hosting.config** file, you must assign a generic partition to use for content that does not belong to any of the origin servers or domains listed. If all partitions for a particular origin server become corrupt, Content Gateway uses the generic partition to store content for that origin server.

The generic partition must have the following format:

```
hostname=* partition=partition_numbers
```

where **partition_numbers** is a comma-separated list of generic partitions.

Examples

The following example configures the proxy to store content from the domain **mydomain.com** in partition 1 and content from **www.myhost.com** in partition 2. The proxy stores content from all origin servers in partitions 3 and 4.

```
domain=mydomain.com partition=1
hostname=www.myhost.com partition=2
hostname=* partition=3,4
```

icp.config

The **icp.config** file defines ICP peers (parent and sibling caches).



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 **icp.config** file contains the name and configuration information for a single ICP peer in the following format:

```
host:host_IP:peer_type:proxy_port:icp_port:MC_on:MC_IP:
MC_TTL:
```

Field	Description
<i>host</i>	The hostname of the ICP peer. This field is optional; if you do not specify the hostname of the ICP peer, you must specify the IP address.
<i>host_IP</i>	The IP address of the ICP peer. This field is optional; if you do not specify the IP address of the ICP peer, you must specify the hostname.
<i>ctype</i>	Use the following options: <ul style="list-style-type: none"> • 1 to indicate an ICP parent cache • 2 to indicate an ICP sibling cache
<i>proxy_port</i>	The port number of the TCP port used by the ICP peer for proxy communication.
<i>icp_port</i>	The port number of the UDP port used by the ICP peer for ICP communication.
<i>multicast_member</i>	Multicast on/off. <ul style="list-style-type: none"> • 0 if multicast is disabled • 1 if multicast is enabled
<i>multicast_ip_str</i>	The multicast IP address.
<i>multicast_ttl</i>	The multicast time to live. Use the following options: <ul style="list-style-type: none"> • 1 if IP multicast datagrams will not be forwarded beyond a single subnetwork • 2 to allow delivery of IP multicast datagrams to more than one subnet (if there are one or more multicast routers attached to the first hop subnet)

Examples

The following example configuration is for three nodes: the local host, one parent, and one sibling.

```
localhost:0.0.0.0:3:8080:3130:0:0.0.0.0:0:
host1:123.12.1.23:1:8080:3131:0:0.0.0.0:0:
host2:123.12.1.24:2:8080:3131:0:0.0.0.0:0:
```

ip_allow.config

The **ip_allow.config** file controls client access to the proxy cache. You can specify ranges of IP addresses that are allowed to use the Content Gateway server as a Web proxy cache.



Important

After you modify the file, run `content_line -x` from the Content Gateway **bin** directory (default location is **/opt/WCG/bin**) to apply the changes. 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 cache.

The action `ip_allow` allows the specified clients to access the proxy cache.

The action `ip_deny` denies the specified clients to access the proxy cache.

By default, the **ip_allow.config** file contains the following line, which allows all clients to access the proxy cache. 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 cache:

```
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 cache:

```
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 cache:

```
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 readdress 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, page 180](#)).

**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 `content_line -x` from the Content Gateway **bin** directory (default location is **/opt/WCG/bin**) to apply the changes. 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 **log_hosts.config** file has the following format:

```
hostname
```

where *hostname* is the hostname of the origin server.

**Note**

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

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, page 307](#).
- ◆ `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, page 308](#).
- ◆ `LogObject` specifies an object that contains a particular format, a local filename, filters, and collation servers. See [LogObject, page 309](#).



Note

The **logs_xml.config** file ignores extra white space, blank lines, and all comments.

LogFormat

The following table lists the LogFormat specifications.

Field	Allowed Inputs
<code><Name = "valid_format_name"/></code>	Required. Valid format names include any name except squid, common, extended, or extended2, which are pre-defined formats. There is no default for this tag.
<code><Format = "valid_format_specification"/></code>	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 Custom logging fields, page 279 . The specified field can be of two types: Simple: for example, %<cqu> A field within a container, such as an HTTP header or a Content Gateway statistic. Fields of this type have the syntax: '%<{field}container>'
<code><Interval = "aggregate_interval_secs"/></code>	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: <ul style="list-style-type: none"> • COUNT • SUM • AVG • FIRST • LAST

LogFilter

The following table lists the LogFilter specifications.

Field	Allowed Inputs
<code><Name = "valid_filter_name"/></code>	Required. All filters must be uniquely named.
<code><Condition = "valid_log_field valid_operator valid_comparison_value"/></code>	<p>Required. This field contains the following elements:</p> <p><code>valid_log_field</code> - the field that will be compared against the given value. For more information, see Logging format cross-reference, page 282.</p> <p><code>valid_operator_field</code> - 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.</p> <p><code>valid_comparison_value</code> - 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.</p> <p>Note: There are no negative comparison operators. If you want to specify a negative condition, use the Action field to REJECT the record.</p>
<code><Action = "valid_action_field"/></code>	Required. ACCEPT or REJECT. This instructs Content Gateway to either accept or reject records satisfying the condition of the filter.

LogObject

The following table lists the `LogObject` specifications.

Field	Allowed Inputs
<code><Format = "valid_format_name"/></code>	Required. Valid format names include the predefined logging formats: <code>squid</code> , <code>common</code> , <code>extended</code> , and <code>extended2</code> , as well as any previously-defined custom log formats. There is no default for this tag.
<code><Filename = "file_name"/></code>	<p>Required. The filename to which the given log file is written on the local file system or on a remote collation server. No local log file will be created if you fail to specify this tag. All filenames are relative to the default logging directory.</p> <p>If the name does not contain an extension (for example, <code>squid</code>), the extension <code>.log</code> is automatically appended to it for ASCII logs and <code>.blog</code> for binary logs. (See <code><Mode = "valid_logging_mode"/></code> below.)</p> <p>If you do not want an extension to be added, end the filename with a single dot (<code>.</code>): for example, <code>squid.</code></p>

Field	Allowed Inputs
<pre><Mode = "valid_logging_mode"/></pre>	<p>Valid logging modes include <code>ascii</code>, <code>binary</code>, and <code>ascii_pipe</code>. The default is <code>ascii</code>.</p> <p>Use <code>ascii</code> to create event log files in human-readable form (plain ASCII).</p> <p>Use <code>binary</code> 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 <code>logcat</code> utility to translate binary log files to ASCII format before you can read them.</p> <p>Use <code>ascii_pipe</code> 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.</p> <p>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.</p>
<pre><Filters = "list_of_valid_filter_names"/></pre>	<p>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.</p>
<pre><Protocols = "list_of_valid_protocols"/></pre>	<p>A comma-separated list of the protocols this object should log. Valid protocol names include <code>HTTP</code> and <code>ICP</code>.</p>
<pre><ServerHosts = "list_of_valid_servers"/></pre>	<p>A comma-separated list of valid hostnames. This tag indicates that only entries from the named servers will be included in the file.</p>
<pre><CollationHosts = "list_of_valid_hostnames"/></pre>	<p>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, <code>host:port</code>).</p>

Field	Allowed Inputs
<code><Header = "header"/></code>	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.
<code><RollingEnabled = "truth value"/></code>	Enables or disables log file rolling for the <i>LogObject</i> . This setting overrides the value for the configuration setting Log Rolling: Enabled/Disabled in Content Gateway Manager or <i>proxy.config.log2.rolling_enabled</i> 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 <i>LogObject</i> .
<code><RollingIntervalSec = "seconds"/></code>	Specifies the seconds between log file rolling for the <i>LogObject</i> . This setting overrides the value for the configuration setting Log Rolling: Interval in Content Gateway Manager or <i>proxy.config.log2.rolling_interval_sec</i> in the records.config file. This option allows you to specify different rolling intervals for different <i>LogObjects</i> .
<code><RollingOffsetHr = "hour"/></code>	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 Content Gateway Manager or <i>proxy.config.log2.rolling_offset_hr</i> 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 port 4000 of the collation host `logs.company.com` and to port 5000 of the collation host `209.131.52.129`.

```
<LogObject>
<Format = "minimal"/>
<Filename = "minimal"/>
<ServerHosts = "company.com, server.somewhere.com"/>
<Protocols = "http"/>
<CollationHosts =
"logs.company.com:4000, 209.131.52.129:5000"/>
</LogObject>
```

WELF (WebTrends Enhanced Log Format)

Content Gateway supports WELF, the WebTrends Enhanced Log Format, 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 Content Gateway Manager.



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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
```

where *ipaddress* is the IP address or range of IP addresses allowed to access Content Gateway Manager.

action must specify either `ip_allow` to allow access to Content Gateway Manager or `ip_deny` to deny access.

By default, the `mgmt_allow.config` file contains the following line, which allows all remote hosts to access 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 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 Content Gateway Manager:

```
src_ip=123.12.3.000-123.12.3.123 action=ip_allow
```

The following example configures Content Gateway to deny the IP address 123.45.67.8 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

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](#), page 74.



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 **parent.config** file must contain a parent caching rule. Content Gateway recognizes three space-delimited 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

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
scheme	A request URL protocol; one of the following: <ul style="list-style-type: none"> • HTTP • FTP
method	A request URL method; one of the following: <ul style="list-style-type: none"> • get • post • put • trace
user_agent	A request header User-Agent value.

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: <ul style="list-style-type: none">• <code>true</code> - Content Gateway goes through the parent cache list in a round-robin based on client IP address.• <code>strict</code> - Content Gateway machines serve requests strictly in turn. For example, machine <code>proxy1</code> serves the first request, <code>proxy2</code> serves the second request, and so on.• <code>false</code> - round-robin selection does not occur.
go_direct	One of the following values: <ul style="list-style-type: none">• <code>true</code> - requests bypass parent hierarchies and go directly to the origin server.• <code>false</code> - requests do not bypass parent hierarchies.

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="p1.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
```

Every line in the **parent.config** file must contain either a `parent=` or `go_direct=` directive.

partition.config

The **partition.config** file lets you manage your cache space more efficiently by creating cache partitions of different sizes. You can further configure these partitions to store data from certain origin servers and domains in the [hosting.config](#) file. This

allows you to take better advantage of caching of frequently visited sites where the content changes infrequently.

**Important**

The partition configuration must be the same on all nodes in a cluster.

You must stop Content Gateway before you change the cache partition size.

Format

For each partition you want to create, enter a line with the following format:

```
partition=partition_number scheme=protocol_type  
size=partition_size
```

where:

partition_number is a number between 1 and 255 (the maximum number of partitions is 255).

protocol_type is **http**.

**Note**

Only HTTP is supported at this time. Streaming media content—**mixt**—is not supported.

partition_size is the amount of cache space allocated to the partition. This value can be either a percentage of the total cache space or an absolute value. The absolute value must be a multiple of 128 MB, where 128 MB is the smallest value. If you specify a percentage, the size is rounded down to the closest multiple of 128 MB. Each partition is striped across several disks to achieve parallel I/O. For example, if there are four disks, a 1 GB partition will have 256 MB on each disk (assuming each disk has enough free space available).

**Note**

If you do not allocate all the disk space in the cache, the extra disk space is not used. You can use the extra space later to create new partitions without deleting and clearing the existing partitions.

Examples

The following example partitions the cache evenly:

```
partition=1 scheme=http size=50%  
partition=2 scheme=http size=50%
```

records.config

The **records.config** file is a list of configurable variables used by Content Gateway.

Many of the values are set when you set configuration options in Content Gateway Manager, or through the command-line interface. Some configuration 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 Content Gateway Manager to configure Content Gateway.



Important

After you modify this file, to apply the changes run `content_line -x` from the Content Gateway **bin** directory (default location is **/opt/WCG/bin**). 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
```

where *DATATYPE* is INT (an integer), STRING (a string), or 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.arm.enabled** is a yes/no flag. A value of 0 (zero) disables the option. A value of 1 enables the option.

```
CONFIG proxy.config.arm.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 tables describe the configuration variables listed in the **records.config** file.

System variables

Local manager

Virtual IP manager

Alarm configuration

ARM (transparency configuration)

Load shedding configuration (ARM)

Authentication basic realm

LDAP

RADIUS authentication

NTLM

Transparent authentication

HTTP engine

Parent proxy configuration

Cache control

Heuristic expiration

Dynamic content and content negotiation

Anonymous FTP password

Cached FTP document lifetime

FTP transfer mode

FTP engine

Customizable user response pages

SOCKS processor

Net subsystem

Cluster subsystem

Cache

DNS

HostDB

Logging configuration

URL remap rules

ICP configuration

Scheduled update configuration

WCCP configuration

ARM (security configuration)

SSL Decryption

ICAP

Connectivity, analysis, and boundary conditions

System variables

Configuration Variable Data Type	Data Type	Default Value	Description
proxy.config.proxy_name	STRING		Specifies the name of the Content Gateway node.
proxy.config.bin_path	STRING	/opt/WCG/bin	Specifies the location of the Content Gateway bin directory. This is the directory in which the Content Gateway binary files are placed by the installer.
proxy.config.proxy_binary	STRING	content_gateway	Specifies the name of the executable that runs the content_gateway process.
proxy.config.proxy_binary_opts	STRING	-M	Specifies the command-line options for starting content_gateway .
proxy.config.manager_binary	STRING	content_manager	Specifies the name of the executable that runs the content_manager process.
proxy.config.cli_binary	STRING	content_line	Specifies the name of the executable that runs the content_line interface.
proxy.config.watch_script	STRING	content_cop	Specifies the name of the executable that runs the content_cop process.
proxy.config.env_prep	STRING	example_prep.sh	Specifies the script that is executed before the content_manager process spawns the content_gateway process.

Configuration Variable Data Type	Data Type	Default Value	Description
proxy.config.config_dir	STRING	config	Specifies the directory, relative to bin_path (above), that contains the Content Gateway configuration files.
proxy.config.temp_dir	STRING	/tmp	Specifies the directory used for Content Gateway temporary files
proxy.config.alarm_email	STRING	websense	Specifies 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	LOG_DAEMON	Specifies the facility used to record system log files. See Working With Log Files, page 167 .
proxy.config.cop.core_signal	INT	3	Specifies 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	45	Specifies the interval, in seconds, between heartbeat tests performed by content_cop to test the health of the content_manager and content_gateway processes. Note: Do not change the value of this variable.
proxy.config.cop.linux_min_swapfree_kb	INT	10240	This variable is not used.
proxy.config.cop.linux_min_memfree_kb	INT	10240	This variable is not used.
proxy.config.output.logfile	STRING	content_gateway.out	Specifies 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.

Configuration Variable Data Type	Data Type	Default Value	Description
<code>proxy.config.snapshot_dir</code>	STRING	snapshots	Specifies 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.
<code>proxy.config.attach_debugger_script</code>	STRING	attach_debugger	This variable should be used only on the direction of Websense Technical Support. If set, when the content_gateway process resets, a debug script (in /opt/WCG/bin) is run.

Local manager

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.lm.sem_id</code>	INT	11452	Specifies the semaphore ID for the local manager. Note: Do not change the value of this variable.
<code>proxy.local.cluster.type</code>	INT	3	Sets the clustering mode: <ul style="list-style-type: none"> • 1 = full-clustering mode • 2 = management-only mode • 3 = no clustering
<code>proxy.config.cluster.rsport</code>	INT	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.
<code>proxy.config.cluster.mcport</code>	INT	8088	Specifies the multicast port. The multicast port is used for node identification. All nodes in a cluster must use the same multicast port.
<code>proxy.config.cluster.mc_group_addr</code>	STRING	224.0.1.37	Specifies the multicast address for cluster communications. All nodes in a cluster must use the same multicast address.
<code>proxy.config.cluster.mc_ttl</code>	INT	1	Specifies the multicast Time-To-Live for cluster communications.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.cluster.log_bogus_mc_msgs</code>	INT	1	Enables (1) or disables (0) logging of bogus multicast messages.
<code>proxy.config.admin.html_doc_root</code>	STRING	ui	Specifies the document root for Content Gateway Manager.
<code>proxy.config.admin.web_interface_port</code>	INT	8081	Specifies the Content Gateway Manager port.
<code>proxy.config.admin.autoconf_port</code>	INT	8083	Specifies the autoconfiguration port.
<code>proxy.config.admin.overseer_port</code>	INT	8082	Specifies the port used for retrieving and setting statistics and configuration variables.
<code>proxy.config.admin.admin_user</code>	STRING	admin	Specifies the administrator ID that controls access to Content Gateway Manager.
<code>proxy.config.admin.admin_password</code>	STRING		Specifies the encrypted administrator password that controls access to Content Gateway Manager. You cannot edit the password; however, you can specify a value of <code>NULL</code> to clear the password. <i>See How do you access Content Gateway Manager if you forget the master administrator password?, page 396.</i>
<code>proxy.config.admin.basic_auth</code>	INT	1	Enables (1) or disables (0) basic user authentication to control access to Content Gateway Manager. Note: If basic authentication is <i>not</i> enabled, any user can access Content Gateway Manager to monitor and configure Content Gateway.
<code>proxy.config.admin.use_ssl</code>	INT	1	Enables the Content Gateway Manager SSL option for secure communication between a remote host and the Content Gateway Manager.
<code>proxy.config.admin.ssl_cert_file</code>	STRING	server.pem	Specifies the filename of the SSL certificate installed on the Content Gateway system for secure communication between a remote host and Content Gateway Manager.

Configuration Variable	Data Type	Default Value	Description
proxy.config.admin.number_config_bak	INT	3	Specifies the maximum number of copies of rolled configuration files to keep.
proxy.config.admin.user_id	STRING	Websense	Specifies the non-privileged user account designated to Content Gateway.
proxy.config.admin.ui_refresh_rate	INT	30	Specifies the refresh rate for the display of statistics in the Monitor pages of Content Gateway Manager.
proxy.config.admin.log_mgmt_access	INT	0	Enables (1) or disables (0) logging of all Content Gateway Manager transactions to the lm.log file.
proxy.config.admin.log_resolve_hostname	INT	1	When enabled (1), the hostname of the client connecting to Content Gateway Manager is recorded in the lm.log file. When disabled (0), the IP address of the client connecting to Content Gateway Manager is recorded in the lm.log file.
proxy.config.admin.subscription	STRING	NULL	Not used.
proxy.config.admin.supported_cipher_list	STRING	RC4-MD5, RC4-SHA, 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 allowed when a browser establishes a secure connection with Content Gateway Manager. No validation is performed on the string. The first good value is used. If there is no good value, the browser is not allowed to connect to the manager and an error is returned.
proxy.config.lm.display_reset_alarm	INT	0	When enabled (1), email is sent to the administrator (proxy.config.alarm_email) whenever Content Gateway resets. Default is 0.

Process manager

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.process_manager.mgmt_port</code>	INT	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	Default Value	Description
<code>proxy.config.vmap.enabled</code>	INT	0	Enables (1) or disables (0) the virtual IP option.

Alarm configuration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.alarm.bin</code>	STRING	<code>example_alarm_bin.sh</code>	Specifies the name of the script file that can execute certain actions when an alarm is signaled. The default file is a sample script named example_alarm_bin.sh located in the bin directory. You must edit the script to suit your needs.
<code>proxy.config.alarm.abs_path</code>	STRING	NULL	Specifies the full path to the script file specified by proxy.config.alarm.bin (prior entry).

ARM (transparency configuration)

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.arm.enabled</code>	INT	1	Enables (1) or disables (0) the ARM, which is used for transparent proxy caching, IP spoofing, and ARM security. See Enabling the ARM option, page 44 .
<code>proxy.config.arm.ignore_ifp</code>	INT	1	When NAT 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 NAT rule.
<code>proxy.config.arm.always_query_dest</code>	INT	0	<p>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.</p> <p>When enabled, IP addresses are logged, instead of domain names.</p> <p>When disabled, domain names are logged. See Reducing DNS lookups, page 61 for additional information.</p> <p>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.</p>
<code>proxy.config.http.outgoing_ip_spoofing_enabled</code>	INT	0	<p>Enables (1) or disables (0) the IP spoofing option, which allows Content Gateway to establish connections to origin servers with the client IP address instead of the Content Gateway IP address.</p> <p>Note: The variable <code>proxy.config.arm.enabled</code> must be enabled for the IP spoofing option to work.</p> <p>See IP spoofing, page 62.</p>

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.arm.bypass_dynamic_enabled</code>	INT	0	Enables (1) or disables (0) the adaptive bypass option to bypass the proxy and go directly to the origin server when clients or servers cause problems. See Dynamic bypass rules, page 58 .
<code>proxy.config.arm.bypass_use_and_rules_bad_client_request</code>	INT	0	Enables (1) or disables (0) dynamic source/destination bypass in the event of non-HTTP traffic on port 80. Note: The variable proxy.config.arm.bypass_on_bad_client_request must also be enabled for this option to work.
<code>proxy.config.arm.bypass_use_and_rules_400</code>	INT	0	Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 400 error. Note: The variable proxy.config.arm.bypass_on_400 must also be enabled for this option to work.
<code>proxy.config.arm.bypass_use_and_rules_401</code>	INT	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.
<code>proxy.config.arm.bypass_use_and_rules_403</code>	INT	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.
<code>proxy.config.arm.bypass_use_and_rules_405</code>	INT	0	Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 405 error. Note: The variable proxy.config.arm.bypass_on_405 must also be enabled for this option to work.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.arm.bypass_use_and_rules_406</code>	INT	0	Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 406 error. Note: The variable proxy.config.arm.bypass_on_406 must also be enabled for this option to work.
<code>proxy.config.arm.bypass_use_and_rules_408</code>	INT	0	Enables (1) or disables (0) dynamic generation of source/destination bypass rules when an origin server returns a 408 error. Note: The variable proxy.config.arm.bypass_on_408 must also be enabled for this option to work.
<code>proxy.config.arm.bypass_use_and_rules_500</code>	INT	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.
<code>proxy.config.arm.bypass_on_bad_client_request</code>	INT	0	Enables (1) or disables (0) dynamic destination bypass in the event of non-HTTP traffic on port 80.
<code>proxy.config.arm.bypass_on_400</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 400 error.
<code>proxy.config.arm.bypass_on_401</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 401 error.
<code>proxy.config.arm.bypass_on_403</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 403 error.
<code>proxy.config.arm.bypass_on_405</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 405 error.
<code>proxy.config.arm.bypass_on_406</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 406 error.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.arm.bypass_on_408</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 408 error.
<code>proxy.config.arm.bypass_on_500</code>	INT	0	Enables (1) or disables (0) dynamic generation of destination bypass rules when an origin server returns a 500 error.

Load shedding configuration (ARM)

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.arm.loadshedding.max_connections</code>	INT	1000000	Specifies the maximum number of client connections allowed before the proxy starts forwarding incoming requests directly to the origin server.

Authentication basic realm

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.proxy.authenticate.basic.realm</code>	STRING	NULL	Specifies the authentication realm name. If the default of NULL is specified, Content Gateway is used.
<code>proxy.config.auth_type</code>	INT	3	Specifies the type of client authentication. <ul style="list-style-type: none"> • 1 = LDAP • 2 = RADIUS • 3 = NTLM

LDAP

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ldap.auth.enabled</code>	INT	0	Enables (1) or disables (0) LDAP proxy authentication. See Using LDAP proxy authentication , page 152.
<code>proxy.config.ldap.cache.size</code>	INT	5000	Specifies 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.
<code>proxy.config.ldap.cache.storage_size</code>	INT	24582912	Specifies 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.
<code>proxy.config.ldap.auth.ttl_value</code>	INT	3000	Specifies the amount of time (in minutes) that entries in the cache remain valid.
<code>proxy.config.ldap.auth.purge_cache_on_auth_fail</code>	INT	1	When enabled (1), configures Content Gateway to delete the authorization entry for the client in the LDAP cache if authorization fails.
<code>proxy.config.ldap.proc.ldap.server.name</code>	STRING	NULL	Specifies the LDAP server name.
<code>proxy.config.ldap.proc.ldap.server.port</code>	INT	389	Specifies the LDAP server port.
<code>proxy.config.ldap.proc.ldap.base.dn</code>	STRING	NULL	Specifies the LDAP Base Distinguished Name (DN). Obtain this value from your LDAP administrator.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ldap.proc.ldap.uid_filter</code>	STRING	sAMAccountName	Specifies the LDAP login name/ID. Use this as a filter to search the full DN database. For eDirectory or other directory services, enter uid in this field.
<code>proxy.config.ldap.secure.bind.enabled</code>	INT	0	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.
<code>proxy.config.ldap.proc.ldap.server.bind_dn</code>	STRING	NULL	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 no value is specified for this field, the proxy attempts to bind anonymously.
<code>proxy.config.ldap.proc.ldap.server.bind_pwd</code>	STRING	NULL	Specifies a password for the user identified by the proxy.config.ldap.proc.ldap.server.bind_dn variable.

RADIUS authentication

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.radius.auth.enabled</code>	INT	0	Enables (1) or disables (0) RADIUS proxy authentication.
<code>proxy.config.radius.proc.radius.primary_server.name</code>	STRING	NULL	Specifies the hostname or IP address of the primary RADIUS authentication server.
<code>proxy.config.radius.proc.radius.primary_server.auth_port</code>	INT	1812	Specifies the RADIUS server port that Content Gateway uses to communicate with the RADIUS server.
<code>proxy.config.radius.proc.radius.primary_server.shared_key</code>	STRING	NULL	Specifies the key used for encoding with the first RADIUS authentication server.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.radius. proc.radius. secondary_server. name</code>	STRING	NULL	Specifies the hostname or IP address of the secondary RADIUS authentication server.
<code>proxy.config.radius. proc.radius. secondary_server. auth_port</code>	INT	1812	Specifies the port that the proxy uses to communicate with the secondary RADIUS authentication server.
<code>proxy.config.radius. proc.radius. secondary_server. shared_key</code>	STRING	NULL	Specifies the key used for encoding with the secondary RADIUS authentication server.
<code>proxy.config.radius. auth.min_timeout</code>	INT	10	Specifies the amount of time the connection to the RADIUS server can remain idle before Content Gateway closes the connection.
<code>proxy.config.radius. auth.max_retries</code>	INT	10	Specifies the maximum number of times Content Gateway tries to connect to the RADIUS server.
<code>proxy.config.radius. cache.size</code>	INT	1000	Specifies the number of entries allowed in the RADIUS cache. The minimum value is 256 entries.
<code>proxy.config.radius. cache.storage_size</code>	INT	15728640	Specifies 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.
<code>proxy.config.radius. auth.ttl_value</code>	INT	60	Specifies the number of minutes that Content Gateway stores username and password entries in the RADIUS cache.

NTLM

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ntlm.auth.enabled</code>	INT	0	Enables (1) or disables (0) NTLM proxy authentication.
<code>proxy.config.ntlm.dc.list</code>	STRING	NULL	Specifies the hostnames of the domain controllers. You must separate each entry with a comma. The format is: <code>host_name[:port]</code> <code>[%netbios_name]</code> or <code>IP_address[:port]</code> <code>[%netbios_name]</code> If you are using Active Directory 2008, you must include the netbios_name or use SMB port 445.
<code>proxy.config.ntlm.dc.load_balance</code>	INT	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.
<code>proxy.config.ntlm.dc.max_connections</code>	INT	10	Specifies the maximum number of connections Content Gateway can have open to the domain controller.
<code>proxy.config.ntlm.cache.enabled</code>	INT	1	Enables (1) or disables (0) the NTLM cache. 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.

Configuration Variable	Data Type	Default Value	Description
proxy.config.ntlm.cache.ttl_value	INT	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	5000	Specifies the number of entries allowed in the NTLM cache.
proxy.config.ntlm.cache.storage_size	INT	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_exception.list	STRING	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.
proxy.config.ntlm.fail_open	INT	1	Enables (1) or disables (0) whether client requests are allowed to proceed when authentication fails due to: <ul style="list-style-type: none"> • no response from the domain controller • badly formed messages from the client • invalid SMB responses Note: Password authentication failures are always failures.

Transparent authentication

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.transparent_auth_hostname</code>	STRING	NULL	Specify 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. For additional information, see Transparent proxy authentication , page 151.
<code>proxy.config.http.transparent_auth_type</code>	INT	1	Specify: <ul style="list-style-type: none"> • 0 to associate a session ID with the username after the user session is authenticated. This setting is required to uniquely identify users who share a single IP address, such as in proxy-chaining or network address translation. • 1 to associate 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 .
<code>proxy.config.http.transparent_auth_session_time</code>	INT	15	Specify 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	Default Value	Description
<code>proxy.config.http.server_port</code>	INT	8080	Specifies the port that Content Gateway uses when acting as a Web proxy server for Web traffic or when serving Web traffic transparently.
<code>proxy.config.http.server_port_attr</code>	STRING	X	Specifies the server port options. You can specify one of the following: <ul style="list-style-type: none"> • C=SERVER_PORT_COMPRESSED • X=SERVER_PORT_DEFAULT • T=SERVER_PORT_BLIND_TUNNEL
<code>proxy.config.http.server_other_ports</code>	STRING	NULL	Specifies the ports other than the port specified by the variable proxy.config.http.server_port to bind for incoming HTTP requests.
<code>proxy.config.http.ssl_ports</code>	STRING	443 563 8081 8071 9443 9444	Specifies the range of ports used for tunneling. Content Gateway allows tunnels only to the specified ports.
<code>proxy.config.http.insert_request_via_str</code>	INT	1	Specify one of the following: <ul style="list-style-type: none"> • 0 = no extra information is added to the string. • 1 = all extra information is added. • 2 = some extra information is added.
<code>proxy.config.http.insert_response_via_str</code>	INT	1	Specify one of the following: <ul style="list-style-type: none"> • 0 = no extra information is added to the string. • 1 = all extra information is added. • 2 = some extra information is added.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.enable_url_expandomatic</code>	INT	1	Enables (1) or disables (0) .com domain expansion, which configures Content Gateway to attempt 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 host , Content Gateway redirects the request to www.host.com .
<code>proxy.config.http.no_dns_just_forward_to_parent</code>	INT	0	When enabled (1), and if HTTP parent caching is enabled, Content Gateway does no DNS lookups on request hostnames.
<code>proxy.config.http.uncacheable_requests_bypass_parent</code>	INT	0	When enabled (1), Content Gateway bypasses the parent proxy for a request that is not cacheable.
<code>proxy.config.http.keep_alive_enabled</code>	INT	1	Enables (1) or disables (0) the use of keep-alive connections to either origin servers or clients.
<code>proxy.config.http.chunking_enabled</code>	INT	1	Specifies whether Content Gateway will generate a chunked response: <ul style="list-style-type: none"> • 0 = Never • 1 = Always

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.send_http11_requests</code>	INT	3	<p>Configures Content Gateway to use HTTP Version 1.1 when communicating with origin servers. You can specify one of the following values:</p> <ul style="list-style-type: none"> • 1 = Content Gateway always uses HTTP 1.1 when communicating with origin servers. • 2 = Content Gateway uses HTTP 1.1 if the origin server has previously used HTTP 1.1. • 3 = Content Gateway uses HTTP 1.1 if the client request is HTTP 1.1 and the origin server has previously used HTTP 1.1. <p>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.</p>
<code>proxy.config.http.send_http11_asfirstrequest</code>	INT	1	<p>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.</p>
<code>proxy.config.http.share_server_sessions</code>	INT	1	<p>Enables (1) or disables (0) the re-use of server sessions.</p> <p>Note: When IP spoofing is enabled, Content Gateway automatically disables this variable.</p>
<code>proxy.config.http.ftp_enabled</code>	INT	1	<p>Enables (1) or disables (0) Content Gateway from serving FTP requests sent via HTTP.</p>
<code>proxy.config.http.record_heartbeat</code>	INT	0	<p>Enables (1) or disables (0) content_cop heartbeat logging.</p>
<code>proxy.config.http.large_file_support</code>	INT	1	<p>When enabled (1), Content Gateway supports downloading of files larger than 2 GB.</p>

Parent proxy configuration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.parent_proxy.routing_enable</code>	INT	0	Enables (1) or disables (0) the HTTP parent caching option. See Hierarchical Caching, page 73 .
<code>proxy.config.http.parent_proxy.retry_time</code>	INT	300	Specifies the amount of time allowed between connection retries to a parent cache that is unavailable.
<code>proxy.config.http.parent_proxy.fail_threshold</code>	INT	10	Specifies the number of times the connection to the parent cache can fail before Content Gateway considers the parent unavailable.
<code>proxy.config.http.parent_proxy.total_connect_attempts</code>	INT	4	Specifies 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).
<code>proxy.config.http.parent_proxy.per_parent_connect_attempts</code>	INT	2	Specifies the total number of connection attempts allowed per parent if multiple parents are used.
<code>proxy.config.http.parent_proxy.connect_attempts_timeout</code>	INT	30	Specifies the timeout value, in seconds, for parent cache connection attempts.
<code>proxy.config.http.forward.proxy_auth_to_parent</code>	INT	0	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.
<code>proxy.config.http.child_proxy.read_auth_from_header</code>	INT	0	When Content Gateway is the parent proxy, read X-Authenticated-User and X-Forwarded-For fields from incoming request headers. 1 = enabled 0 = disabled

Configuration Variable	Data Type	Default Value	Description
proxy.local.http.parent_proxy.disable_ssl_connect_tunneling	INT	0	When enabled (1), HTTPS requests bypass the parent proxy.
proxy.local.http.parent_proxy.disable_unknown_connect_tunneling	INT	0	When enabled (1), non-HTTPS tunnel requests bypass the parent proxy.

HTTP connection timeouts (secs)

Configuration Variable	Data Type	Default Value	Description
proxy.config.http.keep_alive_no_activity_timeout_in	INT	60	Specifies 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	60	Specifies 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	120	Specifies how long Content Gateway keeps connections to clients open if a transaction stalls.
proxy.config.http.transaction_no_activity_timeout_out	INT	120	Specifies how long Content Gateway keeps connections to origin servers open if the transaction stalls.
proxy.config.http.transaction_active_timeout_in	INT	0	Specifies 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	Default Value	Description
<code>proxy.config.http.transaction_active_timeout_out</code>	INT	0	Specifies 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.
<code>proxy.config.http.accept_no_activity_timeout</code>	INT	86400	Specifies the timeout interval in seconds before Content Gateway closes a connection that has no activity.
<code>proxy.config.http.background_fill_active_timeout</code>	INT	60	Specifies how long Content Gateway continues a background fill before giving up and dropping the origin server connection.
<code>proxy.config.http.background_fill_completed_threshold</code>	FLOAT	0.50000	Specifies 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	Default Value	Description
<code>proxy.config.http.connect_attempts_max_retries</code>	INT	6	Specifies the maximum number of connection retries Content Gateway makes when the origin server is not responding.
<code>proxy.config.http.connect_attempts_max_retries_dead_server</code>	INT	2	Specifies the maximum number of connection retries Content Gateway makes when the origin server is unavailable.
<code>proxy.config.http.connect_attempts_rr_retries</code>	INT	2	Specifies 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.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.connect_attempts_timeout</code>	INT	60	Specifies the timeout value in seconds for an origin server connection.
<code>proxy.config.http.down_server.cache_time</code>	INT	30	Specifies how long in seconds Content Gateway remembers that an origin server was unreachable.
<code>proxy.config.http.down_server.abort_threshold</code>	INT	10	Specifies 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.

Congestion control

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.congestion_control.enabled</code>	INT	0	Specifies whether Content Gateway applies congestion control. Congestion control is performed based on rules defined in the congestion.config file. See congestion.config , page 381. 1 = enabled 0 = disabled

Negative response caching

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.negative_caching_enabled</code>	INT	0	<p>When enabled (1), Content Gateway caches negative responses, such as <i>404 Not Found</i>, 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.</p> <p>Content Gateway caches the following negative responses:</p> <ul style="list-style-type: none"> 204 No Content 305 Use Proxy 400 Bad Request 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
<code>proxy.config.http.negative_caching_lifetime</code>	INT	1800	Specifies how long Content Gateway keeps the negative responses as valid in cache.

Proxy users variables

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.anonymize_remove_from</code>	INT	0	When enabled (1), Content Gateway removes the From header that accompanies transactions to protect the privacy of your users.
<code>proxy.config.http.anonymize_remove_referer</code>	INT	0	When enabled (1), Content Gateway removes the Referer header that accompanies transactions to protect the privacy of your site and users.

Configuration Variable	Data Type	Default Value	Description
proxy.config.http.anonymize_remove_user_agent	INT	0	When enabled (1), Content Gateway removes the User-agent header that accompanies transactions to protect the privacy of your site and users.
proxy.config.http.anonymize_remove_cookie	INT	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	1	When enabled (1), Content Gateway removes Client-IP headers for more privacy.
proxy.config.http.anonymize_insert_client_ip	INT	0	When enabled (1), Content Gateway inserts Client-IP headers to retain the client's IP address.
proxy.config.http.append_xforwards_header	INT	0	When enabled (1), Content Gateway appends X-Forwards headers to outgoing requests.
proxy.config.http.anonymize_other_header_list	STRING	NULL	Specifies the headers that Content Gateway will remove from outgoing requests.
proxy.config.http.snarf_username_from_authorization	INT	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	0	When enabled (1), Content Gateway adds the client IP address to the X-Forwarded-For header.

Security

Configuration Variable	Data Type	Default Value	Description
proxy.config.http.push_method_enabled	INT	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.

Cache control

Configuration Variable	Data Type	Default Value	Description
proxy.config.http.cache.http	INT	1	Enables (1) or disables (0) caching of HTTP requests.
proxy.config.http.cache.ftp	INT	1	Enables (1) or disables (0) caching of FTP requests sent via HTTP.
proxy.config.http.cache.ignore_client_no_cache	INT	0	When enabled (1), Content Gateway ignores client requests to bypass the cache.
proxy.config.http.cache.ims_on_client_no_cache	INT	0	When enabled (1), Content Gateway issues a conditional request to the origin server if an incoming request has a no-cache header.
proxy.config.http.cache.ignore_server_no_cache	INT	0	When enabled (1), Content Gateway ignores origin server requests to bypass the cache.
proxy.config.http.cache.cache_responses_to_cookies	INT	3	Specifies how cookies are cached: <ul style="list-style-type: none"> • 0 = do not cache any responses to cookies • 1 = cache for any content-type • 2 = cache only for image types • 3 = cache for all but text content-types

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.cache.ignore_authentication</code>	INT	0	When enabled (1), Content Gateway ignores WWW-Authentication headers in responses. WWW-Authentication headers are removed and not cached.
<code>proxy.config.http.cache.cache_urls_that_look_dynamic</code>	INT	0	Enables (1) or disables (0) caching of URLs that look dynamic.
<code>proxy.config.http.cache.enable_default_vary_headers</code>	INT	0	Enables (1) or disables (0) caching of alternate versions of HTTP objects that do not contain the Vary header.
<code>proxy.config.http.cache.when_to_revalidate</code>	INT	0	Specifies when to revalidate content: <ul style="list-style-type: none"> • 0 = Use cache directives or heuristic (the default value). • 1 = Stale if heuristic. • 2 = Always stale (always revalidate). • 3 = Never stale. • 4 = Use cache directives or heuristic (0) unless the request has an If-Modified-Since header. If the request has an If-Modified-Since header, Content Gateway always revalidates the cached content and uses the client's If-Modified-Since header for the proxy request.
<code>proxy.config.http.cache.when_to_add_no_cache_to_msie_requests</code>	INT	0	Specifies when to add no-cache directives to Microsoft Internet Explorer requests. You can specify the following: <ul style="list-style-type: none"> • 0 = no-cache not added to MSIE requests. • 1 = no-cache added to IMS MSIE requests. • 2 = no-cache added to all MSIE requests.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.cache.required_headers</code>	INT	0	Specifies the type of headers required in a request for the request to be cacheable. <ul style="list-style-type: none"> • 0 = no required headers to make document cacheable. • 1 = at least Last-Modified header required. • 2 = explicit lifetime required, Expires or Cache-Control.
<code>proxy.config.http.cache.max_stale_age</code>	INT	604800	Specifies the maximum age allowed for a stale response before it cannot be cached.
<code>proxy.config.http.cache.range.lookup</code>	INT	1	When enabled (1), Content Gateway looks up range requests in the cache.

Heuristic expiration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.cache.heuristic_min_lifetime</code>	INT	3600	Specifies the minimum amount of time that a document in the cache can be considered fresh.
<code>proxy.config.http.cache.heuristic_max_lifetime</code>	INT	86400	Specifies the maximum amount of time that a document in the cache can be considered fresh.
<code>proxy.config.http.cache.heuristic_lm_factor</code>	FLOAT	0.10000	Specifies the aging factor for freshness computations.
<code>proxy.config.http.cache.fuzz.time</code>	INT	240	Specifies the interval in seconds before the document stale time that the proxy checks for an early refresh.
<code>proxy.config.http.cache.fuzz.probability</code>	FLOAT	0.00500	Specifies the probability that a refresh is made on a document during the specified fuzz time.

Dynamic content and content negotiation

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.cache.vary_default_text</code>	STRING	NULL	Specifies the header on which Content Gateway varies for text documents; for example, if you specify user-agent , the proxy caches all the different user-agent versions of documents it encounters.
<code>proxy.config.http.cache.vary_default_images</code>	STRING	NULL	Specifies the header on which Content Gateway varies for images.
<code>proxy.config.http.cache.vary_default_other</code>	STRING	NULL	Specifies the header on which Content Gateway varies for anything other than text and images.

Anonymous FTP password

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.ftp.anonymous_passwd</code>	STRING	<i>the value of the administrator's email as supplied during installation</i>	Specifies 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.

Cached FTP document lifetime

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.ftp.cache.document_lifetime</code>	INT	259200	Specifies the maximum amount of time that an FTP document can stay in the cache.

FTP transfer mode

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.http.ftp.binary_transfer_only</code>	INT	0	<p>When enabled (1), all FTP documents requested from HTTP clients are transferred in binary mode only.</p> <p>When disabled (0), FTP documents requested from HTTP clients are transferred in ASCII or binary mode, depending on the document type.</p>

Customizable user response pages

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.body_factory.enable_customizations</code>	INT	0	<p>Specifies whether customizable response pages are enabled or disabled and which response pages are used:</p> <ul style="list-style-type: none"> • 0 = disable customizable user response pages • 1 = enable customizable user response pages in the default directory only • 2 = enable language-targeted user response pages
<code>proxy.config.body_factory.enable_logging</code>	INT	0	<p>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.</p>
<code>proxy.config.body_factory.template_sets_dir</code>	STRING	<code>config/body_factory</code>	<p>Specifies the customizable response page default directory.</p>
<code>proxy.config.body_factory.response_suppression_mode</code>	INT	0	<p>Specifies when Content Gateway suppresses generated response pages:</p> <ul style="list-style-type: none"> • 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	Default Value	Description
FTP over HTTP			
proxy.config.ftp.data_connection_mode	INT	1	Specifies the FTP connection mode: <ul style="list-style-type: none"> • 1 = PASV then PORT • 2 = PORT only • 3 = PASV only
proxy.config.ftp.control_connection_timeout	INT	300	Specifies how long Content Gateway waits for a response from the FTP server.
proxy.config.ftp.rc_to_switch_to_PORT	STRING	NULL	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	0	Enables (1) or disables (0) processing of FTP requests from FTP clients.
proxy.config.ftp.cache_enabled	INT	1	Enables (1) or disables (0) FTP documents to be put in the cache. If this option is disabled, Content Gateway always serves FTP documents from the FTP server. Note: Only FTP objects retrieved anonymously (user = anonymous) are cached.
proxy.config.ftp.logging_enabled	INT	1	Enables (1) or disables (0) logging of FTP transactions.
proxy.config.ftp.proxy_server_port	INT	2121	Specifies the port used for FTP connections.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ftp.open_lisn_port_mode</code>	INT	1	<p>Specifies how FTP opens a listening port for a data transfer:</p> <ul style="list-style-type: none"> • 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.
<code>proxy.config.ftp.min_lisn_port</code>	INT	32768	<p>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.</p>
<code>proxy.config.ftp.max_lisn_port</code>	INT	65535	<p>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.</p>
<code>proxy.config.ftp.server_data_default_pasv</code>	INT	1	<p>Specifies the default method used to set up server side data connections:</p> <ul style="list-style-type: none"> • 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).

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ftp.different_client_port_ip_allowed</code>	INT	0	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.
<code>proxy.config.ftp.try_pasv_times</code>	INT	1024	Specifies the number of times Content Gateway can try to open a listening port when the FTP client sends a PASV.
<code>proxy.config.ftp.try_port_times</code>	INT	1024	Specifies the maximum number of times Content Gateway can try to open a listening port when sending a PORT to the FTP server.
<code>proxy.config.ftp.try_server_ctrl_connect_times</code>	INT	6	Specifies the maximum number of times Content Gateway can try to connect to the FTP server's control listening port.
<code>proxy.config.ftp.try_server_data_connect_times</code>	INT	3	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.
<code>proxy.config.ftp.try_client_data_connect_times</code>	INT	3	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.
<code>proxy.config.ftp.client_ctrl_no_activity_timeout</code>	INT	900	Specifies the inactivity timeout, in seconds, for the FTP client control connection.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ftp.client_ctrl_active_timeout</code>	INT	14400	Specifies the active timeout, in seconds, for the FTP client control connection.
<code>proxy.config.ftp.server_ctrl_no_activity_timeout</code>	INT	120	Specifies the inactivity timeout, in seconds, for the FTP server control connection.
<code>proxy.config.ftp.server_ctrl_active_timeout</code>	INT	14400	Specifies the active timeout, in seconds, for the FTP server control connection.
<code>proxy.config.ftp.client_data_no_activity_timeout</code>	INT	120	Specifies the maximum time, in seconds, that a client FTP data transfer connection can be idle before it is aborted.
<code>proxy.config.ftp.client_data_active_timeout</code>	INT	14400	Specifies the maximum time, in seconds, of an FTP data transfer connection from a client.
<code>proxy.config.ftp.server_data_no_activity_timeout</code>	INT	120	Specifies the maximum time, in seconds, that a server FTP data transfer connection can be idle before it is aborted.
<code>proxy.config.ftp.server_data_active_timeout</code>	INT	14400	Specifies the maximum time, in seconds, of an FTP data transfer connection from a server.
<code>proxy.config.ftp.pasv_accept_timeout</code>	INT	120	Specifies the timeout value for a listening data port in Content Gateway (for PASV, the client data connection).
<code>proxy.config.ftp.port_accept_timeout</code>	INT	120	Specifies the timeout value for a listening data port in Content Gateway (for PORT, the server data connection).
<code>proxy.config.ftp.share_ftp_server_ctrl_enabled</code>	INT	1	Enables (1) or disables (0) sharing the server control connections among multiple anonymous FTP clients.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ftp.share_only_after_session_end</code>	INT	1	<p>Specifies how an FTP server control connection is shared between different FTP client sessions:</p> <ul style="list-style-type: none"> • 1 = the FTP server control connection can be used by another FTP client session <i>only</i> 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 <i>only</i> if the FTP client session is not actively using the FTP server connection: for example, if the request is a cache hit or during an idle session.
<code>proxy.config.ftp.server_ctrl_keep_alive_no_activity_timeout</code>	INT	90	Specifies the timeout value when the FTP server control connection is not used by any FTP clients.
<code>proxy.config.ftp.reverse_ftp_enabled</code>	INT	0	Not supported.
<code>proxy.config.ftp.login_info_fresh_in_cache_time</code>	INT	604800	Specifies how long the 220/230 responses (login messages) can stay fresh in the cache.
<code>proxy.config.ftp.directory_listing_fresh_in_cache_time</code>	INT	86400	Specifies how long directory listings can stay fresh in the cache.
<code>proxy.config.ftp.file_fresh_in_cache_time</code>	INT	259200	Specifies how long FTP files can stay fresh in the cache.
<code>proxy.config.ftp.simple_directory_listing_cache_enabled</code>	INT	1	Enables (1) or disables (0) caching of directory listings without arguments: for example, dir/ls .
<code>proxy.config.ftp.full_directory_listing_cache_enabled</code>	INT	1	Enables (1) or disables (0) caching of directory listings with arguments: for example, <code>ls -al</code> , <code>ls *.txt</code> .

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.ftp.file_fresh_mdtm_checking_enabled</code>	INT	0	When enabled (1), before serving the file from cache, check with the FTP server to confirm that file in cache is the most recent.
<code>proxy.config.ftp.data_source_port_20_enabled</code>	INT	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	Default Value	Description
<code>proxy.config.socks.socks_needed</code>	INT	0	Enables (1) or disables (0) the SOCKS option. See Configuring SOCKS firewall integration , page 146.
<code>proxy.config.socks.socks_version</code>	INT	4	Specifies the SOCKS version.
<code>proxy.config.socks.default_servers</code>	STRING	<code>s1.example.com:1080;socks2:4080</code>	Specifies the names and ports of the SOCKS servers with which Content Gateway communicates.
<code>proxy.config.socks.accept_enabled</code>	INT	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.
<code>proxy.config.socks.accept_port</code>	INT	1080	Specifies the port on which Content Gateway accepts SOCKS traffic.

Net subsystem

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.net.connections_throttle</code>	INT	45000	Specifies 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	Default Value	Description
<code>proxy.config.cluster.cluster_port</code>	INT	8086	Specifies the port used for cluster communication.
<code>proxy.config.cluster.ethernet_interface</code>	STRING	<i>your_interface</i>	Specifies the network interface used for cluster traffic. All nodes in a cluster must use the same network interface.

Cache

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.cache.permit.pinning</code>	INT	0	Enables (1) or disables (0) the cache pinning option, which lets you keep objects in the cache for a specified time. You set cache pinning rules in the cache.config file (see cache.config , page 294).
<code>proxy.config.cache.ram_cache.size</code>	INT	-1	Specifies the size of the RAM cache, in bytes. -1 means that the RAM cache is automatically sized at approximately 1 MB per GB of disk.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.cache.limits.http.max_alts</code>	INT	3	Specifies the maximum number of HTTP alternates that Content Gateway can cache.
<code>proxy.config.cache.max_doc_size</code>	INT	0	Specifies the maximum size of documents in the cache (in bytes): 0 = there is no size limit.

DNS

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.dns.search_default_domains</code>	INT	1	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 .
<code>proxy.config.dns.splitDNS.enabled</code>	INT	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, page 149
<code>proxy.config.dns.splitdns.def_domain</code>	STRING	NULL	Specifies 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.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.dns.url_expansions</code>	STRING	NULL	Specifies a list of hostname extensions that are 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 (.).) Note: If the variable proxy.config.http.enable_url_expandomatic is set to 1 (the default value), 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.
<code>proxy.config.dns.lookup_timeout</code>	INT	20	Specifies the DNS lookup timeout duration in seconds. When the timeout period expires, the lookup attempt is terminated.
<code>proxy.config.dns.retries</code>	INT	5	Specifies the number of times a DNS lookup is retried before giving up.

HostDB

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.hostdb.size</code>	INT	200000	Specifies the maximum number of entries allowed in the host database.
<code>proxy.config.hostdb.ttl_mode</code>	INT	0	Specifies the host database time to live mode. You can specify one of the following: <ul style="list-style-type: none"> • 0 = obey • 1 = ignore • 2 = min(X,ttl) • 3 = max(X,ttl)
<code>proxy.config.hostdb.timeout</code>	INT	1440	Specifies the foreground timeout, in minutes.
<code>proxy.config.hostdb.strict_round_robin</code>	INT	0	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	Default Value	Description
<code>proxy.config.log2.logging_enabled</code>	INT	1	Enables and disables event logging: <ul style="list-style-type: none"> • 0 = logging disabled • 1 = log errors only • 2 = log transactions only • 3 = full logging (errors + transactions) See Working With Log Files , page 167.
<code>proxy.config.log2.max_secs_per_buffer</code>	INT	5	Specifies the maximum amount of time before data in the buffer is flushed to disk.
<code>proxy.config.log2.max_space_mb_for_logs</code>	INT	20480	Specifies the amount of space allocated to the logging directory, in megabytes.
<code>proxy.config.log2.max_space_mb_for_orphan_logs</code>	INT	25	Specifies the amount of space allocated to the logging directory, in megabytes, if this node is acting as a collation client.
<code>proxy.config.log2.max_space_mb_headroom</code>	INT	100	Specifies 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.
<code>proxy.config.log2.hostname</code>	STRING	localhost	Specifies the hostname of the machine running Content Gateway.
<code>proxy.config.log2.logfile_dir</code>	STRING	<code>install_dir/logs</code>	Specifies the full path to the logging directory.

Configuration Variable	Data Type	Default Value	Description
proxy.config.log2.logfile_perm	STRING	rw-r--r--	<p>Specifies the log file permissions. The standard UNIX file permissions are used (owner, group, other). Valid values are:</p> <ul style="list-style-type: none"> • - = no permission • r = read permission • w = write permission • x = execute permission <p>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.</p> <p>Permissions for existing log files are not changed when the configuration is changed.</p> <p>Linux only.</p>
proxy.config.log2.custom_logs_enabled	INT	0	<p>When enabled (1), supports the definition and generation of custom log files according to the specifications in logs_xml.config. See logs_xml.config, page 306.</p>
proxy.config.log2.xml_logs_config	INT	1	<p>Specifies the size, in MB, which when reached causes the log files to roll. See Rolling event log files, page 176.</p>
proxy.config.log2.squid_log_enabled	INT	0	<p>Enables (1) or disables (0) the squid log file format.</p>
proxy.config.log2.squid_log_is_ascii	INT	1	<p>Specifies the squid log file type:</p> <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary
proxy.config.log2.squid_log_name	STRING	squid	<p>Specifies the squid log filename.</p>
proxy.config.log2.squid_log_header	STRING	NULL	<p>Specifies the squid log file header text.</p>
proxy.config.log2.common_log_enabled	INT	0	<p>Enables (1) or disables (0) the Netscape common log file format.</p>
proxy.config.log2.common_log_is_ascii	INT	1	<p>Specifies the Netscape common log file type:</p> <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary

Configuration Variable	Data Type	Default Value	Description
proxy.config.log2.common_log_name	STRING	common	Specifies the Netscape common log filename.
proxy.config.log2.common_log_header	STRING	NULL	Specifies the Netscape common log file header text.
proxy.config.log2.extended_log_enabled	INT	1	Enables (1) or disables (0) the Netscape extended log file format.
proxy.config.log2.extended_log_is_ascii	INT	1	Specifies the Netscape extended log file type: <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary
proxy.config.log2.extended_log_name	STRING	extended	Specifies the Netscape extended log filename.
proxy.config.log2.extended_log_header	STRING	NULL	Specifies the Netscape extended log file header text.
proxy.config.log2.extended2_log_enabled	INT	0	Enables (1) or disables (0) the Netscape Extended-2 log file format.
proxy.config.log2.extended2_log_is_ascii	INT	1	Specifies the Netscape Extended-2 log file type: <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary
proxy.config.log2.extended2_log_name	STRING	extended2	Specifies the Netscape Extended-2 log filename.
proxy.config.log2.extended2_log_header	STRING	NULL	Specifies the Netscape Extended-2 log file header text.
proxy.config.log2.separate_icp_logs	INT	0	When enabled (1), configures Content Gateway to store ICP transactions in a separate log file.
proxy.config.log2.separate_host_logs	INT	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 HTTP host log splitting , page 180).

Configuration Variable	Data Type	Default Value	Description
<code>proxy.local.log2.collation_mode</code>	INT	0	Specifies the log collation mode: <ul style="list-style-type: none"> • 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 on sending XML-based custom formats to the collation server, see logs_xml.config , page 306.
<code>proxy.config.log2.collation_host</code>	STRING	NULL	Specifies the hostname of the log collation server.
<code>proxy.config.log2.collation_port</code>	INT	8085	Specifies the port used for communication between the collation server and client.
<code>proxy.config.log2.collation_secret</code>	STRING	foobar	Specifies the password used to validate logging data and prevent the exchange of unauthorized information when a collation server is being used.
<code>proxy.config.log2.collation_host_tagged</code>	INT	0	When enabled (1), configures Content Gateway to include the hostname of the collation client that generated the log entry in each entry.
<code>proxy.config.log2.collation_retry_sec</code>	INT	5	Specifies the number of seconds between collation server connection retries.
<code>proxy.config.log2.rolling_enabled</code>	INT	1	Enables (1) or disables (0) log file rolling. See Rolling event log files , page 176.
<code>proxy.config.log2.rolling_interval_sec</code>	INT	21600	Specifies the log file rolling interval, in seconds. The minimum value is 300 (5 minutes). The maximum value is 86400 seconds (one day).
<code>proxy.config.log2.rolling_offset_hr</code>	INT	0	Specifies the file rolling offset hour. The hour of the day that starts the log rolling period.
<code>proxy.config.log2.rolling_size_mb</code>	INT	10	Specifies the size, in megabytes, which when reached causes the current file to be closed and a new file to be created.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.log2.auto_delete_rolled_files</code>	INT	1	Enables (1) or disables (0) automatic deletion of rolled files.
<code>proxy.config.log2.sampling_frequency</code>	INT	1	Configures Content Gateway to log only a sample of transactions rather than every transaction. You can specify the following values: <ul style="list-style-type: none">• 1 = log every transaction• 2 = log every second transaction• 3 = log every third transaction and so on...

URL remap rules

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.url_remap.default_to_server_pac</code>	INT	0	<p>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.</p> <p>For this type of redirection to work, the variable <code>proxy.config.reverse_proxy.enabled</code> must be set to 1.</p>
<code>proxy.config.url_remap.default_to_server_pac_port</code>	INT	-1	<p>Sets the PAC port so that PAC requests made to the Content Gateway proxy service port are redirected to this port.</p> <p>-1 specifies that the PAC port will be set to the autoconfiguration port (the default autoconfiguration port is 8083). This is the default setting.</p> <p>This variable can be used together with the <code>proxy.config.url_remap.default_to_server_pac</code> variable 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.</p>
<code>proxy.config.url_remap.remap_required</code>	INT	0	<p>Set this variable to 1 if you want Content Gateway to serve requests only from origin servers listed in the mapping rules of the <code>remap.config</code> file. If a request does not match, the browser will receive an error.</p>
<code>proxy.config.url_remap.pristine_host_hdr</code>	INT	0	<p>Set this variable to 1 if you want to retain the client host header in a request during remapping.</p>

ICP configuration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.icp.enabled</code>	INT	0	Sets ICP mode for hierarchical caching: <ul style="list-style-type: none"> • 0 = disables ICP. • 1 = allows Content Gateway to receive ICP queries only. • 2 = allows Content Gateway to send and receive ICP queries. See ICP cache hierarchies , page 75.
<code>proxy.config.icp.icp_interface</code>	STRING	<i>your_interface</i>	Specifies the network interface used for ICP traffic.
<code>proxy.config.icp.icp_port</code>	INT	3130	Specifies the UDP port that you want to use for ICP messages.
<code>proxy.config.icp.multicast_enabled</code>	INT	0	Enables (1) or disables (0) ICP multicast.
<code>proxy.config.icp.query_timeout</code>	INT	2	Specifies the timeout used for ICP queries.

Scheduled update configuration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.update.enabled</code>	INT	0	Enables (1) or disables (0) the Scheduled Update option.
<code>proxy.config.update.force</code>	INT	0	Enables (1) or disables (0) a force immediate update. When enabled, Content Gateway overrides the scheduling expiration time for all scheduled update entries and initiates updates until this option is disabled.
<code>proxy.config.update.retry_count</code>	INT	10	Specifies the number of times Content Gateway retries the scheduled update of a URL in the event of failure.

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.update.retry_interval</code>	INT	2	Specifies the delay in seconds between each scheduled update retry for a URL in the event of failure.
<code>proxy.config.update.concurrent_updates</code>	INT	100	Specifies the maximum simultaneous update requests allowed at any time. This option prevents the scheduled update process from overburdening the host.

WCCP configuration

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.wccp.enabled</code>	INT	0	Enables (1) or disables (0) WCCP.
<code>proxy.config.wccp.version</code>	INT	1	Specifies the version of WCCP being used: <ul style="list-style-type: none"> • 1 = Version 1.0 • 2 = Version 2.0

WCCP 1.0 variables

<code>proxy.config.wccp.router_ip</code>	STRING	NULL	Specifies the IP address of the router sending traffic to Content Gateway.
<code>proxy.config.wccp.ethernet_interface</code>	STRING	<i>your_interface</i>	Specifies the Ethernet interface used to talk to the WCCP v1 router. Note: The Content Gateway installation script detects your Ethernet interface and sets this variable appropriately. If your system has multiple network interfaces, check that this variable specifies the correct interface.

WCCP 2.0 variables

<code>proxy.config.wccp2.security_enabled</code>	INT	0	Enables (1) or disables (0) security so that the router and the proxy can authenticate each other. (If you enable security in Content Gateway, you must also enable security on the router. See your Cisco router documentation.)
--	-----	---	---

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.wccp2.password</code>	STRING	NULL	Specifies the password used for authentication. This must be the same password configured on the router. It must be at least seven characters long.
<code>proxy.config.wccp2.multicast_enabled</code>	INT	0	Enables (1) or disables (0) multicast mode.
<code>proxy.config.wccp2.multicast_address</code>	STRING	NULL	Specifies the IP multicast address.
<code>proxy.config.wccp2.multicast_ttl</code>	INT	1	Specifies the number of hops multicast traffic must travel to reach its destination. Assign a value to this variable if you anticipate that traffic will travel more than 3 hops. In addition, assign a value to this variable if your WCCP router is not directly connected to Content Gateway. The default is 1.
<code>proxy.config.wccp2.rev_encapsulation</code>	INT	0	Enables (1) or disables (0) packet encapsulation mode, which enables Content Gateway to send encapsulated returned (bypassed) packets to the router.
<code>proxy.config.wccp2.fwd_method</code>	INT	1	Specifies the preferred forward method. If the device doesn't support the specified forward method, Content Gateway negotiates to use a supported method. 1 = GRE 2 = L2
<code>proxy.config.wccp2.assign_method</code>	INT	1	Specifies the preferred assignment method. 1 = HASH 2 = MASK
<code>proxy.config.wccp2.return_method</code>	INT	1	Specifies the preferred return method. If the device doesn't support the specified return method, Content Gateway negotiates to use a supported method. 1 = GRE 2 = L2

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.wccp2.number_of_routers</code>	INT	0	If multicast is not enabled, the routers on your network are not automatically discovered. You must specify the number of routers that direct traffic to Content Gateway. WCCP v2 supports a maximum of 32 routers.
<code>proxy.config.wccp2.router0_ip</code> ... <code>proxy.config.wccp2.router10_ip</code>	STRING	NULL	If multicast is not enabled, the routers on your network are not automatically discovered. You must specify the IP address of each router that directs traffic to the Content Gateway server.
<code>proxy.config.wccp2.svc_HTTP</code>	INT	1	When enabled (1), a WCCP v2 router can route HTTP requests to Content Gateway.
<code>proxy.config.wccp2.svc_FTP</code>	INT	0	When enabled (1), a WCCP v2 router can route FTP requests from FTP clients to Content Gateway.
<code>proxy.config.wccp2.svc SOCKS</code>	INT	0	When enabled (1), a WCCP v2 router can route SOCKS traffic to Content Gateway.
<code>proxy.config.wccp2.svc_DNS</code>	INT	0	When enabled (1), a WCCP v2 router can route DNS requests to Content Gateway.
<code>proxy.config.wccp2.svc_HTTPS</code>	INT	0	When enabled (1), enables the HTTPS service group.
<code>proxy.local.wccp2.ethernet_interface</code>	STRING	<i>your_interface</i>	Specifies the ethernet interface used to talk to the WCCP v2 router. Note: The Content Gateway installation script detects your Ethernet interface and sets this variable appropriately. If your system has multiple network interfaces, check that this variable specifies the correct interface.

ARM (security configuration)

Configuration Variable	Data Type	Default Value	Description
proxy.config.arm.security_enabled	INT	0	Enables (1) or disables (0) ARM security. See Controlling host access to the proxy server , page 140.

SSL Decryption



Note

All SSL decryption setting should be made in the Content Gateway Manager. None of the variables in the table below should be modified directly in records.config.

Configuration Variable	Data Type	Default Value	Description
proxy.config.ssl_decryption.use_decryption	INT	0	When enabled (1), Content Gateway performs SSL decryption.
proxy.config.ssl_decryption_ports	INT	443	Specifies the HTTPS ports. Content Gateway allows SSL decryption and policy lookup only to the specified ports.
proxy.config.ssl_decryption.ui_enabled	INT	0	When enabled (1), the SSL configuration tab is displayed in the Content Gateway Manager.
proxy.config.ssl_management_port	INT	8071	The management port on which the SSL Manager listens.
proxy.config.ssl_inbound_port	INT	8070	The port on which SSL Manager listens for inbound (client-facing) traffic.
proxy.config.ssl_outbound_port	INT	8090	The port SSL Manager uses for outbound (Internet-facing) traffic.
proxy.config.ssl_outbound_ip	STRING	127.0.0.1	The IP address of the SSL Manager inbound and outbound proxy.

Configuration Variable	Data Type	Default Value	Description
proxy.config.ssl_forward_to_inbound	INT	1	Do not change. When SSL Manager is enabled, causes SSL traffic to be forwarded to the correct proxy port.
proxy.config.administrator_id	STRING	NULL	Do not change. Holds the encrypted administrator ID. The variable is used by SSL Manager.

ICAP

Configuration Variable	Data Type	Default Value	Description
proxy.config.icap.enabled	INT	0	Enables (1) or disables (0) ICAP support with Websense Data Security Suite (DSS). See Working With Websense Data Security, page 105 .
proxy.config.icap.ICAPUri	STRING	NULL	Specifies the Uniform Resource Identifier for the ICAP service. Obtain the identifier from your DSS administrator. Enter the URI in the following format: <pre>icap://hostname:port/path</pre> For <i>hostname</i> , enter the IP address or hostname of the DSS Protector appliance. The default ICAP port is 1344. <i>Path</i> is the path of the ICAP service on the host machine. For example: <pre>icap://ICAP_machine:1344/opt/icap_services</pre> You do not need to specify the port if you are using the default ICAP port 1344.
proxy.config.icap.FailOpen	INT	1	Set to: <ul style="list-style-type: none"> • 0 to send a block page if Content Gateway encounters an error while communicating with Data Security Suite. • 1 to allow traffic

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.icap.BlockHugeContent</code>	INT	0	Set to: <ul style="list-style-type: none"> • 0 to send a block page if a file larger than the size limit specified in DSS is sent. The default size limit in DSS is 12 MB. • 1 to allow traffic
<code>proxy.config.icap.AnalyzeSecureContent</code>	INT	1	Set to: <ul style="list-style-type: none"> • 0 if decrypted traffic should be sent directly to its destination. • 1 if decrypted traffic should be sent to Data Security Suite for analysis.
<code>proxy.config.icap.AnalyzeFTP</code>	INT	1	When enabled (1), send native FTP upload file transfers to ICAP server for analysis.

Data Security

Configuration Variable	Data Type	Default Value	Description
<code>proxy.config.dss.enabled</code>	INT	0	Enables (1) or disables (0) support for on-box Data Security. See Working With Websense Data Security , page 105.
<code>proxy.config.dss.AnalyzeFTP</code>	INT	1	When enabled (1), send native FTP upload file transfers to the on-box Data Security policy engine for analysis.
<code>proxy.config.dss.AnalyzeSecureContent</code>	INT	1	Set to: <ul style="list-style-type: none"> • 0 if decrypted traffic should be sent directly to its destination. • 1 if decrypted traffic should be sent to Data Security Suite for analysis.
<code>proxy.config.dss.analysis_timeout</code>	INT	10000	Specifies the maximum length of time, in milliseconds, that a single file analysis can take before analysis is aborted.

Connectivity, analysis, and boundary conditions

Configuration Variable	Data Type	Default Value	Description
wtg.config.subscription_key	STRING	NULL	Holds the Websense Security Gateway or Websense Security Gateway Anywhere subscription key value.
wtg.config.download_server_ip	STRING	download.websense.com	Holds the hostname or IP address of the Websense download server.
wtg.config.download_server_port	INT	80	Holds the port number of the Websense download server.
wtg.config.policy_server_ip	STRING		Holds the IP address of the Websense Policy Server.
wtg.config.policy_server_port	INT	55806	Holds the port number of the Websense Policy Server.
wtg.config.wse_server_ip	STRING		Holds the IP address of the Websense Filtering Service.
wtg.config.wse_server_port	INT	15868	Holds the port number of the Websense Filtering Service WISP interface.
wtg.config.ssl_bypassed_categories	STRING	NULL	This variable takes 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 Manager to specify categories to bypass SSL decryption.
wtg.config.ssl_decryption_bypass_ip_based	INT	0	Specifies that the SSL category bypass process use only the IP address (not the hostname) when performing a category lookup. 0 = disabled 1 = enabled
wtg.config.fail_open	INT	1	Specifies whether Content Gateway will permit or block the request when Websense Web filtering (Filtering Service) is unavailable. Set to: <ul style="list-style-type: none"> • 0 to send a block page • 1 to permit the request

Configuration Variable	Data Type	Default Value	Description
<code>wtg.config.fail_open_low_memory</code>	INT	0	Specifies whether Content Gateway will permit or block traffic when there is a system low memory condition. Set to: <ul style="list-style-type: none"> • 0 to block traffic • 1 to permit traffic, bypassing scanning
<code>wtg.config.archive_depth</code>	INT	5	Specifies the maximum depth of analysis performed on archive files.
<code>wtg.config.max_decompressions</code>	INT	10	Specifies the maximum number of total decompressions to be performed on archive files (per transaction). The value should not exceed 25.
<code>wtg.config.max_subsamples</code>	INT	10000	Specifies the maximum number of discrete files within an archive file that Content Gateway may decompress and analyze to classify a given transaction.
<code>wtg.config.zipbomb_action</code>	INT	1	For internal use. Indicates zip bomb analysis status. Do not change the value of this variable.
<code>wtg.config.min_mem_avail</code>	INT	100	Specifies in megabytes, the minimum amount of available memory which when reached causes Content Gateway to request more memory from the system.
<code>wtg.config.max_mem_allowed</code>	INT	0	Specifies in megabytes, the maximum amount of memory, which when consumed, causes Content Gateway to perform more extensive memory monitoring.

remap.config

The `remap.config` file contains mapping rules that Websense 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, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 space-delimited fields: type, target, and replacement. The following table describes the format of each field.

Field	Description
<i>type</i>	Enter one of the following: <ul style="list-style-type: none"> <code>redirect</code>—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. <code>redirect_temporary</code>—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).
<i>target</i>	Enter the origin or <i>from</i> URL. You can enter up to four components: <pre>scheme://host:port/path_prefix</pre> <i>scheme</i> can be <code>http</code> , <code>https</code> , or <code>ftp</code> .
<i>replacement</i>	Enter the destination or <i>to</i> URL. You can enter up to four components: <pre>scheme://host:port/path_prefix</pre> <i>scheme</i> can be <code>http</code> , <code>https</code> , or <code>ftp</code> .



Note

The scheme type (HTTP, HTTPS, FTP) of the target and replacement must match.

Examples

The following section shows example mapping rules in the **remap.config** file.

Redirect mapping rules

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 the following information:

- ◆ The SOCKS servers that Websense Content Gateway passes through to access specific origin servers, and the order in which the proxy goes through the SOCKS server list



Note

You can specify your *default* SOCKS servers in the Content Gateway Manager or by editing the configuration variable:

```
proxy.config.socks.socks.default_servers
```

However, the **socks.config** file lets you perform additional SOCKS configuration; you can send requests to specific origin servers through specific SOCKS servers.

- ◆ The origin servers you want Content Gateway to access directly *without* going through the SOCKS server
- ◆ The user name and password that Content Gateway uses to connect to a SOCKS server (SOCKS version 5 only)



Important

After you modify this file, you must restart the proxy.

Format

To specify the SOCKS servers Content Gateway must go through to reach specific origin servers, you must add a rule to the **socks.config** file with the following format:

```
dest_ip=ipaddress parent=server_name:port
[round_robin=value]
```

where:

ipaddress is the origin server IP address or range of IP addresses separated by – or /.

server_name is the hostname of the SOCKS server.

port is the port number through which the proxy communicates with the SOCKS server.

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 the origin servers you want Content Gateway to access directly *without* going through the SOCKS server, enter a rule in the `socks.config` file in the following format:

```
no_socks ipaddress
```

where *ipaddress* is a comma-separated list of the IP addresses or IP address ranges associated with the origin servers you want Content Gateway to access directly. Do not specify the all networks broadcast address: 255.255.255.255.

To specify the user name and password Content Gateway uses for authentication with the SOCKS version 5 server, enter a rule in the `socks.config` file in the following format:

```
auth u username password
```

where *username* is the user name, and *password* is the password used for authentication.

**Note**

Each rule in the `socks.config` file 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 `socks1` on port 1080 and `socks2` on port 4080. Because the optional specifier `round_robin` is set to `strict`, the proxy sends the first request to `socks1`, the second request to `socks2`, the third request to `socks1`, and so on.

```
dest_ip=123.14.15.1 - 123.14.17.4
parent=socks1:1080;socks2:4080 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
```

The following example configures Content Gateway to use the user name `content_gateway` and the password `secret` for authentication with the SOCKS version 5 server:

```
auth u content_gateway secret
```

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](#), page 149.



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. 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 **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
<i>dest_domain</i>	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.
<i>dest_host</i>	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.
<i>url_regex</i>	A valid URL regular expression. This specifies that the DNS server selection be based on a regular expression.
<i>dns_server</i>	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.
<i>def_domain</i>	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 .
<i>search_list</i>	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 .

Examples

Consider the following DNS server selection specifications:

```
dest_domain=internal.company.com named=255.255.255.255:212
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**.

storage.config

The **storage.config** file lists all the files, directories, or hard disk partitions that make up the cache.

**Important**

After you modify this file, you must restart the proxy.

Format

The format of the **storage.config** file is:

pathname size

where *pathname* is the name of a partition, directory, or file, and *size* is the size of the named partition, directory, or file, in bytes. You must specify a size for directories or files. For raw partitions, size specification is optional.

You can use any partition of any size. For best performance, the following guidelines are recommended:

- ◆ Use raw disk partitions.
- ◆ For each disk, make all partitions the same size.
- ◆ For each node, use the same number of partitions on all disks.

Specify pathnames according to your operating system requirements. See the following examples.

**Important**

In the **storage.config** file, a formatted or raw disk must be at least 2 GB. The recommended disk cache size is 147 GB.

update.config

The **update.config** file controls how Websense Content Gateway performs a scheduled update of specific local cache content. The file contains a list of URLs specifying objects that you want to schedule for update.

A scheduled update performs a local HTTP `GET` on the objects at the specific time or interval. You can control the following parameters for each specified object:

- ◆ The URL
- ◆ URL-specific request headers, which overrides the default
- ◆ The update time and interval

- ◆ The recursion depth



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (default location is `/opt/WCG/bin`) to apply the changes. When you apply the changes to a node in a cluster, Content Gateway applies the changes to all nodes in the cluster.

Scheduled update supports the following tag/attribute pairs when performing recursive URL updates:

- ◆ ``
- ◆ ``
- ◆ ``
- ◆ `<body background=" ">`
- ◆ `<frame src=" ">`
- ◆ `<iframe src=" ">`
- ◆ `<fig src=" ">`
- ◆ `<overlay src=" ">`
- ◆ `<applet code=" ">`
- ◆ `<script src=" ">`
- ◆ `<embed src=" ">`
- ◆ `<bgsound src=" ">`
- ◆ `<area href=" ">`
- ◆ `<base href=" ">`
- ◆ `<meta content=" ">`

Scheduled update is designed to operate on URL sets consisting of hundreds of input URLs (expanded to thousands when recursive URLs are included); it is *not* intended to operate on massively large URL sets, such as those used by Internet crawlers.

Format

Each line in the `update.config` file uses the following format:

```
URL\request_headers\offset_hour\interval\recursion_depth\
```

The following table describes each field.

Field	Allowed Inputs
<i>URL</i>	HTTP and FTP-based URLs.
<i>request_headers</i>	(Optional.) A list of headers (separated by semi-colons) passed in each GET request. You can define any request header that conforms to the HTTP specification. The default is no request header.
<i>offset_hour</i>	The base hour used to derive the update periods. The range is 00-23 hours.
<i>interval</i>	The interval, in seconds, at which updates should occur, starting at offset hour.
<i>recursion_depth</i>	The depth to which referenced URLs are recursively updated, starting at the given URL.

Examples

The following example illustrates an HTTP scheduled update:

```
http://www.company.com\User-Agent: noname user
agent\13\3600\5\
```

This example specifies the URL and request headers, an offset hour of 13 (1 p.m.), an interval of one hour, and a recursion depth of 5. This would result in updates at 13:00, 14:00, 15:00, and so on. To schedule for an update to occur only once a day, use an interval value of 24 hours x 60 minutes x 60 seconds = 86400.

The following example illustrates an FTP scheduled update:

```
ftp://anonymous@ftp.company.com/pub/misc/
test_file.cc\18\120\0\
```

This example specifies the FTP request, an offset hour of 18 (6 p.m.), and an interval of every two minutes. The user must be *anonymous* and the password must be specified by the variable *proxy.config.http.ftp.anonymous_passwd* in the **records.config** file.

congestion.config

The **congestion.config** file defines what to do when congestion control is on and a server is congested. You edit the file by going to **Configure > Networking > Connection Management > Congestion Control** and modifying the **Congestion Rules** on the screen.

Related topics:

[Connection Management, page 265](#)

Format

Each line in the **congestion.config** file consists of one or more sets of tag value pairs, preceded by one or more destination specifiers.

The format for each line is:

```
<dest_specifier_tag>=<value> [secondary_specifier_tag]  
<rule1_tag>=<value> ... <rule2_tag>=<value>
```

For example:

```
dest_domain=xyz.com max_connection_failures=3  
dest_host=gumby port=80 fail_window=100 max_connection=10
```

Destination specifier tags

The following tags specify the destination host to be used for congestion control. You must include exactly one primary destination specifier on each line of the **congestion.config** file. You can include any number of secondary specifiers, but you cannot duplicate them on the same line.

Tag	Type	Description
dest_domain	Primary	The domain name of the destination host. For example: xyz.com
dest_host	Primary	The host name of the destination host. For example: gumby.
dest_ip	Primary	The IP address of the destination host. For example: 123.45.6.7.
dest_regex	Primary	A regular expression describing the destination host. For example: gumby pokey.
port	Secondary	The port number to use to connect to the destination host. For example: 8877.
Prefix	Secondary	The prefix of the destination host. For example: //NIM1/.

Congestion control rule variables

The following variables or tags specify the rules to use for congestion control. Requests to a server that do not have an applicable rule in **congestion.config** use the values for the “origins server connect attempts” variables in **records.config**.

Variable	Input Parameter	Default Value	Description
max_connection_failures	INT	5	Specifies the maximum number of times that a connection can fail before no further attempts are made.
fail_window	INT	120	The period, in seconds, after which a server is marked congested if more than the specified number connection failures are experienced.
proxy_retry_interval	INT	10	The period, in seconds, after which the proxy should retry a congested server. If a server is marked congested, the proxy does not forward requests to it until Proxy Retry After Time (PRAT), which is current time + the proxy_retry_interval.
client_wait_interval	INT	300	The period, in seconds, the client should wait before retrying the request during the PRAT interval. The Client Retry After Time (CRAT) equals PRAT - current time + the client_wait_interval + a random integer from 0 to the wait_interval_alpha).
wait_interval_alpha	INT	30	The maximum client_wait_interval.
live_os_conn_timeout	INT	60	The number of seconds for each connection attempt before a timeout. If the proxy does not succeed, then one connection failure is counted towards the server.
live_os_conn_retries	INT	2	For each request to a live (non-congested) server, the proxy tries at most this number of times to connect to the server.
dead_os_conn_timeout	INT	15	The timeout, in seconds, before the proxy considers a connection attempt failed.
dead_os_conn_retries	INT	1	The number of times the proxy tries to connect to a congested server after its PRAT time. A congested server will stay dead if the proxy cannot make a successful connection; otherwise, the server becomes live again
max_connection	INT	-1	The maximum number of connections allowed to the proxy. “-1” means unlimited. Content Gateway temporarily marks a server as congested if a max_connection number to the server is reached. If a new client requests a new connection to the server, the client receives the message, “503 Retry-After”. There is no PRAT on the servers that reach their max_connection.

Variable	Input Parameter	Default Value	Description
error_page	Page URL	“congestion #retryAfter”	The error that is returned when a connection is found to be congested.
congestion_scheme	per_ip per_host	per_ip	<p>The congestion scheme to use to identify the server. Options include “per_ip” or “per_host”.</p> <p>If a domain has 2 IP address and per_ip is used, then each IP address has its own number of connection failures, and each IP address is marked either congested or uncongested.</p> <p>If per_host is used, then one connection failure is counted for the entire host. If the host is marked congested, both IPs are marked congested.</p>

F

Error Messages

Websense Content Gateway error messages

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.websense.com and then navigate to Support and Knowledge Base.

Process fatal

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.
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 the request to the parent proxy.

Warnings

Message	Description
<i>Logfile error: error_number</i>	Generic logging error.
Bad cluster major version range <i>version1-version2</i> for node <i>IP address</i> connect failed	Incompatible software versions causing a problem.
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 <i>IP address</i> , closing	The specified client is not allowed to connect to Content Gateway. The client IP address is not listed in the ip_allow.config file.

Message	Description
Could not rename log <i>filename</i> to <i>rolled filename</i>	System error when renaming log file during roll.
Did <i>this_amount</i> of backup still to do <i>remaining_amount</i>	Congestion is approaching.
Different clustering minor versions <i>version 1</i> , <i>version 2</i> for node <i>IP address</i> continuing	Incompatible software versions causing a problem.
log format symbol <i>symbol_name</i> not found	Custom log format references a field symbol that does not exist. See Event Logging Formats, page 279 .
missing field for field marker	Error reading a log buffer.
Unable to accept cluster connections on port: <i>cluster_port_number</i>	Contact Websense Technical Support. Go to www.websense.com/support/ for Technical Support contact information
Unable to open log file <i>filename</i> , <i>errno=error_number</i>	Cannot open the log file.
Error accessing disk <i>disk_name</i>	Content Gateway might have a cache read problem. You might have to replace the disk.
Too many errors accessing disk <i>disk_name</i> : declaring disk bad	Content Gateway is not using the cache disk because it encountered too many errors. The disk might be corrupt and might have to be replaced.
No cache disks specified in storage.config file: cache disabled	The Content Gateway storage.config file does not list any cache disks. Content Gateway is running in proxy-only mode. You must add the disks you want to use for the cache to the storage.config file (see storage.config, page 379).

Alarm messages

The following table describes alarm messages that you may see in Content Gateway Manager.

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.

Message	Description/Solution
[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 Rolling event log files, page 176 .
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 Rolling event log files, page 176 .
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 Custom logging fields, page 279 , 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 Working With Websense Data Security, page 105 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 Websense Data Security Suite server is entered correctly in Content Gateway Manager. See Working With Websense Data Security, page 105 for information on the format of the URI.
Content Gateway is not able to communicate with the DSS server. Please try again.	Ensure that the Websense Data Security Suite server is up and running, and accepting connections on the port specified in the <i>proxy.config.icap.ICAPUri</i> variable. Contact your Websense Data Security Suite 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. Investigate the status of the domain controller.

HTML messages sent to clients

Websense 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 [Standard HTTP response messages, page 391](#). 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.

Title	HTTP Code	Description	Customizable Filename
Access Denied	403	You are not allowed to access the document at location <i>URL</i> .	access#denied
Bad HTTP request for FTP Object	400	Bad HTTP request for FTP object.	ftp#bad_request
Cache Read Error	500	Error reading from cache. Please retry request.	cache#read_error
Connection Timed Out	504	Server has not sent any data for too long a time.	timeout#inactivity
Content Length Required	400	Could not process this request because no Content-Length was specified.	request#no_content_length

Title	HTTP Code	Description	Customizable Filename
Cycle Detected	400	Your request is prohibited because it would cause an HTTP proxy cycle.	request#cycle_detected
Forbidden	403	<i>port_number</i> 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
FTP Authentication Required	401	You need to specify a correct user name and password to access the requested FTP document <i>URL</i> .	ftp#auth_required
FTP Connection Failed	502	Could not connect to the server <i>server_name</i> .	connect#failed_connect
FTP Error	502	The FTP server <i>server_name</i> returned an error. The request for document <i>URL</i> failed.	ftp#error
Host Header Required	400	An attempt was made to transparently proxy your request, but this attempt failed because your browser did not send an HTTP <code>Host</code> header. Manually configure your browser to use <code>https://proxy_name:proxy_port</code> as an HTTP proxy. See your browser's documentation for details. Alternatively, end users can upgrade to a browser that supports the HTTP <code>Host</code> header field.	interception#no_host
Host Header Required	400	Your browser did not send a <code>Host</code> HTTP header field and therefore the virtual host being requested could not be determined. To access this Web site correctly, you will need to upgrade to a browser that supports the HTTP <code>Host</code> header field.	request#no_host
HTTP Version Not Supported	505	The origin server <i>server_name</i> is using an unsupported version of the HTTP protocol.	response#bad_version

Title	HTTP Code	Description	Customizable Filename
Invalid HTTP Request	400	Could not process this <i>client_request</i> HTTP method request for <i>URL</i> .	request#syntax_error
Invalid HTTP Response	502	The host <i>server_name</i> did not return the document <i>URL</i> correctly.	response#bad_response
Malformed Server Response	502	The host <i>server_name</i> did not return the document <i>URL</i> correctly.	response#bad_response
Malformed Server Response Status	502	The host <i>server_name</i> did not return the document <i>URL</i> correctly.	response#bad_response
Maximum Transaction Time exceeded	504	Too much time has passed transmitting document <i>URL</i> .	timeout#activity
No Response Header From Server	502	The host <i>server_name</i> did not return the document <i>URL</i> correctly.	response#bad_response
Not Cached	504	This document was not available in the cache, and you (the client) accept cached copies only.	cache#not_in_cache
Not Found on Accelerator	404	The request for <i>URL</i> on host <i>server_name</i> was not found. Check the location and try again.	urlrouting#no_mapping
NULL	502	The host <i>hostname</i> did not return the document <i>URL</i> correctly.	response#bad_response
Proxy Authentication Required	407	Please log in with user name and password.	access#proxy_auth_required
Server Hangup	502	The server <i>hostname</i> closed the connection before the transaction was completed.	connect#hangup
Temporarily Moved	302	The document you requested, <i>URL</i> , has moved to a new location. The new location is <i>new_URL</i> .	redirect#moved_temporarily
Transcoding Not Available	406	Unable to provide the document <i>URL</i> in the format requested by your browser.	transcoding#unsupported
Tunnel Connection Failed	502	Could not connect to the server <i>hostname</i> .	connect#failed_connect

Title	HTTP Code	Description	Customizable Filename
Unknown Error	502	The host <i>hostname</i> did not return the document <i>URL</i> correctly.	response#bad_response
Unknown Host	500	Unable to locate the server named <i>hostname</i> . 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
Unsupported URL Scheme	400	Cannot perform your request for the document <i>URL</i> because the protocol scheme is unknown.	request#scheme_unsupported

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	OK
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

Message	Description
501	Not Implemented
502	Bad Gateway
504	Gateway Timeout

G

The req_ca.cnf File

Create a **req_ca.cnf** file and copy the code below into that file. See [Creating a sub-certificate authority](#), page 117 for information on the **req_ca.cnf** file.

```
#
# Configuration file for generating a CA Request
#
HOME = .
RANDFILE = $ENV::HOME/.rnd
#
[ policy_match ]
countryName = match
stateOrProvinceName = match
organizationName = match
organizationalUnitName = optional
commonName = supplied
emailAddress = optional
#
[ policy_anything ]
countryName = optional
stateOrProvinceName = optional
localityName = optional
organizationName = optional
organizationalUnitName = optional
commonName = supplied
emailAddress = optional
#####
#
[ req ]
default_bits = 1024
default_keyfile = privkey.pem
distinguished_name = req_distinguished_name
string_mask = nombstr
req_extensions = v3_req # The extensions to add to a certificate
request
[ req_distinguished_name ]
countryName = Country Name (2 letter code)
countryName_default = US
countryName_min = 2
countryName_max = 2
stateOrProvinceName = State or Province Name (full name)
stateOrProvinceName_default = Some-State
localityName = Locality Name (eg, city)
0.organizationName = Organization Name (eg, company)
0.organizationName_default = Internet Widgits Pty Ltd
#organizationalUnitName = Organizational Unit Name (eg, section)
```

```
#organizationalUnitName_default =
commonName = Common Name (Name of Sub-CA)
commonName_max = 64
emailAddress = Email Address
emailAddress_max = 64
[ v3_req ]
# Extensions to add to a certificate request to make it a CA
basicConstraints=CA:TRUE
nsCertType = sslCA
keyUsage = cRLSign, keyCertSign
```

H

FAQs and Troubleshooting Tips

Frequently Asked Questions (FAQs)

- ◆ *How do disk I/O errors affect the cache and what does Content Gateway do when a cache disk fails?*, page 395
- ◆ *If a client disconnects during the time that Content Gateway is downloading a large object, is any of the object saved in the cache?*, page 396
- ◆ *Can Content Gateway cache Java applets, JavaScript programs, or other application files like VBScript?*, page 396
- ◆ *How do you access Content Gateway Manager if you forget the master administrator password?*, page 396
- ◆ *How do you apply changes to the logs_xml.config file to all nodes in a cluster?*, page 397
- ◆ *In Squid- and Netscape-format log files, what do the cache result codes mean?*, page 397
- ◆ *What does the cctx field record in a custom log file?*, page 399
- ◆ *Does Content Gateway refresh entries in its host database after a certain period of time if they have not been used?*, page 399
- ◆ *Can you improve the look of your custom response pages by using images, animated gifs, and Java applets?*, page 399
- ◆ *How do you configure Content Gateway to serve only transparent requests?*, page 400

See *Troubleshooting tips*, page 401 for additional information.

How do disk I/O errors affect the cache and what does Content Gateway do when a cache disk fails?

If a disk drive fails five successive I/O operations, Content Gateway considers the drive inaccessible and removes the whole disk from the cache. Normal cache operation continues on all other Content Gateway disk drives.

If a client disconnects during the time that Content Gateway is downloading a large object, is any of the object saved in the cache?

When a client disconnects during an HTTP or FTP operation, Content Gateway continues to download the object from the origin server for up to 10 seconds. If the transfer from the origin server completes successfully within 10 seconds after the client disconnect, Content Gateway stores the object in the cache. If the origin server download does not complete successfully within 10 seconds, Content Gateway disconnects from the origin server and deletes the object from the cache. Content Gateway does not store partial documents in the cache.

Can Content Gateway cache Java applets, JavaScript programs, or other application files like VBScript?

Content Gateway can store and serve Java applets, JavaScript programs, VBScripts, and other executable objects from its cache according to the freshness and cacheability rules for HTTP objects.

Content Gateway does not execute the applets, scripts, or programs. These objects run only when the client system that sent the request loads them.

How do you access Content Gateway Manager if you forget the master administrator password?

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

If you forget the administrator password and cannot access 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 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.

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is `/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 (default location is **/opt/WCG/bin**), run `content_line -x` to apply the changes.
5. Log on to Content Gateway Manager. When prompted for the user name and password, enter the administrator ID and leave the password entry blank.
Because you have already cleared the password in the **records.config** file, you do not need a password to log on as the administrator.
6. Navigate to the **Configure > My Proxy > UI Setup > Login** tab.
7. In the **Administrator** section, leave the Old Password field empty. Type the new password in the **New Password** field, and then retype the new password in the **New Password (Retype)** field.
8. Click **Apply**.
The next time you access Content Gateway Manager, you must use the new password.

How do you apply changes to the `logs_xml.config` file to all nodes in a cluster?

After you modify the `logs_xml.config` file on one Content Gateway node, enter the following command from the Content Gateway **bin** directory (default location is **/opt/WCG/bin**):

```
content_line -x
```

Content Gateway applies the changes to all nodes in the cluster. The changes take effect immediately.

In Squid- and Netscape-format log files, what do the cache result codes mean?

The following table describes the 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 <code>if-modified-since</code> request to the origin server and the origin server sent a <code>304 not-modified</code> response. The proxy sent the cached object to the client.

Cache Result Code	Description
TCP_REF_FAIL_HIT	Indicates that the object was in the cache but was stale. Content Gateway made an <code>if-modified-since</code> 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 <code>if-modified-since</code> 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 <code>no-cache</code> header. The proxy obtained the requested object from the origin server and sent a copy to the client. Content Gateway deletes any previous copy of the object from the cache.
TCP_IMS_HIT	Indicates that the client issued an <code>if-modified-since</code> request and the object was in the cache and fresher than the IMS date, or an <code>if-modified-since</code> 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 <code>if-modified-since</code> request and the object was either not in cache or was stale in cache. The proxy sent an <code>if-modified-since</code> 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.

What does the `cqtx` field record in a custom log file?

The `cqtx` field records the complete client request text (minus headers) in the log file. For example, `get http://www.company.com HTTP/1.0`.

Does Content Gateway refresh entries in its host database after a certain period of time if they have not been used?

By default, the Content Gateway host database observes the time-to-live (ttl) values set by name servers. You can reconfigure Content Gateway to a different value.

1. Open the **records.config** file located in the Content Gateway **config** directory (default location is `/opt/WCG/config`).
2. Edit the following variable:

Variable	Description
<code>proxy.config.hostdb.ttl_mode</code>	<p>Set to</p> <p>0 - to obey the ttl values set by the name servers</p> <p>1 - to ignore the ttl values set by name servers and use the value set by the Content Gateway configuration variable proxy.config.hostdb.timeout. Set this variable to a value appropriate for your environment.</p> <p>2 - to use the lower of the two values (the one set by the name server or the one set by Content Gateway)</p> <p>3 - to use the higher of the two values (the one set by the name server or the one set by Content Gateway)</p>

3. Save and close the file.
4. From the Content Gateway **bin** directory (default location is `/opt/WCG/bin`), run `content_line -x` to apply the configuration changes.

Can you improve the look of your custom response pages by using images, animated gifs, and Java applets?

Content Gateway can respond to clients only with 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 do 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.

How do you configure 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 from the **ip_allow.config** file by specifying ranges of IP addresses that are allowed to use the proxy cache. If Content Gateway receives a request from an IP address not listed in a range specified in the file, it discards the request. See [ip_allow.config](#), page 303.
- ◆ 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 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 (where *ipaddress* is the IP address of your Content Gateway system and *port_number* is a port number on which no service is listening):

```
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
```

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 two 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
```

Troubleshooting tips

- ◆ *When you clear the cache, the statistics in Content Gateway Manager do not reset to zero, page 401*
- ◆ *The throughput statistic is inaccurate in Content Gateway Manager, page 401*
- ◆ *You are unable to execute Content Gateway commands, page 402*
- ◆ *You observe inconsistent behavior when one node obtains an object from another node in the cluster, page 402*
- ◆ *Web browsers may display an error document with a data missing message, page 403*
- ◆ *Content Gateway does not resolve any Web sites, page 403*
- ◆ *Maximum document size exceeded message in the system log file, page 404*
- ◆ *DrainIncomingChannel message in the system log file, page 404*
- ◆ *No cop file message in the system log file, page 404*
- ◆ *Warning in system log file when editing vadders.config (Linux), page 405*
- ◆ *Non transparent requests fail after enabling always_query_destination, page 405*
- ◆ *Content Gateway is running but no log files are created, page 406*
- ◆ *Content Gateway error indicates too many network connections, page 406*
- ◆ *Low memory symptoms, page 407*
- ◆ *Connection timeouts with the origin server, page 407*
- ◆ *IBM Web servers do not work with Content Gateway, page 407*
- ◆ *Content Gateway does not start (or stop), page 408*

When you clear the cache, the statistics in Content Gateway Manager do not reset to zero

The clear cache command (`content_gateway -Cclear`) does not reset the statistics to zero in Content Gateway Manager. To reset the statistics, follow the procedure below.

1. Access Content Gateway Manager.
2. Add **clear_stats.html** to the URL in your browser window and press **Return**. For example:

```
http://proxy1:8081/clear_stats.html  
Content Gateway Manager statistics should all show zero.
```

The throughput statistic is inaccurate in Content Gateway Manager

Content Gateway updates the throughput statistic after it has transferred 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 inaccuracy is more noticeable with a light load. A heavier load yields a more accurate statistic.

You are unable to execute Content Gateway commands

Commands do not execute under the following conditions:

- ◆ If the `content_manager` process is not running.
Check if the `content_manager` process is running by entering the following command:

```
ps aux | grep content_manager
```

or

```
./WCGAdmin status
```

If the `content_manager` process is not running, enter the following command from the Content Gateway **bin** directory (default location is **/opt/WCG/bin**) to start it:

```
./content_manager
```



Important

If you must stop Content Gateway, it is recommended that you restart it using **./WCGAdmin**. Stop it with **./WCGAdmin stop** and start it with **./WCGAdmin start** to ensure that all the processes stop and start correctly. See [Getting Started, page 9](#).

- ◆ If you are not executing the command from **\$WCGHome/bin**.
If the Content Gateway **bin** directory is not in your path, prepend the commands with `./` (for example, `./content_line -h`).
- ◆ If multiple Content Gateway installations are present and you are not executing the command from the active path specified in **/etc/content_gateway**.

Always change to the correct directory by issuing the command:

```
cd `cat /etc/content_gateway`/bin
```

You observe inconsistent behavior when one node obtains an object from another node in the cluster

As part of the system preparation process, you must synchronize the clocks on all the nodes in your cluster. Minor time differences cause no problems, but differences of more than a few minutes can affect Content Gateway operation.

It is recommended that you run a clock synchronization daemon such as `xntpd`. You can obtain the latest version of `xntpd` from the following URL:

<http://www.ntp.org>

Web browsers may display an error document with a data missing message

A message similar to the following displays in Web browsers:

```
Data Missing
This document resulted from a POST operation and has
expired from the cache. If you wish you can repost the
form data to re-create the document by pressing the
reload button.
```

Web browsers maintain their local cache in memory and/or disk on the client system. Browser messages about documents that have expired from cache see the browser local cache, *not* the Content Gateway cache. There is no Content Gateway message or condition that can cause such messages to appear in a Web browser.

For information about browser cache options and effects, see the browser documentation.

Content Gateway does not resolve any Web sites

The browser indicates that it is contacting the host and then times out with the following message:

```
The document contains no data; Try again later, or contact
the server's Administrator....
```

Make sure that the system is configured correctly and that Content Gateway can read the name resolution file:

- ◆ Check if the server can resolve DNS lookups by issuing the `nslookup` command. For example:

```
nslookup www.myhost.com
```
- ◆ Check if the `/etc/resolv.conf` file contains the valid IP address of your DNS server(s).
- ◆ On some systems, if the `/etc/resolv.conf` file is unreadable or has no name server entry, the operating system will use localhost as a name server. However, Content Gateway does not use this convention. If you want to use localhost as a name server, you must add a name server entry for 127.0.0.1 or 0.0.0.0 in the `/etc/resolv.conf` file.
- ◆ Check that the Content Gateway user account has permission to read the `/etc/resolv.conf` file. Change the file permissions to `rw-r--r-- (644)`.



Important

If the IP addresses in `/etc/resolv.conf` change, Content Gateway must be restarted.

Maximum document size exceeded message in the system log file

The following message appears in the system log file.

```
WARNING: Maximum document size exceeded
```

A requested object was larger than the maximum size allowed in the proxy cache. Content Gateway provided proxy service for the oversized object but did not cache it.

You can set the object size limit for the cache by modifying the **Maximum Object Size** field on the **Configure > Subsystems > Cache > General** tab. If you do not want to limit the size of objects in the cache, set the document size to 0 (zero).

DrainIncomingChannel message in the system log file

The following messages appear in the system log file:

```
Feb 20 23:53:40 louis content_manager[4414]: ERROR ==>
[drainIncomingChannel] Unknown message: 'GET http://
www.telechamada.pt/ HTTP/1.0'
Feb 20 23:53:46 louis last message repeated 1 time
Feb 20 23:53:58 louis content_manager[4414]: ERROR ==>
[drainIncomingChannel] Unknown message: 'GET http://
www.ip.pt/ HTTP/1.0'
```

These error messages indicate that a browser is sending HTTP requests to one of the Content Gateway cluster ports, either `rsport` (default port 8087) or `mcpport` (default port 8088). Content Gateway discards the request. This error does not cause any Content Gateway problems. The browser must be reconfigured to use the correct proxy port.



Note

Content Gateway clusters work best when configured to use a separate network interface and cluster on a private subnet so that client machines have no access to the cluster ports.

No cop file message in the system log file

The following message appears repeatedly in the system log file:

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

The file **config/internal/no_cop** 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` or the `stop_content_gateway` commands. 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` or the `start_content_gateway` command.

The Content Gateway installation script creates a **no_cop** file so that Content Gateway does not start automatically. After you have completed installation and configuration, and have rebooted the operating system, use the `./WCGAdmin start` or the `start_content_gateway` command to start Content Gateway. See [Getting Started, page 9](#), for information on starting and stopping Content Gateway.

Warning in system log file when editing `vaddrs.config` (Linux)

If you edit the `vaddrs.config` file on a Linux system as a non-root user, Content Gateway issues a warning message in the system log file similar to the following:

```
WARNING: interface is ignored: Operation not permitted.
```

You can ignore this message. Content Gateway does apply your configuration edits.



Important

It is recommended that you always configure virtual IP addresses from Content Gateway Manager. Editing the `vaddrs.config` file can lead to unpredictable results.

Non transparent requests fail after enabling `always_query_destination`

The variable `proxy.config.arm.always_query_dest` in the `records.config` file configures Content Gateway in transparent mode to ignore host headers and always ask for the IP address of the origin server. When you enable this variable, Content Gateway obtains the origin server's IP address from the existing NAT map list rather than trying to resolve the destination host name with a DNS lookup. As a result, logged URLs contain only IP addresses, not host names. To log domain names, set `proxy.config.arm.always_query_dest` to 0. However, setting `proxy.config.arm.always_query_dest` to 0 does not reduce the number of DNS lookups.

However, explicit requests (non transparent requests, including requests on port 80) fail, as there is no matching map in the NAT list.



Note

The `always_query_destination` option works only on the primary proxy port.

Content Gateway is running but no log files are created

Content Gateway writes event log files only when there is information to record. If Content Gateway is idle, there may be no log files.

Ensure that you are looking in the correct directory. By default, Content Gateway creates log files in its **logs** directory. Check the location of the log files in Content Gateway Manager by examining the **Log Directory** field on the **Configure > Subsystems > Logging > General** tab. Alternatively, you can check the value of the variable `proxy.config.log2.logfile_dir` in the **records.config** file.

Check that the log directory has read/write permissions for the Content Gateway user account. If the log directory does not have the correct permissions, the `content_gateway` process is unable to open or create log files.

Check that logging is enabled. In Content Gateway Manager, examine the **Logging** area on the **Configure > Subsystems > Logging > General** tab. Alternatively, you can check the value of the variable `proxy.config.log2.logging_enabled` in the **records.config** file.

Check that a log format is enabled. In Content Gateway Manager, check that a standard format is enabled on the **Configure > Subsystems > Logging > Formats** tab or that the custom format is enabled on the **Custom** tab. In the **records.config** file, you select standard formats or the custom format by editing variables in the **Logging Config** section.

Content Gateway error indicates too many network connections

By default, Content Gateway supports 8000 network connections: half of this number is allocated for client connections and half for origin server connections. A connection throttle event occurs when client or origin server connections reach 90% of half the configured limit (3600 by default). When a connection throttle event occurs, Content Gateway continues processing all existing connections but does not accept new client connection requests until the connection count falls below the limit.

Connection throttle events can occur under the following conditions:

- ◆ If there is a *connection spike* - if thousands of client requests all reach the proxy at the same time. Such events are typically transient and require no corrective action.
- ◆ If there is a *service overload* - if client requests continuously arrive faster than the proxy can service them. Service overloads often indicate network problems between Content Gateway and origin servers or indicate that Content Gateway needs more memory, CPU, cache disks, or other resources to handle the client load.

Examine the Performance graphs to determine the nature of the connection throttle. In particular, check the Client Connections, TCP Connections, and Client Ops Per Second graphs. You can also check error messages in the system log file, error log file, or event log files.

If necessary, you can reset the maximum number of connections supported by the proxy on the **Configure > Networking > Connection Management > Throttling** tab

or by editing the value of `proxy.config.net.connections_throttle` in the `records.config` file. Do not increase the connection throttle limit unless the system has adequate memory to handle the client connections required. A system with limited RAM might need a throttle limit lower than the default value.



Important

Do not set this variable below the minimum value of 100.

Low memory symptoms

Under heavy load, the Linux kernel can run out of RAM. The low memory condition can cause slow performance and a variety of system problems. RAM exhaustion can occur even if the system has plenty of free swap space.

Symptoms of extreme memory exhaustion include the following messages in the system log files (`/var/log/messages`):

```
WARNING: errno 105 is ENOBUFS (low on kernel memory),
consider a memory upgrade
kernel: eth0: can't fill rx buffer (force 0)!
kernel: recvmmsg bug: copied E01BA916 seq E01BAB22
```

To avoid memory exhaustion, add more RAM to the system or reduce the load on Content Gateway.

Connection timeouts with the origin server

Some origin servers take longer than 30 seconds to post HTTP requests, which results in proxy connection timeouts. To prevent such connection timeouts, in Content Gateway Manager go to the **Configure > Protocols > HTTP > Timeouts** tab, and in the **Active Timeout** section, change the value of **Origin Server Response** to 60 seconds or more.

IBM Web servers do not work with Content Gateway

IBM Web servers do not support the TLS (Transport Layer Security) protocol. For IBM Web servers to work with Content Gateway, you must edit the value of a configuration variable.

1. Open the `records.config` file located in the Content Gateway `config` directory (default location is `/opt/WCG/config`).
2. Edit the following configuration variable:

Variable	Description
<code>proxy.config.ssl.TLSv1</code>	Set this variable to 0 (zero).

3. Save and close the file.
4. From the Content Gateway **bin** directory (default location is `/opt/WCG/bin`), run `content_line -x` to apply the changes.

Content Gateway does not start (or stop)

Content Gateway starts automatically upon installation. If you must stop the product, the preferred method to stop and restart Content Gateway is to use the `./WCGAdmin start` and `./WCGAdmin stop` commands.

Starting or stopping Content Gateway

1. Become root:

```
su
```
2. Change to the Content Gateway **bin** directory. The default location is `/opt/WCG`
3. Start the proxy:

```
./WCGAdmin start
```

Stop the proxy:

```
./WCGAdmin stop
```

Glossary

alternates

Different versions of the same Web object. Some origin servers answer requests to the same URL with a variety of objects. The content of these objects can vary, depending on whether a server delivers content for different languages, targets different browsers with different presentation styles, or delivers variable content at different times of the day.

ARM

Adaptive Redirection Module. Used in transparent proxy caching, ARM is a component that redirects intercepted client traffic destined for an origin server to the Content Gateway application. Before the traffic is redirected by the ARM, it is intercepted by an **L4 switch** or router.

cache

Stores copies of frequently accessed objects close to users and serves them to users when requested. See also **object store**.

cache hierarchy

Levels of caches that communicate with each other. All cache hierarchies recognize the concepts of **parent cache** and **child cache**.

cache hit

An object in the cache that can be served directly to the client.

cache miss

An object that is not in the cache or that is in the cache but no longer valid. In both cases, the proxy must get the object from the **origin server**.

caching web proxy server

A Web proxy server with local cache storage that allows the proxy to fulfill client requests locally, using a cached copy of the origin server's previous response.

CGI

Common Gateway Interface. A set of rules that describe how an origin server and another piece of software (a *CGI program*) located on the same machine communicate.

cgi-bin

The most common directory name on an origin server in which **CGI** programs are stored.

child cache

A cache lower in a **cache hierarchy** for which Content Gateway is a parent. See also **parent cache**.

cluster

A group of Content Gateway nodes that share configuration information and can act as a single large virtual cache.

Configure mode

One of two modes in **Content Gateway Manager**. Configure mode lets you configure the Content Gateway system. See also **Monitor mode**.

content_cop

A Content Gateway process that monitors the health of the **content_gateway** and **content_manager** processes by periodically issuing heartbeat requests to fetch synthetic Web pages.

content_gateway

A Content Gateway process that is the cache processing engine of the Content Gateway product. `content_gateway` is responsible for accepting connections, processing requests, and serving documents from the **cache** or **origin server**.

Content Gateway Manager

Content Gateway's browser-based interface consisting of a series of Web pages that enable you to monitor performance and change configuration settings.

content_manager

A Content Gateway process and the command and control facility. `content_manager` is responsible for launching, monitoring, and reconfiguring the **content_gateway** process. It is also responsible for the administration user interface, the proxy auto-configuration port, and the statistics interface, cluster administration, and **virtual IP failover**.

cookie

A piece of information sent by an origin server to a Web browser. The browser software saves the information and sends it back to the server whenever the browser makes additional requests from the server. Cookies enable origin servers to keep track of users.

DNS

Domain Name Service. Content Gateway includes a fast, asynchronous DNS resolver to streamline conversion of host names to IP addresses.

explicit proxy caching

A Content Gateway configuration option, in which client software (typically a browser) must be specifically configured to send Web requests to the Content Gateway proxy.

FTP

File Transfer Protocol. A protocol based on TCP/IP for reliable file transfer.

FULL CLUSTERING

A Content Gateway cluster distributes its cache across its nodes into a single, virtual object store, rather than replicating the cache node by node. See also [management-only clustering](#).

HTTP

Hypertext Transfer Protocol. The client/server protocol upon which the World Wide Web is based.

ICP

Internet Cache Protocol. A protocol for proxy caches to exchange information about their content.

IP

Internet Protocol. The lowest-layer protocol under TCP/IP responsible for end-to-end forwarding and long packet fragmentation control.

ISP

Internet Service Provider. An organization that provides access to the Internet.

JavaScript

A scripting language designed to give Web pages the ability to interact with the people viewing them. Examples of such interaction include performing actions in response to mouse movement or mouse clicks and validation of what has been entered into forms.

L4 switch

An Ethernet switch that can control network traffic flow using Level 4 rules. The switch can intercept desired client protocol packets and direct them to a proxy for transparent operation.

management-only clustering

A Content Gateway option, in which all nodes in a cluster automatically share configuration information. See also [FULL CLUSTERING](#).

Monitor mode

One of two modes in [Content Gateway Manager](#). Monitor mode lets you view statistics about Content Gateway performance and Web traffic. See also [Configure mode](#).

MRTG

Multi Router Traffic Grapher. A graphing tool provided with Content Gateway that creates the Performance graphs that allow you to monitor Content Gateway performance.

Netscape log format

A standard access log format. Using the Netscape log format, you can analyze Content Gateway access log files with off-the-shelf log analysis scripts. See also [Squid log format](#).

object store

A custom high-speed database, on which Content Gateway stores all cached objects.

origin server

The Web server that contains the original copy of the requested information.

PAC file

Proxy Auto-Configuration file. A JavaScript function definition that a browser calls to determine how requests are handled.

parent cache

A cache higher up in a [cache hierarchy](#), to which the proxy can send requests.

proxy server

See [web proxy server](#).

router

A device that handles the connection between two or more networks. Routers look at destination addresses of the packets passing through them and decide which route to send them on.

SOCKS

A circuit-level proxy protocol that provides a tunneling mechanism for protocols that cannot be proxied conveniently.

Squid log format

A standard access log format. Using the Squid log format, you can analyze Content Gateway event log files with off-the-shelf log analysis scripts. See also [Netscape log format](#).

SSL

Secure Sockets Layer. A protocol that enables encrypted, authenticated communications across the Internet. Used mostly in communications between origin servers and Web browsers.

syslog

The UNIX system logging facility.

TCP

Transmission Control Protocol. An Internet standard transport layer protocol. TCP provides reliable end-to-end communication by using sequenced data sent by IP.

transparent proxy caching

A configuration option that enables Content Gateway to intercept and respond to Internet requests without requiring users to reconfigure their browser settings. It does this by intercepting traffic destined for an origin server and redirecting that traffic through the proxy cache.

URL

Uniform Resource Locator. The address that defines the route to a file on the Web or other Internet facility.

virtual IP failover

An option available to clustered Content Gateway servers, in which WCG maintains a pool of virtual IP addresses that it assigns to the nodes of a cluster. If a node fails, the remaining nodes mask the fault and take over the failed node's virtual interface.

WCCP

Web Cache Control Protocol. A protocol used by Cisco IOS-based routers to redirect traffic during transparent proxy caching.

web proxy server

A proxy server that forwards client requests to **origin servers**. The proxy server may deny requests according to filter rules or security limitations.

web server

A computer that provides World Wide Web services on the Internet. See also **origin server**.

WPAD

Web Proxy Auto-Discovery. A protocol that allows clients to automatically locate a Web proxy, providing the benefits of a proxy without the need for explicit client configuration.



Websense® Content Gateway Online Help

©1996-2010, Yahoo, Inc., and Websense, Inc.

All rights reserved.

10240 Sorrento Valley Rd., San Diego, CA 92121, USA

Published September 8, 2010

Printed in the United States of America

R050410750

This document contains proprietary and confidential information of Yahoo, Inc. and Websense, Inc. The contents of this document may not be disclosed to third parties, copied, or duplicated in any form, in whole or in part, without prior written permission of Websense, Inc.

Websense and Websense Enterprise are registered trademarks of Websense, Inc., in the United States and certain international markets. Websense has numerous other unregistered trademarks in the United States and internationally. All other trademarks are the property of their respective owners.

Every effort has been made to ensure the accuracy of this manual. However, Websense Inc., and Yahoo, Inc. make no warranties with respect to this documentation and disclaim any implied warranties of merchantability and fitness for a particular purpose. Websense Inc. 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.

Traffic Server is a trademark or registered trademark of Yahoo! Inc. in the United States and other countries.

Red Hat is a registered trademark of Red Hat Software, Inc.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, Windows NT, and Active Directory are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation in the United States and in other countries.

UNIX is a registered trademark of AT&T.

All other trademarks are property of their respective owners.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States.

Contractor/manufacturer is Websense, Inc., 10240 Sorrento Valley Parkway, San Diego, CA 92121.

Portions of Websense Content Gateway include third-party technology used under license. Notices and attribution are included below.

Portions of Websense Content Gateway include the following technology:**OpenSSL 0.9.6**

The OpenSSL is an open source toolkit licensed under the GNU General Public License. Copyright © 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUAL-

ITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

Netscape Directory SDK 4.0 for C

Netscape Directory SDK 4.0 for C is available without license fee under the terms of the Netscape ONE SDK End User License Agreement.

Each of the Components is provided on an "AS IS" basis, without warranty of any kind, including without limitation the warranties of merchantability, fitness for a particular purpose and non-infringement. The entire risk as to the quality and performance of the Components is borne by you. Should the Components prove defective or inaccurate, as the case may be, you and not Netscape or its suppliers assume the entire cost of any service and repair. In addition, the security mechanisms, if any, implemented by the Components have inherent limitations, and you must determine that each of the Components sufficiently meets your requirements. This disclaimer of warranty constitutes an essential part of the agreement. SOME JURISDICTIONS DO NOT ALLOW EXCLUSIONS OF AN IMPLIED WARRANTY, SO THIS DISCLAIMER MAY NOT APPLY TO YOU AND YOU MAY HAVE OTHER LEGAL RIGHTS THAT VARY BY JURISDICTION.

Tcl 8.3

Tcl software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files. The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

libdb

LIBDB Copyright © 1991, 1993 The Regents of the University of California. All rights reserved. This product includes software developed by the University of California, Berkeley and its contributors.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER ARISING IN CONTRACT, STRICT LIABILITY OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

INN

Copyright © 1991, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001 The Internet Software Consortium and Rich Salz. This code is derived from software contributed to the Internet Software Consortium by Rich Salz. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the Internet Software Consortium and its contributors. 4. Neither the name of the Internet Software Consortium nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INTERNET SOFTWARE CONSORTIUM AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE INTERNET SOFTWARE CONSORTIUM OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF AD-

WARRANTY OF THE POSSIBILITY OF SUCH DAMAGE.

MRTG

Multi Router Traffic Grapher (MRTG) is freely available under the terms of the GNU General Public License. Copyright © 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

Libregx

Copyright © 1992, 1993, 1994, 1997 Henry Spencer. All rights reserved. This software is not subject to any license of the American Telephone and Telegraph Company or of the Regents of the University of California.

libmagic

Copyright (c) Ian F. Darwin 1986, 1987, 1989, 1990, 1991, 1992, 1994, 1995.

Software written by Ian F. Darwin and others; maintained 1994- Christos Zoulas.

This software is not subject to any export provision of the United States Department of Commerce, and may be exported to any country or planet.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice immediately at the beginning of the file, without modification, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Index

A

- absolute freshness limit, setting, 17
- access
 - web site, 128
- access control list, 140
- access log files, 135
- accessing Content Gateway Manager, 10, 142
- Adaptive Redirection Module. See ARM.
- adding certificate authorities, 119, 120
- adding nodes to a cluster, 68
- adding virtual IP addresses, 71
- administration tools, 4
- administrator ID, 10
- administrator ID, changing, 142
- administrator ID, setting, 142
- administrator password, 142
- administrator password default
 - administrator ID, 10
- aging factor
 - modifying, 17
- alarm messages, 386
- alarm script file, 98
- alarms, 97
 - clearing, 98
 - email notification, 98
 - viewing, 97
- Alarms button, 194
- allow certificates, 120
- allowing requests, 297
- alternates for caching, 30
- ARM, 3, 43, 59
 - bypass and WCCP, 45
 - enabling, 44
 - redirection rules, 44
 - security, 44, 60, 140
 - static bypass rules, 58
- ARM security
 - enabling, 140
- arm_security.config file, 140, 141, 287
- ASCII log files, 174
- ASCII_PIPE mode, 172, 310
- auth.config file, 289
- authentication
 - timeout period, 151
 - transparent, 151
- autodeleting log files, 169

B

- backing up certificates, 121
- binary log files, 174

- browsers
 - supported by Content Gateway Manager, 10
- bypass rules
 - deny, 293
 - dynamic, 58
 - static, 59
 - viewing, 60
- bypass.config file, 291
 - examples, 294
 - format, 293
- bypassing certificate validation, 126
- bypassing parent proxies, 74, 314

C

- cache
 - changing capacity, 77, 78
 - child, 73
 - clearing, 82
 - contents, 77, 78, 379
 - hit, 16
 - miss, 16
 - partitioning, 80
 - scheduling updates to, 21, 379
 - statistics, 94
- cache affinity, 45, 48
- cache capacity, 379
- cache pinning, 23
- cache request overview, 15
- cache space
 - managing, 316
- cache statistics, 194
- cache.config file, 17, 294
- cache-control headers, 18
- cached objects
 - expiration date, 16
 - freshness, 16
 - FTP, 16
 - HTTP, 16
- caching, 16
- caching alternates, 30
- caching cookied content, 29
- caching dynamic content, 29
- caching FTP objects, 31
- certificate authorities
 - adding, 119, 120
- certificate revocation lists
 - updating, 127
- certificate status, 119
- certificate validation, bypassing, 126
- certificates, 115, 120
 - backing up, 121
 - bypassing validation, 126
 - deleting, 120
 - deny, 120

- generating, 115
 - importing, 116
 - managing, 119
 - restoring, 121
 - revocation status, 127
 - slow, 120
 - sub-certificate authority, 117
 - changing, 142
 - changing cache capacity, 77, 78
 - changing certificate status, 119, 120
 - changing records.config variables, 89
 - changing size
 - RAM cache, 82
 - changing status of, 120
 - changing virtual IP addresses, 71
 - child cache, 73
 - clearing alarms, 98
 - clearing the cache, 82
 - client access control lists, 58
 - client access to proxy cache, 58, 139, 303
 - client no-cache directives, 25
 - clustering
 - adding nodes, 68
 - management-only, 2, 66
 - modes, 2
 - collating event log files, 182
 - command line interface
 - variables, 212
 - command-line interface, 12
 - commands, 211
 - commands
 - content_line -h, 12
 - WCGAdmin start, 9
 - configuration
 - remote, 143
 - configuration files, 89
 - congestion.config, 381
 - filter.config, 297
 - configuration information, sharing, 66
 - configuration options, 147
 - changing in the records.config file, 89
 - configuration snapshots
 - deleting, 92
 - restoring, 92
 - taking, 91
 - configuration variables (records.config), 318
 - Configure mode
 - Content Gateway Manager, 85
 - configuring proxy authentication
 - NTLM, 159
 - RADIUS, 156
 - configuring service groups, 49
 - configuring Websense Content Gateway, 12, 85, 89
 - via the command line, 89
 - congestion.config file, 381
 - Content Gateway Manager, 94, 143
 - accessing, 10
 - alarms, 97
 - Alarms button (Monitor), 194
 - Configure mode, 10, 85
 - controlling access, 141
 - logging on, 10
 - Monitor mode, 93
 - Performance button (Monitor), 96
 - starting, 10
 - starting Monitor mode, 93
 - supported browsers, 10
 - user accounts, 142
 - viewing statistics, 93, 94
 - Content Routing
 - statistics, 197
 - Content Routing button
 - Content Gateway Manager, 95
 - content_cop process, 4
 - content_gateway process, 4
 - content_line -h command, 12
 - content_manager process, 4
 - contents, 135
 - controlling
 - access to Content Gateway Manager, 141, 313
 - client access to proxy cache, 139
 - host access to Websense Content Gateway machine, 140
 - controlling host access to Content Gateway Manager, 143
 - controlling information leakage, 106
 - cookies. See caching cookie'd content
 - custom logging, 171
 - fields, 279
 - customer support, 7
- ## D
- Date header, 16
 - default, 10
 - deleting certificates, 120
 - deleting configuration snapshots, 92
 - deny bypass rules, 59, 293
 - deny certificates, 120
 - denying requests, 297
 - deployment options, 1
 - directory listings
 - full, 32
 - disabling
 - caching directory listings, 32
 - FTP over HTTP caching, 32
 - FTP proxy object caching, 32
 - HTTP caching, 28
 - logging, 168

- disk usage
 - restricting, 316
- displaying cache statistics, 94, 194
- DNS
 - resolver, 3
- DNS servers
 - specifying, 149, 377
- dynamic bypass rules, 58
 - deny bypass, 59, 293
 - setting, 59
- dynamic bypass statistics, viewing, 59
- dynamic content
 - caching, 29

E

- editing virtual IP addresses, 71
- emailing alarms, 98
- enabling ARM security, 140
- encryption, 396
- error log files, 167
- error messages, 385
 - HTML, 388
- event log entries, examples, 187
- event log files
 - collating, 182
 - converting binary to ASCII, 175
 - managing, 169
 - splitting, 179
 - statistics, 185
 - summary logs, 173
- expiration date, 16
- Expires header, 16
- explicit proxy
 - HTTPS PAC file, 113
 - SSL, 112
- explicit proxy caching, 2, 15

F

- files
 - arm_security.config, 140, 141, 287
 - auth.config, 289
 - bypass.config, 291
 - cache.config, 17, 294
 - hosting.config, 300
 - icp.config, 302
 - ip_allow.config, 139, 303
 - ipnat.conf, 304
 - log_hosts.config, 181
 - logs_xml.config, 171
 - mgmt_allow.config, 143, 313
 - parent.config, 74, 314
 - partition.config, 316
 - records.config, 17, 318

- socks.config, 375
- splittedns.config, 149, 377
- storage.config, 77, 78, 379
- update.config, 379
- filter.config file, 297
 - examples, 299
 - format, 298
- filtering rules, 297
- force immediate update option, 23
- forcing object caching, 29
- freshness computations, 17
- FTP client application, 39
- FTP objects
 - caching, 31
 - freshness, 20
- full directory listings, 32

G

- getting started, 9
- Graphs button
 - Content Gateway Manager, 94
 - statistics, 194
- Graphs button (Monitor), 194

H

- header requirements, 18
- headers
 - cache-control, 18
 - Expires, 16
 - Last-Modified, 16
 - max-age, 16
 - WWW-Authenticate, 27
- headroom limit (logging), 169
- hierarchical caching, 2, 73
 - HTTP hierarchies, 73
 - ICP hierarchies, 75
 - parent failover, 74
- host access, 143
- host access to Content Gateway Manager, 143
- host access, controlling, 140
- host log splitting, 180
- hosting.config file, 300
- HTML error messages, 388
- HTTP
 - alternates, 30
 - cache hierarchies, 73, 74, 314
 - host, separate logs, 180
- HTTP object freshness, 16
- HTTP response messages, 391

I

- ICAP, 106
- ICAP (Internet Content Adaptation Protocol)

- protocols supported, 106
- ICAP Service URI, 109
- ICP, 75
 - log filename, 179
 - peer, 75
 - separate logs, 179
- ICP cache hierarchies, 75
- icp.config file, 302
- immediate update, 23
- inbound traffic
 - SSL, 122
- incidents, 128
- increasing cache capacity, 78
- information leakage, controlling, 106
- interception strategies, 45
- internal root CA, 115
 - backup, 119
- IP spoofing, 62
- ip_allow.config file, 139, 303
 - examples, 303
 - format, 303
- ipnat.conf file, 304

J

- Java, 10
- JavaScript, 10

K

- keeping header information, 297

L

- Last-Modified header, 16
- LDAP authentication rules, 297
- LDAP proxy authentication, 152
- listing cache contents, 77, 78, 379
- listing commands, 12
- log collation, 182
- log collation server, 182
- log files
 - autodeleting, 169
- log files, contents, 135
- log formats, 170
- log_hosts.config file, 181
- logcat application, 175
- LogFilter specification, 308
- LogFormat specification, 308
- logging
 - access logs, 133
 - activity logs, 133
 - aggregate summaries, 173
 - ASCII_PIPE, 172, 310
 - choosing log file formats, 170
 - collating log files, 182

- converting binary files to ASCII, 175
- custom logging fields, 279
- disabling, 168
- example log entries, 187
- file splitting, 179
- headroom limit, 169
- length of time to keep files, 134
- managing log files, 169
- Netscape Common formats, 283
- Netscape Extended formats, 283
- Netscape Extended-2 formats, 284
- offset hour, 178
- rolling intervals, 178
- size of log files, 134
- Squid formats, 282
- SSL Manager, 133
- stand-alone collator (SAC), 184
- statistics, 185
- time stamps, 178
- WELF, 313

- logging on to Content Gateway Manager, 10
- LogObject specification, 309
- logs_xml.config file, 171

M

- management-only clustering, 2, 66
- managing certificates, 119
- max-age header, 16
- messages
 - certificate validation failure, 136
 - connection failure, 137
- mgmt_allow.config file, 143, 313
- modifying aging factor, 17
- monitor
 - remote, 143
- Monitor mode, 93
- My Proxy
 - statistics, 191
- My Proxy button
 - Monitor tab, 94

N

- naming rolled log files, 177
- Netscape Common logging formats, 283
- Netscape Extended logging formats, 283
- Netscape Extended-2 logging formats, 284
- Networking
 - statistics, 203
- Networking button
 - Content Gateway Manager Monitor tab, 96
- nodes
 - adding to a cluster, 68
- NTLM authentication rules, 297

NTLM proxy authentication, 158, 159

O

object caching, forcing, 29
 object freshness
 aging factor, 17
 object store, 77
 offset hour, 178
 Online certification status protocol, 127
 origin server, 15
 orphan log files, 182
 outbound traffic
 SSL, 122
 outgoing content, examining, 106
 overwriting dynamic bypass rules, 294

P

PAC file
 HTTPS, 113
 SSL Manager, 112
 parent cache, 73
 parent failover, 74
 parent proxies
 bypassing, 74, 314
 parent.config file, 74, 314
 partition.config file, 316
 partitioning the cache, 80
 partitions, 379
 passphrase, 117
 password, 10, 142
 password encryption, 396
 password, setting administrator, 142
 Performance graphs, 5
 pin-in-cache, 296
 print_bypass utility, 60
 processes (Websense Content Gateway), 4
 Protocol
 statistics, 194
 Protocols button
 Content Gateway Manager Monitor tab, 94
 proxy authentication
 LDAP, 152
 NTLM, 158
 RADIUS, 155
 proxy cache
 client access to, 139
 controlling client access to, 303
 proxy caching
 cache-control headers, 18
 client no-cache directives, 25
 cookied content, 29
 disabling HTTP caching, 28
 dynamic content, 29

 explicit, 15
 FTP object freshness, 20
 header requirements, 18
 HTTP alternates, 30
 revalidating HTTP objects, 19
 scheduling cache updates, 21
 server no-cache directives, 27
 transparent, 15
 whether to cache, 24
 WWW-Authenticate headers, 27
 PUSH, 298

R

RADIUS proxy authentication, 155, 156
 RAM cache, 77, 82
 raw disk, 379
 Read authentication from child proxy, 222
 re-authentication, 151
 records.config file, 17
 records.config file, 318
 redirecting requests (ARM), 43
 reducing cache capacity, 79
 regular expressions, 285
 remote monitoring and configuration, 143
 restoring certificates, 121
 restoring configuration snapshots, 92
 restoring Websense Content Gateway configurations, 90, 92
 restricting access to Content Gateway Manager, 141
 restricting host access to Websense Content Gateway machine, 140
 revalidation, 19
 revocation list
 restoring, 121
 revocation status, 127
 rolled log files, 177
 rolling intervals, 178
 root CA
 internal, 115
 backup, 119
 routers
 configuring, 49
 rules
 filtering, 297

S

SAC (stand-alone collator), 184
 sample records.config file, 90
 saving configurations, 91
 scheduling cache updates, 21
 scheduling updates, 379
 script file for alarms, 98

- Secure Sockets Layer, 143
 - Security
 - statistics, 198
 - security, 139
 - Content Gateway Manager
 - access, 141
 - host access, 140, 287
 - options, 6, 139
 - proxy authentication, 150
 - SOCKS, 146
 - split DNS, 149
 - SSL for secure administration, 143
 - Security button
 - Content Gateway Manager Monitor tab, 95
 - server no-cache directives, 27
 - service group ID number, 49
 - service groups, 51
 - configuring, 49
 - disabling WCCP processing, 51
 - enabling WCCP processing, 50
 - guidelines for configuring, 50
 - setting absolute freshness limit, 17
 - setting administrator ID and password, 141, 142
 - setting administrator password, 142
 - sharing configuration information, 66
 - simple directory listings, 32
 - snapshots
 - deleting, 92
 - restoring, 92
 - taking, 91
 - SOCKS, 147
 - proxy options, 147
 - SOCKS servers
 - specifying, 375
 - socks.config file, 375
 - split DNS, 149
 - splitdns.config file, 149, 377
 - splitting event log files, 179
 - spoofing, 62
 - Squid logging formats, 282
 - SSL, 143
 - certificates, 143
 - enabling (Content Gateway Manager), 144
 - inbound traffic, 122
 - outbound traffic, 122
 - SSL Manager
 - enabling, 114
 - stand-alone collators, 184
 - starting, 9
 - Content Gateway Manager Configure mode, 85
 - Content Gateway Manager Monitor mode, 93
 - starting Content Gateway Manager, 10
 - starting Websense Content Gateway, 9
 - static bypass rules, 59, 294
 - statistics
 - Content Routing, 197
 - My Proxy, 191
 - Networking, 203
 - Protocol, 194
 - Subsystems, 200
 - viewing from Content Gateway Manager, 94
 - viewing from the command line, 96
 - viewing in Content Gateway Manager, 93
 - status
 - changing certificate, 120
 - status, certificate, 119
 - storage.config file, 77, 78, 379
 - format, 379
 - stripping header information, 297
 - Subsystems
 - statistics, 200
 - Subsystems button
 - Content Gateway Manager Monitor tab, 95
 - summary log files, 173
- ## T
- technical support, 7
 - time in cache, 16
 - time stamps (log files), 177
 - timeout period for user authentication, 151
 - traffic analysis options, 5
 - traffic graphs, see Performance graphs, 5
 - transaction logging, 5
 - transparent authentication, 151
 - transparent proxy
 - caching, 15
 - interception strategies, 45
 - transparent proxy caching, 43
 - L4 switch, 45
 - policy-based routing, 56
 - software solutions, 56
 - WCCP, 46
- ## U
- update.config file, 22, 379
 - updates
 - scheduling, 379
 - URL regular expressions, 285
 - url_regix, 285
 - user accounts, 142
- ## V
- validation, bypassing certificate, 126
 - variables
 - records.config file, 89, 318
 - verifying that Websense Content Gateway is running, 11

- verifying URLs, 23
- viewing alarms, 97
- viewing bypass rules, 60
- viewing certificates
 - certificates
 - viewing, 120
- viewing dynamic bypass statistics, 59
- viewing logging statistics, 185
- viewing statistics
 - from Content Gateway Manager, 94
 - from the command line, 96
- virtual IP addresses, 70
 - adding, 71
 - editing, 71
- virtual IP failover, 2, 70

W

- WCCP, 46
 - load balancing, 48
 - service groups, 51
- WCCP 2.0
 - enabling, 53
 - security, 51
- WCCP processing

- disabling, 51
- enabling, 50
- WCCP2
 - configuring service groups, 49
- WCCP2 routers
 - configuring, 49
- WCGAdmin start command, 9
- web proxy caching, 1, 15
- web site access, 128
- Websense Content Gateway, 9
 - verifying, 11
- Websense Content Gateway components, 3
- Websense Content Gateway configurations
 - saving, 90
- Websense Content Gateway Manager, 194
 - Monitor mode, 10
- Websense Content Gateway processes, 4
- WELF, 313
- WWW-Authenticate headers, 27

X

- X-Authenticated-User, 222
- X-Forwarded-For, 222
- XML custom log formats, 171, 306