



Content Gateway Manager Help

Websense® Content Gateway

v7.8.x

Websense Content Gateway Online Help

February 2014

R250214782

Copyright © 1996-2014 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 TRITON 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 Unpublished rights reserved under the Copyright Laws of the United States.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 | 2 |
| | SSL inspection | 2 |
| | As a Web proxy cache | 3 |
| | In a cache hierarchy | 4 |
| | In a managed cluster | 4 |
| | As a DNS proxy cache | 4 |
| | Components | 4 |
| | Cache | 4 |
| | RAM cache | 5 |
| | Adaptive Redirection Module | 5 |
| | Host database | 5 |
| | DNS resolver | 6 |
| | Processes | 6 |
| | Administration tools | 7 |
| | Proxy traffic analysis features | 7 |
| | Online Help | 8 |
| | Technical Support | 8 |
| Topic 2 | Getting Started | 11 |
| | Accessing the Content Gateway manager | 11 |
| | Configuring Content Gateway for two-factor authentication | 13 |
| | Accessing the Content Gateway manager if you forget the master administrator password | 15 |
| | Entering your subscription key | 16 |
| | Providing system information | 17 |
| | Verifying that the proxy is processing Internet requests | 17 |
| | Using the command-line interface | 18 |
| | Starting and stopping Content Gateway on the Command Line | 19 |
| | The no_cop file | 20 |
| Topic 3 | Web Proxy Caching | 21 |
| | Cache requests | 21 |
| | Ensuring cached object freshness | 22 |
| | HTTP object freshness | 22 |
| | FTP object freshness | 26 |
| | Scheduling updates to local cache content | 27 |
| | Configuring the Scheduled Update option | 28 |
| | Forcing an immediate update | 28 |
| | Pinning content in the cache | 29 |

| | | |
|----------------|--|-----------|
| | Setting cache pinning rules | 29 |
| | Enabling cache pinning | 30 |
| | To cache or not to cache? | 30 |
| | Caching HTTP objects | 30 |
| | Client directives | 31 |
| | Origin server directives | 32 |
| | Configuration directives | 34 |
| | Forcing object caching | 36 |
| | Caching HTTP alternates | 36 |
| | Configuring how Content Gateway caches alternates. | 36 |
| | Limiting the number of alternates for an object | 37 |
| | Caching FTP objects. | 37 |
| | Disabling FTP over HTTP caching. | 38 |
| Topic 4 | Explicit Proxy | 39 |
| | Manual browser configuration | 39 |
| | Using a PAC file. | 40 |
| | Sample PAC file | 41 |
| | Using WPAD | 42 |
| | Configuring FTP clients in an explicit proxy environment | 43 |
| Topic 5 | Transparent Proxy and ARM | 47 |
| | The ARM | 48 |
| | Transparent interception strategies | 49 |
| | Transparent interception with a Layer 4 switch | 49 |
| | Transparent interception with WCCP v2 devices | 50 |
| | Transparent interception and multicast mode | 66 |
| | Transparent interception with policy-based routing | 67 |
| | Transparent interception with software-based routing | 68 |
| | Configuring Content Gateway to serve only transparent requests | 69 |
| | Interception bypass | 70 |
| | Dynamic bypass rules | 71 |
| | Static bypass rules. | 72 |
| | Viewing the current set of bypass rules | 73 |
| | Connection load shedding | 73 |
| | Reducing DNS lookups | 74 |
| Topic 6 | Additional Proxy Configuration | 77 |
| | IP spoofing | 77 |
| | Range-based IP spoofing. | 78 |
| | IP spoofing and the flow of traffic | 79 |
| | Configuring IP spoofing | 80 |
| | Support for IPv6 | 82 |

| | | |
|-----------------|---|------------|
| | IPv6 configuration summary | 83 |
| Topic 7 | Clusters | 85 |
| | Management clustering | 86 |
| | Changing clustering configuration | 86 |
| | Adding nodes to a cluster | 89 |
| | Deleting nodes from a cluster | 91 |
| | Virtual IP failover | 91 |
| | What are virtual IP addresses? | 92 |
| | Enabling and disabling virtual IP addressing | 92 |
| | Adding and editing virtual IP addresses | 92 |
| Topic 8 | Hierarchical Caching | 95 |
| | HTTP cache hierarchies | 95 |
| | Parent failover | 96 |
| | Configuring Content Gateway to use an HTTP parent cache | 96 |
| Topic 9 | Configuring the Cache | 99 |
| | Changing cache capacity | 99 |
| | Querying cache size | 100 |
| | Increasing cache capacity | 100 |
| | Reducing cache capacity | 101 |
| | Partitioning the cache | 102 |
| | Making changes to partition sizes and protocols | 102 |
| | Partitioning the cache according to origin server or domain | 102 |
| | Configuring cache object size limit | 103 |
| | Clearing the cache | 104 |
| | Changing the size of the RAM cache | 104 |
| Topic 10 | DNS Proxy Caching | 107 |
| | Configuring DNS proxy caching | 108 |
| Topic 11 | Configuring the System | 111 |
| | Content Gateway manager | 111 |
| | Using Configure mode | 111 |
| | Command-line interface | 112 |
| | Configuration files | 113 |
| | Saving and restoring configurations | 114 |
| | Taking configuration snapshots | 115 |
| | Restoring configuration snapshots | 115 |
| | Deleting configuration snapshots | 116 |
| Topic 12 | Monitoring Traffic | 117 |
| | Viewing statistics | 117 |
| | Using Monitor mode | 117 |

| | | |
|-----------------|--|------------|
| | Viewing statistics from the command line | 121 |
| | Working with alarms | 121 |
| | Clearing alarms | 122 |
| | Configuring Content Gateway to email alarm messages | 123 |
| | Using a script file for alarms | 123 |
| | Using Performance graphs | 123 |
| | Creating SSL-related reports | 124 |
| | Certificate Authorities | 124 |
| | Incidents | 126 |
| Topic 13 | Working With Websense Data Security | 129 |
| | Threats dashboard with Web Security Gateway | 129 |
| | Web DLP and the Threats dashboard with Web Security Gateway Anywhere | 129 |
| | How Web DLP works | 130 |
| | Data Security components on-box with Content Gateway | 131 |
| | Data Security over ICAP | 131 |
| | Registering and configuring Data Security | 132 |
| | Registration and configuration details | 133 |
| | Configuration options | 134 |
| | Configuring the ICAP client | 135 |
| | ICAP failover and load balancing | 136 |
| Topic 14 | Working With Encrypted Data | 139 |
| | Running in explicit proxy mode | 141 |
| | Enabling SSL support | 143 |
| | Initial SSL configuration tasks | 144 |
| | Certificates | 145 |
| | Internal Root CA | 145 |
| | Importing your Root CA | 146 |
| | Creating a new Root CA | 147 |
| | Creating a subordinate CA | 148 |
| | Backing up your internal Root CA | 153 |
| | Managing certificates | 154 |
| | View a certificate | 154 |
| | Delete a certificate | 154 |
| | Change the allow/deny status of a certificate | 155 |
| | Adding new certificate authorities | 155 |
| | Backing up certificates | 156 |
| | Restoring certificates | 156 |
| | Decryption and Encryption | 156 |
| | SSL configuration settings for inbound traffic | 156 |
| | SSL configuration settings for outbound traffic | 158 |

| | |
|--|------------|
| Validating certificates | 159 |
| Configuring validation settings | 159 |
| Bypassing verification | 162 |
| Keeping revocation information up to date | 163 |
| Certificate revocation lists | 163 |
| Online certification status protocol (OCSP) | 163 |
| Managing HTTPS website access | 164 |
| Viewing incidents | 164 |
| Changing the status of an incident | 166 |
| Deleting an incident | 167 |
| Changing the text of a message | 167 |
| Viewing incident details | 167 |
| Adding websites to the Incident List | 168 |
| Client certificates | 169 |
| When a client certificate is requested | 170 |
| Importing client certificates | 170 |
| When a client certificate is always required: the Hostlist | 170 |
| Deleting client certificates | 171 |
| Customizing SSL connection failure messages | 171 |
| Certificate validation failed | 172 |
| SSL connection failure | 172 |
| Topic 15 Security | 175 |
| Controlling client access to the proxy | 175 |
| Controlling access to the Content Gateway manager | 176 |
| Setting the administrator ID and password | 176 |
| Creating a list of user accounts | 177 |
| Controlling host access to the Content Gateway manager | 178 |
| Using SSL for secure administration | 178 |
| FIPS 140-2 Mode | 179 |
| Filtering Rules | 181 |
| Creating filtering rules | 181 |
| Configuring SOCKS firewall integration | 185 |
| Configuring SOCKS servers | 186 |
| Setting SOCKS proxy options | 187 |
| Setting SOCKS server bypass | 188 |
| Using the Split DNS option | 188 |
| Content Gateway user authentication | 189 |
| Selecting the authentication method | 191 |
| Supported domain controllers and directories | 191 |
| Best practices when using Windows Active Directory | 191 |
| Backup domain controllers | 191 |

| | | |
|-------------------|--|------------|
| | Transparent user authentication | 191 |
| | Browser limitations | 192 |
| | Global authentication options | 193 |
| | Surrogate credentials | 196 |
| | Integrated Windows Authentication | 197 |
| | Legacy NTLM authentication | 202 |
| | LDAP authentication | 204 |
| | RADIUS authentication | 208 |
| | Rule-Based Authentication | 211 |
| | Mac and iPhone/iPad authentication | 234 |
| Topic 16 | Working With Log Files. | 241 |
| | Event log files | 242 |
| | Managing event log files | 243 |
| | Choosing the logging directory | 243 |
| | Controlling logging space | 243 |
| | Event log file formats | 245 |
| | Using standard formats | 245 |
| | Custom format | 246 |
| | Choosing binary or ASCII | 249 |
| | Using logcat to convert binary logs to ASCII | 250 |
| | Rolling event log files | 251 |
| | Rolled log filename format | 252 |
| | Rolling intervals | 253 |
| | Setting log file rolling options | 253 |
| | Splitting event log files | 254 |
| | HTTP host log splitting | 254 |
| | Setting log splitting options | 255 |
| | Collating event log files | 256 |
| | Configuring Content Gateway to be a collation server | 257 |
| | Configuring Content Gateway to be a collation client | 258 |
| | Using a stand-alone collator | 258 |
| | Viewing logging statistics | 259 |
| | Viewing log files | 260 |
| | Example event log file entries | 261 |
| | Squid format | 261 |
| | Netscape examples | 262 |
| | Cache result codes in Squid- and Netscape-format log files | 264 |
| Appendix A | Statistics. | 267 |
| | My Proxy | 267 |
| | Summary | 267 |
| | Node | 269 |

| | | |
|-------------------|--|------------|
| | Graphs | 270 |
| | Alarms | 270 |
| | Diagnostics | 271 |
| | Protocols | 273 |
| | HTTP | 273 |
| | FTP | 275 |
| | Security | 275 |
| | Integrated Windows Authentication | 276 |
| | LDAP | 278 |
| | Legacy NTLM | 278 |
| | SOCKS | 279 |
| | Data Security | 279 |
| | Subsystems | 280 |
| | Cache | 280 |
| | Clustering | 282 |
| | Logging | 282 |
| | Networking | 282 |
| | System | 283 |
| | ARM | 283 |
| | ICAP | 284 |
| | WCCP | 285 |
| | DNS Proxy | 286 |
| | DNS Resolver | 286 |
| | Virtual IP | 287 |
| | Client Connection Status | 287 |
| | Performance | 287 |
| | SSL | 290 |
| | SSL Key Data | 290 |
| | CRL Statistics | 291 |
| | Reports | 291 |
| Appendix B | Commands and Variables | 293 |
| | Websense Content Gateway commands | 293 |
| | Websense Content Gateway variables | 294 |
| | Statistics | 295 |
| Appendix C | Configuration Options | 299 |
| | My Proxy | 299 |
| | Basic | 300 |
| | Subscription | 304 |
| | UI Setup | 305 |
| | Snapshots | 308 |
| | Logs | 310 |

| | |
|--|------------|
| Protocols | 311 |
| HTTP | 311 |
| HTTP Responses | 322 |
| HTTP Scheduled Update | 323 |
| HTTPS | 325 |
| FTP | 326 |
| Content Routing | 327 |
| Hierarchies | 328 |
| Mapping and Redirection | 330 |
| Browser Auto-Config | 332 |
| Security | 332 |
| Connection Control | 333 |
| FIPS Security | 333 |
| Data Security | 334 |
| Access Control | 335 |
| SOCKS | 351 |
| Subsystems | 354 |
| Cache | 354 |
| Logging | 356 |
| Networking | 360 |
| Connection Management | 360 |
| ARM | 362 |
| WCCP | 368 |
| DNS Proxy | 372 |
| DNS Resolver | 373 |
| ICAP | 376 |
| Virtual IP | 377 |
| Health Check URLs | 378 |
| SSL | 380 |
| Appendix D | |
| Event Logging Formats | 381 |
| Custom logging fields | 381 |
| Logging format cross-reference | 384 |
| Squid logging formats | 385 |
| Netscape Common logging formats | 385 |
| Netscape Extended logging formats | 386 |
| Netscape Extended-2 logging formats | 386 |
| Appendix E | |
| Configuration Files | 389 |
| Specifying URL regular expressions (url_regex) | 390 |
| Examples | 391 |
| auth_domains.config | 391 |
| Format | 391 |

| | |
|--------------------------------------|-----|
| auth_rules.config | 393 |
| Format | 394 |
| bypass.config | 395 |
| Format | 396 |
| Dynamic deny bypass rules | 396 |
| Examples | 397 |
| cache.config | 397 |
| Format | 398 |
| Examples | 399 |
| filter.config | 400 |
| Format | 401 |
| Examples | 402 |
| hosting.config | 403 |
| Format | 403 |
| Examples | 404 |
| ip_allow.config | 405 |
| Format | 405 |
| Examples | 405 |
| ipnat.conf | 406 |
| Format | 406 |
| Examples | 406 |
| log_hosts.config | 406 |
| Format | 407 |
| Examples | 407 |
| logs_xml.config | 408 |
| Format | 408 |
| Examples | 413 |
| WELF (WebTrends Enhanced Log Format) | 415 |
| mgmt_allow.config | 415 |
| Format | 415 |
| Examples | 416 |
| parent.config | 416 |
| Format | 417 |
| Examples | 418 |
| partition.config | 419 |
| Format | 419 |
| Examples | 420 |
| records.config | 420 |
| Format | 420 |
| Examples | 420 |
| Configuration variables | 421 |

| | |
|---|-----|
| System variables | 422 |
| Local manager | 425 |
| Process manager | 427 |
| Virtual IP manager | 427 |
| Alarm configuration | 428 |
| ARM | 428 |
| Load shedding configuration (ARM) | 432 |
| Authentication basic realm | 433 |
| LDAP | 434 |
| RADIUS authentication | 436 |
| NTLM | 437 |
| Integrated Windows Authentication | 440 |
| Transparent authentication | 441 |
| HTTP engine | 442 |
| Parent proxy configuration | 445 |
| HTTP connection timeouts (secs) | 446 |
| Origin server connection attempts | 447 |
| Negative response caching | 449 |
| Proxy users variables | 449 |
| Security | 451 |
| Cache control | 451 |
| Heuristic expiration | 453 |
| Dynamic content and content negotiation | 454 |
| Anonymous FTP password | 454 |
| Cached FTP document lifetime | 455 |
| FTP transfer mode | 455 |
| Customizable user response pages | 456 |
| FTP engine | 457 |
| SOCKS processor | 462 |
| Net subsystem | 462 |
| Cluster subsystem | 463 |
| Cache | 463 |
| DNS | 464 |
| DNS proxy | 465 |
| HostDB | 466 |
| Logging configuration | 467 |
| URL remap rules | 471 |
| Scheduled update configuration | 472 |
| SNMP configuration | 473 |
| Plug-in configuration | 473 |
| WCCP configuration | 473 |
| FIPS (Security Configuration) | 473 |
| SSL Decryption | 474 |

| | |
|---|------------|
| ICAP | 479 |
| Data Security | 481 |
| Connectivity, analysis, and boundary conditions | 481 |
| remap.config | 484 |
| Format | 485 |
| Examples | 485 |
| socks.config | 486 |
| Format | 486 |
| Examples | 487 |
| socks_server.config | 488 |
| Format | 488 |
| Examples: | 488 |
| splitdns.config | 489 |
| Format | 489 |
| Examples | 490 |
| storage.config | 491 |
| Format | 491 |
| update.config | 491 |
| Format | 492 |
| Examples | 493 |
| wccp.config | 493 |
| Appendix F Error Messages | 495 |
| Websense Content Gateway error messages | 495 |
| Process fatal errors | 495 |
| Warnings | 496 |
| Alarm messages | 498 |
| HTML messages sent to clients | 500 |
| Standard HTTP response messages | 504 |
| Index | 505 |

1

Overview

Help | Content Gateway | Version 7.8.x

Websense® Content Gateway is the Web proxy component of Websense Web Security Gateway and Web Security Gateway Anywhere.

Content Gateway works together with Websense Web Security to protect users and networks from malicious and unwanted content by performing advanced content analysis precisely when it is needed—as the content flows through the proxy—using the results of analysis to apply appropriate Web Security policy. This on-demand analysis protects users and networks at the same time that it makes rapidly changing, *Web 2.0* sites safe for your organization and users. Depending on Content Gateway configuration, advanced analysis is applied to HTTP, HTTPS, and FTP channels.

The precise application of advanced analysis is configured by the administrator for each Web Security Gateway (Anywhere) deployment.

Content Gateway can also be configured to function as a high-performance Web proxy cache that caches 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:

- ◆ [SSL inspection](#)
- ◆ [As a Web proxy cache](#)
- ◆ [In a cache hierarchy](#)
- ◆ [In a managed cluster](#)
- ◆ [As a DNS proxy cache](#)

Content Gateway can also be configured to:

- ◆ Ensure that clients are authenticated before they access content. Content Gateway supports Integrated Windows Authentication, legacy NTLM (NTLMSSP), LDAP, and RADIUS. See, [Content Gateway user authentication, page 189](#).
- ◆ Control client access to the proxy. See, [Controlling client access to the proxy, page 175](#).
- ◆ Use different DNS servers, depending on whether the proxy needs to resolve host names located inside or outside a firewall. This enables you to keep your internal network configuration secure while providing transparent access to external sites on the Internet. See, [Using the Split DNS option, page 188](#).

- ◆ Use the on-box Data Security policy engine or the ICAP interface to enable sites using Websense Data Security to examine outbound material such as Web postings, and block or allow based on company policy. See *Working With Websense Data Security*, page 129.
- ◆ Control access to the Content Gateway manager using:
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access
 - User accounts that define which users can access the manager and which activities they can perform (for example, view statistics only or view statistics and configure Content Gateway).
- ◆ Integrate into your firewall and control traffic through a SOCKS server. See *Security*, page 175.

Related topics:

- ◆ *Deployment options*, page 2
- ◆ *Components*, page 4
- ◆ *Proxy traffic analysis features*, page 7
- ◆ *Online Help*, page 8
- ◆ *Technical Support*, page 8

Deployment options

Help | Content Gateway | Version 7.8.x

SSL inspection

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

Content Gateway includes a complete set of certificate-handling capabilities. See [Working With Encrypted Data, page 139](#)

**Note**

Content Gateway does not cache HTTPS content.

**Important**

Even when HTTPS is **not** enabled, Content Gateway performs HTTPS URL filtering. This means that for every HTTPS request, a URL lookup is performed and policy is applied.

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

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

As a Web proxy cache

When Content Gateway is deployed as a Web proxy cache, user requests for Web content pass through Content Gateway on their 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 is typically deployed to receive client requests in one of the 2 following ways:

- ◆ As an *explicit proxy* in which the user's browser or client software is configured to send requests directly to Content Gateway. See [Explicit Proxy, page 39](#).
- ◆ As a *transparent proxy* in which user requests are transparently routed to Content Gateway on their way to the destination server. The user's client software

(typically a browser) is unaware that it is communicating with a proxy. See [Transparent Proxy and ARM](#), page 47.

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, either to other Content Gateway servers or to other caching products. See [Hierarchical Caching](#), page 95.

In a managed cluster

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

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

If the virtual IP failover option is enabled, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes of the cluster. Content Gateway can detect node failures (such as power supply or CPU failures) and reassign IP addresses of the failed node to the operational nodes. See [Virtual IP failover](#), page 91, for details.

If Content Gateway is configured as a transparent proxy with WCCP, failover is handled by WCCP and virtual IP failover should not be used. See [WCCP load distribution](#), page 53.

For complete information, see [Clusters](#), page 85.

As a DNS proxy cache

As a DNS proxy cache, Content Gateway can resolve DNS requests for clients. This offloads remote DNS servers and reduces response times for DNS lookups. See [DNS Proxy Caching](#), page 107.

Components

Help | Content Gateway | Version 7.8.x

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 99](#).

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 104](#).

Adaptive Redirection Module

The Adaptive Redirection Module (ARM) provides several essential functions. One is to send device notifications for cluster communication interface failover. Another is to inspect incoming packets before the IP layer sees them and readdress them to Content Gateway for processing.

The ARM is always active.

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 transit properly through a proxy.

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 73](#).

Host database

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

- ◆ DNS information (for fast conversion of host names to IP addresses)
- ◆ The HTTP version of each host (so advanced protocol features can be used with hosts running modern servers)

- ◆ Host reliability and availability information (to avoid waits for non-functional servers)

DNS resolver

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



Important

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

Processes

Help | Content Gateway | Version 7.8.x

Content Gateway has 4 primary processes:

| Process name | Description |
|------------------|--|
| content_gateway | Accepts connections, processes protocol requests, and serves documents from the cache or origin server. |
| content_manager | <p>Launches, monitors, and reconfigures the content_gateway process.</p> <p>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.</p> <p>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.</p> |
| content_cop | <p>Monitors the health of content_gateway and content_manager.</p> <p>The content_cop process periodically (several times each minute) queries content_gateway and content_manager by issuing heartbeat requests to fetch synthetic Web pages. If no response is received within the timeout interval or if an incorrect response is received, content_cop restarts content_manager and content_gateway.</p> |
| analytics_server | Manages the requests made and processes spawned for Content Classification Analytics. |

Administration tools

Help | Content Gateway | Version 7.8.x

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

Sometimes it is convenient or necessary to use the Content Gateway command-line interface. You can execute individual commands or script a series of commands in a shell. This method is only partially available when Content Gateway is installed on a Websense appliance. Use the Content Gateway manager and the Appliance manager **Command Line Utility** instead.

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

See:

[Content Gateway manager](#), page 111

[Command-line interface](#), page 112

[Configuration files](#), page 113

Proxy traffic analysis features

Help | Content Gateway | Version 7.8.x

Content Gateway provides options for network traffic analysis and monitoring:

- ◆ *Manager statistics and graphs* show network traffic information. View graphs and statistics from the Content Gateway manager, or collect and process statistics using the command-line interface.
- ◆ A variety of *Performance* graphs show historical information about virtual memory usage, client connections, document hit rates, and so on. View *Performance* graphs in the Content Gateway manager.
- ◆ *Manager alarms* are presented in the Content Gateway manager. Content Gateway signals an alarm for any detected failure condition. You can configure Content Gateway to send email or page support personnel when an alarm occurs. Content Gateway also sends select alarms to Web Security manager, where they are referred to as **alerts**. Summary alert messages are displayed on the Web Security **Status > Today** page. The full alert message is displayed on the **Alerts**

page. Web Security administrators can configure which Content Gateway conditions cause alert messages to be sent, and which methods (email or SNMP) are used to send the alert.

- ◆ *Transaction logging* lets you record information in a log file about every request the proxy receives and every error it detects. Use the logs to determine how many people use the proxy, how much information each person requested, and which pages are most popular. You can see why a transaction was in error 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 117. For logging options, see [Working With Log Files](#), page 241.

Online Help

Help | Content Gateway | Version 7.8.x

Click on **Get Help!** on any page in the Content Gateway manager to get 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 online help, or to access [Release Notes](#), installation and deployment information, FAQs, tips, and other technical information, go to the [Websense Technical Library](#).

Technical Support

Help | Content Gateway | Version 7.8.x

Technical information about Websense products is available 24 hours a day at:

<http://support.websense.com>

In the Support site you will find:

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

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

2

Getting Started

Help | Content Gateway | Version 7.8.x

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

Refer to the following procedures to get started:

- ◆ [Accessing the Content Gateway manager, page 11](#)
- ◆ [Entering your subscription key, page 16](#)
- ◆ [Verifying that the proxy is processing Internet requests, page 17](#)
- ◆ [Using the command-line interface, page 18](#)
- ◆ [Starting and stopping Content Gateway on the Command Line, page 19](#)

Accessing the Content Gateway manager

Help | Content Gateway | Version 7.8.x

The web browser-based Content Gateway manager is the management console for Content Gateway.

The Content Gateway manager is supported on:

- ◆ Microsoft Internet Explorer 8, 9, 10, and 11
- ◆ Mozilla Firefox versions 5 and later
- ◆ Google Chrome 13 and later; use of Chrome with some extensions (add-ons) may result in unexpected behavior

Use of other browsers and versions may result in unexpected behavior.

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

There are 3 ways to access the Content Gateway manager:

- ◆ From the Content Gateway button in Web Security manager.* For configuring access from Web Security manager, see Web Security Help.
- ◆ By entering the IP address and port of the Content Gateway host system in your browser. See below.

- ◆ When Content Gateway is a module on a Websense appliance, by opening the appliance Logon portal and clicking the Content Gateway button.

*When two-factor authentication (certificate or RSA SecurID) is configured in the TRITON Unified Security Center, the only way to access the Content Gateway manager is through the Web Security manager. See [Configuring Content Gateway for two-factor authentication](#), page 13.

**Note**

When single sign-on is used, the browser must be configured to allow pop-ups on the Content Gateway IP address.

To access the Content Gateway manager directly:

1. Open your Web browser and enter:

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

where *nodename* is the IP address of Content Gateway and *adminport* is the port number assigned to the Content Gateway manager (default: 8081).

For more information on using HTTPS to start the Content Gateway manager, see [Using SSL for secure administration](#), page 178.

2. Log on to the 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 the Content Gateway manager](#), page 176.

The 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 117, for additional information on the Monitor tab and [Configuring the System](#), page 111 for information on the configuration options in the Content Gateway manager.

Security certificate alerts

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

permanently accept the certificate within the browser. See your browser documentation for details.

**Note**

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

Windows 7 considerations

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

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

Logging off of the manager

How you log on to Content Gateway manager affects the log off behavior.

If you log on to Content Gateway manager from Web Security manager using **single sign-on**, when you log off of Content Gateway manager your session is closed.

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

Configuring Content Gateway for two-factor authentication

Help | Content Gateway | Version 7.8.x

Two-factor (certificate) authentication:

- ◆ Is configured for and applies to TRITON Unified Security Center logon only.
- ◆ Requires administrators to provide 2 forms of identification to log on.
- ◆ Can be made to apply to the Content Gateway manager by forcing administrators to log on to TRITON Unified Security Center before accessing the Content Gateway manager.
- ◆ Requires single sign-on to be configured for administrators allowed access to the Content Gateway manager.
- ◆ **Requires that the password logon capability be disabled on Content Gateway** (see below), preventing administrators not configured for single sign-on from accessing the Content Gateway manager. If Content Gateway is deployed on an appliance, password access is disabled using an appliance manager command. See the V-Series or X-Series manager Help.

For more information about configuring two-factor authentication, see “Configuring Certificate Authentication” in TRITON console Help.

Disabling and enabling Content Gateway password logon

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



Important

If Content Gateway is installed on a Websense appliance, see Appliance Manager Help for details.

To disable password logon:

1. Set up single sign-on in the Web Security manager.
2. If two-factor authentication will be used, set up two-factor authentication in the TRITON console.
3. Log on to the Content Gateway host system and acquire root privileges.
4. Change directory to “/etc” and check to see if there is a “websense” subdirectory. If not, create one (“mkdir websense”).
5. Change directory to “websense” (path is now “/etc/websense”) and check to see if the file “password-logon.conf” exists.
6. If not, create it (“touch password-logon.conf”).
7. Edit “password-logon.conf”.
8. Add the line, or modify the existing line to:

```
password-logon=disabled
```
9. Write and exit the file.

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

To re-enable password logon for all administrators:

1. Log on to the Content Gateway host system and acquire root privileges.
2. Change directory to “/etc/websense”.
3. Edit “password-logon.conf” and change:

```
password-logon=disabled
```

to:

```
password-logon=enabled
```
4. Write and exit the file.

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

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

Help | Content Gateway | Version 7.8.x



Note

The following procedure applies to Content Gateway standalone (software) installations.

If Content Gateway is running on a Websense appliance, the password is reset on the **Administration > Account Management** page of the Appliance manager.

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

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

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



Note

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

3. Save and close the file.
4. From the Content Gateway **bin** directory (**/opt/WCG/bin**), run **./content_line -x** to apply the changes.
5. Log on to the 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 the Content Gateway manager, you must use the new password.

Entering your subscription key

Help | Content Gateway | Version 7.8.x

Related topic:

- ◆ [Providing system information, page 17](#)

When Content Gateway is deployed with Web Security Gateway or Web Security Gateway Anywhere, there is no need to enter your subscription key in the Content Gateway manager. The key is shared automatically when it is specified in Web Security.



Note

The Web Security instance that is used is determined by the Policy Server that is configured. The configured Policy Server IP address is shown in the Content Gateway manager on the **Monitor > My Proxy > Summary** page when the **More Details** view is selected.

To configure Policy Server:

- ◆ In the V-Series or X-Series manager go to **Configuration > Web Security Components**.
- ◆ On a software install, edit `/opt/WCG/websense.ini` and set the value of **PolicyServerIP**. Then stop and start Content Gateway processes:

```
/opt/WCG/WCGAdmin stop
```

```
/opt/WCG/WCGAdmin start
```

When Content Gateway is deployed with only Websense Data Security you must enter your subscription key in the Content Gateway manager.

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

Providing system information

Help | Content Gateway | Version 7.8.x

If Content Gateway is the proxy integration for Web Security Gateway or Web Security Gateway Anywhere, the Policy Server IP address and port were specified during installation.

To complete configuration of Policy Server and Filtering Service timeout conditions and action (permit or block traffic), perform the following:

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

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

3. In the **Action for Communication Errors** section, select to permit or block traffic if a communication timeout condition occurs. When a timeout occurs, Content Gateway applies the setting and regularly polls the services to detect their return to service.
4. Beginning with 7.8.3, use the **Scanning Data Files Update** section to configure a delay for the download of the security scanning data files used by Content Gateway analysis. Select a **Delay time** from the drop-down provided.

Keep in mind that the longer the delay, the higher the security risk. The **Suspend updates** option is not recommended for use for an extended period of time. Selecting it will prompt an alarm as a reminder that data file downloads have been suspended. It is recommended that you not clear the alarm until **Delay time** has been reset.

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

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

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

This feature is typically used for a backup system.

5. Click **Apply**.


Verifying that the proxy is processing Internet requests

Help | Content Gateway | Version 7.8.x

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

1. Open the Content Gateway manager. See [Accessing the Content Gateway manager](#), page 11.
2. Go to the **Monitor > My Proxy > Summary** page to view subscription detail, scanning data file status, and node details, including the number of objects served, the hit rate, and other basic proxy service information.
3. Navigate to **Monitor > Protocol > HTTP > General** to display the General HTTP Statistics table.
4. Note the current **Total Document Bytes** statistic in the **Client** section of the table.

Check the value of this statistic.



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

5. Set your browser to the proxy port.
6. Browse the Internet.
7. Recheck the **Total Document Bytes** statistic.

This value increases as the proxy processes HTTP requests.

Using the command-line interface

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

**Note**

The command-line interface is not available when Content Gateway is installed on a Websense appliance. Use the Content Gateway manager and the V-Series or X-Series manager Command Line Utility instead.

You can execute individual commands or script multiple commands in a shell. See [Websense Content Gateway commands, page 293](#).

1. Become root:

```
su
```

2. Change to the Content Gateway **bin** directory (`/opt/WCG/bin`). Run Content Gateway commands from this directory.

Commands take the form:

```
content_line -command argument
```

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
```

Starting and stopping Content Gateway on the Command Line

Help | Content Gateway | Version 7.8.x

To stop or start Content Gateway from the command line:

1. Become root:

```
su
```

2. Change to the Content Gateway installation directory (`/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
```

After you have installed Content Gateway, open the Content Gateway manager (the management interface) to verify that the proxy is running. See [Accessing the Content Gateway manager](#), page 11 and [Verifying that the proxy is processing Internet requests](#), page 17.

The no_cop file

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

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

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

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

3

Web Proxy Caching

Help | Content Gateway | Version 7.8.x

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 deployment**, in which the user's client software (typically a browser) is unaware that it is communicating with a proxy, and **explicit proxy deployment**, in which the user's client software is configured to send requests directly to the proxy.

Cache requests

Related topics:

- ◆ [Ensuring cached object freshness, page 22](#)
- ◆ [Scheduling updates to local cache content, page 27](#)
- ◆ [Pinning content in the cache, page 29](#)
- ◆ [To cache or not to cache?, page 30](#)
- ◆ [Caching HTTP objects, page 30](#)
- ◆ [Forcing object caching, page 36](#)
- ◆ [Caching HTTP alternates, page 36](#)
- ◆ [Caching FTP objects, page 37](#)

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

1. Content Gateway receives a user request for a Web object.

2. Using the Web address, the proxy tries to locate the requested object in its object store (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 22](#)). 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.
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.

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.

Content Gateway does not store partial documents in the cache. Should a client disconnect while an HTTP or FTP download is underway, Content Gateway continues the download for up to 10 seconds after the disconnect. If the transfer completes successfully, Content Gateway stores the object in the cache. If the download does not complete, Content Gateway disconnects from the origin server and deletes the object from the cache.

Ensuring cached object freshness

Help | Content Gateway | Version 7.8.x

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 22](#).
- ◆ FTP objects stay in the cache for a specified time period. See [FTP object freshness, page 26](#).

HTTP object freshness

Help | Content Gateway | Version 7.8.x

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 23.

The computed freshness limit is bound by minimum and maximum boundaries. See [Setting an absolute freshness limit](#), page 24.
- ◆ 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 24.
- ◆ 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 397.

Modifying the aging factor for freshness computations

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 31).
- ◆ 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

Help | Content Gateway | Version 7.8.x

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 22](#). 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 397](#).

FTP object freshness

Help | Content Gateway | Version 7.8.x

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 only the FTP objects requested from HTTP clients.

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 397).

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**.

Scheduling updates to local cache content

Help | Content Gateway | Version 7.8.x

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.

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 28 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 28.

Configuring the Scheduled Update option

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 29](#).
- ◆ Enable the cache pinning option. See [Enabling cache pinning, page 30](#).

Setting cache pinning rules

Help | Content Gateway | Version 7.8.x

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 311.
4. Click **Add** to add the rule to the list, and then click **Apply**.
 5. Click **Close**.

Enabling cache pinning

Help | Content Gateway | Version 7.8.x

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

To cache or not to cache?

Help | Content Gateway | Version 7.8.x

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 30.
- ◆ For FTP objects, the proxy responds to caching directives you specify through configuration options and files. See [Caching FTP objects](#), page 37.

Caching HTTP objects

Help | Content Gateway | Version 7.8.x

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 31
- ◆ [Origin server directives](#), page 32

- ◆ [Configuration directives](#), page 34

Client directives

Help | Content Gateway | Version 7.8.x

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 31.

- ◆ **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 35.

- ◆ **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

Help | Content Gateway | Version 7.8.x

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 the Content Gateway manager (in the **Behavior** section of the **Configure > Protocols > HTTP > Cacheability** tab).

Origin server directives

Help | Content Gateway | Version 7.8.x

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 34.

- ◆ **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 33.

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

Configuring the proxy to ignore server no-cache headers

Help | Content Gateway | Version 7.8.x

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 |
|---|---|
| <code>proxy.config.http.cache.ignore_server_no_cache</code> | 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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 35](#).
- ◆ 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 35](#).
- ◆ Cache objects served in response to the **Cookie:** header. See [Caching cookied objects, page 35](#).

- ◆ Observe never-cache rules in the **cache.config** file. See [cache.config](#), page 397.

Disabling HTTP object caching

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 cookie content of all types.
 - Select **No Cache on Cookies** to *not* cache cookie content of any type.
3. Click **Apply**.

Forcing object caching

Help | Content Gateway | Version 7.8.x

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 **ttl-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 311](#).
4. Click **Add**, and then click **Apply**.
5. Click **Close**.

Caching HTTP alternates

Help | Content Gateway | Version 7.8.x

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 38](#)).
- ◆ Set never cache rules in the **cache.config** file (see [cache.config, page 397](#)).
- ◆ 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 31](#)).

Caching is not supported for FTP objects requested from FTP clients.

Disabling FTP over HTTP caching

Help | Content Gateway | Version 7.8.x

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**.

4

Explicit Proxy

Help | Content Gateway | Version 7.8.x

If Internet requests are not transparently routed to Content Gateway via a Layer 4 switch or router (see [Transparent Proxy and ARM, page 47](#)), traffic must be **explicitly** routed to Content Gateway by configuring the client's Internet browser. (This is sometimes referred to as an *explicit proxy deployment*.)

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

- ◆ By directly configuring their browsers to send requests directly to the proxy. See [Manual browser configuration, page 39](#).
- ◆ By configuring their browsers to download proxy configuration instructions from a PAC (Proxy Auto-Config) file. See [Using a PAC file, page 40](#).
- ◆ 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 42](#).

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 43](#).

Manual browser configuration

Help | Content Gateway | Version 7.8.x

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

- ◆ The proxy's hostname or IP address.



Important

If Integrated Windows Authentication is configured for user authentication, the Fully Qualified Domain Name must be used. Specifying the IP address will result in authentication failure. See [Integrated Windows Authentication, page 197](#).

- ◆ The proxy server port. The Content Gateway default proxy server port is 8080.

**Important**

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

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

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

For Microsoft Internet Explorer version 7.0 and greater, 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 4.0 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

Help | Content Gateway | Version 7.8.x

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 the HTTPS protocol option is enabled, see [Running in explicit proxy mode, page 141](#), 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 41](#). See, also, the article titled “PAC File Best Practices” in the [Websense Technical Library](#).
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 hostname **proxy1** and Content Gateway uses the default port 8083 to serve the file, users must specify the following URL in the proxy configuration settings:

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

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

**Note**

By default, Microsoft Internet Explorer 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. Log on to the Content Gateway manager and go to **Configure > Content Routing > Browser Auto-Config > WPAD** to display the **wpad.dat** file.
3. The WPAD Settings area displays the **wpad.dat** file:
 - If you copied an existing **wpad.dat** file into the Content Gateway **config** directory, the file contains your proxy configuration settings. Check the settings and make changes if necessary.
 - If you did not copy an existing **wpad.dat** file into the Content Gateway **config** directory (`/opt/WCG/config`), the WPAD Settings area is empty. Enter a script that will provide the proxy server configuration settings. A sample script is provided in [Sample PAC file, page 41](#) (a **wpad.dat** file can contain the same script as a **proxy.pac** file).
4. Click **Apply**.
5. Navigate to **Configure > Networking > ARM**.

6. In the **Network Address Translation (NAT)** section, click **Edit File** to add a special remap rule to the **ipnat.conf** file.
7. Enter information in the fields provided, and then click **Add**:
 - 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 **Destination IP** field, enter the IP address of the Content Gateway server that will be resolved to the WPAD server name by the local name servers.
 - In the **Destination Port** field, enter **80**.
 - In the **Redirected Destination IP** field enter the same IP address you entered in the **Destination IP**.
 - In the **Redirected Destination Port** field, enter **8083**.
8. Click **Add**.
9. Use the arrow keys on the left side to move the new rule to the first line in the file.
10. Click **Apply**, and then click **Close**.
11. Click **Restart** on the **Configure > My Proxy > Basic > General**.

Configuring FTP clients in an explicit proxy environment

Help | Content Gateway | Version 7.8.x

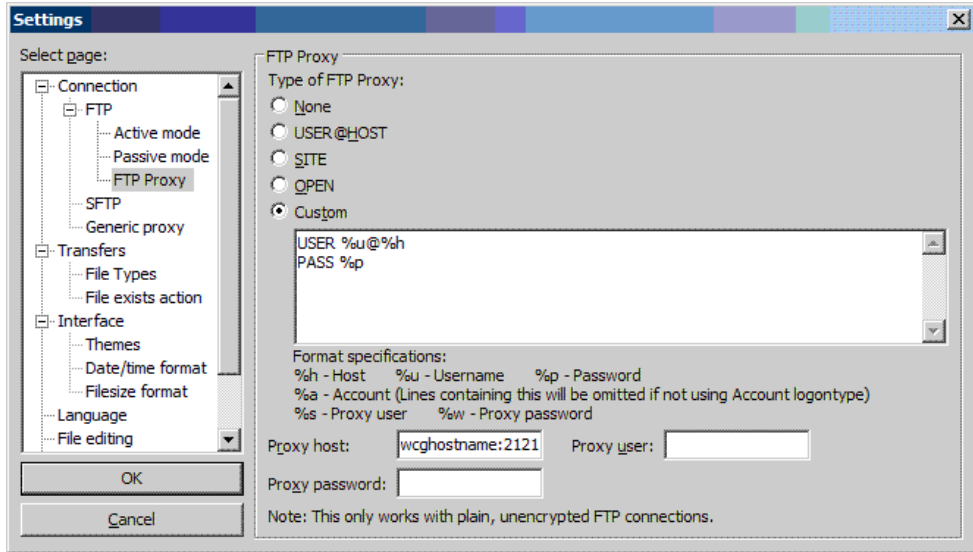
When Content Gateway is configured to proxy FTP traffic (see [FTP, page 326](#)), FTP client applications, such as FileZilla or WS_FTP, should be configured to send FTP requests to the proxy. When so configured, the user works with the FTP client application as if no proxy were present.

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**.

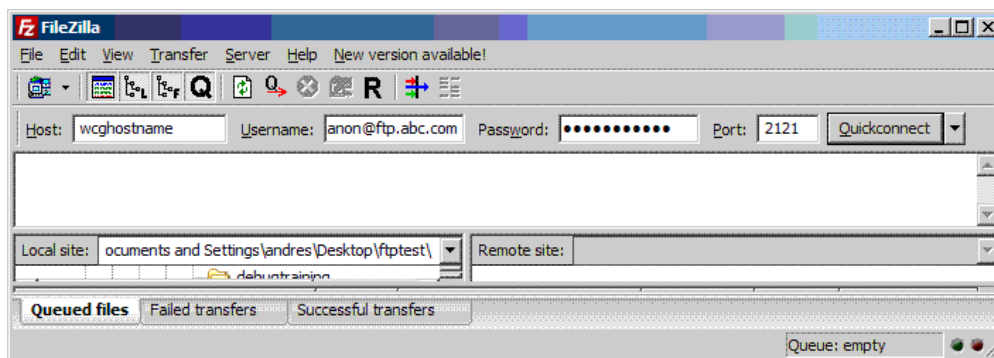
The user then enters FTP connection information in the usual way, as if no proxy were present. For example:

Host: ftp.example.com
 Username: anon
 Password: 123abc

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

Host: Content Gateway proxy hostname
 Username: anon@ftp.example.com
 Password: 123abc

Port: 2121



5

Transparent Proxy and ARM

Help | Content Gateway | Version 7.8.x

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

In a transparent proxy deployment:

1. The proxy intercepts client requests to origin servers via a switch or router. See [Transparent interception strategies, page 49](#).
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. (The ARM is always enabled.)
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.



Important

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

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

servers from being redirected to the proxy. See [Interception bypass](#), page 70.

Related topics:

- ◆ [Transparent interception strategies](#), page 49
- ◆ [Interception bypass](#), page 70
- ◆ [Connection load shedding](#), page 73
- ◆ [Reducing DNS lookups](#), page 74
- ◆ [IP spoofing](#), page 77
- ◆ [Support for IPv6](#), page 82

The ARM

Help | Content Gateway | Version 7.8.x

The Content Gateway ARM inspects incoming packets before the IP layer sees them and readdresses the packets to Content Gateway for processing.

The ARM can make two changes to an incoming packet's address. It can change 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 always active.

For the proxy to serve HTTP, HTTPS, FTP, or DNS requests transparently, you must check the redirection rules in the **ipnat.conf** file and edit them if necessary. If you are using WCCP for transparent interception, there must be a redirection rule for every port in every active service group. Rules for standard ports are included by default. To view and work with ARM redirection rules, follow these steps.

1. Log on to the Content Gateway manager and go to **Configure > Networking > ARM > General**.

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

- 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 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 362.

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

Transparent interception strategies

Help | Content Gateway | Version 7.8.x

Websense Content Gateway supports the following transparent interception solutions:

- ◆ A Layer 4 switch. See [Transparent interception with a Layer 4 switch](#), page 49.
- ◆ A router or switch that supports WCCP v2. Cisco IOS-based routers are the most common. See [Transparent interception with WCCP v2 devices](#), page 50.
- ◆ Policy-based routing. See [Transparent interception and multicast mode](#), page 66.
- ◆ Software routing. See [Transparent interception with software-based routing](#), page 68.


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

To configure Content Gateway to serve only transparent traffic, see [Configuring Content Gateway to serve only transparent requests](#), page 69.

Transparent interception with a Layer 4 switch

Help | Content Gateway | Version 7.8.x

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.

Transparent interception with WCCP v2 devices

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [WCCP load distribution](#), page 53
- ◆ [Configuring WCCP v2 routers](#), page 55
- ◆ [Enabling WCCP v2 in Content Gateway](#), page 60
- ◆ [ARM bypass and WCCP](#), page 52

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

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

A list of *WCCP v2 supported features* follows the setup outline.



Important

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

Following is a WCCP v2 setup outline.

1. Install and configure your WCCP v2 devices.
On each WCCP v2 device:
 - Configure the service groups.
 - Configure password security, if needed.
 - Configure multicast communication, if needed.
 See *Configuring WCCP v2 routers*, page 55.
2. Configure Content Gateway to work with your WCCP devices.
 - Define matching service groups.
In addition to network interface, protocols, ports, authentication (if used), and multicast communication (if used), also configure:
 - The IP addresses of the WCCP v2 devices.
 - The Packet Forward Method and Packet Return Method.
 - If Content Gateway is deployed in a cluster, **assignment method** load distribution, if desired.
 - Create ARM NAT rules for non-standard ports.
See *Enabling WCCP v2 in Content Gateway*, page 60 and *The ARM*, page 48.
3. Validate the configuration with test traffic.

WCCP v2 supported features

Content Gateway supports the following WCCP v2 features:

- ◆ Multiple routers in a proxy cluster
- ◆ Multiple ports per service group
- ◆ Multiple service groups per protocol. Sometimes it is necessary or convenient to have different service groups for different WCCP devices. For example, for Cisco ASA firewall, different service groups are required for each WCCP device in the network.
- ◆ Dynamic load distribution in a cluster through **assignment method** HASH or MASK. See *WCCP load distribution*, page 53.


- ◆ Packet Return Method and Packet Forward Method negotiation
- ◆ MD5 password security per service group
- ◆ Multicast mode

In a Content Gateway cluster, it is recommended that you **not** enable virtual IP failover in WCCP environments. WCCP v2 and the Content Gateway configuration handles node failures and restarts. (See [WCCP load distribution, page 53](#) and [Virtual IP failover, page 91](#).)

Content Gateway also supports cache affinity. If a node becomes unavailable and then recovers, the node’s cache does not need to be repopulated.

How WCCP v2 interception works:

1. WCCP v2 devices send HTTP, HTTPS, FTP, and DNS traffic, per the configuration of the service group, to the proxy server or cluster of servers.
2. The ARM readdresses traffic. For example, HTTP traffic on port 80 is readdressed to Content Gateway port 8080.
3. The proxy processes the request as usual, sending the response back to the client.
4. The ARM readdresses the proxy port in the response header to port 80 (undoing the readdressing it did on the way to the proxy). As a result, the user sees the response as if it had been sent directly from the origin server.



ARM bypass and WCCP

Help | Content Gateway | Version 7.8.x

If Content Gateway has an ARM bypass rule (discussed in [Interception bypass](#), page 70), Content Gateway forwards particular client requests directly to the origin server, bypassing the proxy.

Bypassed requests are unchanged by the ARM; they retain their client source IP address.

With WCCP v2, you can exclude certain router interfaces from redirection. **Content Gateway ARM bypass rules work only if you exclude the router interface that Content Gateway is connected to from WCCP redirection.** You do this on the router by selecting the interface connected to Content Gateway and issuing the router configuration command **ip wccp redirect exclude in**. This causes the router to exclude traffic inbound on the specified interface from all redirection rules.

WCCP load distribution

Help | Content Gateway | Version 7.8.x

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

- ◆ Load distribution is configured in Content Gateway Manager and is pushed to the WCCP devices.
- ◆ Load distribution is configured **per service group**.

For each service group:

- Participating cluster members must be registered to the service group. (The WCCP device makes no decisions about load balancing.)
- The HASH or MASK assignment method is selected. HASH is typically used with the GRE forward/return method, and MASK with the L2 forward/return method.



Important

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

- One or more **distribution attributes** are selected. Typically the destination IP address is used.
- If load is to be distributed to different cluster members in different proportions, a **weight** value is set on each cluster member. These values determine the proportion of requests each will receive relative to other members of the cluster. This option is only useful if the **Synchronize in the Cluster** option is disabled. See [Configuring service groups in the Content Gateway manager](#), page 61.

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

- There are multiple Content Gateway servers with different performance capabilities, for example a V-Series V10000 G2 and a V10000 G3.
- The Internet traffic profile doesn't lend itself to even distribution due to preferences for specific origin servers (and therefore destination IP addresses).

How dynamic redistribution works:

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

For configuration steps, see [Configuring service groups in the Content Gateway manager, page 61](#).

How the weight value supports asymmetric load distribution:



Important

Weight is only useful if the **Synchronize in the Cluster** option is **disabled**. See [Configuring service groups in the Content Gateway manager, page 61](#).

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

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

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

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

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

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

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

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

Configuring WCCP v2 routers

Help | Content Gateway | Version 7.8.x

It is strongly recommended that you consult the documentation and the manufacturer's support site for information regarding configuration and performance of your WCCP v2 device.

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

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

To prepare WCCP v2 devices for use with the proxy:

1. Configure one or more service groups for the protocols you intend to use. A service group can handle one or multiple protocols. See *Configuring service groups on the WCCP device*, page 55.
2. Configure the router to enable WCCP processing for these service groups. See *Enabling WCCP processing for a service group*, page 56.
3. Optionally, enable router security. Router security must also be enabled for the service group in Content Gateway. See *Enabling WCCP v2 security on the router*, page 59.



Note

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

4. When you are done configuring the router, you must also configure and enable WCCP in the Content Gateway manager. See *Enabling WCCP in the Content Gateway manager*, page 60.

Configuring service groups on the WCCP device

Help | Content Gateway | Version 7.8.x

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

A service group can intercept:

- one or more protocols
- on one or more ports

Service groups are assigned a unique integer identifier (ID) from 0 to 255.

Service groups IDs are user defined; they do not have a default port or traffic type.

The following table illustrates a set of service group definitions that are often found in networks. If you are configuring for IP spoofing, see the table in *IP spoofing*, page 77, for common reverse service groups IDs.

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

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

The best practice is to configure the router(s) first and Content Gateway second.

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

1. To see what has been configured on the router for WCCP, enter:


```
show running-config | include wccp
```
2. To enable WCCP v2, enter:


```
ip wccp version 2
```
3. If you used another proxy cache with your router prior to Content Gateway, disable the service ID that was previously used. For example, if you have a Cisco router, disable the service ID **web-cache** by issuing this command:


```
no ip wccp web-cache
```
4. 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.

Enabling WCCP processing for a service group

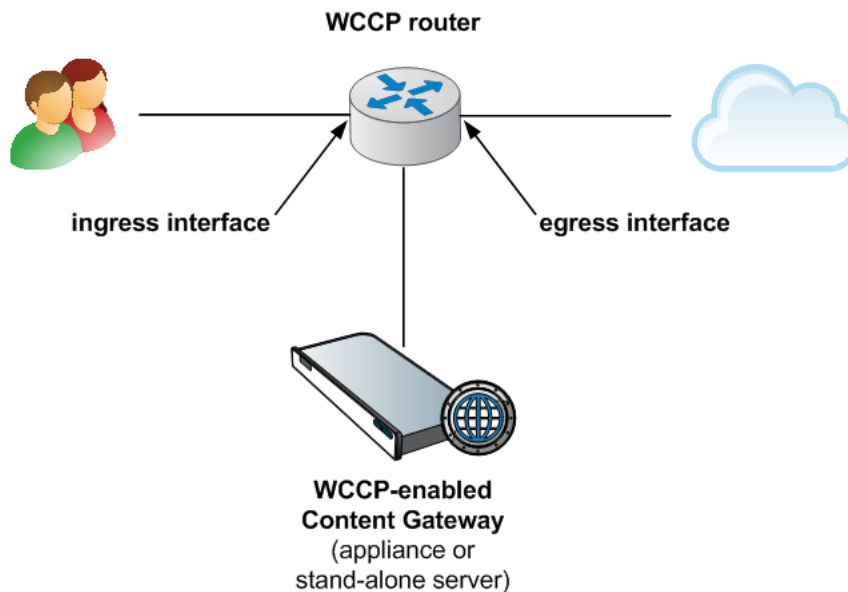
Help | Content Gateway | Version 7.8.x

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

WCCP v2 routers contain multiple network interfaces, including:

- ◆ one or more interfaces that receive inbound (ingress) client traffic
- ◆ one or more interfaces connected to Content Gateway

- ◆ an interface dedicated to outbound (egress) traffic that is aimed at the Internet



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

1. Turn on the WCCP feature:

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



Note

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

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

First, select the interface to configure:

```
interface <type> <number>
```

Second, establish your redirection rules:

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

Examples for inbound redirection:

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

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

```
ip wccp 0 redirect in
```

To turn on redirection of HTTPS destination port traffic:

```
ip wccp 70 redirect in
```

To turn on redirection of FTP destination port traffic enter:

```
ip wccp 5 redirect in
```

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

```
ip wccp 20 redirect in
```

Examples for *outbound* (egress) redirection:

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

First, select the interface to configure:

```
interface <type> <number>
```

Second, establish your redirection rules:

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

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

```
ip wccp 0 redirect out
```

To turn on redirection for HTTPS:

```
ip wccp 70 redirect out
```

To turn on redirection for FTP enter:

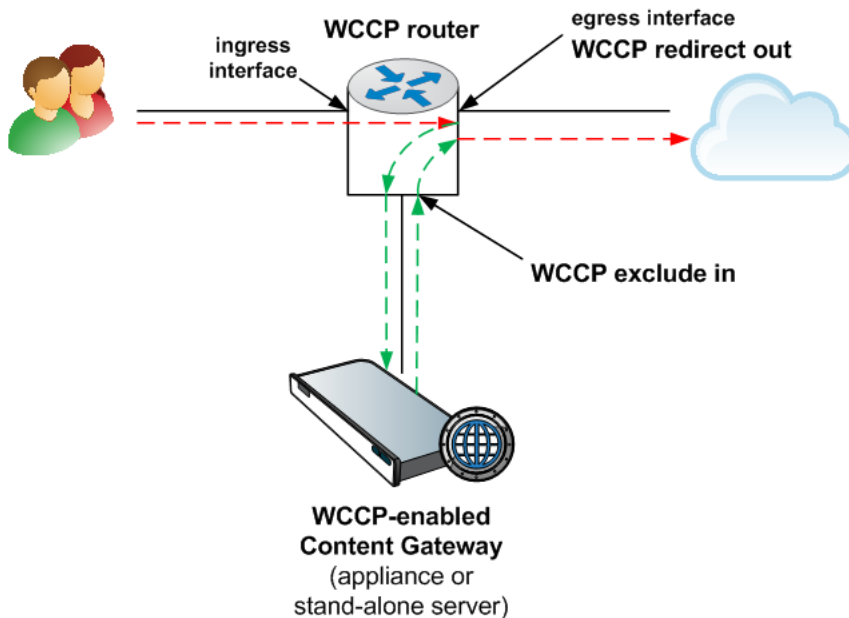
```
ip wccp 5 redirect out
```

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

```
interface <type> <number>
```
 - b. Exclude Content Gateway outbound traffic on the interface from all redirection rules on the router:

```
ip wccp redirect exclude in
```


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



Disabling WCCP processing for a service group

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

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

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

1. Telnet to the router and switch to Enable mode.

2. At the prompt, enter the following command to configure the router from the terminal:
`configure terminal`
3. If you defined a password when you enabled WCCP on the router, skip to step 4. Otherwise, enter the following command for each service group that the router intercepts:
`hostname(config)# ip wccp service_group password 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 the Content Gateway configuration for this service group.
4. Exit and save the router configuration.

Enabling WCCP v2 in Content Gateway

[Help](#) | [Content Gateway](#) | Version 7.8.x

Related topics:

- ◆ [Configuring WCCP v2 routers, page 55](#)
- ◆ [Configuring service groups on the WCCP device](#)
- ◆ [Enabling WCCP processing for a service group](#)
- ◆ [Enabling WCCP v2 security on the router, page 59](#)

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

1. [Enabling WCCP in the Content Gateway manager](#)
2. [Configuring service groups in the Content Gateway manager](#)
3. Restarting Content Gateway



Important

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

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

Enabling WCCP in the Content Gateway manager

[Help](#) | [Content Gateway](#) | Version 7.8.x

1. Go to **Configure > My Proxy > Basic > General**.

2. In the **Networking** section of the **Features** table, locate **WCCP**, click **On**, and then **Apply**. Do **not** restart Content Gateway.

Configuring service groups in the Content Gateway manager

Help | Content Gateway | Version 7.8.x

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

To define service groups, go to **Configure > Networking > WCCP**.

- a. The **Service Groups** table displays the list of configured service groups and a subset of their configuration settings.

Entries are stored in the **wccp.config** file.

The **Refresh** button rereads **wccp.config**, refreshing the table.

To add, modify, delete, or reorder service groups, click **Edit File**.

- b. **Synchronize in the Cluster:** If Content Gateway is configured in a cluster, enable (default) or disable the **Synchronize in the Cluster** option. (The value of this option is always synchronized in the cluster.)

When this option is enabled, the WCCP configuration (stored in **wccp.config**) is synchronized in the cluster and configuration changes can be made on any node in the cluster.

When this option is disabled, the WCCP configuration is not synchronized in the cluster and changes to the WCCP configuration must be made individually on each node. A common use case for this is to control which service groups are enabled/disabled on each node, and/or to use proportional load distribution using **weight**.

If after being disabled this option is enabled, the configuration on the node on which the administrator enables the option is used to initially synchronize the cluster.

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

Configuring a service group (editing **wccp.config**)

1. On **Configure > Networking > WCCP**, click **Edit File** to open **wccp.config** in the editor.

Defined service groups are summarized at the top of the page.

Click an entry in the list to view its complete details, modify, or reposition it.

When an entry is selected, the down and up arrows to the left of the list reposition the entry in the list.

Click “X” to delete a selected entry.

2. **Service Group Information**

- a. **Service Group Status:** To enable a service group, select **Enabled**. A service group can be defined but not active.
- b. **Service Group Name:** Specify a unique service group name. The service group name is an aid to administration.
- c. **Service Group ID:** Specify a WCCP service group identification number from 0-255. This ID must match a corresponding service group ID configured on the router. See *Configuring service groups on the WCCP device*.
- d. **Protocol:** Specify the network protocol applicable to the service group, either TCP or UDP.
- e. **Ports:** Specify the ports that this service group will use. You can specify up to 8 ports in a comma-separated list.

**Important**

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

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

3. Mode Negotiation

The **Packet Forward Method** determines how traffic is transmitted from the WCCP router to the proxy.

The **Packet Return Method** specifies the method used to return traffic back to the WCCP router.

Typically the router supports only one method.

Typically, the forward and return methods match.

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

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

Packet Forward Method: Select L2 or GRE.

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

**Important**

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

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

Packet Return Method: Select L2 or GRE.



Important

GRE cannot be used with WCCP multicast mode.



Important

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

4. Advanced Settings

- a. **Assignment Method:** Specify the parameters used to distribute intercepted traffic among multiple nodes in a cluster. For a description of the WCCP load distribution feature, see [WCCP load distribution, page 53](#).

HASH applies a hash operation to the selected distribution attributes.

- With HASH, more than one distribution attribute can be selected.
- The result of the hash operation determines the cluster member that receives the traffic.

MASK applies a mask operation to the selected distribution attribute.

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

The following distribution attributes can be selected:

- Destination IP address
- Destination Port
- Source IP address
- Source Port

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

- b. **Weight:** Only useful when **Synchronize in the Cluster** is disabled.

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

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

**Important**

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

**Note**

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

For more information about load distribution, see [WCCP load distribution](#), page 53.

- c. **Reverse Service Group ID:** For IP spoofing. Allows you to specify a reserve service group ID.

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

**Note**

Only HTTP and HTTPS are supported for IP Spoofing.

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

**Note**

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

5. Router Information

**Note**

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

- a. **Security:** To use optional WCCP authentication, select **Enabled** and enter the same password used for service group authentication on the router. See [Enabling WCCP v2 security on the router, page 59](#).
- b. **Multicast:** To run in multicast mode, select **Enabled** and enter the multicast IP address. The multicast IP address must match the multicast IP address specified on the router. See [Transparent interception and multicast mode, page 66](#).



Important

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

- c. **WCCP Routers:** Specify up to 10 **WCCP Router IP Addresses**. These routers must be configured with a corresponding service group.

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

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

The **Local GRE Tunnel Endpoint IP Address:**

- Must be IPv4
- Must be unique and not assigned to any device
- Must be a routable IP address
- Should reside on the same subnet as the proxy. If it is not, you must define a route for it.
- Is not intended to be a client-facing proxy IP address
- Is bound to the physical interface specified for the service group (on a V-Series appliance, eth0 = P1; eth1 = P2)

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

You can use “ping” to test connectivity to the router.

- From Content Gateway, ping each router defined in the service group (in the Router IP Address field).

- If ping doesn't return a response, you need to define a **GRE Tunnel Next Hop** to that router. Intervening routers must have a route to the WCCP router, or a next hop.



Note

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

6. Click **Add** to add a new entry, or click **Set** to save changes to the selected entry.
7. Click **Apply** and then **Close** to close the editor. Navigating away from the page before clicking **Apply** results in the loss of all changes.
8. Restart the proxy to cause the changes to take effect. Go to **Configure > My Proxy > Basic > General** and click **Restart**.



Note

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

Transparent interception and multicast mode

Help | Content Gateway | Version 7.8.x

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



Important

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

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

1. Telnet to the router and switch to Enable mode.

2. At the prompt, enter the following command to configure the router from the terminal:

```
configure terminal
```
3. At the prompt, enter the following command for each service group that the router intercepts:

```
hostname(config)# ip wccp service_group group-address  
multicast_address
```

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.
4. 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.
5. At the prompt, enter the following command for each service group that the router intercepts:

```
hostname(config-if)# ip wccp service_group group-listen
```
6. Exit and save the router configuration.

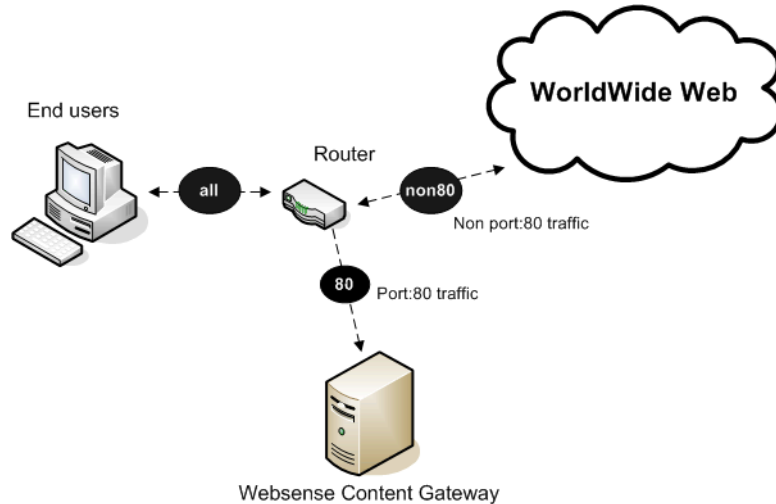
Transparent interception with policy-based routing

Help | Content Gateway | Version 7.8.x

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 91.



Transparent interception with software-based routing

Help | Content Gateway | Version 7.8.x

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 91](#).) The Content Gateway cluster failover mechanism is similar to the Hot Standby Router Protocol (HSRP).

Configuring Content Gateway to serve only transparent requests

Help | Content Gateway | Version 7.8.x

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

- ◆ You can control client access to Content Gateway by specifying ranges of IP addresses that are allowed to connect to the proxy. If Content Gateway receives a request from an IP address not listed in a specified range, it discards the request. See [Controlling client access to the proxy, page 175](#).
- ◆ 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 (**Configure > Networking > ARM > General**) 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
```

Interception bypass

Help | Content Gateway | Version 7.8.x

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

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

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

Content Gateway follows 2 types of bypass rules:

- ◆ *Dynamic* (also called adaptive) bypass rules are generated dynamically if you configure Content Gateway to bypass the cache when it detects non-HTTP traffic on port 80 or when it encounters certain HTTP errors. See [Dynamic bypass rules, page 71](#).
- ◆ *Static* bypass rules must be manually configured in the **bypass.config** file. See [Static bypass rules, page 72](#).



Note

Do not confuse ARM bypass rules with client access control lists. Bypass rules are created in response to interoperability problems. Client access control is simply restriction of the client IP addresses that can access the proxy, as described in [Controlling client access to the proxy, page 175](#).

Dynamic bypass rules

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Setting dynamic bypass rules, page 72](#)
- ◆ [Viewing dynamic bypass statistics, page 72](#)

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 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 itself 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 73](#).

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 395.

Setting dynamic bypass rules

Help | Content Gateway | Version 7.8.x

By default, Content Gateway is not configured to bypass itself 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 > Networking > ARM > Dynamic Bypass**.
2. Enable the **Dynamic Bypass** option.
3. In the **Behavior** section, select the dynamic bypass rules you want to use.
4. Click **Apply**.
5. Click **Restart** on the **Configure > My Proxy > Basic > General** tab.

Viewing dynamic bypass statistics

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

You can configure 3 types of static bypass rules:

- ◆ **Source bypass**, in which Content Gateway bypasses a particular source IP address or range of IP addresses. 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.

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 395).

Viewing the current set of bypass rules

Help | Content Gateway | Version 7.8.x

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 and then change directory to the Content Gateway **bin** directory (`/opt/WCG/bin`).
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.

Connection load shedding

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

In version 7.8.2 and later, when Always Query Destination is enabled, the value defined for the variable `proxy.config.arm.use_hostname_for_wisp_and_reporting` determines whether IP address or hostname is captured for reporting purposes.



Important

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

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

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

To enable Always Query Destination:

1. Open the **records.config** file in the Content Gateway **config** directory (**/opt/WCG/config**).
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
```


6

Additional Proxy Configuration

Help | Content Gateway | Version 7.8.x

Explicit and transparent proxy deployments can be used with other optional features.

Related Topics:

[IP spoofing, page 77](#)

[Support for IPv6, page 82](#)

IP spoofing

Help | Content Gateway | Version 7.8.x

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

IP spoofing configures the proxy to use:

- ◆ The IP address of the client when communicating with the origin server (basic IP spoofing), or
- ◆ A specified IP address when communicating with the origin server (*Range-based IP spoofing*)

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

IP spoofing features and restrictions:

- ◆ Prior to 7.8.3, IP spoofing is supported for transparent proxy requests only.
- ◆ IP spoofing is supported for HTTP and HTTPS traffic only.
- ◆ When IP spoofing is enabled, it is applied to both HTTP and HTTPS. It cannot be configured for only one protocol.
- ◆ HTTPS traffic is spoofed whether SSL support is enabled or not.

- ◆ IP spoofing relies on the ARM.
- ◆ IP spoofing is **not** supported with edge devices such as a Cisco ASA or PIX firewall. When this is attempted, requests made by Content Gateway using the client IP address are looped back to Content Gateway.
- ◆ IP spoofing, supported with IPv6 in 7.8.4 and later, requires all IP addresses in the same routing path use the same format. That is, all IP addresses must be either IPv6 or IPv4. A combination of IPv6 and IPv4 addresses is not supported.
- ◆ Prior to 7.8.4, IP spoofing does not support IPv6.



Warning

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

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

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



Important

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

Range-based IP spoofing

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

Among other uses, range-based IP spoofing facilitates:

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

spoofing for some clients, and standard proxy IP address substitution for some clients. The latter is done by creating a group that specifies the proxy IP address.




Important

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

Beginning with 7.8.4, IP Spoofing is supported for IPv6. However, range-based IP Spoofing is not supported for IPv6.

IP spoofing and the flow of traffic

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

6. The switch redirects the origin server response to the proxy, completing the proxy-to-origin server TCP connection.
7. A proxy response to the client is generated and returned to the client on the proxy-to-client TCP connection.



Note

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

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

Following is a set of suggested definitions.

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

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

Configuring IP spoofing

Basic IP spoofing

From Websense Content Manager:

1. Go to **Configure > Networking > ARM > General**.
2. Under **IP Spoofing**, select **Enabled**.
3. Click **Apply**.

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



Warning

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

For information about configuring WCCP routers, see [Configuring WCCP v2 routers](#), page 55.

Range-based IP spoofing

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

To create the range-based IP spoofing table:

- Go to **Configure > Networking > ARM > General**.
- Under **IP Spoofing**, select **Enabled**. Basic IP spoofing must be enabled to enable range-based IP spoofing.
- Under **Range Based IP Spoofing**, select **Enabled**.
- In the **Client IP Addresses** field, enter a comma separated list of individual IP addresses and/or IP address ranges.

In a range, the first IP address is separated from last with a hyphen. For example:
10.100.100.0-10.100.100.254

CIDR notation is allowed. Do not use spaces.

- In the **Specified IP Address** field, enter a single IP address.
- Click **Apply** to add the entry to the table.
Warning: If any of the formatting is invalid, all of the data in that row is cleared.
- To add a new row to the table, click **Add Row**.
- To put new entries into effect, click **Apply** and then restart Content Gateway.

To remove an entry from the IP spoofing table:

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

Support for IPv6

Help | Content Gateway | Version 7.8.x

TRITON Enterprise, including the Content Gateway proxy component, provides support for IPv6.



Important

Prior to 7.8.4, IPv6 support is provided for explicit proxy deployments only. Beginning with 7.8.4, in transparent proxy deployments, IPv6 support requires WCCP v2.01. If you use a Cisco router, it must be version 15.4(1) or later.

Content Gateway support for IPv6 includes:

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

Limits and restrictions:

- ◆ IPv6-only internal networks are not supported
- ◆ IPv4 must be used to communicate among all TRITON components, including other members of a Content Gateway cluster (multicast address)
- ◆ With all user authentication, the Domain Controller(s) must be reachable on an IPv4 address
- ◆ Prior to 7.8.4, the ARM does not support IPv6 addresses, including for redirection rules (`ipnat.config`) and static bypass rules (`bypass.config`)

- ◆ The parent proxy in a chain cannot be IPv6
- ◆ Prior to 7.8.4, IP spoofing is not supported. While IPv6 spoofing is available in 7.8.4 and later, range-based IP Spoofing is still not available for IPv6.
- ◆ SOCKS proxy is not supported
- ◆ IPv6 support, added in 7.8.4, is not available for FTP passive mode with the transparent proxy.
- ◆ IPv6 only clients do not display a block page correctly. The user is blocked from the site as expected but will receive a browser error rather than a block page. Dual-stack IPv6 clients receive the normal block page.



Warning

When the operating system is Red Hat Enterprise Linux 6, update 3, or CentOS 6, update 3, do not specify an IPv6 nameserver in `/etc/resolv.conf`. If an entry is included, Content Gateway will reset every time it attempts to start. If your deployment requires an IPv6 nameserver, upgrade your operating system to Red Hat Enterprise Linux 6, update 4, or CentOS 6, update 4.

IPv6 proxy statistics:

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

Effect of IPv6 on Event logs:

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

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

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

IPv6 configuration summary

IPv6 support is disabled by default.

If Content Gateway is deployed on a Websense Appliance, first enable IPv6 in the Appliance Manager on the **Configuration > Network Interfaces > IPv6** tab.

IPv6 is enabled in the Content Gateway manager in the **Network** section of the **Configure > My Proxy > Basic** page. When it is enabled, support is enabled for all functional areas as enumerated in the preceding section.

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

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

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

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

7

Clusters

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Changing clustering configuration](#), page 86
- ◆ [Adding nodes to a cluster](#), page 89
- ◆ [Deleting nodes from a cluster](#), page 91
- ◆ [Virtual IP failover](#), page 91

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

- ◆ Content Gateway detects the addition and deletion of nodes in the cluster and can detect when a node is down.
- ◆ You can add or delete a node from a cluster at any time.
- ◆ When you remove a node from the cluster, Content Gateway removes all references to the missing node.
- ◆ Restarting a node in the cluster causes all nodes in the cluster to restart.
- ◆ When the [Virtual IP failover](#) feature is enabled, the live nodes in a cluster can assume a failed node's traffic.
- ◆ Nodes in a cluster automatically share configuration information.



Note

Filtering Service and Policy Service IP addresses are not propagated around the cluster.

In transparent proxy deployments with WCCP, the service group Enabled/Disabled state and Weight settings are not propagated. See [Transparent interception with WCCP v2 devices](#), page 50.

When SSL support is enabled, the Dynamic Incident List is not propagated around the cluster.

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

**Important**

It is recommended that a dedicated network interface be used for Content Gateway cluster communication, **except** when the host is a V-Series appliance, in which case the P1 (eth0) interface is recommended.

**Important**

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

Management clustering

Help | Content Gateway | Version 7.8.x

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

**Note**

The number of nodes in a cluster can be 2 or more.

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

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

Changing clustering configuration

Help | Content Gateway | Version 7.8.x

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

1. In the Content Gateway manager, go to **Configure > My Proxy > Basic > Clustering**.
2. In the **Cluster Type** area, select the clustering mode:
 - Select **Management Clustering** to include this proxy in a cluster.
 - Select **Single Node** if this node is not part of a cluster.
3. In the **Interface** area, enter the name of the network interface. This is the interface used by Content Gateway to communicate with other nodes in the cluster, for example: eth1.

It is recommended that you use a dedicated secondary interface.

Node configuration information is multicast, in plain text, to other Content Gateway nodes on the same subnet. Therefore, Websense recommends that clients be located on a separate subnet from Content Gateway nodes (multicast communications for clustering are not routed).

On V-Series appliances, P1 (eth0) is the recommended interface. However, you may also use P2 (eth1) if you want to isolate cluster management traffic.

4. In the **Cluster Multicast Group Address** area, enter the multicast group address that all members of the cluster share. The default is 224.0.1.37.

**Warning**

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

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

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

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

If Content Gateway is on an appliance:

- ◆ Log on to the appliance manager, go to **Administration > Toolbox**, and open the **Command Line Utility**.
- ◆ Select the **Websense Content Gateway Module** and then the command `'content-line -s'`.
- ◆ Specify the **Variable Name** `proxy.config.cluster.mc_group_addr` and give **Value** the multicast IP address.
- ◆ Check each member of the cluster to ensure that they are all using the same multicast IP address.
- ◆ Restart the node.

If Content Gateway is installed on a separate server:

- ◆ Log on to the Linux host and go to `/opt/WCG/config`.
- ◆ Edit (vi) `records.config`, find `proxy.config.cluster.mc_group_addr` and assign it the value of the multicast IP address.
- ◆ Check each member of the cluster to ensure that they are all using the same multicast IP address.
- ◆ Restart the node.

-
5. Click **Apply**.

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

**Important**

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

Adding nodes to a cluster

Help | Content Gateway | Version 7.8.x

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

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

**Important**

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

1. Install the appropriate hardware and connect it to your network.
2. Install the Content Gateway software using the appropriate procedure for installing a cluster node. See the [Content Gateway installation procedure in the Deployment and Installation Center](#).
3. During the installation procedure, make sure that the following is true:
 - The cluster name that you assign to the new node is the same as the cluster name for the existing nodes.
 - The port assignments for the new node are the same as the port assignments used by the other nodes.
 - You have added multicast addresses and multicast route settings.
4. Restart Content Gateway. See *Starting and stopping Content Gateway on the Command Line*, page 19.

To add an existing Content Gateway installation to the cluster:

1. In the Content Gateway manager, go to **Configure > My Proxy > Basic > General** and set **Proxy Name** to the name of the cluster.
2. Go to **Configure > My Proxy > Basic > Clustering**.

3. Set **Interface** to the interface used by the cluster. All members must use the same interface.
4. Set the **Multicast Group Address** to the address being used by the cluster.
5. In the **Type** area, select **Management Clustering**.
6. Click **Apply**.
7. Click **Restart** on **Configure > My Proxy > Basic > General**.

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

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

| Variable | Description |
|--|--|
| <i>proxy.local.cluster.type</i> | Specify the clustering mode: 2 = management mode 3 = no clustering |
| <i>proxy.config.proxy_name</i> | Specify the name of the Content Gateway cluster. All nodes in a cluster must use the same name. |
| <i>proxy.config.cluster.mc_group_addr</i> | Specify the multicast address for cluster communications. All nodes in a cluster must use the same multicast address. |
| <i>proxy.config.cluster.rsport</i> | 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. |
| <i>proxy.config.cluster.mcport</i> | Specify the multicast port. The multicast port is used for node identification. All nodes in a cluster must use the same multicast port. The default port is 8088. |
| <i>proxy.config.cluster.ethernet_interface</i> | Specify the network interface for cluster traffic. All nodes in a cluster must use the same network interface. |

3. Save and close the file.

4. Restart Content Gateway (`/opt/WCG/WCGAdmin restart`).

Deleting nodes from a cluster

Help | Content Gateway | Version 7.8.x

On the node you want to remove from the cluster:

1. Go to **Configure > My Proxy > Basic > Clustering**.
2. In the **Cluster Type** area, select **Single Node**.
3. Click **Apply**.
4. If you are permanently removing the node from the cluster, it is a best practice to change the proxy name to a name other than the cluster name.
Go to **Configure > My Proxy > Basic > General** and change the **Proxy Name** to the system hostname or another meaningful value.
5. Restart the proxy.

Virtual IP failover

Help | Content Gateway | Version 7.8.x

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 address reassignment is done through a process called **ARP rebinding**.

What are virtual IP addresses?

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Enabling and disabling virtual IP addressing, page 92](#)
- ◆ [Adding and editing virtual IP addresses, page 92](#)

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

Help | Content Gateway | Version 7.8.x

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**.

Adding and editing virtual IP addresses

Help | Content Gateway | Version 7.8.x

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. Go 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 **Close**.

8

Hierarchical Caching

Help | Content Gateway | Version 7.8.x

Websense Content Gateway can participate in *HTTP cache hierarchies*, [page 95](#), in which requests not fulfilled in one cache can be routed to other regional caches, taking advantage of the contents and proximity of nearby caches.

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

HTTP cache hierarchies

Help | Content Gateway | Version 7.8.x

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 96.

**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 416).

**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

Help | Content Gateway | Version 7.8.x

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 416. 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

Help | Content Gateway | Version 7.8.x

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 [Hierarchies](#), page 328.
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.

9

Configuring the Cache

Help | Content Gateway | Version 7.8.x

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 parts of Web pages, providing optimum bandwidth savings. Using object management, the object store can cache alternate versions of the same object, varying on language or encoding type, and can store small and large documents, minimizing wasted space. When the cache is full, Content Gateway removes stale data.

Fault tolerance: Content Gateway can tolerate disk failures on cache disks. If a disk drive fails five successive I/O operations, Content Gateway marks the disk as down, removes the drive from the cache, and sends an alarm message to the Content Gateway manager, indicating which disk failed. Normal cache operation continues on the remaining cache disks. 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 99](#).
- ◆ Partition the cache by reserving cache disk space for specific protocols and origin servers and domains. See [Partitioning the cache, page 102](#).
- ◆ Specify a size limit for objects allowed in the cache. See [Configuring cache object size limit, page 103](#).
- ◆ Delete all data in the cache. See [Clearing the cache, page 104](#).
- ◆ Change the size of the RAM cache. See [Changing the size of the RAM cache, page 104](#).

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 104](#).

Changing cache capacity

Help | Content Gateway | Version 7.8.x

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.

The minimum disk cache size is 2 GB.

Related topics:

- ◆ [Querying cache size, page 100](#)
- ◆ [Increasing cache capacity, page 100](#)
- ◆ [Reducing cache capacity, page 101](#)

Querying cache size

Help | Content Gateway | Version 7.8.x

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 (**/opt/WCG/bin**).

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

Increasing cache capacity

Help | Content Gateway | Version 7.8.x

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 on the Command Line, page 19](#).
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 in the Content Gateway **config** directory (**/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 491.
4. Restart Content Gateway.

Reducing cache capacity

Help | Content Gateway | Version 7.8.x

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 491.
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

Help | Content Gateway | Version 7.8.x

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

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

HTTP is the only protocol supported.

Help | Content Gateway | Version 7.8.x

Making changes to partition sizes and protocols

Help | Content Gateway | Version 7.8.x

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 104](#)) before making many changes to cache partition sizes.

Partitioning the cache according to origin server or domain

Help | Content Gateway | Version 7.8.x

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 the Content Gateway manager.

In the Content Gateway manager:

1. Configure the cache partitions according to size and protocol, as described in [partition.config](#), page 419.

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 354.
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**.

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

```
WARNING: Maximum document size exceeded
```

Clearing the cache

Help | Content Gateway | Version 7.8.x

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 on the Command Line](#), page 19.
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

Help | Content Gateway | Version 7.8.x

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**.

The default size is 104857600 (100 MB).

**Note**

A value of “-1” directs Content Gateway to automatically size the RAM cache to be approximately 1 MB per 1 GB of disk cache.

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

10

DNS Proxy Caching

Help | Content Gateway | Version 7.8.x

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

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



Important

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

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

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



Note

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

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

Configuring DNS proxy caching

Help | Content Gateway | Version 7.8.x

To configure Content Gateway as a DNS proxy cache:

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



Important

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

In the Content Gateway manager:

1. Go to **Configure > Networking > ARM > General**.
2. In the **Network Address Translation (NAT)** section, click **Edit File** to open the file editor for the **ipnat.conf** file.
3. Enter information in the fields provided:
 - In the **Ethernet Interface** field, enter the Content Gateway ethernet interface to which client DNS requests are routed. For example, eth0.
 - In the **Connection Type** drop-down list, select **udp**.
 - In the **Destination IP** field, enter **0.0.0.0** to accept DNS requests from all clients.
 - In the **Destination CIDR field** (optional), enter the CIDR mask value. If you have specified 0.0.0.0 in the Destination IP field, enter '0' here.
 - In the **Destination Port** field, enter the port on which DNS requests are sent to Content Gateway. The default port is 53.
 - In the **Redirected Destination IP** field, enter the IP address of Content Gateway.
 - In the **Redirected Destination Port** field, enter the port that Content Gateway uses to communicate with the DNS server. The default port is 5353.
 - In the **User Protocol** drop-down list, select **dns**.
4. Click **Add**, then click **Apply**, and then click **Close**. Postpone the prompted restart until step 8.

5. Go to **My Proxy > Basic** and in the **Features** table, enable **DNS Proxy** in the **Networking** section and click **Apply**. Postpone the prompted restart until step 8.
6. Go to **Networking > DNS Proxy**.
7. In the **DNS Proxy Port** field, enter the DNS proxy port. The default port is 5353.
8. Click **Apply** and restart Content Gateway.
9. Configure your layer 4 switch or WCCP v2 router to send DNS traffic to the Content Gateway DNS port (default: 53).

11

Configuring the System

Help | Content Gateway | Version 7.8.x

Websense Content Gateway provides several options for configuring the system:

- ◆ [Content Gateway manager, page 111](#)
- ◆ [Command-line interface, page 112](#)
- ◆ [Configuration files, page 113](#)
- ◆ [Saving and restoring configurations, page 114](#)

Many configuration changes require a restart Content Gateway. When such a change is made in Content Gateway manager, a message informs you that a restart is needed. When such a change is made on the command line or in a configuration file there is no notification. It is recommended that you consult Content Gateway documentation to confirm whether a restart is needed.

Content Gateway manager

Help | Content Gateway | Version 7.8.x

The Content Gateway manager provides a Web-based user interface for configuring the Content Gateway web proxy.



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 112](#) and [Configuration files, page 113](#).

For instructions on logging on to the Content Gateway manager, see [Accessing the Content Gateway manager, page 11](#).

Using Configure mode

By default, the Content Gateway manager opens in Monitor mode.

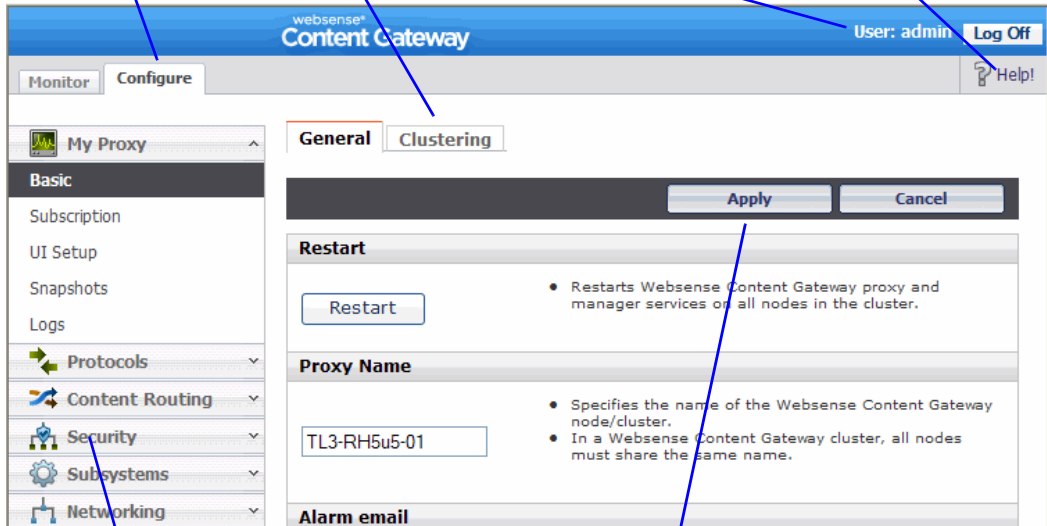
Click the **Configure** tab to display the Configure mode buttons.

Click to display the Configure buttons.

Click a tab to display more options.

Shows the current user logged on to Content Gateway manager.

Click Help! to display the online Help system.



Click a button to display its configuration options.

Click Apply to save the configuration changes on the current tab.

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

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

Command-line interface

Help | Content Gateway | Version 7.8.x

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

1. Log on to a Content Gateway node as root, and then change directory ('cd') to the Content Gateway **bin** directory (/opt/WCG/bin).
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 421).

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
```

Configuration files

Help | Content Gateway | Version 7.8.x

You can change some Content Gateway configuration options by editing specific variables in the **records.config** file, located in **/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 (**/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 389](#).

Saving and restoring configurations

Help | Content Gateway | Version 7.8.x

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



Note

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

This section describes how to perform the following tasks:

- ◆ Take a snapshot of the current configuration. See [Taking configuration snapshots, page 115](#).

- ◆ Restore previously taken configuration snapshots. See [Restoring configuration snapshots](#), page 115.
- ◆ Delete configuration snapshots stored on the Content Gateway node. See [Deleting configuration snapshots](#), page 116.

Taking configuration snapshots

Help | Content Gateway | Version 7.8.x

You can save all of the current configuration settings on your Content Gateway system using a facility in the Content Gateway manager.

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

1. Go 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.
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. Go 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

Help | Content Gateway | Version 7.8.x

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. Go 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. Go 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

Help | Content Gateway | Version 7.8.x

1. Go 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.

12

Monitoring Traffic

Help | Content Gateway | Version 7.8.x

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

- ◆ Statistics that show Content Gateway performance and network traffic information. See [Viewing statistics, page 117](#). The command-line interface provides an alternative method of viewing this information. See [Viewing statistics from the command line, page 121](#).
- ◆ Alarms that signal detected failure conditions. See [Working with alarms, page 121](#).
- ◆ Performance graphs that show historical Content Gateway performance and network traffic information. See [Using Performance graphs, page 123](#).
- ◆ Reports for SSL traffic, including the status of certificate authorities and incidents. See [Creating SSL-related reports, page 124](#).

Viewing statistics

Help | Content Gateway | Version 7.8.x

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

For instructions on logging on to the Content Gateway manager, see [Accessing the Content Gateway manager, page 11](#).

Using Monitor mode

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

All statistics displayed in Monitor mode are described in detail in [Statistics, page 267](#).

My Proxy

Click **My Proxy** to see statistics about Content Gateway.

- ◆ Click **Summary** to see a concise view of your Content Gateway system. The top portion of the page displays information about the features of your 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**.

**Important**

The graph is displayed in your browser using a Java applet. You should have the latest version of Java installed on your PC (at least version 1.7). To validate your access to Content Gateway statistics, you will be prompted for Content Gateway logon credentials.

- ◆ Click **Alarms** to view the alarms that Content Gateway has signaled. See [Working with alarms](#), page 121.
- ◆ Beginning with 7.8.4, select **Diagnostics** to run automatic or manual diagnostic tests to help determine the cause for a problem you might be having. Run the diagnostics listed on the **Automatic** tab to test network connections. Enter parameters for the diagnostic commands provided on the **Manual** tab to execute tests typically run from the command line.

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.



Note

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

Security

The Security button provides information about proxy authentication, and SOCKS server 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 **Integrated Windows Authentication (IWA)** to see the negotiated requests counters, the NTLM request counters and the Basic authentication request counters. The IWA tab appears only if you have enabled the IWA 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, DNS proxy, 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.
- ◆ Click **WCCP** to see WCCP v2 fragmentation statistics and the configuration of every WCCP service group enabled on the Content Gateway node. The WCCP tab appears only if you have enabled WCCP in the Features table on the **Configure > My Proxy > Basic > General** tab.
- ◆ Click **DNS Proxy** to see the total number of DNS requests served by Content Gateway, and the number of cache hits and misses. The DNS Proxy button appears only if you have enabled the DNS Proxy option 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.
- ◆ Click **Client Connections** to see client connection totals for current connections, unique clients since last restart, clients that have exceeded the connection limit, and client for which connections were closed.

Performance

The Performance button displays historical performance graphs. See [Using Performance graphs, page 123](#).

SSL

The SSL button provides information and statistics on SSL key data, certificate revocation and OCSP counts, and SSL logs.

- ◆ Click the **SSL Key Data** button to see information about SSL service health (Is alive), connection statistics, and SSL session cache hits and misses.

Click **CRL Statistics** to see CRL and OCSP statistics.

Click Reports to generate Certificate Authority reports and Incident reports.

Viewing statistics from the command line

Help | Content Gateway | Version 7.8.x

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

You can also configure, stop, and restart Content Gateway from the command line. See [Command-line interface](#), page 112, and [Websense Content Gateway variables](#), page 294.

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 Content Gateway node.
3. From the Content Gateway **bin** directory (/opt/WCG/bin), enter the following command:

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

where **variable** is the variable that holds the information you want. For a list of the variables you can specify, see [Websense Content Gateway variables](#), page 294.

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

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

Working with alarms

Help | Content Gateway | Version 7.8.x

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 **Content Gateway subscription download failed: error connecting** 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.

| Node | Alarm | Clear |
|--------------|--|--------------------------|
| d1-rhe5u3-01 | [Tue Jan 31 14:13:53 2012] After several attempts, Content Gateway failed to connect to the Policy Server. Please troubleshoot the connection. | <input type="checkbox"/> |



Note

Content Gateway also sends select alarms to the Web Security manager, where they are referred to as **alerts**. Summary alert messages are displayed on the Web Security **Status > Today** page. Web Security administrators can configure which Content Gateway conditions cause alert messages to be sent, and which methods (email or SNMP) are used to send the alert, on the **Settings > Alerts** pages.

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 498](#), provides a description of some of the alarm messages that Content Gateway generates.



Important

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

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

Configuring Content Gateway to email alarm messages

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@example.com`
3. Click **Apply**.

Using a script file for alarms

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

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

Using Performance graphs

Help | Content Gateway | Version 7.8.x

The Performance graphing tool (Multi Router Traffic Grapher) allows you to monitor 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.

Go to **Monitor > Performance** to access performance graphs.



Important

To run Multi Router Traffic Grapher (the Performance graphing tool), you must have Perl v5.005 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 (/opt/WCG/bin).
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 SSL-related reports

Help | Content Gateway | Version 7.8.x

You can request a report detailing the status of certificate authorities (see [Certificate Authorities](#), page 124) or listing incidents (see [Incidents](#), page 126).

Reports can be either in HTML or comma-separated format. The comma-separated reports appear as Excel spreadsheets.

Certificate Authorities

Help | Content Gateway | Version 7.8.x

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 CSV, 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 163.
5. Click **Generate Report** to generate the report.

HTML output looks like this:

| Certificate Authorities | | Incidents | | | |
|---|------------|------------|-----------|------------|-----------------|
| Validation Reports | | | | | |
| Certificate Authority | Count good | Percentage | Count bad | Percentage | Last Access Dat |
| Go Daddy Class 2 Certification Authority | 519 | 26.04 % | 0 | 0.00 % | 2014-01-09 14: |
| Go Daddy Secure Certification Authority | 519 | 26.04 % | 0 | 0.00 % | 2014-01-09 14: |
| VeriSign Class 3 International Server CA - G3 | 69 | 3.46 % | 0 | 0.00 % | 2014-01-09 15: |
| GeoTrust Global CA | 2 | 0.10 % | 0 | 0.00 % | 2014-01-10 08: |
| GeoTrust SSL CA | 1 | 0.05 % | 0 | 0.00 % | 2014-01-10 08: |
| Entrust.net Certification Authority (2048) | 1 | 0.05 % | 0 | 0.00 % | 2014-01-10 08: |
| DigiCert High Assurance EV Root CA | 84 | 4.21 % | 0 | 0.00 % | 2014-01-10 09: |
| GlobalSign Organization Validation CA | 2 | 0.10 % | 0 | 0.00 % | 2014-01-10 09: |
| GlobalSign Root CA | 2 | 0.10 % | 0 | 0.00 % | 2014-01-10 09: |
| Thawte SSL CA | 3 | 0.15 % | 0 | 0.00 % | 2014-01-10 09: |
| Thawte Premium Secure CA | 2 | 0.10 % | 0 | 0.00 % | 2014-01-10 09: |

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

| | A | B | C | D | E | F | G |
|----|---|-------|--------|-----------|-------|----------------|-----------|
| 1 | CSV Report of EVA - Certificate Authorities | | | | | | |
| 2 | | | | | | | |
| 3 | Certificate Authority | Count | goc | Percentag | Count | bac | Percentag |
| 4 | Go Daddy Class 2 Certification Authority | 519 | 26.04% | 0 | 0.00% | 1/9/2014 14:28 | |
| 5 | Go Daddy Secure Certification Authority | 519 | 26.04% | 0 | 0.00% | 1/9/2014 14:28 | |
| 6 | VeriSign Class 3 International Server CA - G3 | 69 | 3.46% | 0 | 0.00% | 1/9/2014 15:13 | |
| 7 | GeoTrust Global CA | 2 | 0.10% | 0 | 0.00% | 1/10/2014 8:30 | |
| 8 | GeoTrust SSL CA | 1 | 0.05% | 0 | 0.00% | 1/10/2014 8:30 | |
| 9 | Entrust.net Certification Authority (2048) | 1 | 0.05% | 0 | 0.00% | 1/10/2014 8:30 | |
| 10 | DigiCert High Assurance EV Root CA | 84 | 4.21% | 0 | 0.00% | 1/10/2014 9:50 | |
| 11 | GlobalSign Organization Validation CA | 2 | 0.10% | 0 | 0.00% | 1/10/2014 9:51 | |
| 12 | GlobalSign Root CA | 2 | 0.10% | 0 | 0.00% | 1/10/2014 9:51 | |
| 13 | Thawte SSL CA | 3 | 0.15% | 0 | 0.00% | 1/10/2014 9:52 | |
| 14 | Thawte Premium Server CA | 3 | 0.15% | 0 | 0.00% | 1/10/2014 9:52 | |
| 15 | DigiCert High Assurance CA-3 | 4 | 0.20% | 0 | 0.00% | 1/10/2014 9:52 | |
| 16 | Entrust.net Secure Server Certification Authority | 4 | 0.20% | 0 | 0.00% | 1/10/2014 9:52 | |
| 17 | GTE CyberTrust Global Root | 33 | 1.66% | 0 | 0.00% | 1/31/2014 4:48 | |
| 18 | VeriSign Class 3 Public Primary Certification Authority | 15 | 0.75% | 0 | 0.00% | 1/31/2014 4:48 | |
| 19 | VeriSign Class 3 Extended Validation SSL CA | 9 | 0.45% | 0 | 0.00% | 1/31/2014 4:48 | |
| 20 | AddTrust External CA Root | 7 | 0.35% | 0 | 0.00% | 1/31/2014 4:49 | |
| 21 | UTN-USERFirst-Hardware | 7 | 0.35% | 0 | 0.00% | 1/31/2014 4:49 | |
| 22 | Class 3 Public Primary Certification Authority | 289 | 14.50% | 0 | 0.00% | 1/31/2014 4:49 | |
| 23 | Equifax Secure Certificate Authority | 420 | 21.07% | 0 | 0.00% | 2/13/2014 8:54 | |
| 24 | | | | | | | |



Note

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

Incidents

Help | Content Gateway | Version 7.8.x

To generate a report of SSL 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 support was enabled
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 HTTPS website access, page 164](#).
5. Click **Generate Report** to generate the report.

HTML output looks like this:

Validation Reports

HTML Report of EVA - Incidents

Profile: default_default

| 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:

Validation Reports

CSV Report of EVA - Incidents

Profile: default_default

| Hostname | Count | Percentage | last modification |
|----------------------------|-------|------------|-------------------|
| data.coremetrics.com:443 | 12 | 7.84% | ##### |
| tc.bankofamerica.com | 2 | 1.31% | ##### |
| *.coremetrics.com | 2 | 1.31% | ##### |
| egov.ins.usdoj.gov | 4 | 2.61% | ##### |
| egov.immigration.gov:443 | 2 | 1.31% | ##### |
| *.usps.com | 2 | 1.31% | ##### |
| urs.microsoft.com | 19 | 12.42% | ##### |
| revoked.microdasys.net | 9 | 5.88% | ##### |
| revoked.microdasys.net:443 | 11 | 7.19% | ##### |
| www.microdasys.net | 3 | 1.96% | ##### |



Note

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

13

Working With Websense Data Security

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Registering and configuring Data Security, page 132](#)
- ◆ [Configuring the ICAP client, page 135](#)
- ◆ [ICAP failover and load balancing, page 136](#)

Websense Content Gateway works together with Websense Data Security components to support:

- ◆ The Threats dashboard with Web Security Gateway
- ◆ Web data loss prevention (DLP) and the Threats dashboard with Web Security Gateway Anywhere

Threats dashboard with Web Security Gateway

When Content Gateway is deployed with Web Security Gateway, several Data Security components are installed on the Content Gateway and TRITON management servers to support the Web Security Threats dashboard (see [Web Security Help](#)). Components include the Data Security Policy Engine (on the Content Gateway machine), and the Data Security Forensics Repository (on the TRITON management server).

Content Gateway registers with these components when it's first configured and then checks registration status upon restart, automatically re-registering if necessary.

Web DLP and the Threats dashboard with Web Security Gateway Anywhere

When Content Gateway is deployed with Web Security Gateway Anywhere (or with Web Security Gateway and a full Data Security subscription), capabilities include

forensics data in the Threats dashboard and data loss prevention (DLP) over Web channels such as HTTP, HTTPS, FTP, and FTP over HTTP. (A full Data Security deployment can extend Web DLP to include channels such as mobile devices, removable media, and printers. For a complete description of Websense Data Security, visit the Data Security product page at www.websense.com.)

Web DLP, as well as extended Data Security configurations, require separate installation of Data Security and other Data Security components. Before configuring Content Gateway to work with Data Security, see the deployment and installation information hosted in the [Websense Technical Library](#).

Content Gateway supports 2 methods of working with Data Security:

- ◆ Preferred: Data Security components installed with Content Gateway.
- ◆ Over ICAP using Data Security components located on a separate host. This intended for use with Data Security Suite versions 7.1 and earlier.

Only one method can be used at a time.

How Web DLP works

In addition to the Web DLP data flow described below, enabling a special analytic engine causes outbound traffic to be analyzed for data theft. In the Web Security manager, see the **Outbound security** options on **Scanning > Scanning Options**.

Web DLP data flow works as follows:

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 Data Security Web DLP 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.

**Note**

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

- ◆ Content Gateway forwards the block page to the sender in a 403 Forbidden message.
- ◆ The block page must be larger than 512 bytes or some user agents (e.g., Internet Explorer) will substitute a generic error message.
- ◆ The block page can be customized. See [Modifying Data Endpoint Confirm And Block Messages](#).

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

Transaction details are logged by Data Security, per its configuration.

Data Security components on-box with Content Gateway

When Content Gateway is installed, a small number of Data Security components are installed on the same box. Content Gateway registers with these components when it's first configured and then checks the registration status whenever it's restarted, automatically re-registering if necessary. For more information about Data Security registration, see [Registering and configuring Data Security](#), page 132.

After policies have been created and deployed in the Data security manager, Content Gateway sends content, such as postings and uploads, to Data Security for analysis and policy enforcement.

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
- ◆ 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 279.

Data Security over ICAP

When the Data Security policy engine is located on a separate host, Content Gateway can communicate with Data Security over ICAP v1.0. For configuration details, see [Configuring the ICAP client](#), page 135. Note that integration with on-box components is the preferred deployment.

Registering and configuring Data Security

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Configuring the ICAP client, page 135](#)

For an introduction to Websense Data Security, see [Working With Websense Data Security, page 129](#).

Registration and configuration summary:

- ◆ Registration with on-box Data Security components is automatic. No configuration is required.
Threat dashboard forensics data is collected automatically by Websense Web Security.
If registration fails, an alarm displays.
- ◆ Registration with off-box Data Security Management Server is automatic after **Configure > My Proxy > Basic > Data Security > Integrated on-box** is enabled and Content Gateway is restarted.
Content Gateway queries the TRITON console for the presence of Data Security Management Server.



Important

Content Gateway and Data Security Management Server system times should be synchronized to within a few minutes.

Registration is tested and retried, if needed, every time Content Gateway is started.

If automatic registration fails, an alarm displays.



Important

Data Security and Content Gateway communicate over several ports. If IPTables are configured on the Content Gateway host system, these ports must be open in IPTables. See these Technical Library articles: [Content Gateway Ports](#) and [Configuring IPTables for Websense Content Gateway](#).

- ◆ Web DLP policies are configured in the Data Security manager in the **System Modules** section. You must **deploy** the Data Security policies to put them into effect. See Data Security Help for details.

- ◆ View registration status in the Content Gateway manager on the **Monitor > Summary** page by clicking **More Detail** and checking the list at the bottom of the **Subscription Details** section.
- ◆ Registration success and failure information is logged in: /opt/WCG/logs/dss_registration.log

Registration and configuration details

Whether you are deploying Web Security Gateway or Web Security Gateway Anywhere, registration with the Forensics Repository is automatic. There is no additional configuration.

If you are deploying Web Security Gateway Anywhere to use Web DLP, you must enable the Data Security integration in the Content Gateway manager:

- ◆ Go to **Configure > My Proxy > Basic** and enable **Data Security > Integrated on-box**. If this option is **not** enabled, registration is with the Forensics Repository only.



Important

Before enabling **Data Security > Integrated on-box**, ensure that the Content Gateway and Data Security Management Server computers are running and accessible, and that their system clocks are synchronized within a few minutes.

After **Data Security > Integrated on-box** is enabled, registration with **Data Security Management Server** is automatic and is performed, if needed, every time that Content Gateway starts. To perform registration, Content Gateway queries Websense Web Security Policy Broker for needed information, including IP address and cluster ID.

Registration status can be viewed in the Content Gateway manager on the **Monitor > Summary** page by clicking **More Detail** and reviewing the list at the bottom of the **Subscription Details** section.

Once registered, Content Gateway uses the Web DLP policy engine for malware detection. Go to the Data Security manager to configure and deploy Web DLP policies. You must deploy Web DLP in the Data Security manager.

If automatic registration fails, an alarm displays.

Manual registration

After **Data Security > Integrated on-box** is enabled and Content Gateway has been restarted, you can attempt a manual registration by going to **Configure > Security > Data Security** (see below).

Restarting Content Gateway always checks the registration status and initiates an auto-registration attempt, if needed.

Registration success and failure information is logged in: /opt/WCG/logs/dss_registration.log



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.

Manual registration with Data Security Management Server:

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. Ensure that **Data Security > Integrated on-box** is enabled. In the Content Gateway manager select **Configure > Basic > General**. In the list of **Features**, under **Networking** locate **Data Security**, select **On**, then select **Integrated on-box**, and then click **Apply**.
3. Next to **Integrated on-box**, click the **Not registered** link. This opens the **Configure > Security > Data Security** registration screen.
4. Enter the IP address of the **Data Security Management Server**.
5. Enter a user name and password for logging onto the Data Security manager. The user must be a Data Security administrator with Deploy Settings privileges.
6. Click **Register**. If registration is successful, a message confirms the result and prompts you to restart Content Gateway.

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

Configuration options

When registration is successful, on the **Configure > Security > Data Security** page set the following options:

1. **Analyze FTP Uploads**: Select this option to send FTP uploads to Data Security for analysis and policy enforcement.
2. **Analyze HTTPS Content**: Select this option to send decrypted HTTPS posts to Data Security for analysis and policy enforcement. The HTTPS protocol option must be enabled on Content Gateway. See .



Note

For these options to have any effect, Content Gateway must be configured to proxy FTP and HTTPS traffic.

3. Click **Apply** to save your settings and then restart Content Gateway.
4. Go to the Data Security manager to configure the Data Security Content Gateway module. See “Configuring the Web Content Gateway module” in *Data Security Manager Help*.

Data Security and Content Gateway communicate over several ports. If IPTables are configured on the Content Gateway host system, these ports must be open in IPTables. See these Technical Library articles: [Content Gateway Ports](#) and [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
-

Configuring the ICAP client

Help | Content Gateway | Version 7.8.x

ICAP can be used with any version of Websense Data Security, however **the direct interface is recommended when the policy engine is on-box with Content Gateway**. See [Registering and configuring Data Security, page 132](#).

ICAP **must** be used for inter-operation with Data Security Suite versions 7.1 and earlier.



Note

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

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

See [ICAP failover and load balancing](#), below.

To configure integration with ICAP, log on to the Content Gateway manager and go to **Configure > My Proxy > Basic > General**.

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 primary ICAP service, followed by a comma (no space) and the URI of the secondary ICAP service. A secondary ICAP service is optional.

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 for analysis or sent directly to the destination. The HTTPS protocol option must be enabled to send traffic to Data Security. See [Working With Encrypted Data](#), page 139.
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. See [FTP](#), page 326.
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 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 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.

ICAP failover and load balancing

Help | Content Gateway | Version 7.8.x

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

Load balancing between 2 ICAP servers is also an option.

Time to failover

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

Failure conditions leading to failover

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

Excluded failure conditions

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

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

Recovery Conditions

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

Content Gateway tests for recovery conditions for each down ICAP server at a specified interval. If load balancing is disabled, requests continue to be sent to a secondary ICAP server until the primary comes back online. If load balancing is enabled, Content Gateway starts sending requests to a server (round-robin) as soon as it is marked up.

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

Recovery actions

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

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

Fail open

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

Configuration settings

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

| Configuration Variable | Data Type | Default Value | Description |
|---------------------------------|-----------|---------------|---|
| proxy.config.icap.ICAPUri | STRING | (empty) | A comma-separated list of ICAP URIs. For example: icap://1.2.3.4:1344/reqmod, icap://4.3.2.1:1344/reqmod |
| proxy.config.icap.ActiveTimeout | INT | 5 | The read/response timeout in seconds. The activity is considered a failure if the timeout is exceeded. |
| proxy.config.icap.RetryTime | INT | 5 | The recovery interval, in seconds, to test whether a down server is back up. |
| proxy.config.icap.FailOpen | INT | 1 | Set to: <ul style="list-style-type: none"> • 1 to allow traffic when the ICAP servers are down • 0 to send a block page if the ICAP servers are down |
| proxy.config.icap.LoadBalance | INT | 1 | Set to: <ul style="list-style-type: none"> • 1 to distribute requests to all available servers • 0 to distribute requests to only the primary server. |

14

Working With Encrypted Data

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Running in explicit proxy mode, page 141](#)
- ◆ [Initial SSL configuration tasks, page 144](#)
- ◆ [Enabling SSL support, page 143](#)
- ◆ [Certificates, page 145](#)
- ◆ [Internal Root CA, page 145](#)
- ◆ [Managing certificates, page 154](#)
- ◆ [SSL configuration settings for inbound traffic, page 156](#)
- ◆ [SSL configuration settings for outbound traffic, page 158](#)
- ◆ [Validating certificates, page 159](#)
- ◆ [Managing HTTPS website access, page 164](#)
- ◆ [Client certificates, page 169](#)
- ◆ [Customizing SSL connection failure messages, page 171](#)

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

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

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

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

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

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

The advanced proxy support is simply called HTTPS support or SSL support. How it works and how it's configured is described in the following sections.

In the Content Gateway manager, SSL support is enabled on the **Configure > My Proxy > Basic > General** page in the **Protocols** area with the **HTTPS** option.



Important

Even when HTTPS support is **not** enabled and HTTPS is not decrypted, Content Gateway performs URL filtering. This means that for every HTTPS request received from a client, a URL lookup is performed and policy is applied.

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

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




Note

Content Gateway does not cache HTTPS content.

When HTTPS is enabled, each HTTPS request consists of two separate sessions:

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

Different certificates are required for each session.



For additional information on SSL, TLS, and SSL/TLS certificates, search the Internet or consult any of the commercially available books.

Running in explicit proxy mode

Help | Content Gateway | Version 7.8.x

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 use as a basis for building a custom PAC file.

1. On the **Configure > My Proxy > Basic > General** tab, ensure that HTTPS is enabled. If it is disabled, set it to **On**, click **Apply**, and **Restart** Content Gateway.
2. Go 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. This template is for basic testing only. Further modify this file to meet all of your organization's needs.

```
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 hostname **proxy1** and Content Gateway uses the default port 8083 to serve the file, users must specify the following URL in the proxy configuration settings:

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

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

For Microsoft Internet Explorer version 7.0 and later:

1. Go 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. Go 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 your browser documentation for details.

Enabling SSL support

Help | Content Gateway | Version 7.8.x

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



Note

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

2. Click **Apply** and then **Restart**.
3. Enter the name of the SSL certificate file. See [Creating a subordinate CA](#), page 148.

Use the **Configure > Protocols > HTTPS** page to specify:

- The HTTPS port (default is 8080)
 - Skype tunneling (explicit proxy only)
 - Tunneling when a request returns an **Unknown protocol** error
1. The **HTTPS Proxy Server Port** is the port used for client to Content Gateway connections. The default is 8080. If traffic is transparent on 443, a default ARM NAT rule readdresses the requests to 8080. See **Configure > Networking > ARM: Network Address Translation**.
 2. If Content Gateway is an **explicit proxy** and you want to allow Skype traffic, enable the **Tunnel Skype** option. This option is necessary because, although Skype presents an SSL handshake, Skype data flow does not conform to the SSL standard. Unless the traffic is tunneled, the connection is dropped.

To complete the configuration, in the Web Security manager ensure that filtering policies that apply to users of Skype allow “Internet telephony”. This is required for users of Skype whether HTTPS support is enabled or not.

Also, if not prevented, after the initial handshake Skype will route traffic over a non-HTTP port. To force Skype traffic to go through Content Gateway, a GPO should be used as described in the [Skype IT Administrators Guide](#).



Important

There is no need to set this option if HTTPS support is not enabled.

This option is not valid and has no effect when Content Gateway is a transparent proxy.

3. To tunnel HTTPS requests when the SSL handshake results in an unknown protocol error, enable **Tunnel Unknown Protocols**.

By default, Content Gateway will not try to tunnel non-ssl traffic. Beginning with 7.8.3, a new variable is available that will enable tunneling of non-ssl traffic when

Content Gateway is an **explicit proxy**. Beginning with 7.8.4, the same variable is used for transparent proxy as well.

Add the following to the records.config file (in `/opt/WCG/config`, by default) to turn on tunneling of non-ssl traffic.

```
CONFIG proxy.config.ssl_decryption_bypass.tunnel_non-ssl_traffic INT 1
```

Reset the value to 0 to disable the feature and turn off tunneling of non-ssl traffic.

A restart of Content Gateway is required for this setting to take affect.



Warning

Tunneled connections are not decrypted or inspected.

Also, for 7.8.1 only, the setting of this option persists even when the HTTPS feature is disabled on **Configure > My Proxy > Basic > General**. Because this page is only displayed when HTTPS is enabled, you should disable this option before disabling HTTPS.

Web Security behavior varies based on the type of proxy deployment.

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

Initial SSL configuration tasks

Help | Content Gateway | Version 7.8.x

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

1. Create an internal root CA (certificate authority). In order to sign SSL traffic, Content Gateway requires an internal SSL Certificate Authority that has the ability to sign SSL certificates. This is for traffic between the browser and Content Gateway. See [Internal Root CA, page 145](#).
2. Add this CA to the certificate tree. Servers, such as destination servers, check this tree to ensure that they can trust users because they have certificates from an authority listed here. The certificates listed on the certificate tree are certificate authorities you empower (trust) to verify the validity of individual websites. Any site signed by a certificate authority in the certificate tree with the “allow” status is allowed through Content Gateway. See [Managing certificates, page 154](#)

3. Customize pages that browser users will see. See [Customizing SSL connection failure messages](#), page 171. Among the pages that can be customized are a connect failure and certificate verification failure page.

Certificates

Help | Content Gateway | Version 7.8.x

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

- ◆ It must be current (not expired or revoked). See [Validating certificates](#), page 159.
- ◆ It must be issued by a trusted CA (certificate authority). See [Managing certificates](#), page 154
- ◆ The URL and the certificate owner must match. See [Configuring validation settings](#), page 159.

HTTPS connections between the client browser and Content Gateway require a certificate issued by an internal CA. See [Internal Root CA](#), page 145.

Connections between Content Gateway and the origin server require a certificate signed by one of the certificate signing authorities listed in the Certificate Authority Tree on the **Configure > SSL > Certificates > Certificate Authorities** tab. See [Managing certificates](#), page 154.

Internal Root CA

Help | Content Gateway | Version 7.8.x

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

- ◆ You must have an internal Root CA to complete an inbound connection.
- ◆ You can either import or create the internal Root CA.

- ◆ The internal Root CA is stored in the SSL configuration database.



Important

Back up the existing internal Root CA before importing or creating a new one. This enables you to return to an earlier version, if necessary. See [Backing up your internal Root CA, page 153](#) for details.

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



Important

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

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

There are three options for creating an internal Root CA:

- ◆ Leverage an existing corporate CA and import it into Content Gateway. See [Importing your Root CA, page 146](#).
- ◆ Create a new Root CA and make that CA available to browsers. See [Creating a new Root CA, page 147](#).
- ◆ Create a subordinate CA. This leverages a corporate CA, but can also be revoked by the corporate CA. See [Creating a subordinate CA, page 148](#).

Importing your Root CA

Help | Content Gateway | Version 7.8.x

If your organization already has a Root CA, you can import it. This certificate must be trusted by all browsers in your organization. Be sure to back up any new internal Root CA that you import. See [Backing up your internal Root CA, page 153](#) for details.

1. Go 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.



Important

The certificate and private key format must match.

Additionally, the private key format must match the format required by the importing node (unencrypted or encrypted).

To verify the certificate and private key format, view the files in a text editor. Use **Backup Root CA** to export the CA from the database.

For information on converting the private key format, see:

- ◆ [Preparing an Internal Root CA for importing into a FIPS 140-2 enabled node](#)
- ◆ [Converting an RSA key type to a PKCS#8 key type](#)
- ◆ [Converting an encrypted private key to an RSA key](#)

4. Enter and confirm the passphrase.
5. Click **Import Root CA**. The imported CA is stored in the SSL configuration database.

Creating a new Root CA

Help | Content Gateway | Version 7.8.x

Related topic:

- ◆ [Creating a subordinate CA, page 148](#)

If you do not already have a Root CA, fill in the fields on this tab to create one.

The process uses **openssl pkcs#8**.

Be sure to back up any new Root CAs that you create. See [Backing up your internal Root CA, page 153](#) 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**, 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 subordinate CA

Help | Content Gateway | Version 7.8.x

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

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

Preparation

- ◆ **If you are not the Enterprise domain administrator, you will need to work with that person to get the correct domain permissions to generate a sub CA.**
- ◆ Install the **OpenSSL 1.0.1e** toolkit (www.openssl.org) on a Windows or Linux computer.

Creating a Certificate Signing Request (CSR)

1. Create a CSR with OpenSSL.
In a Windows Command Prompt or on the Linux command line, create a CSR with the following **openssl** command:

```
openssl req -new -newkey rsa:2048 -keyout wcg.key -out
wcg.csr
```

```
[root@JS-WCG ~]# openssl req -new -newkey rsa:2048 -keyout wcg.key -out wcg.csr
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'wcg.key'
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [GB]:US
State or Province Name (full name) [Berkshire]:California
Locality Name (eg, city) [Newbury]:San Diego
Organization Name (eg, company) [My Company Ltd]:Websense, INC.
Organizational Unit Name (eg, section) []:Technical Support
Common Name (eg, your name or your server's hostname) []:10.212.4.164
Email Address []:

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
[root@JS-WCG ~]# █
```

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

The openssl command generates 2 files:

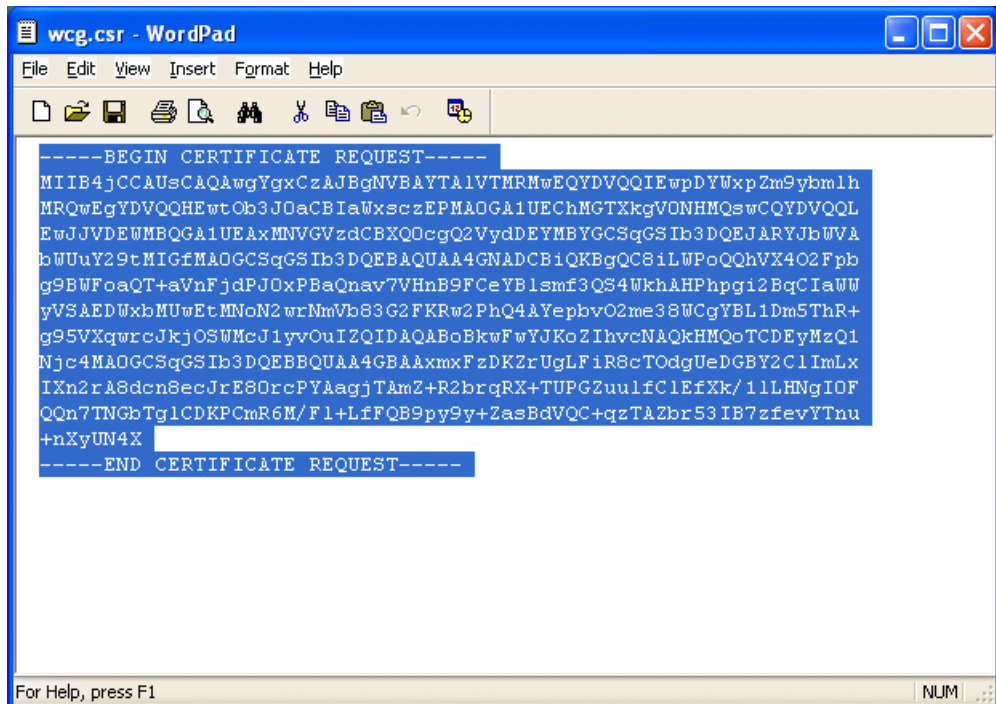
- **wcg.csr** -- the CSR that will be signed by the Certificate Authority to create the final certificate
- **wcg.key** -- the private key

3. If you created the CSR on a Linux system, copy it to your Windows host with WinSCP or some other file transfer utility.

Signing the request

You must sign the request with Microsoft Certificate Services.

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

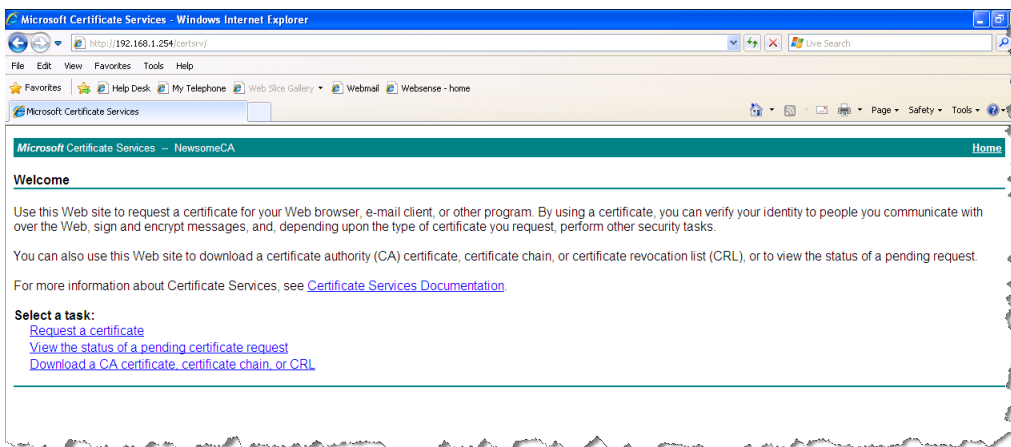


2. In **Internet Explorer**, go to the **Microsoft CA server**.

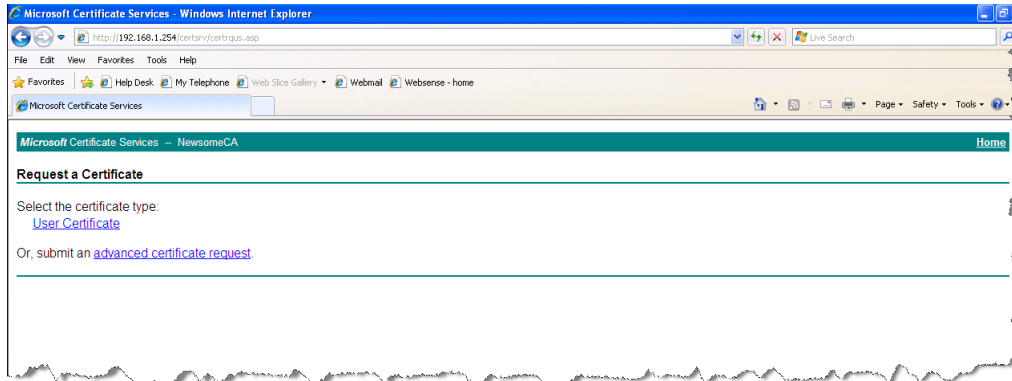
Enter the following URL:

http://<CA_server_IP_address>/certsrv

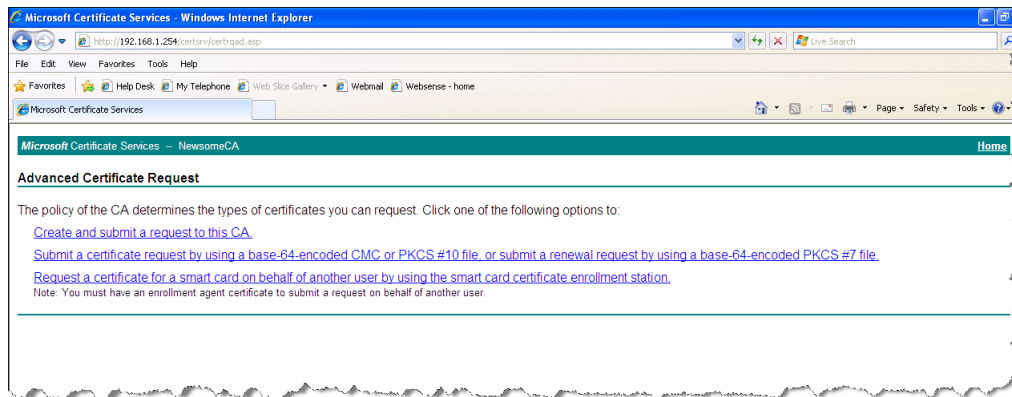
The **Certificate Services** applet starts.



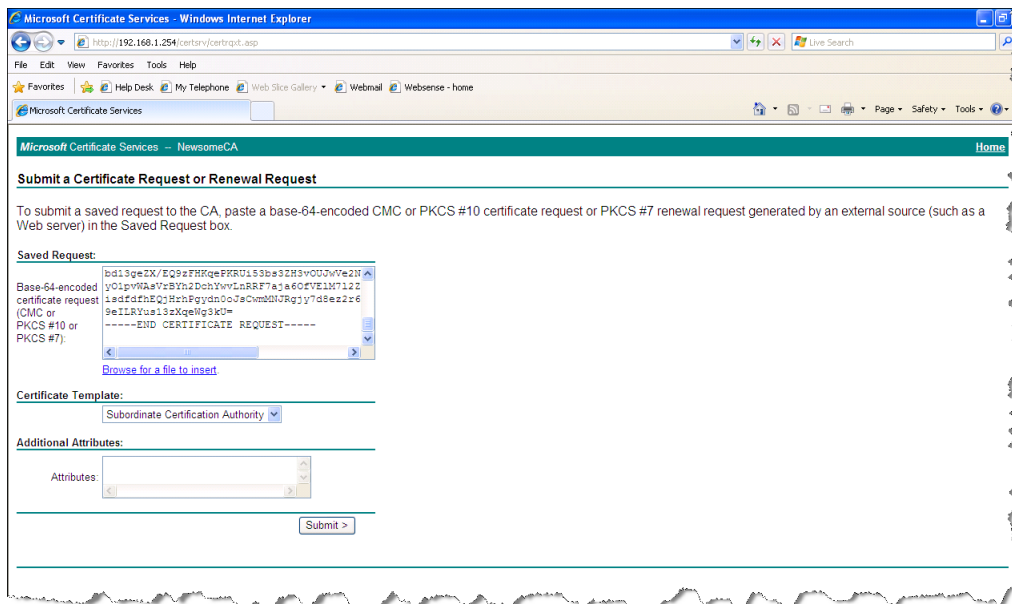
3. On the **Welcome** screen, below the **Select a task** heading, select **Request a certificate**. The **Request a certificate** page displays.



4. Select to submit an **advanced certificate request**.




- On the **Advanced Certificate Request** screen, select **Submit a certificate request by using a base-64-encoded CMC**. The **Submit a Certificate Request or Renewal Request** screen displays.



- On the **Submit a Certificate Request or Renewal Request** screen, paste the content of the **wcg.csr** file (previously placed on the clipboard) in the **Certificate Template** drop down window and click **Submit**.

The certificate is issued and the **Certificate Issued** screen displays. If, instead, the **Certificate Pending** screen displays, you do not have sufficient privileges to create a sub CA. Contact your Enterprise domain administrator to complete the certificate creation process and then proceed to step 7.



- Select the **Base 64 encoded** radio button and then select **Download certificate**. Save the certificate to your desktop. Later you will import it into Content Gateway.

With the base 64 encoded certificate on your desktop, along with the private key created during the CSR generating process, you are ready to import both into Content Gateway.

Importing the sub-CA into Content Gateway

1. Open the Content Gateway manager and go to **Configure > SSL > Internal Root CA > Import Root CA**.

The screenshot shows the 'Import Root CA' configuration page in the Content Gateway manager. The page has a blue header with 'websense Content Gateway' and 'User: admin Log'. A sidebar on the left contains navigation options: Monitor, Configure, My Proxy, Protocols, Content Routing, Security, Subsystems, Networking, and SSL. The main content area is titled 'Import Root CA' and includes buttons for 'Create Root CA' and 'Backup Root CA'. Below this, there is a section for 'Internal Root CA' with a 'Certificate' field, a 'Private key' field, a 'Passphrase' field, and a 'Confirm passphrase' field. Each of the first two fields has a 'Browse...' button and a red warning message: 'Please use only base64-encoded certificates.' At the bottom of the form is an 'Import Root CA' button.

2. **Browse** to select the certificate. The certificate must be in X.509 format and base-64-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**.
6. Restart Content Gateway.

Backing up your internal Root CA

Help | Content Gateway | Version 7.8.x

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. Go 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

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Adding new certificate authorities, page 155](#)
- ◆ [Backing up certificates, page 156](#)
- ◆ [Restoring certificates, page 156](#)

Content Gateway initially populates the Certificate Authority Tree (trusted certificate store) with the list qualified by Mozilla for Firefox ([see this mozilla.org page](#)). The CA tree is listed on the **Configure > SSL > Certificates > Certificate Authorities** tab. Content Gateway trusts origin servers that offer these certificates.

In the list a small “i” appears before the names of certificates that can be validated via CRL (certificate revocation lists) or OCSP (online certification status protocol). See [Keeping revocation information up to date, page 163](#) for information about checking the revocation status of a certificate. Content Gateway checks the revocation status of certificates used for both inbound and outbound traffic.

Click on the name of a certificate authority to:

- ◆ [View a certificate, page 154](#)
- ◆ [Delete a certificate, page 154](#)
- ◆ [Change the allow/deny status of a certificate, page 155](#)

View a certificate

1. Go 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. Go 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. Go 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

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Backing up certificates, page 156](#)
- ◆ [Restoring certificates, page 156](#)

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 156](#). If you want to back up your entire Content Gateway configuration, see [Saving and restoring configurations, page 114](#).

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 155](#) for additional information.

Backing up certificates

Help | Content Gateway | Version 7.8.x

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.

Use the page **Configure > SSL > Certificates > Backup Certificates** to back up certificates.

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

To back up your entire Content Gateway configuration, see [Saving and restoring configurations](#), page 114.

Restoring certificates

Help | Content Gateway | Version 7.8.x

Use **Configure > SSL > Certificates > Restore Certificates** to restore certificates. The certificate database is propagated around the cluster.

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

Help | Content Gateway | Version 7.8.x

[SSL configuration settings for inbound traffic](#), page 156

[SSL configuration settings for outbound traffic](#), page 158

SSL configuration settings for inbound traffic

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [SSL configuration settings for outbound traffic](#), page 158

Use **Configure > SSL > Decryption / Encryption > Inbound** to configure SSL and TLS settings and ciphers for inbound traffic.

- Under **Protocol Settings**, indicate which protocols you want Content Gateway to support. Supported protocols are:
 - SSLv2
 - SSLv3 (enabled by default)
 - TLSv1 (enabled by default)



Note

TLSv1.1 and TLSv1.2 are also supported.

Both are enabled by default for inbound connections, but not for outbound (to destination server) connections.

Each is enabled/disabled with **records.config** variables:

`proxy.config.ssl.client.TLSv11 INT 1 --(0 = disabled)`

`proxy.config.ssl.client.TLSv12 INT 1 --(0 = disabled)`

When Content Gateway is on a Websense appliance, use the Appliance manager Toolbox Command Line Utility to set the value.

When Content Gateway is installed as software, use “`content_line -s`” on the Linux command line to set the value.

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

You must select at least one protocol.

These settings override the settings for these protocols in the users’ browsers.

You can select different protocols for outbound traffic.

- The cipher list describes available algorithms and level of encryption between the client and Content Gateway.

The **Default** setting indicates to use all available ciphers **except** the eNULL, ADH, and EXP suites.

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

Additional cipher settings are:

- **High** encryption cipher suites: 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 (clients connections to Content Gateway), consider using Low encryption to improve performance.

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

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

SSL configuration settings for outbound traffic

Help | Content Gateway | Version 7.8.x

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

1. Under **Protocol Settings**, indicate which protocols you want Content Gateway to support. Supported protocols are:
 - SSLv2
 - SSLv3 (enabled by default)
 - TLSv1 (enabled by default)



Note

TLSv1.1 and TLSv1.2 are also supported.

Both are disabled by default for outbound connections.

They are enabled/disabled with records.config variables:

```
proxy.config.ssl.client.TLSv11 INT 0 --(0 = disabled)
```

```
proxy.config.ssl.client.TLSv12 INT 0 --(0 = disabled)
```

When Content Gateway is on a Websense appliance, use the Appliance manager Toolbox Command Line Utility to set the value.

When Content Gateway is installed as software, use “content_line -s” on the Linux command line to set the value.

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

You must select at least one protocol.

You can select different protocols for inbound traffic.

2. Select **Session Cache** if you want to cache keys until the time specified in **Session Cache Timeout** expires. This can improve performance. If keys are not cached, each request is negotiated again.

Setting the **Session Cache Timeout** to 0 (zero) causes session caching to be disabled.
3. Indicate, in seconds, how long keys should be kept in the cache. The default is 300 seconds (5 minutes).
4. The cipher list describes available algorithms and level of encryption between Content Gateway and the origin server.

The **Default** setting indicates to use all available ciphers **except** the eNULL, ADH, and EXP suites.

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

Additional cipher settings are:

- **High** encryption cipher suites: 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, consider using one of the higher encryption levels to improve security.

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

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

Validating certificates

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Bypassing verification, page 162](#)
- ◆ [Keeping revocation information up to date, page 163](#)

SSL certificate verification is an important component of SSL security. It is through certificate exchange and verification that the client, in this case Content Gateway, and the origin server verify that each is who it says it is.

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

Use the tabs on **Configure > My Proxy > SSL > Validation** to enable and configure the CVE.

For information about options when verification fails and you prefer to trust the site, see [Bypassing verification, page 162](#).

For a comprehensive discussion of the use and best practices of the CVE, see [SSL Certificate Verification Engine v7.8](#).

Configuring validation settings

1. Go to **Configure > SSL > Validation > General**.
2. **Enable the certificate verification engine:** This option enables and disables the certificate verification engine.

Certificate verification is **disabled** by default. This prevents the Content Gateway administrator and network users from being surprised by the effects of certificate verification when HTTPS is initially enabled (on **Configuration > My Proxy > Basics > General**).

If this option is not selected, certificate validation does not occur.



Important

If you disable the CVE, you need to provide settings on only the following pages:

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

3. **Deny certificates where the common name does not match the URL:** When enabled, two checks are made:
 - First, the certificate's Common Name is checked for an exact match of the destination URL.
 - If the first check fails, the certificate's Subject Alternative Name (SAN) list is checked for an exact match of the destination URL.

Checks are case **insensitive**.

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

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

4. **Allow wildcard certificates:** This is a sub-option of **When Deny Certificates where the common name does not match the URL**. When enabled, this option allows matches with Common Names that include the "*" (wildcard) character in the name.

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

Use of the wildcard means that individual servers within the domain are not verified; they are included as a result of the wildcard.

Allowing wildcard certificates eases the strict matching burden when a Common Name match is required. It is also helpful for domains that have multiple subdomains like google.com or yahoo.com. It also introduces some risk that a fraudulent or undesirable variation of a domain may go unblocked.

5. **No expired or not yet valid certificates:** When enabled, denies access to sites that offer an expired or not yet valid certificate. This is a basic check that is important because many malicious sites operate with expired certificates. If this option is not selected, access to those sites is permitted.

**Note**

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

6. **Verify entire certificate chain:** When enabled, verifies expiration and revocation status of all certificates between the site certificate and the root Certificate Authority as specified in the certification path of the certificate. This is an important check.
7. **Check certificate revocation by CRL:** Certificate revocation lists (CRLs) are used to check a certificate's revocation status. CRLs list certificates that have been issued and subsequently revoked by the CA.

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

If this option is enabled, it is recommended that you verify that the daily CRL update feature is enabled. Go to the **Revocation Settings** tab and enable the check box in **CRL Settings**.

If this option is **not** used, it is recommended that you disable the daily CRL update feature. Go to the **Revocation Settings** tab and disable the check box in **CRL Settings**.

8. **Check certificate revocation by OCSP:** Online Certificate Status Protocol (OCSP) is an alternate way to check a certificate's revocation status. While OCSP is beneficial, it is not used as widely as CRLs and therefore is not as reliable. Also, it is a real-time, Internet-hosted check that can introduce some request handling latency.

**Note**

It is recommended that you use OCSP in addition to, rather than instead of, CRLs. See [Keeping revocation information up to date](#), page 163 for more information on CRLs and OCSP.

9. **Block certificates with Unknown OCSP state:** When OCSP revocation checking is enabled, enable this option to block certificates that return the "Unknown" status.
10. **Preferred method for revocation check:** When both CRL and OCSP revocation checking are enabled, use this option to indicate which method to apply first. The default is CRL.

11. **Block certificates with no CRL URI and with no OCSP URI:** When CRL checking, OCSP checking, or both are enabled, use this option to block certificates that do not have the expected, associated URIs. For example, if only CRL checking is enabled and the certificate doesn't have a CRL URI, if this option is enabled the connection is blocked. When both CRL and OCSP checking are enabled, the block occurs only if both CRL and OCSP lack a URI.

You can view URI information in the certificate when you select to view the certificate in your browser. See [View a certificate, page 154](#) for details.

Because many certificates do not include CRL or OCSP information, this option can result in a high number of verification failures. Often the failures are reported as "Unknown revocation state" errors.

This can result in a highly restrictive security policy, with many access denials.

As with all verification failures, you can allow for exceptions using the Incident List. See [Managing HTTPS website access, page 164](#).

Bypassing verification

Help | Content Gateway | Version 7.8.x

Use the **Configure > SSL > Validation > Verification Bypass** page to allow users to visit a site when certificate verification fails.

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. This is referred to as **verification bypass**. If this check box is not selected, users do not have the option to browse to the site.
2. If verification bypass is enabled, you can specify a period of time, in minutes, that the user is allowed to bypass a particular site without having to click-through the warning again. Specify that period in the **Time before the user is notified again for the site** entry field. The default is 6 minutes.
3. Select **Enable the SSL session cache for bypassed certificates** to store information about bypassed certificates in cache and reuse the connections.
 - If this option is selected, performance is better, but 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, but performance is not as fast.

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 HTTPS website access, page 164](#).

Keeping revocation information up to date

Help | Content Gateway | Version 7.8.x

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 163) and through OCSP (see [Online certification status protocol \(OCSP\)](#), page 163).

Certificate revocation lists

Use the **Configure > SSL > Validation > Revocation Settings** page to configure how Content Gateway keeps revocation information current. By default, Content Gateway 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.



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.

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

Use **Configure > SSL > Validation > Revocation Settings** to configure how Content Gateway 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 HTTPS website access

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Viewing incidents, page 164](#)
- ◆ [Changing the status of an incident, page 166](#)
- ◆ [Deleting an incident, page 167](#)
- ◆ [Changing the text of a message, page 167](#)
- ◆ [Viewing incident details, page 167](#)
- ◆ [Adding websites to the Incident List, page 168](#)

These tabs can help you manage access to websites and can aid the HelpDesk in troubleshooting access issues.

When a client receives an access denial message because the website does not comply with security policy, Content Gateway generates an *incident*. See [Viewing incidents, page 164](#).

If you want to specify how Content Gateway treats a particular site, you can add that to the Incident List as well. See [Adding websites to the Incident List, page 168](#).

Additional information on troubleshooting can be found in [SSL Certificate Verification Engine v7.8](#).

Viewing incidents

Help | Content Gateway | Version 7.8.x

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

A separate Incident List is kept for every node in a cluster. Incidents that are added or modified by the administrator, are copied around the cluster (synchronized).

Unexpected incidents that, by default, result in an access denial message, are not synchronized in the cluster.

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

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

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 |
|-------|--|
| Node | The name of the Content Gateway node on which the list entry is located. |
| ID | Assigned by the system, this is the incident ID number, also called the Ticket ID. The HelpDesk 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. |

| Field | Description |
|---------|--|
| Status | <p>Determines how Content Gateway will treat this website in the future. Four conditions are possible:</p> <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 website is blocked, unless Verification Bypass is configured, in which case the block page includes a “Visit site anyway” button. See Bypassing verification, page 162. • Tunnel The site is tunneled. Traffic is not decrypted and Content Gateway does not check the certificate. Tunneling can be used to bypass inspection of trusted sites and improve performance. Note: Tunnel by URL does not work with all transparent proxy traffic. See Adding websites to the Incident List, page 168. <p>You can change the status of a site via the drop-down box in the Action column.</p> |
| 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 websites to the Incident List, page 168 . |
| URL | The URL of a site whose certificate could not be validated. |
| Message | Enables you to edit the error message. See Changing the text of a message, page 167 for information on customizing error messages. The pencil and the magnifying glass each represent links. See Viewing incident details, page 167 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 167 . |

Changing the status of an incident

Help | Content Gateway | Version 7.8.x

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

1. Go 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 165](#) for an explanation of these options.
 - Tunnel
 - Block
 - Blacklist
 - Allow
3. Click **OK**. The icon in the Status column changes to reflect the new status.

Deleting an incident

Help | Content Gateway | Version 7.8.x

1. Go 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 164](#).
3. In the Action column, select **Delete** from the Action drop-down list, and then click **OK**.

If it is necessary or convenient, the entire Incident List can be deleted using a sqlite3 command:

```
sqlite3 /opt/WCG/config/new_scip3.db "delete from
certificate_acl;"
```

Changing the text of a message

Help | Content Gateway | Version 7.8.x

1. Go to **Configure > SSL > Incidents > Incident List**.
2. Locate the incident you want to examine more closely. See [Viewing incidents, page 164](#).
3. Click the pencil to open a window where you can change the text of this error message. For example, the HelpDesk 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

Help | Content Gateway | Version 7.8.x

1. Go to **Configure > SSL > Incidents > Incident List**.
2. Locate the incident you want to examine more closely. See [Viewing incidents, page 164](#).
3. Click the magnifying glass to see additional details about the incident, such as the:
 - **Description** – the message that appears in the incident listing

- **Created** – The time the incident was created
- **Modified** – The time the incident was modified
- **Access attempts** – How many times users have attempted to access this site

Adding websites to the Incident List

Help | Content Gateway | Version 7.8.x

Use **Configure > SSL > Incidents > Add Website** 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 164](#).

1. Enter the URL of the site you are adding to the Incident List.



Note

When specifying an IPv6 address, **do not** enclose the address in square brackets ([]).

2. Select either **By Certificate** or **By URL**.
 - **By Certificate** provides greater security. If you add a site by certificate, clients cannot bypass the policy by using the IP address rather than the URL. When you select **By Certificate**, Content Gateway retrieves the server certificate and adds the site to the Incident List. See [Viewing incidents, page 164](#).
If sites are blocked by certificates, wildcard certificates are not accepted, even if the common name is recognized.
 - Select **By URL** to tunnel, allow, or blacklist the site.
3. In the Action drop-down list, specify if the site should be added with **Tunnel**, **Allow**, or **Blacklist** status. See [The incident report, page 165](#) for details.

- **Tunnel:** (Valid for **By URL** only) The site is tunneled. Traffic is not decrypted and Content Gateway does not check the certificate.



Important

Tunnel by URL does not work for all transparent proxy requests.

It works under these conditions:

- ◆ When the client application uses TLS and includes an SNI (server name indication), Content Gateway checks the Incident list for the hostname in the SNI.
- ◆ When there is no SNI, Content Gateway connects to the origin server to retrieve the certificate. If the Common Name is a unique FQDN, Content Gateway looks it up in the Incident list. If the Common Name contains a "*" (wildcard), or is not a unique FQDN, Content Gateway looks for the IP address in the Incident list.

Alternatively, use ARM *Static bypass rules*.

- **Allow:** Users can access the site even if the certificate is not valid. Traffic is decrypted, and certificate checking is disabled.
- **Blacklist:** The site is completely blocked. Users cannot access this site even if the Verification Bypass is configured.

4. Click **Apply**.

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 CVE disabled. (See [Configuring validation settings, page 159](#).) This enables you to improve performance by tunneling trusted sites and blocking those you know should not be accessed. See [The incident report, page 165](#) for information about assigning a status, such as tunneling, to a site and incident.

Client certificates

Help | Content Gateway | Version 7.8.x

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

Related topics:

- ◆ [Importing client certificates, page 170](#)
- ◆ [When a client certificate is always required: the Hostlist, page 170](#)
- ◆ [Deleting client certificates, page 171](#)

When a client certificate is requested

1. Go to **Configure > SSL > Client Certificates > General**.
2. Select **Tunnel** or **Create incident** to specify how Content Gateway 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 165](#) for a listing of possible dispositions.
3. Click **Apply**.

Importing client certificates

Help | Content Gateway | Version 7.8.x

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



Important

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

Help | Content Gateway | Version 7.8.x

Use **Configure > SSL > Client Certificates > Hostlist** to list destination servers that always require a client certificate. Be sure to import the certificate before adding it to the Hostlist. See [Importing client certificates, page 170](#).

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**.



Important

For browsers that don't send a Server Name Indicator (SNI), such as Internet Explorer version 8 and earlier, create an entry for both the destination IP address and the hostname.

Deleting client certificates

Help | Content Gateway | Version 7.8.x

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

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

Customizing SSL connection failure messages

Help | Content Gateway | Version 7.8.x

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 172](#).
- ◆ There is a connection failure. See [SSL connection failure, page 172](#).

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 Content Gateway server |
| %u | Complete URL |
| \$\$DETAILS | Detailed error message |
| \$\$TICKET_ID | The ID number of the incident. |

Certificate validation failed

Help | Content Gateway | Version 7.8.x

Use **Configure > SSL > Customization > Certificate Failure** 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.

There is a known problem in Internet Explorer 10 that sometimes results in the wrong block page being displayed. To work around the problem, click on Preview repeatedly until the correct page is displayed, or disable TLS 1.0 in Internet Explorer 10.

1. Edit the HTML code in the window to reflect your message. See [Customizing SSL connection failure messages, page 171](#) 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.

SSL connection failure

Help | Content Gateway | Version 7.8.x

Use **Configure > SSL > Customization > Connect Error** to customize the message users receive when Content Gateway 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.

There is a known problem in Internet Explorer 10 that sometimes results in the wrong block page being displayed. To work around the problem, click on Preview repeatedly until the correct page is displayed, or disable TLS 1.0 in Internet Explorer 10.

1. Edit the text in the window to reflect your message. See [Customizing SSL connection failure messages, page 171](#) 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.

15

Security

Help | Content Gateway | Version 7.8.x

Websense Content Gateway allows you to establish secure communication between the proxy and other computers on the network. You can:

- ◆ Control which clients are allowed to access the proxy. See [Controlling client access to the proxy, page 175](#).
- ◆ Control access to the Content Gateway manager using:
 - Administrator accounts (see [Setting the administrator ID and password, page 176](#) and [Creating a list of user accounts, page 177](#)).
 - SSL (Secure Sockets Layer) protection for encrypted, authenticated access (see [Using SSL for secure administration, page 178](#)).
- ◆ Create filtering rules to control access to the Internet, specify special authentication requirements, and control other traffic transiting the proxy. See [Filtering Rules, page 181](#).
- ◆ Configure Content Gateway integration into your firewall and control traffic through one or more SOCKS servers. See [Configuring SOCKS firewall integration, page 185](#).
- ◆ Configure Content Gateway to use multiple DNS servers to match your site's security configuration. See [Using the Split DNS option, page 188](#).
- ◆ Configure Content Gateway to perform user authentication. The proxy supports Integrated Windows Authentication (with Kerberos), legacy NTLM (NTLMSSP), LDAP, and RADIUS user authentication. There is also support for multiple authentication methods with multiple authentication realms. See [Content Gateway user authentication, page 189](#).

Controlling client access to the proxy

Help | Content Gateway | Version 7.8.x

You can configure Content Gateway to allow only certain clients to use the proxy.

To allow access, specify client IP addresses and IP address ranges in **ip_allow.config**.

To deny access, do not include those client IP addresses in the file.

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**. 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 access to the Content Gateway manager

Help | Content Gateway | Version 7.8.x

You can restrict access to the Content Gateway manager to ensure that only authenticated users can change configuration options and view performance and network traffic statistics.

You can:

- ◆ Set the master administrator ID and password. A user who logs on to the Content Gateway manager with the administrator ID has access to all Content Gateway manager activities. See [Setting the administrator ID and password](#), page 176.
- ◆ Create and maintain a list of user accounts that determines who can log on to the Content Gateway manager and which activities they can perform. See [Creating a list of user accounts](#), page 177.
- ◆ Create an access control list of IP addresses that defines which machines can access the Content Gateway manager. See [Controlling host access to the Content Gateway manager](#), page 178.
- ◆ Use SSL for secure administration. See [Using SSL for secure administration](#), page 178.
- ◆ Require administrators to log on to TRITON Unified Security Center, with or without two-factor authentication, and then use the Web Security manager Content Gateway access page to log on to the Content Gateway manager. See, [Accessing the Content Gateway manager](#), page 11

Setting the administrator ID and password

Help | Content Gateway | Version 7.8.x

During installation, you assign a password that controls administrative access to the Content Gateway manager. A user who logs on to the Content Gateway manager using the correct ID and password can view all the statistics on the Monitor tab and change any configuration options on the Configure tab.

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 the Content Gateway manager unless you have set up a list of IP addresses that are denied access (see [Controlling host access to the Content Gateway manager, page 178](#)).
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 [Accessing the Content Gateway manager if you forget the master administrator password, page 15](#).
5. Click **Apply**.

Creating a list of user accounts

Help | Content Gateway | Version 7.8.x

If a single administrator ID and password for the 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 the 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 the 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 the Content Gateway manager

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

In addition to using an administrator ID and user accounts, you can control which hosts have access to the 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 305](#).
4. Click **Apply**, and then click **Close**.

Using SSL for secure administration

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

Websense supports the Secure Sockets Layer protocol (SSL) to provide protection for remote administrative monitoring and configuration using the 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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

You can obtain an SSL certificate from a recognized certificate authority (for example, VeriSign). Install the certificate in the Content Gateway **config** directory (**/opt/WCG/bin**). You must either rename the certificate to the default filename **private_key.pem**, or specify the name of the certificate in the Content Gateway manager (follow the procedure in [Enabling SSL, page 178](#)).

Enabling SSL

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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 file name only if the certificate file does not use the default name **private_key.pem**.

4. Click **Apply**.

FIPS 140-2 Mode

Help | Content Gateway | Version 7.8.x

FIPS (Federal Information Processing Standard) 140-2 is a U.S. government security standard for hardware and software cryptography modules. Modules validated against the standard assure government and other users that the cryptography in the system meets the standard.

The cryptographic libraries used in version 7.8 of Websense Web Security Gateway (Anywhere), including the Content Gateway component, have passed FIPS 140-2 validation. To see a listing of the validation, go to the 2012 list of [Validated FIPS 140 1 and FIPS 140-2 Cryptographic Modules](#) and search for “Websense”. For more information about the NIST FIPS 140-2 program, see [Cryptographic Module Validation Program \(CMVP\) validation page](#).

By default, Content Gateway does not operate in FIPS 140-2 mode. Content Gateway still uses the FIPS-validated libraries, but it also allows cryptographic algorithms that are not supported by the FIPS 140-2 standard.

You can configure Content Gateway to enforce FIPS 140-2 on HTTPS connections.

When FIPS is enabled:

- ◆ HTTPS connections use TLSv1
- ◆ HTTPS connections use FIPS 140-2 approved algorithms

- ◆ Content Gateway generates SHA-256 certificates in response to origin server certificate requests

**Warning**

Once the FIPS 140-2 option is enabled, you cannot disable it without a complete reinstall of Content Gateway. If Content Gateway is on an appliance, the appliance must be reimaged.

**Important**

Where Web Security Gateway interfaces with some TRITON Enterprise components, there may be a FIPS 140-2 boundary. These include:

- ◆ In Web Security Gateway Anywhere, traffic that flows through the cloud does not use FIPS 140-2.
 - ◆ ThreatScope traffic does not use FIPS 140-2.
 - ◆ Websense Data Security does not use FIPS 140-2.
 - ◆ TRITON Mobile Security does not use FIPS 140-2.
 - ◆ When RSA SecurID is configured for the TRITON console logon, the connection to RSA SecurID is not FIPS 140-2.
-

**Important**

Due to a system limitation, FIPS 140-2 mode cannot be used with IWA fallback to NTLM or Legacy NTLM user authentication.

To enable FIPS 140-2 on HTTPS connections:

1. In the Content Gateway manager go to **Configure > Security > FIPS Security**.
2. Review the warning, select **Enabled**, and click **Apply**.
3. If you are sure that you want to enable FIPS, restart Content Gateway.
4. If you do not want to enable FIPS, select **Disable** and click **Apply**.

**Note**

Even after FIPS 140-2 mode is enabled, by default SHA-1 certificates continue to be used for logon to TRITON management consoles. To learn about how to create and install stronger SHA certificates, see: [How do I create a stronger SHA certificate for Websense management consoles?](#)

Filtering Rules

Help | Content Gateway | Version 7.8.x

Content Gateway supports the ability to create rules that inspect requests for certain parameters and, when matched, apply a specified action. Rules can be created to:

- ◆ Deny or allow URL requests
- ◆ Insert custom headers
- ◆ Allow specified applications, or requests to specified websites to bypass user authentication
- ◆ Keep or strip header information from client requests
- ◆ Prevent specified applications from transiting the proxy



Note

To create rules for IWA, NTLM, and LDAP user authentication, see [Rule-Based Authentication, page 211](#). To get started with Content Gateway user authentication options, see [Content Gateway user authentication, page 189](#).

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.

Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 ThreatScope destinations.

After adding, deleting, or modifying a rule, restart Content Gateway.

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.

add_hdr — causes a custom header-value pair to be inserted. Requires that **Custom Header** and **Header Value** are specified. Provides support for destination hosts that require a specific header-value pair. For an example, see [Creating an add_hdr rule to allow Google enterprise gmail](#), below.

**Note**

The “radius” rule type is **not** supported.

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. If the rule applies to only inbound traffic on a specific port, enter a value for **Proxy Port**.
6. If the rule type is **add_hdr**, specify the **Custom Header** and **Header Value**. The **Custom Header** and **Header Value** must be values that the destination host expects. See the example for Google Business Gmail below.
7. Provide values for any required or desired **Secondary Specifiers**. They include:
 - Time** — Specifies a time range, such as 08:00-14:00.
 - Prefix** — Specifies a prefix in the path part of a URL.
 - Suffix** — Specifies a file suffix in the URL.
 - Source IP address** — Specifies a single client IP address, or an IP address range of clients.
 - Port** — Specifies the port in a requested URL.
 - Method** — Specifies a request URL method:
 - get
 - post
 - put

- trace

Scheme — Specifies the protocol of a requested URL. Options are:

- HTTP
- HTTPS
- FTP (for FTP over HTTP only)

User-Agent — Specifies a request header User-Agent value. This is a regular expression (regex).

You can use the User-Agent field to create application filtering rules that:

- Allow applications that don't properly handle authentication challenges to bypass authentication
- Block particular client-based applications from accessing the Internet

See the Websense knowledge base article titled “When authentication prevents devices, browsers, and custom applications from working with the proxy” for more information and several examples.

8. When you have finished defining the rule, click **Add** to add the rule and then **Apply** to save the rule.
9. When you are done adding rules, click **Apply** to save all the changes and then click **Close** to close the edit window.

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.

Creating an add_hdr rule to allow Google enterprise gmail

Google provides a mechanism in the form of a custom header in the request, that allows Google to recognize and allow or block access to enterprise gmail and other Google Apps for Business.

To make Google's solution work for enterprise gmail:

1. In the Web Security manager **allow** the Web Security category **Internet Communication > General Email**.
2. In the Content Gateway manager enable **HTTPS** (SSL decryption). If your site does not already use SSL support, acquaint yourself with the feature before enabling it.
3. In the Content Gateway manager, on the **Configure > Security > Access Control** page, open **filter.config** and create an **add_hdr** rule.

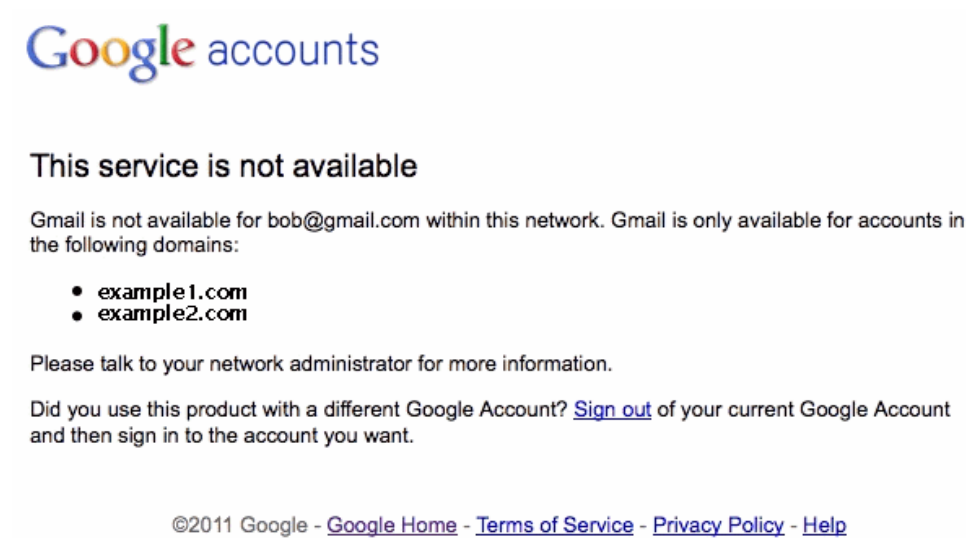


Note

The **add_hdr** rule type can be used with any site that uses a custom header-value pair to accomplish special handling.

- a. Select **add_hdr**.
- b. For **Primary Destination Type** select **dest_domain**.
- c. For **Primary Destination Value** specify “mail.google.com”.
- d. In the **Custom Header** field, specify “X-GoogApps-Allowed-Domains”.
- e. In the **Header Value** field, specify your domain, or a list of domains separated by commas. For example: www.example1.com,www.example2.com
- f. Optionally, in the **Source IP** field specify the source IP address or range of source IP addresses to which this rule will be applied. For example: 10.10.20.30 or 10.10.1.1-10.30.40.50
- g. Click **Add** to add the rule.
- h. Click **Apply** to save all the changes, and then click **Close** to close the edit window.

When a user attempts to access Google services from an unauthorized account, Google displays a block page similar to this:



For Google’s description of the filtering solution, see the article [Block access to consumer accounts and services while allowing access to Google Apps for your organization](#).

Configuring SOCKS firewall integration

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Configuring SOCKS servers, page 186](#)
- ◆ [Setting SOCKS proxy options, page 187](#)
- ◆ [Setting SOCKS server bypass, page 188](#)

SOCKS is commonly used as a network firewall, allowing hosts behind a SOCKS server to gain full access to the Internet while preventing unauthorized access from the Internet to hosts inside the firewall.

When Content Gateway receives a request for content that is not in the cache, 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. If caching is enabled, Content Gateway caches the content before sending it to the client.

- ◆ Content Gateway can act as a SOCKS client, where it receives and serves HTTP or FTP requests as usual.
- ◆ Content Gateway can act as a SOCKS proxy, relaying requests to and from the SOCKS server (usually on port 1080).
- ◆ When Content Gateway is installed on a V-Series appliance it can act as a SOCKS server, providing all of the services of a SOCKS server. (When Content Gateway is **not** installed on an appliance, it cannot act as a SOCKS server.)



Note

Content Gateway does not perform authentication with the client. However, Content Gateway can perform user name and password authentication with a SOCKS server running SOCKS version 5.

Configuring SOCKS servers

Help | Content Gateway | Version 7.8.x

Content Gateway can be configured to work with one or more SOCKS servers in your network. When Content Gateway is installed on a V-Series appliance, a SOCKS server is included with the module.



Note

When Content Gateway is **not** installed on a V-Series appliance, no SOCKS server is provided with Content Gateway.

To configure SOCKS servers:

1. Enable the SOCKS feature.
 - a. Navigate to **Configure > My Proxy > Basic > General**.
 - b. In the **Security** section of the **Features** table, click **SOCKS On**, and click **Apply**.
 - c. Restart Content Gateway.
2. Specify the SOCKS version.
 - a. Go to **Configure > Security > SOCKS > General**.
 - b. Select the SOCKS version running on your SOCKS servers and click **Apply**.
3. To configure the V-Series on-appliance SOCKS server:
 - a. Select the **Server** tab.
 - b. In the **On-Appliance SOCKS Server** area, select **Enabled** and click **Apply**.
An entry for the server is created in the `socks_server.config` file.
 - c. To change the default entry, in the **SOCKS Server** area click **Edit File**. In the editor, select the **On-Appliance-SOCKS-Server** rule.
You can change the port, whether it will be the default SOCKS server, and whether server authentication is applied.
You cannot change the server name or the IP address, which is always the loopback address.
After you make the needed changes, click **Set**.
4. To configure use of other SOCKS servers in your network:
 - a. Select the **Server** tab and in the **SOCKS Server** area click **Edit File**.
 - b. Enter a SOCKS server name.
 - c. Enter the SOCKS server IP address or a domain name that is resolvable by the DNS server inside your network.
 - d. Select whether it will be the default SOCKS server.
 - e. If authentication will be used, provide a SOCKS user name and password.
 - f. Click **Set** to add the server to the list.

You can always return to the editor, select the rule, make changes, and click **Set** to save them.

5. If there are multiple SOCKS servers, after they have been added, or while they are being added, you can arrange them in precedence-order by selecting an entry and moving it up or down the list with the up and down arrows.
6. Click **Apply** to accept your changes, and **Close** to close the editor.
7. In the **SOCKS Server Rules** area you can create rules for specific routing and bypass by destination IP address. See, [Setting SOCKS server bypass](#), page 188.
8. To review configuration options that apply to all SOCKS servers, select the **Options** tab.
 - a. Review and adjust the **Server Connection Timeout** value. It specifies how many seconds Content Gateway waits attempting to connect to a SOCKS server before timing out.
 - b. Review and adjust the **Connection Attempts Per Server** value. It specifies how many times Content Gateway attempts to connect to a given SOCKS server before marking the server as unavailable.
 - c. Review and adjust the **Server Pool Connection Attempts** value. It specifies how many times Content Gateway attempts to connect to a given SOCKS server in the pool before giving up.
9. When SOCKS server configuration is complete, click **Apply** and then go to **Configure > My Proxy > General** and restart Content Gateway.

To remove a server from the list:

1. In the **SOCKS Server** area click **Edit File**.
2. In the list, select the entry you want to delete and click **X**, to the left of the list.
3. Click **Apply** and then **Close**, when you're ready to exit the editor.
4. When configuration is complete, go to **Configure > My Proxy > General** and restart Content Gateway.

Setting SOCKS proxy options

Help | Content Gateway | Version 7.8.x

To configure Content Gateway as a SOCKS proxy, you must enable the SOCKS proxy option and specify the port on which Content Gateway accepts SOCKS traffic from SOCKS clients.

As a SOCKS proxy, Content Gateway can receive SOCKS packets (usually on port 1080) from the client and forward requests directly to the SOCKS server.



Note

You must set SOCKS proxy options in addition to enabling the SOCKS option and specifying SOCKS server information described in [Configuring SOCKS servers](#), page 186.

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 SOCKS server bypass

Help | Content Gateway | Version 7.8.x

You can configure Content Gateway to bypass SOCKS servers and access certain origin servers directly.

1. Navigate to **Configure > Security > SOCKS > Server**. In the **SOCKS Server Rules** area click **Edit File** to open **socks.config**.
2. To modify an existing rule, select it from the list, make your changes, and click **Set**.
3. To create a new rule, specify the parameters and click **Add**.
 - a. Select a **Rule Type**:
 - Route through SOCKS server**
 - Do not route through SOCKS server**
 - b. Specify a destination IP address or range of addresses. Never specify the all networks broadcast address: 255.255.255.255
 - c. Select the SOCKS servers to be used for the traffic.
 - d. Select whether the traffic will be distributed to the specified SOCKS servers in round robin fashion.
 - e. Click **Add** to add the rule.
4. Click **Apply** and then **Close**.
5. Click **Restart** on **Configure > My Proxy > Basic > General**.

Using the Split DNS option

Help | Content Gateway | Version 7.8.x

You can configure Content Gateway to use multiple DNS servers, depending on your security requirements. For example, you can configure Content Gateway to look to one set of DNS servers to resolve host names on your internal network, while allowing DNS servers outside the firewall to resolve hosts on the Internet. This maintains the security of your intranet, while continuing to provide direct access to sites outside your organization.

To configure Split DNS, you must perform the following tasks:

- ◆ Specify the rules for performing DNS server selection based on the destination domain, the destination host, or a URL regular expression.
- ◆ Enable the Split DNS option.

In the Content Gateway manager:

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**.

Content Gateway user authentication

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Browser limitations](#), page 192
- ◆ [Global authentication options](#), page 193
- ◆ [Integrated Windows Authentication](#), page 197
- ◆ [Legacy NTLM authentication](#), page 202
- ◆ [LDAP authentication](#), page 204
- ◆ [RADIUS authentication](#), page 208
- ◆ [Rule-Based Authentication](#), page 211
- ◆ [Mac and iPhone/iPad authentication](#), page 234

Content Gateway supports several methods of authenticating users before their requests are allowed to proceed. These methods can be used together with Websense Web Security user identification (XID) features to provide fallback should user authentication fail or become unavailable.

In both explicit and transparent proxy modes, Content Gateway supports user authentication with:

- ◆ [Integrated Windows Authentication](#) (Kerberos with SPNEGO to NTLM)

- ◆ *Legacy NTLM authentication* (NTLMSSP)
- ◆ *LDAP authentication*
- ◆ *RADIUS authentication*

Content Gateway also supports combinations of Integrated Windows Authentication (IWA), Legacy NTLM, and LDAP using:

- ◆ *Rule-Based Authentication, page 211*

Rule-Based Authentication summary

Rule-Based Authentication is an ordered list of authentication rules. When a request is processed, the list is traversed top to bottom and the first match is applied.

Rules specify:

1. How to match a client.
By:
 - IP address
 - Inbound proxy port (explicit proxy only)
 - User-Agent value
 - A combination of the above
2. The domain or ordered list of domains to authenticate against. With a list, the first successful authentication is remembered and used in subsequent authentications for that user.
3. Whether a customizable web portal page should be used for authentication.

Multiple Realm Networks: Rule-Based Authentication supports multiple realm network structures in which Windows Active Directory domains do not have mutual trust relationships and therefore require that each domain's members be authenticated by a domain controller within their domain. In this environment rules are created that specify:

1. Members of the realm (untrusted domain) by IP address or proxy port
2. The realm (domain) they belong to

Authenticating when domain membership is unknown: Some organizations do not always know what domain a user belongs to. For example, this can happen when organizations are rapidly acquiring new businesses. The unknown domain membership problem can be handled in rule-based authentication by creating a rule (or rules) for IP address lists or ranges that also specifies an ordered list of domains to attempt to authenticate against. The first successful authentication is remembered and used in later authentications.

Authentication based on User-Agent value: One or more User-Agent values can be specified in an authentication rule. Often this is a list of browsers. When the User-Agent value matches a rule, authentication is performed against the specified domain(s). If the User-Agent value doesn't match any rule, and no rule matches based on other values, no authentication is performed (this is always true; if no rule matches, no authentication is performed).

Selecting the authentication method

The authentication method is selected in the **Authentication** section of the **Configure > My Proxy > Basic** page. Configuring authentication for rule-based authentication begins with selecting **Rule-Based Authentication**.

Supported domain controllers and directories

- ◆ Windows NT domain controllers
- ◆ Windows 2003, 2008, 2012 Active Directory
- ◆ Novell eDirectory 8.7 and 8.8 (LDAP only)
- ◆ Oracle DSEE 11g, Sun Java 7 and 6.2 (LDAP only)

Best practices when using Windows Active Directory

If you have only one Active Directory domain, or if all of your Active Directory domains share inbound and outbound trust relationships, the best option is to deploy Integrated Windows Authentication. However, if you want to control authentication based on User-Agent values, you must use Rule-Based Authentication.

If you have multiple domains or realms and user authentication is a requirement, you must use Rule-Based Authentication. For details, see [Rule-Based Authentication, page 211](#).

If user identification is sufficient, you can use one of the Web Security user identification options. See the section titled *User Identification* in Web Security Help.

Backup domain controllers

For Integrated Windows Authentication and Legacy NTLM, Content Gateway supports the specification of backup domain controllers for failover. If the primary domain controller (DC) does not respond to proxy requests, Content Gateway contacts the next DC in the list (the backup domain controller). For the next request, the proxy tries to contact the primary DC again and then contacts the backup DC if the connection fails.

Transparent user authentication

Content Gateway supports both transparent (Single Sign-On) and interactive (prompted) authentication. Transparent authentication is supported with Integrated Windows Authentication and Legacy NTLM. Some browsers provide only limited support. See [Browser limitations, page 192](#).

On Windows networks, Single Sign-On allows users to sign on only once so that they can transparently access all authorized network resources. Therefore, if a user has already logged on to the Windows network successfully, the credentials specified

during Windows logon are used for proxy authentication and the user is not prompted again for a username and password.

[Interactive authentication is supported in networks that are not configured for Single Sign-On and for use with browsers that don't support Single Sign-On. With interactive authentication, users are prompted for credentials before they can access content through Content Gateway.

Browser limitations

Help | Content Gateway | Version 7.8.x

Not all web browsers support transparent user authentication.



Note

Please see the [Version 7.8 Content Gateway Release Notes](#) for the most up-to-date information.

The following table indicates how a browser responds to an authentication request when Integrated Windows Authentication (IWA) is configured.

| Browser/ Operating System | Internet Explorer (v10 and 11 tested) | Firefox | Chrome | Opera (v12.02 tested) | Safari (v6.02 tested) |
|---|---------------------------------------|--|---|--|--|
| Windows | Performs transparent authentication | Performs transparent authentication (v28 tested) | Performs transparent authentication (v33 and 34 tested) | Falls back to NTLM and prompts for credentials | Falls back to NTLM and prompts for credentials |
| Mac OS X | Not applicable | Performs transparent authentication (v28 tested) | Performs transparent authentication (v33 and 34 tested) | Not tested. | Performs transparent authentication |
| Red Hat Enterprise Linux, update 6 | Not applicable | Performs transparent authentication (v28 tested) | Explicit proxy: Performs transparent authentication. Transparent proxy: Falls back to NTLM and prompts for credentials. (v33 and 34 tested) | Not tested. | Not applicable |

Global authentication options

Help | Content Gateway | Version 7.8.x

Use the **Configuration > Security > Access Control > Global Authentication Options** page to configure:

- ◆ User authentication *Fail Open*/fail closed behavior
- ◆ *Credential Caching* options
- ◆ The *Redirect Hostname* (required for transparent proxy deployments)

These settings apply to all proxy user authentication configurations, within the parameters stated for each option below.

Whenever changes are made to any of these settings, click **Apply** to save your changes and then restart the proxy to put the changes into effect.

Fail Open

Fail Open specifies whether requests are allowed to proceed to Web Security for processing when user authentication fails.

When Fail Open is enabled and a Web Security XID agent is configured, if authentication fails and the client is identified by the XID agent, user-based policy is applied. If the user cannot be identified and a policy is assigned to the client's IP address, that policy is applied. Otherwise, the Default policy is applied.



Important

The Fail Open setting does not apply when IWA is the authentication method and the client fails to retrieve a kerberos ticket from the domain controller (DC) because the DC is down.

The Fail Open setting does apply with IWA when IWA falls back to NTLM.

Options include:

- ◆ **Disabled** – specifies that requests do not proceed when authentication failures occur.
- ◆ **Enabled only for critical service failures** (default) – specifies that requests proceed if authentication fails due to:
 - No response from the domain controller
 - The client is sending badly formatted messages

- ◆ **Enabled for all authentication failures, including incorrect password** – specifies that requests proceed for all authentication failures, including password failures.



Important

When user authentication is rule-based with a domain list:

- ◆ If **Enabled only for critical service failures** is selected, when a critical service failure occurs fail open is **not** applied. An error always results in fail closed.
 - ◆ If **Enabled for all authentication failures, including incorrect password** is selected, after trying basic credentials with every domain in the list, fail open is applied.
-

Credential Caching

Credential Caching options include:

- ◆ Caching Method
- ◆ Cache Time-To-Live (TTL), in minutes
- ◆ LDAP Specific Settings

Credential caching settings apply to all clients whether Content Gateway is an explicit or transparent proxy.

Credential caching applies to:

- ◆ All authentication methods when Content Gateway is a transparent proxy
- ◆ When Content Gateway is an explicit proxy:
 - NTLM when Integrated Windows Authentication (IWA) falls back to or negotiates NTLM
 - Legacy NTLM

When IWA authenticates with Kerberos, Kerberos handles ticket (credential) caching.

Caching Method options:

Cache using IP address only – specifies that all credentials are cached with IP address surrogates. This is the recommended method when all clients have unique IP addresses.

Cache using Cookies only – specifies that all credentials are cached with cookie surrogates. This is recommended when all clients share IP addresses, as with multi-host servers such as Citrix servers, or when traffic is NATed by a device that is forwarding traffic to Content Gateway.

Cache using both IP addresses and Cookies – specifies to use cookie surrogates for the IP addresses listed in the cookie caching list, and to use IP address surrogates for

all other IP addresses. This is recommended when the network has a mix of clients, some with unique IP addresses and some using multi-user hosts or that are subject to NATing.

For a description of surrogate credentials, see [Surrogate credentials](#).



Important

Cookie mode caching:

- ◆ Cookie mode caching does not work with applications that do not support cookies, or with browsers in which cookie support has been disabled.
- ◆ When the browser is Internet Explorer, the full proxy hostname in the form “http://host.domain.com” must be added to the Local intranet zone.
- ◆ When the browser is Chrome, it must be configured to allow third-party cookies or configured for an exception to allow cookies from the proxy hostname in the form “host.domain.com”.
- ◆ When the IP address is set for cookie mode and the request method is CONNECT, no caching is performed.
- ◆ Cookie mode caching is not performed for FTP requests.
- ◆ Cookie mode caching is not supported by Captive Portal with 7.8.3. Support for cookie mode caching was added with 7.8.4.



Note

The user interface setting to disable the NTLM cache for explicit proxy has been removed. Although not recommended, the cache can be disabled for explicit proxy traffic in records.config by setting the value of **proxy.config.ntlm.cache.enabled** to **0** (zero).

Cache Time-To-Live

Cache Time-To-Live (TTL) specifies the duration, in minutes, that an entry in the cache is retained. When the TTL expires, the entry is removed and the next time that that user submits a request, the user is authenticated. If the authentication succeeds, an entry is placed in the cache.

The default TTL is 15 minutes. The range of valid values is 5 to 1440 minutes.

LDAP Specific Settings

When enabled, **Purge LDAP cache on authentication failure** causes the proxy to delete the authorization record for the client from the LDAP cache when an LDAP user authentication failure occurs.

Redirect Hostname

Redirect Hostname specifies an alternate hostname for the proxy.



Note

Redirect Hostname is not used by Integrated Windows Authentication.

By default, authenticating clients are redirected to the hostname of the Content Gateway machine. If clients are unable to resolve that hostname through DNS, or if an alternate DNS name for the proxy is defined, that hostname should be specified in the **Redirect Hostname** field.



Note

To ensure that user authentication for transparent proxy occurs transparently (without prompting the user for credentials), the browser must be configured so that the Redirect Hostname is in its **Intranet Zone**. Typically, this is achieved by ensuring that the Redirect Hostname is in the same domain as the computer on which the browser is running. For example, if the client is **workstation.example.com** and the Redirect Hostname is **proxyhostname.example.com**, the browser allows authentication to occur transparently. Consult your browser documentation.



Note

Content Gateway supports transparent authentication in proxy clusters that use WCCP load distribution. However, the **assignment method distribution attribute** must be the source IP address. For more information see [WCCP load distribution](#), page 53.

Surrogate credentials

Surrogate credentials are entries placed in the credential cache after initial successful authentications.

- ◆ An IP address surrogate ties a credential to an IP address and assumes that the IP address is used by only one user at any given time.

- ◆ A cookie surrogate is tied to a cookie placed on the client's system and depends on client application support for cookies. This method is required when a client IP address is shared by more than one user at a time, as with multi-user hosts such as Citrix servers.

After the initial successful authentication, Content Gateway uses the surrogate credential to respond to subsequent authentication requests on behalf of the user, thus reducing latency and the load on domain controllers and directory services. Credential surrogate entries are deleted when the Time-To-Live expires.

Integrated Windows Authentication

Help | Content Gateway | Version 7.8.x

Integrated Windows Authentication (IWA) is a robust method of authenticating users who belong to shared-trust Windows domains (one or many).

Integrated Windows Authentication:

- ◆ Uses Kerberos and SPNEGO
- ◆ Supports NTLM in both explicit and transparent proxy modes
- ◆ Supports NTLMv2 and NTLMv1 with Session Security
- ◆ Supports Windows Active Directory 2003, 2008, and 2012
- ◆ Can be used with Rule-Based Authentication
- ◆ Supports Internet Explorer 7 and later, Firefox 4 and later, Google Chrome 6 and later, Windows Safari 4 and later, Safari 4 and later on iPad iOS4, and Opera 10 and later
- ◆ Supports UTF-8 user names
- ◆ Supports fall back to prompted authentication

Requires that:

- ◆ Clients be joined to the domain
- ◆ Client browsers specify the Fully Qualified Domain Name (FQDN) of Content Gateway as an intranet site or trusted site
- ◆ In explicit proxy deployments, browsers must specify the FQDN of Content Gateway

If you are using IWA with rule-based authentication, see [Rule-Based Authentication, page 211](#), for configuration steps.

Integrated Windows Authentication: Configuration summary

Follow these steps to configure IWA as the user authentication method for your Content Gateway deployment:

- ◆ In the Content Gateway manager, enable **Integrated Windows Authentication** on the **Configure > My Proxy > Basic** page and click **Apply**.
- ◆ Configure [Global authentication options](#).

- ◆ Join Content Gateway to the Windows domain. See *Configuring Integrated Windows Authentication* for a list of required conditions.

Configuring Integrated Windows Authentication

1. Go to **Configure > My Proxy > Basic > General**. In the **Authentication** section, click **Integrated Windows Authentication On**, and click **Apply**.
2. Configure the *Global authentication options*.
3. Join the Windows domain.

To join the domain:

- Content Gateway must be able to resolve the domain name.
 - Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
 - The correct domain Administrator name and password must be specified.
 - There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
 - If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services must be reachable by Content Gateway on the network.
- a. In the **Domain Name** field, enter the fully qualified domain name.
 - b. In the **Administrator Name** field enter the Windows Administrator user name.
 - c. In the **Administrator Password** field enter the Windows Administrator password.
The name and password are used only during the join and are not stored.
 - d. Select how to locate the domain controller:
 - **Auto-detect using DNS**
 - **DC name or IP address**
If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.
 - e. In the **Content Gateway Hostname** field, confirm that the hostname is the correct hostname and that it is no more than 15 characters (no more than 11 characters on V-Series appliances). If it is longer, it must be shortened if IWA is to be used. The length restriction results from the 15 character limit on NetBIOS hostnames.



Warning

Do not change the hostname after the domain is joined. If the hostname is changed, IWA immediately stops working and will not work again until the domain is unjoined and then re-joined with the new hostname.

- f. Click **Join Domain**. If there is an error, ensure that the conditions outlined above are met and then see *Failure to join the domain*.

**Important**

All clients subject to authentication must be joined to the domain.

Browsers and other proxy clients must be configured to specify the FQDN of Content Gateway as an intranet site or trusted site.

- g. Restart Content Gateway and run some test traffic through the proxy to verify that authentication is working as expected. If there is a problem, see *Troubleshooting Integrated Windows Authentication*.

To unjoin the current domain and join a new domain

1. Navigate to the **Configure > Security > Access Control > Integrated Windows Authentication** tab and click **Unjoin**.
2. To join a new domain, in the **Domain Name** field, enter the fully qualified domain name.
3. In the **Administrator Name** field enter the Windows Administrator user name.
4. In the **Administrator Password** field enter the Windows Administrator password. The name and password are used only during the join and are not stored.
5. Select how to locate the domain controller:
 - **Auto-detect using DNS**
 - **DC name or IP address**If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.
6. Click **Join Domain**.

To change the way the domain controller is found

1. Navigate to the **Configure > Security > Access Control > Integrated Windows Authentication** tab.
2. In the **Domain Controller** section, select how to locate the domain controller:
 - **Auto-detect using DNS**
 - **DC name or IP address**If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.
3. Click **Apply**.

Configuring Integrated Windows Authentication with a load balancer

IWA with a load balancer is supported for v7.8.2 and later.



Important

After upgrade to v7.8.3 check and, if necessary, rejoin IWA domains.

Transparent proxy deployments do not require any special configuration.

Explicit proxy deployments that are behind a load balancer require a custom configuration

With Websense Content Gateway, Integrated Windows Authentication (IWA) uses the Kerberos protocol, with NTLM fallback.

In a load-balanced environment:

- ◆ Clients explicitly point to the Content Gateway cluster via the FQDN, which, when a load balancer is used, must resolve to the load balancer's VIP.
- ◆ Kerberos then returns a ticket for the load balancer's VIP, which the client then sends to Content Gateway.
- ◆ Because the ticket is not issued for the proxy's IP address, but rather for the load balancer's VIP, Content Gateway cannot decrypt the ticket and authentication fails.

To restate the problem, it's not possible to configure clients to request Content Gateway's Kerberos ticket because the client's operating system handles the ticket request based on the FQDN of the proxy, which resolves to the VIP of the load balancer.

Normally, Content Gateway would be configured to share the hostname of the load balancer, but this is not possible when the load balancer requires hostname resolution (as with DNS-based load balancing).

Because it's not possible to stop clients from sending a load-balancer's Kerberos ticket to Content Gateway, the proxies must be configured to accept the load-balancer's ticket, making the Content Gateway nodes appear as the load-balancer within the scope of Kerberos.

Please contact Websense Technical Support for detailed, step-by-step configuration instructions.

Troubleshooting Integrated Windows Authentication

Help | Content Gateway | Version 7.8.x

This section covers 2 common problems:

- ◆ [Failure to join the domain](#)
- ◆ [Failure to authenticate clients](#)

Failure to join the domain

These conditions are required for Content Gateway to join a domain:

- ◆ Content Gateway must be able to resolve the domain name.
- ◆ Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
- ◆ The correct domain Administrator name and password must be specified.
- ◆ There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
- ◆ If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services, must be reachable by Content Gateway on the network.
- ◆ If the Active Directory is configured with multiple Sites, ensure that the subnet that Content Gateway is on is added to one of them.

Troubleshooting:

- ◆ Errors encountered in the join action are reported at the top of the screen (the Integrated Windows Authentication tab).
- ◆ The error message usually includes a link to the failure log where you can get more details.
- ◆ Join failures are logged to `/opt/WCG/logs/smbadmin.join.log`
- ◆ In most cases, the failure message in the log is a standard Samba and Kerberos error message that is easily found with an Internet search.

Failure to authenticate clients

These conditions are required to authenticate clients:

- ◆ Content Gateway clients must be a member of the same domain as that joined by Content Gateway.
- ◆ Client system time must be in sync with the domain controller and Content Gateway to plus or minus 1 minute.
- ◆ Explicit proxy clients must **not** be configured to send requests to the IP address of Content Gateway. Clients must use the Fully Qualified Domain Name (FQDN) of Content Gateway. If the IP address is used, NTLM authentication is always performed.
- ◆ The Content Gateway FQDN must be in DNS and resolvable by all proxy clients.
- ◆ Browsers and other client applications must specify the FQDN of Content Gateway as an intranet site or trusted site.
- ◆ When the Active Directory is configured with multiple Sites, the subnet that Content Gateway is on must be added to one of them. If it's not, the following alarm may be generated when Content Gateway is restarted:

```
Windows domain [domain name] unreachable or bad membership
status
```

Troubleshooting:

In the Content Gateway manager, use the **Diagnostic Test** function on the **Monitor > Security > Integrated Windows Authentication** tab. This Monitor page displays authentication request statistics and provides the diagnostic test function.

The **Diagnostic Test** function performs connectivity and authentication testing and reports errors. It also shows domain controller TCP port connectivity and latency.

Errors and messages are logged to:

- ◆ /var/log/messages
- ◆ content_gateway.out
- ◆ /opt/WCG/logs/smbadmin.log
- ◆ /opt/WCG/logs/smbadmin.join.log

Performance issues:

- ◆ **IWA (Kerberos):** Authentication performance is bound by CPU. There is no communication to the domain controllers for Kerberos authentication.
- ◆ **NTLM and Basic:** Domain controller responsiveness effects performance. The **Monitor > Security > Integrated Windows Authentication** page shows average response time.

Legacy NTLM authentication

Help | Content Gateway | Version 7.8.x

Content Gateway supports the NTLM (NT LAN Manager) authentication protocol as a method of ensuring that users in a Windows network are authenticated before they access the Internet.



Important

This implementation of NTLM support (Legacy NTLM) relies solely on the NTLMSSP protocol. Although it performs reliably as documented in this section, it is highly recommended that the [Integrated Windows Authentication](#) mode be used instead. It provides more robust and secure support for NTLM.



Important

If rule-based authentication will be used, configure Legacy NTLM authentication through the [Rule-Based Authentication](#) option.

However, read this section to become familiar with Legacy NTLM features and restrictions.

When the Legacy NTLM option is enabled, the proxy challenges users who request content for proof of their credentials. The proxy then sends the proof of the user's

credentials directly to the Windows domain controller to be validated. If the credentials are valid, the proxy serves the requested content and stores the credentials in the NTLM cache for future use. If the credentials are not valid, the proxy sends an *authentication failed* message.

Restrictions:

1. **WINS resolution** is not supported. Domain controllers must have host names that can be resolved by a DNS server.
2. **Extended security** is not supported and cannot be enabled on the domain controller.
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 192.

If you are using Legacy NTLM with rule-based authentication, see [Rule-Based Authentication](#), page 211, for configuration steps.

Configuring Legacy NTLM authentication

1. Go to **Configure > My Proxy > Basic > General**.
2. In the **Authentication** section, click **Legacy NTLM On**, and click **Apply**.
3. Configure the [Global authentication options](#).
4. Go to **Configure > Security > Access Control > Legacy NTLM**.
5. In the **Domain Controller Hostnames** field, enter the hostname of the primary domain controller, followed, optionally, by a comma separated list of backup domain controllers. The format of the hostname must be:

```
host_name[:port] [%netbios_name]
```

or

IP_address[:port] [%netbios_name]

**Note**

If you are using Active Directory 2008, you must include the netbios_name or use SMB port 445. If you **do not** use port 445, you must ensure that the Windows Network File Sharing service is running on the Active Directory server. See your Windows Server 2008 documentation for details.

**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. Click **Apply** and restart Content Gateway (**Configure > My Proxy > Basic > General**).

Optionally, you can configure Content Gateway to allow certain clients to access specific sites on the Internet without being authenticated by the NTLM server; See [Access Control, page 335](#)).

LDAP authentication

Help | Content Gateway | Version 7.8.x

Content Gateway supports the LDAP option to ensure that users are authenticated with an LDAP server before accessing content through the proxy.



Important

If rule-based authentication will be used, configure LDAP authentication through the [Rule-Based Authentication](#) option. However, read this section to become familiar with LDAP features and restrictions.

When LDAP is enabled:

- Content Gateway acts as an LDAP client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the LDAP server to check that the credentials are correct.
- If the LDAP server accepts the username and password, the proxy serves the client the requested content and stores the username and password in the credential cache.
- Future authentication requests for that user are served from the cache until the cache entry expires (Time-To-Live value).
- If the LDAP server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

LDAP authentication supports both simple and anonymous bind.

LDAP user authentication can support passwords containing special characters. Configuration is made directly in the **records.config** file. The following parameter must be enabled, and the correct encoding name to which the special characters belong must be configured. Add these entries to **records.config**. Note that the default setting is 0 (feature disabled).

```
// To enable the feature specify 1.
CONFIG proxy.config.ldap.proc.encode_convert INT <1 or 0>
// Specify an encoding name here. For example,
// for German specify "ISO-8859-1".
CONFIG proxy.config.ldap.proc.encode_name STRING <encoding
name>
```

Configuring Content Gateway to be an LDAP client

1. Go to **Configure > My Proxy > Basic > General**.
2. In the **Authentication** section, click **LDAP On**, and then click **Apply**.
3. Configure the [Global authentication options](#).
4. Go to **Configure > Security > Access Control > LDAP**.
5. Enter the hostname of the LDAP server.

6. Enter the port on which Content Gateway communicates with the LDAP server. The default is port 389.

**Note**

When the LDAP directory service is Active Directory, requests from users located outside the global catalog's base domain will fail to authenticate. This is because the default port for LDAP is 389 and requests sent to 389 search for objects only within the global catalog's base domain. To authenticate users from outside the base domain, change the LDAP port to 3268. Requests sent to 3268 search for objects in the entire forest.

7. Enable **Secure LDAP** if you want the proxy to use secure communication with the LDAP server. Secure communication is performed on port 636 or 3269. Change the port value in the previous field, if necessary.
8. Select the type of directory service to set the filter for searching.
 - **Microsoft Active Directory** sets the type to **sAMAccountName** (default).
 - **Other** sets the type to **uid** for eDirectory or other directory services.
9. Enter the **Bind Distinguished Name** (fully qualified name) of a user in the LDAP-based directory service. For example:
`CN=John Smith, CN=USERS, DC=MYCOMPANY, DC=COM`
Enter a maximum of 128 characters in this field.
If no value is specified for this field, the proxy attempts to bind anonymously.
10. Enter a password for the user specified in the previous step.
11. Enter the **Base Distinguished Name** (DN). Obtain this value from your LDAP administrator.
12. Click **Apply**.
13. Click **Restart** on **Configure > My Proxy > Basic > General**.

As optional steps, you can:

- ◆ Change LDAP cache options. See [Setting LDAP cache options, page 206](#).
- ◆ 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 335](#)).

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 **/opt/WCG/config**.

- 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. |

- Save and close the file.
- From the Content Gateway **bin** directory (`/opt/WCG/bin`), run **content_line -L** to restart the proxy on the local node or **content_line -M** to restart the proxy on all the nodes in a cluster.

Configuring secure LDAP

By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology. You can enable LDAP over SSL (LDAPS) by installing a properly formatted certificate from either a Microsoft certification authority (CA) or a non-Microsoft CA.

To use LDAPS with Content Gateway:

- Open the **records.config** file located in `/opt/WCG/config`.
- Add following entry to **records.config**:

```
CONFIG proxy.config.ldap.secure.bind.enabled INT 1
```

3. Navigate to **Configure > Security > Access Control > LDAP** and change the port to 3269.

**Note**

The Directory Service must be configured to support LDAPS authentication. See to the documentation provided by the directory provider for instructions.

RADIUS authentication

Help | Content Gateway | Version 7.8.x

Content Gateway supports the RADIUS option to ensure that users are authenticated with a RADIUS server before accessing content through the proxy.

When the RADIUS option is enabled:

- Content Gateway acts as a RADIUS client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the RADIUS server to check that the credentials are correct.
- If the RADIUS server accepts the username and password, the proxy serves the client with the requested content and stores the username and password entry in the RADIUS cache; all future authentication requests for that user are served from the RADIUS cache until the entry expires.
- If the RADIUS server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

Content Gateway supports a primary RADIUS server and a secondary RADIUS server for failover. If the primary server does not respond to the proxy request within the specified timeout (60 seconds by default), Content Gateway tries to check the username and password again. If a response from the primary RADIUS server is not received after the maximum number of retries (10 by default), the proxy contacts the secondary RADIUS server. If Content Gateway cannot contact the secondary RADIUS server, the user is prompted again for a username and password.

The RADIUS cache is held in memory and stored on disk. Content Gateway updates the data on disk every 60 seconds. In addition, Content Gateway stores username and password entries in the RADIUS cache for 60 minutes. If a password and username entry is expired in the RADIUS cache, Content Gateway contacts the RADIUS server to accept or reject the username and password.

To configure Content Gateway to be a RADIUS client:

- ◆ Enable the RADIUS option.
- ◆ Specify the hostname or IP address of the primary and secondary (optional) RADIUS servers, and the port and shared key that Content Gateway uses to communicate with the RADIUS servers.

See [Configuring Content Gateway to be a RADIUS client](#), page 209.

Configuring Content Gateway to be a RADIUS client

1. Go to **Configure > My Proxy > Basic > General**.
2. In the Authentication section, click **Radius On**, and then click **Apply**.
3. Navigate to **Configure > Security > Access Control > Radius**.
4. Enter the hostname of your primary RADIUS server.
5. Enter the port number through which Content Gateway communicates with the primary RADIUS server.
6. Enter the key used for encoding.
7. If you are using a secondary RADIUS server, enter the hostname, port, and shared key in the appropriate fields of the **Secondary Radius Server (Optional)** area.
8. Click **Apply**.
9. Click **Restart** on **Configure > My Proxy > Basic > General**.



Note

In addition to performing these procedures, you must add the Content Gateway machine as a trusted client on the primary and secondary RADIUS servers and provide the shared key you want to use for the Content Gateway machine (the shared key must be the same one you specify in the procedure below). See your RADIUS server documentation.

Setting RADIUS cache and server timeout options

By default, the RADIUS cache and RADIUS server timeout options are configured as follows:

- ◆ The RADIUS cache is configured to store 1,000 entries and each entry is considered fresh for 60 minutes.
- ◆ Content Gateway can try to re-establish a connection to the RADIUS server if the connection remains idle for 10 seconds and can retry the connection a maximum of 10 times.

Change these default values by editing the **records.config** file.

1. Open the **records.config** file located 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 (`/opt/WCG/bin`), run `content_line -L` to restart Content Gateway on the local node or `content_line -M` to restart WCG on all the nodes in a cluster.

Rule-Based Authentication

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Global authentication options, page 193](#)
- ◆ [Rule-based authentication Domain list, page 216](#)
- ◆ [Creating an authentication rule, page 221](#)
- ◆ [Working with existing authentication rules, page 224](#)
- ◆ [Rule-based authentication use cases, page 225](#)
- ◆ [Authentication based on User-Agent, page 227](#)
- ◆ [Authentication using Captive Portal](#)
- ◆ [Troubleshooting authentication rules, page 231](#)

Using an ordered list of authentication rules, rule-based authentication provides support for multiple realm, multiple domain, and other special authentication requirements. When a request is processed, the rule list is traversed top to bottom, and the first match is applied.

Authentication rules specify:

1. How to match a user.
 - By:
 - IP address
 - Inbound proxy port (explicit proxy only)
 - User-Agent value
 - A combination of the above
2. The domain or ordered list of domains to authenticate against.

With a list of domains, the first successful authentication is cached and used in subsequent authentications. If IP address caching is configured, the IP address is cached. If Cookie Mode is configured, the cookie (user) is cached.
3. Whether a customizable web portal page should be used for authentication.

In rule-based authentication, only the first matching rule is tried. If authentication is unsuccessful, no further authentication is attempted.

Rule-based authentication is designed to meet these special requirements:

- ◆ **Multiple realm networks:** Rule-based authentication supports multiple realm networks in which domains do not share trust relationships and therefore require that each domain's members be authenticated by a domain controller within their domain. In this environment rules are created that specify:
 - Members of the realm (untrusted domain) by IP address or proxy port

- The realm (domain) they belong to

**Note**

Users of Multiple Realm Authentication from v7.7.x and prior, should think of rule-based authentication as:

- ◆ A rename and expansion of the Multiple Realm Authentication feature.
- ◆ The addition of support for ordered domain lists.
- ◆ A reorganization and simplification of how domains are specified for use in rules.
- ◆ A simplification of how credential caching is configured (see [Global authentication options](#), page 193).

For a complete overview, please read the rest of this section.

- ◆ **Authentication when domain membership is unknown:** Some organizations do not always know what domain a user belongs to. For example, this can happen when organizations acquire new businesses and directory services are not mapped or consolidated. The unknown domain membership problem can be handled in rule-based authentication by creating a rule for IP address lists or ranges that specifies an ordered list of domains to attempt to authenticate against. The first successful authentication is remembered and used in later authentications. If authentication is not successful or the browser times out, no authentication is performed.
- ◆ **Authentication based on User-Agent value:** One or more User-Agent value can be specified in an authentication rule. Often this is a list of browsers. When the User-Agent value matches a rule, authentication is performed against the specified domain(s). If the User-Agent value doesn't match any rule and no rule matches based on other values, no authentication is performed (this is always true in rule-based authentication; if no rule matches, no authentication is performed).

For use case examples see [Rule-based authentication use cases](#), page 225.

**Note**

If all the users in your network can be authenticated by domain controllers that share trust relationships, you probably don't need rule-based authentication.

However, the option is well suited to single domain environments that may benefit from multiple rules based on IP addresses, inbound proxy port (explicit proxy), and/or User-Agent values.

Rule-based authentication structure and logic

Structure:

- ◆ A list of domains is created and maintained.
When a domain is added to the list, the authentication method is specified: IWA, Legacy NTLM, or LDAP. RADIUS is not supported.
Only domains on the domain list can be specified in authentication rules.
The domain list is created and maintained on the **Configure > Security > Access Control > Domains** tab. The domain list is stored in the **auth_domains.config** file.
- ◆ Authentication rules identify users (clients) by IP address, inbound proxy port (explicit proxy only), and/or User-Agent values, and attempt to authenticate the user against a specified domain or list of domains.
Authentication rules are defined on the **Configure > Security > Access Control > Authentication Rules** tab. Rules are stored in the **auth_rules.config** file.



Note

Credential caching configuration is performed on the **Configure > Security > Access Control > Global Configuration Options** tab. On that page you specify IP address caching, cookie caching, or both. The setting applies to both transparent proxy and explicit proxy traffic. When both IP address caching and cookie caching are specified, the IP addresses that cookie caching is applied to must be specified.

Logic:

- ◆ One or more rules are defined for clients and domains (**Configure > Security > Access Control > Authentication Rules**).
- ◆ When a request for web content is received:
 - A top-down rule list traversal begins
 - The first match is applied
 - If the rule includes a list of domains, authentication proceeds as follows:
 - The proxy attempts to authenticate with the first domain using the method configured for that domain. For example, if the first domain is IWA, Content Gateway transparently negotiates with the browser for credentials (407 or 401).
 - If authentication fails and Content Gateway hasn't already challenged (prompted) for credentials, it then prompts for credentials.

Exception: When Content Gateway is an explicit proxy and the first and second domains are IWA, there is no prompt for basic credentials. Instead, Content Gateway uses the Kerberos ticket provided by the client

to attempt to authenticate with the second domain. If the attempt fails and the fallback to NTLM authentication fails, the user is prompted for credentials.

When Content Gateway is a transparent proxy the standard behavior applies. This is because when the user is not a member of the first domain, the request for a Kerberos ticket fails because the client does not trust the FQDN sent with the request. The fallback to NTLM authentication also fails and the user is prompted for credentials.

- Content Gateway then uses the basic credentials with each domain, starting with the second, proceeding sequentially until authentication succeeds or the list is exhausted.
- Content Gateway then uses the basic credentials to attempt, again, to authenticate with the first domain.
- If authentication fails with all domains and the **Fail Open (Configuration > Security > Access Control > Global Authentication Options)** setting is:

Enabled only for critical service failures, the proxy assumes that the user mis-entered their credentials, prompts again for basic credentials, and attempts, again, to authenticate sequentially against the list.

Enabled for all authentication failures, including incorrect password, fail open is applied.

- If no rule matches, no authentication is attempted
- ◆ Transactions are logged with the user name used by Filtering Service.
- ◆ Proxy authentication statistics are collected and reported individually for each authentication method. See [Security, page 275](#) (in the Statistics section).



Important

Content Gateway must be configured with a DNS server that can resolve the fully qualified domain name (FQDN) of Content Gateway for every realm used by IWA. If this isn't done, IWA fails to work. How to configure the DNS server is up to the network administrator. One option is to configure a DNS transfer zone (Sub Zone) between the primary DNS server of Content Gateway and the DNS server of each authentication realm (isolated domain).

Rule-based authentication configuration summary

1. If Content Gateway is an explicit proxy and you want to bring traffic in on multiple ports, specify the ports on the **Configure > Protocol > HTTP** tab.



Important

You must also configure your clients to use the correct port.

2. Configure [Global authentication options](#), page 193 (**Configure > Security > Access Control > Global Authentication Options**).
3. Create a domain list (**Configure > Security > Access Control > Domains**).
 - To specify a domain in a rule, it must be a member of the **Domain List**.
 - Active Directory domains used with IWA must be joined.

Handling of unknown users:



Important

In rule-based authentication, Content Gateway may authenticate users that Web Security doesn't know about (are outside User Services primary domain). In these cases, Content Gateway can be configured to send an "alias" user name that Web Security knows about. Or, to apply the Default policy, send no name. This specification is made for each domain in the domains list.

For more information, see [Unknown users and the 'alias' option](#), below.

4. Create authentication rules (**Configure > Security > Access Control > Authentication Rules**).
5. Restart Content Gateway to make the new rules take effect.

Rule-based authentication best practices

- ◆ If you don't need rules, don't use rule-based authentication. Deploying a single authentication method should provide the best performance.
- ◆ Use the fewest number of rules needed to satisfy your requirements.
- ◆ Do not use a domain list in a rule if it's not needed.

When a domain list is used

- ◆ If there is an IWA or NTLM domain, make it first in the list.
- ◆ If there is more than one IWA or NTLM domain, place the domain with the most active members first in the list. In other words, make the first domain the one that will most often authenticate users.
- ◆ Note that if an IWA domain is first in the list and the user is not joined to that domain, the user will be prompted for credentials.
- ◆ Note that if the first domain in the list is LDAP, every user who matches the rule will be prompted for credentials. The credentials provided will be offered to each successive domain.
- ◆ If the domain list includes an IWA domain, the Captive Portal (available beginning with 7.8.3) option is disabled.

Unknown users and the ‘alias’ option

In rule-based authentication it’s possible for Content Gateway to authenticate a user who, when passed to Web Security, is not recognized because the name is not in the User Services directory.

When an authenticated user name is not found by Web Security, the **Default** policy is applied. There are several ways to address this:

- ◆ Change the Web Security User Services configuration so that it can discover and add the names to its directory.
- ◆ Add the unrecognized names to Web Security’s primary domain. The names must match exactly. Define policies for the new names.
- ◆ For users who match a particular authentication rule, pass an alias name and add the alias name to Web Security’s primary domain. The names must match exactly. Define a policy for the alias name.
- ◆ If the existing Web Security Default policy is sufficient, do nothing, or for every user who matches a particular rule, in the rule select to use a blank (empty) alias.

For some illustrative use cases, see [Rule-based authentication use cases](#).

Rule-based authentication Domain list

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

To use rule-based authentication, you create and maintain a **Domain List**. There must be at least one domain on the list before an authentication rule can be defined.

When a domain is added to the list, the authentication method is specified.

When a rule is defined, the domain or domains are selected from the domain list.

Supported domain types include:

- ◆ Active Directory (AD) domains to be used with IWA. These domains must be joined by Content Gateway, as well as by its members (users).
- ◆ Domain Controllers (DC) to be used with Legacy NTLM
- ◆ AD and uid domain controllers and directory servers to be used with LDAP

Domain specification configuration summary:

1. Rule-based authentication must be enabled (**Configure > My Proxy > General**).
2. On **Configure > Security > Access Control > Domains**, click **New Domain**.
3. Select the authentication method.
4. Specify a unique name that will help you recognize the domain and its purpose.
5. Optionally, configure the **Aliasing** option.
6. Specify the domain settings. These vary by authentication method.

See:

- ◆ [Adding an Active Directory domain for use with IWA](#)

- ◆ [Adding an NTLM domain controller for use with Legacy NTLM](#)
- ◆ [Adding a domain \(directory service\) for use with LDAP](#)

Adding an Active Directory domain for use with IWA

Active Directory (AD) domains to be used with IWA must be joined by both Content Gateway and directory members (clients).

To join a domain:

- Content Gateway must be able to resolve the domain name.
- Content Gateway system time must be synchronized with the domain controller's time, plus or minus 1 minute.
- The correct domain Administrator name and password must be specified.
- There must be TCP/UDP connectivity to the domain controller(s) (ports 88, 389, 445).
- If backup domain controllers are configured, they and their Kerberos Distribution Center (KDC) services, must be reachable by Content Gateway on the network.

To specify and join a domain:

1. Go to **Configure > Security > Access Control > Domains** and click **New Domain**.
2. Select **Integrated Windows Authentication** from the **Authentication Method** drop down box.
3. In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose.
4. Optionally, configure the **Aliasing** option. For information, see [Unknown users and the 'alias' option, page 216](#).
5. In the **Domain Name** field, enter the fully qualified domain name. For example, ad1.example.com.
6. In the **Administrator Name** field enter the Windows Administrator user name.
7. In the **Administrator Password** field enter the Windows Administrator password.
The name and password are used only during the join and are not stored.
8. Select how to locate the **domain controller**:
 - **Auto-detect using DNS**
 - **DC name or IP address**
If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.

9. Confirm the **Content Gateway Hostname**.



Warning

Do not change the hostname after the domain is joined. If it is changed, IWA immediately stops working and will not work again until the domain is unjoined and then re-joined with the new hostname.

10. Click **Join Domain**.

The **Joined Domain Connections** section of the **Monitor > Security > Integrated Windows Authentication** page displays a list of joined domains and connections, and provides a diagnostic test function.

For troubleshooting tips, see [Failure to join the domain](#).

To change the way the domain controller is found, and other attributes

1. On the **Domains** page, in the list select the domain you want to change and click **Edit**.
2. In the **IWA Domain Details** section, select how to locate the domain controller:
 - **Auto-detect using DNS**
 - **DC name or IP address**
If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list, no spaces.
3. You can also change the **Aliasing** setting. See [Unknown users and the 'alias' option, page 216](#).
4. Click **Apply**.

Adding an NTLM domain controller for use with Legacy NTLM

Support for Legacy NTLM has these restrictions:

- ◆ **WINS resolution** is not supported. Domain controllers must have hostnames that can be resolved by a DNS server.
- ◆ **Extended security** is not supported and cannot be enabled on the domain controller.
- ◆ **NTLM2 session security** is not supported and cannot be enabled on clients. In the Security Settings area of the Windows operating system, inspect the **Network Security: Minimum session security** settings.
- ◆ **NTLMv2** is not supported with Active Directory 2008.
- ◆ Not all browsers support transparent NTLM authentication. See [Browser limitations, page 192](#).

For a complete description of support for Legacy NTLM, see [Legacy NTLM authentication, page 202](#).

To add an NTLM domain for use in rule-based authentication:

1. Go to **Configure > Security > Access Control > Domains** and click **New Domain**.
2. Select **Legacy NTLM** from the **Authentication Method** drop down box.
3. In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose. After the domain is added, the name cannot be changed.
4. Optionally, configure the **Aliasing** option. For information see: [Unknown users and the 'alias' option, page 216](#).
5. In the **Legacy NTLM Domain Details** section:
 - a. In the **Domain Controller** entry field enter the IP address and port number of the primary domain controller. If no port is specified, Content Gateway uses port 139.
You can also specify secondary domain controllers in a comma-separated list. The supported formats are:

```
host_name[:port][%netbios_name]
```

```
IP_address[:port][%netbios_name]
```

 The **netbios_name** is required with Active Directory 2008.
 - b. Specify whether load balancing should be applied among multiple DCs.

**Note**

Even if load balancing is **not** selected, if multiple domain controllers are specified and the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term fail over provision, until such time that the primary domain controller can accept new connections.

6. Click **Add Domain**.

Adding a domain (directory service) for use with LDAP

When LDAP is used:

- Content Gateway acts as an LDAP client and directly challenges users who request content for a username and password.
- After receiving the username and password, Content Gateway contacts the LDAP server to check that the credentials are correct.
- If the LDAP server accepts the username and password, the proxy serves the client the requested content and stores the username and password in the credential cache.
- Future authentication requests for that user are served from the cache until the cache entry expires (Time-To-Live value).

- If the LDAP server rejects the username and password, the user's browser displays a message indicating that authorization failed and prompts again for a username and password.

LDAP authentication supports both simple and anonymous bind.

To add an LDAP domain to the Domains list:

1. Go to **Configure > Security > Access Control > Domains** and click **New Domain**.
2. Select **LDAP** from the **Authentication Method** drop down list.
3. In the **Domain Identifier** field, enter a unique name that will help you recognize the domain and its purpose. After the domain is added, the name cannot be changed.
4. Optionally, configure the **Aliasing** option. For information see: [Unknown users and the 'alias' option, page 216](#).
5. In the **LDAP Domain Details** section:
 - a. In the **LDAP Server Name** field, enter the fully qualified domain name or IP address of the LDAP server.
 - b. If the LDAP server port is other than the default (389), in the **LDAP Server Port** field, enter the LDAP server port.
 - c. Enter the **LDAP Base Distinguished Name**. Obtain this value from your LDAP administrator.
 - d. Select the **LDAP Server Type** from the drop down list.
 - Select **sAMAccountName** for Active Directory
 - Select **uid** for other directory services
 - e. In the **Bind Domain Name** field, enter the bind distinguished name. This must be a Full Distinguished Name of a user in the LDAP directory service. For example:
`CN=John Smith,CN=USERS,DC=MYCOMPANY,DC=COM`
 - f. In the **Bind Password** field, enter the password for the name given in the **Bind Domain Name** field.
 - g. Enable **Secure LDAP** if you want Content Gateway to use secure communication with the LDAP server. If enabled, set the LDAP port to 636 or 3269.
6. Click **Add Domain**.

To unjoin or remove a domain from the Domain List

On the **Domains** page, select the domain from the list and click **Unjoin** or **Delete**.

A confirmation dialogue displays. Confirm that you want to remove the domain from the list.



Warning

When a domain is removed, it is also removed from any authentication rules that specify it.

If it is the only domain specified in a rule, when the domain is removed the rule is made invalid and, therefore, the rule is removed.

Creating an authentication rule

Help | Content Gateway | Version 7.8.x

Before you create an authentication rule you must:

- ◆ Enable **Rule-Based Authentication** on **Configure > My Proxy > Basic > General**.
- ◆ Configure [Global authentication options, page 193](#)
- ◆ Create a [Rule-based authentication Domain list, page 216](#)

You must also know:

- ◆ The name of the domain(s) to be specified in the rule. This is the unique name that was specified when the domain was added to the Domains list.
- ◆ How to match users.

By:

- IP address – individual addresses or address ranges can be specified
- Inbound proxy port (explicit proxy only)
- User-Agent values
- A combination of the above

To create a rule:



Note

In the Rule editor, after entering all specifiers, click **Add** before clicking **Apply**. If **Apply** is clicked first, or the edit window is closed, all entry fields are cleared.

The size of a rule cannot exceed 512 characters.

1. Go to **Configure > Security > Access Control** and review and adjust the **Global Authentication Options** and **Domains** list.
2. If AD domains are used with IWA, go to **Monitor > Security > Integrated Windows Authentication** and confirm that the IWA domains are joined and that connections are established.

3. Go to **Configure > Security > Access Control > Authentication Rules**. A list of existing authentication rules is displayed at the top of the page.
4. Click **Edit File** to open the rule editor.
5. If some rules have already been defined, note the order of the rules in the list at the top of the page.



Important

Rule order matters. The rule match traversal is performed top-to-bottom. Only the first match is applied.

6. Select **Status Enabled** if you want the rule to be active after the rule is added and Content Gateway is restarted.
7. Enter a unique **Rule Name** (required). A short, descriptive name will help you recognize the rule and its purpose. It is recommended that the name not exceed 50 characters.
8. If the rule applies to specific IP addresses, in the **Source IP Addresses** field, enter a comma-separated list of individual IP addresses and/or IP address ranges. Do not use spaces. For example:

10.4.1.1,10.12.1.1-10.12.254.254

Source IP address ranges can overlap. Overlapping ranges may be useful as a quick way of identifying sub-groups in a large pool. In overlapping ranges, the first match is used.

If this field is empty (undefined), all IP addresses match.

9. If the rule applies to inbound traffic on a specific port, select the **Proxy Port** from the drop down list. This option is valid with explicit proxy only.

Inbound ports are specified on the **Configure > My Proxy > Protocols > HTTP > General** page in the **Secondary HTTP Proxy Server Ports** field. Client applications must be configured to send requests to the desired port.

If undefined, all ports match. Transparent proxy deployments should leave the field undefined.

10. To apply the rule to specific **User-Agent** values, enter POSIX-compliant regular expressions (regex) to match the desired values. To specify a common browser type, select a **Predefined** regex from the drop down list and click **Include**.

If undefined, all User-Agents match.

You can edit the field directly.

Use the “|” character (logical ‘or’) to separate regexes.

The “^” regex operator is not supported.

The regex is validated when the rule is committed to the configuration file, which happens after clicking **Add** or **Set** and then **Apply**. **If the regex is not valid, the rule is deleted and must be recreated with a valid regex.**

For an extended description and examples, see [Authentication based on User-Agent, page 227](#).

11. Specify the domain(s) to authenticate against.

- a. From the **Domains** drop down list, select the applicable domain and click **Include**. Only domains that have been added to the **Domains** list are available (**Configure > Security > Access Control > Domains**).
- b. If an ordered list of domains will be used, select each domain one at a time and click **Include**. Then select domains in the list and use the up and down arrows to achieve the desired order.

**Important**

The *Fail Open*/fail closed setting is applied after every domain in the list is tried.

12. Beginning with 7.8.3, click **Enable** next to Captive Portal to redirect users to a customizable web portal page for authentication.

Beginning with 7.8.4, click:

- **Enabled for HTTPS Authentication page** to redirect users to a customizable web portal page for authentication.

When this selection is enabled, the page will display using HTTPS.

When HTTPS is used, a server certification is generated based on the internal root CA. To use this feature, you must import the internal root CA to ensure there is no certificate error. See *Importing your Root CA* for details.

- **Enabled for HTTP Authentication page** to redirect users to a customizable web portal page for authentication.

With this selection, the page is displayed using the HTTP protocol.

This option is disabled if an IWA domain is included in the domains list.

If this option is enabled and an IWA domain is added to the domains list, an error message will display.

Note that when Content Gateway receives an unauthenticated POST request from a user who matches a Captive Portal rule, it redirects the user to the web portal authentication page and does not record the POST data. After successful authentication, the original POST data must be input again.

See *Authentication using Captive Portal* for additional details.

13. Click **Add** to add the rule.
14. At the top of the page, check and adjust the position of the rule in the rule list. The first rule matched is applied.
15. Click **Apply** and then restart Content Gateway to put the rule into effect.

**Warning**

If a rule has invalid values, a warning message displays that identifies the invalid rule. The rule is not written to the file.

Working with existing authentication rules

Help | Content Gateway | Version 7.8.x

Use the rule editor in the Content Gateway manager. Do not directly edit `auth_rules.config`.

Editing a rule

1. Go to **Configure > Security > Access Control > Authentication Rules** and click **Edit File**.
2. In the table of rules, click on the rule to be changed. Its values populate the fields in the definition area.
3. Make the desired changes, click **Set** and then click **Apply**.



Important

If a field value is not valid, the rule is not committed and the rule entry is discarded. To avoid difficulty in recreating a rule, separately record the field values so that it is easy to correct the bad field value and recreate the rule.

4. Click **Close** to return to the **Authentication Rules** tab and click **Refresh** to see the updated list.
5. **Restart** Content Gateway to put the changes into effect.

Reordering the list of rules

Authentication rules are matched top-down in the list. Only the first match is applied.

1. Go to **Configure > Security > Access Control > Authentication Rules** and click **Edit File**.
2. In the table of rules, click on the rule that you want to reposition and then click the down or up arrow on the left to reposition the rule.
3. When the rules are in the desired order, click **Apply**.
4. Click **Close** to return to the **Authentication Rules** tab and click **Refresh** to see the updated list.
5. **Restart** Content Gateway to put the changes into effect.

Deleting a rule

1. Go to **Configure > Security > Access Control > Authentication Rules** and click **Edit File**.
2. In the table of rules, click on the rule to be deleted and click the “X” button on the left.
3. When you are done deleting rules, click **Apply**.
4. Click **Close** to return to the **Authentication Rules** tab and click **Refresh** to see the updated list.
5. **Restart** Content Gateway to put the changes into effect.

Rule-based authentication use cases

Help | Content Gateway | Version 7.8.x

[Multiple realm use case 1: Domain acquired; explicit proxy, page 225](#)

[Multiple realm use case 2: Internal domain added; explicit proxy, page 226](#)

[Multiple realm use case 3: Temporary domain added; transparent proxy, page 226](#)

[Authentication based on User-Agent, page 227](#)

Multiple realm use case 1: Domain acquired; explicit proxy

This describes a common case in which a second domain is added to an existing, single-domain environment. Content Gateway is an explicit proxy; clients use a PAC file.

An organization—let’s call them Quality Corp—uses a software installation of Content Gateway. They have one domain (QCORP), and one domain controller. They use NTLM to authenticate users.

Quality Corp acquires New Corp who has their own domain (NCORP) and domain controller. They use LDAP to authenticate users.

Quality Corp would like to manage the combined employees in a single domain, but they aren’t ready to make the infrastructure changes. Until they are, they would like to have a separate use policy for New Corp users (i.e., not use the “default” user on the QCORP domain).

Rule-based authentication makes this possible.

To configure the solution, Quality Corp would:

1. Enable Rule-Based Authentication.
2. Add a second, non-default HTTP port (**Configure > Protocols > HTTP > General**). This port will be used by all members of NCORP.
3. Create a PAC file for members of NCORP that causes them to connect to Content Gateway on the new, second port.
4. Create authentication rules, one each for the QCORP and NCORP domains:
 - a. On **Configure > Security > Access Control > Domains**, add the QCORP and NCORP domains to the Domains list.
 - When adding NCORP, use the **Aliasing** option to specify “NCorpUser” for use in policy determination.
 - b. On **Configure > Security > Access Control > Authentication Rules**, create an NCORP rule for connections on the second port. You must know the IP addresses/ranges of New Corp users, and specify the NCORP domain.
 - c. Define the QCORP rule to handle all other connections.
5. In Web Security manager, add “NCorpUser” to the QCORP domain as a valid user and create policy for that user.

At this point, everyone connecting to Content Gateway from NCORP is authenticated against the NCORP domain controller and gets the group policy associated with NCorpUser. Note that no individual user-based policy or features, such as quota time, are possible in this scenario. Transactions are logged as NCorpUser. This is all performed with no effect on the authentication, policy, or logging of users on the QCORP domain.

Multiple realm use case 2: Internal domain added; explicit proxy

This describes a common case in which a second domain is added to an existing, single-domain environment. Content Gateway is an explicit proxy; clients use a PAC file.

An organization—let's call it BigStars—uses a software installation of Content Gateway. They have one domain (BIG), and one domain controller. They use NTLM to authenticate users.

A group in the company converts to Apple computers, which can't be authenticated with NTLM. The IT group installs an LDAP server and creates a new domain—BIGAPL—for the Apple users.

Because this group of users previously existed and was managed on the primary domain (BIG), the IT department expects that both user-based policy and logging still apply.

The Rule-Based Authentication feature makes this possible.

To configure the solution, BigStars would:

1. Verify that every user in BIGAPL is also in BIG with the exact same user name.
2. Enable Rule-Based Authentication.
3. Add a second, non-default HTTP port (**Configure > Protocols > HTTP**). This port will be used by all members of BIGAPL.
4. Create a PAC file for members of BIGAPL that causes them to connect to Content Gateway on the new, second port.
5. Create authentication rules, one each for the BIGAPL and BIG domains.
 - a. On **Configure > Security > Access Control > Domains**, add the BIGAPL and BIG domains to the Domains list.
 - b. On **Configure > Security > Access Control > Authentication Rules**, create a BIGAPL rule for connections on the second port.
 - c. Define the BIG rule to handle all other connections.

At this point, all members of BIGAPL are authenticated with LDAP, but maintain their individual policy as specified by their existing NTLM identities. Logs and reports also refer to that same user.

Multiple realm use case 3: Temporary domain added; transparent proxy

This describes a common case in which a second, special-purpose domain is added to an existing, single-domain environment. Content Gateway is a transparent proxy using WCCP v2.

An organization—let's call it Creative Corp—uses a software installation of Content Gateway. They have one domain (CCORP), and one domain controller. They use NTLM to authenticate users.

Creative Corp is about to launch a new product and wants to make a big splash. They decide to have an open house complete with kiosks, demonstrations, and presenters. The kiosks only need the default Internet policy to properly demonstrate the new product. The IT manager wants to keep the kiosk network as walled off from the corporate intranet as possible. In this scenario, logging individual users isn't a requirement.

The Rule-Based Authentication feature makes this possible.

To configure the solution, Creative Corp would:

1. Build a new, temporary network complete with its own domain controller. Let's call this domain CTEMP.
2. Add one or more users to CTEMP. They can either match one-to-one with existing users on the primary domain, or be one or more generic users for use by the presenters.
3. Redirect Internet traffic on CTEMP to Content Gateway with WCCP v2.
4. Enable Rule-Based Authentication.
5. Create authentication rules, one each for the CTEMP and CCORP domains:
 - a. On **Configure > Security > Access Control > Domains**, add the CTEMP domain, enable Aliasing and leave the name field blank. This will have the result of applying the Default policy to all users of CTEMP.
 - b. Add the CCORP domain to the Domains list.
 - c. On **Configure > Security > Access Control > Authentication Rules**, create a CTEMP rule to apply to all connections coming from the IP address range assigned to the CTEMP domain.
 - d. Define the CCORP rule to handle all other connections.

At this point, anyone using the Internet on one of the kiosks is authenticated against the CTEMP network and has the Default policy applied to their requests.

Authentication based on User-Agent

In an authentication rule, a Request header User-Agent value can be used to determine if user authentication will be performed. This is useful when you want to authenticate users using a known set of client applications, usually browsers, and allow other applications, often a set of applications that don't support authentication, to proceed without authentication. Such rules can also specify IP addresses and, if Content Gateway is an explicit proxy, inbound proxy port.

As with all authentication rules, the first matching rule is applied. (For a complete description of rule-based authentication, see [Rule-Based Authentication](#), page 211.)

When the User-Agent field is used, the critical element is the regular expression (regex) that performs the match.

- ◆ The regex must be POSIX-compliant.
 - The “^” regex operator is not supported.
- ◆ Predefined regexes are provided for the most common browsers.
- ◆ When the field is empty, all User-Agent values match.
- ◆ You can create a custom regex by directly editing the field.
- ◆ Multiple regexes are allowed. They must be separated by a “|” (‘or’ operator).

When you click **Apply** (after Add or Set), the regex is parsed and validated. **If the regex is not valid, the rule is deleted and must be recreated with a valid regex.**

Following are a few examples of custom regexes.

Microsoft Internet Explorer 7, 8, or 9:

```
MSIE ([7-9]{1}[\.0-9]{0})
```

Example User-Agent string:

```
Mozilla/5.0 (Windows; U; MSIE 9.0; Windows NT 9.0; en-US)
```

Microsoft Internet Explorer Mobile, all versions:

```
IEMobile
```

Example User-Agent string:

```
Mozilla/5.0 (compatible; MSIE 9.0; Windows Phone OS 7.5; Trident/5.0; IEMobile/9.0)
```

Apple iPhone, all versions:

```
(iPhone) OS (\d+)_(\d+) (?:_(\d+))?
```

Example User-Agent string:

```
Mozilla/5.0 (iPod; U; CPU iPhone OS 4_3_3 like Mac OS X; ja-jp) AppleWebKit/533.17.9 (KHTML, like Gecko) Version/5.0.2 Mobile/8J2 Safari/6533.18.5
```

Apple iPad, all versions:

```
(iPad).+ OS (\d+)_(\d+) (?:_(\d+))?
```

Example User-Agent string:

```
Mozilla/5.0 (iPad; CPU OS 6_0 like Mac OS X) AppleWebKit/536.26 (KHTML, like Gecko) Version/6.0 Mobile/10A5355d Safari/8536.25
```

Search the Internet for lists of User-Agent strings, example regular expressions, regex checkers, and related resources.

Use case:

This describes a case in which an organization with a single domain wants to authenticate requests from 2 common web browsers. They also want to bypass authentication for web applications that do not support authentication.

An organization—let’s call it Best Corp—uses Content Gateway. They have one domain (BCORP), and one domain controller. They use IWA to authenticate users.

Best Corp wants to ensure that:

- ◆ Requests from common web browsers are authenticated. They control which web browsers are allowed on their computers.
- ◆ Web applications that don't support authentication bypass authentication.

The User-Agent feature of rule-based authentication makes this possible.

To configure the solution, Best Corp:

1. Enables Rule-Based Authentication.
2. Adds the BCORP domain to the Domains list.
3. Creates an IWA rule that:
 - a. Optionally, specifies the supported client IP address ranges.
 - b. Specifies, by User-Agent value, the web browsers to authenticate.
In the **User-Agent** field, they use the **Predefined** drop down list to select and **Add** Internet Explorer and Firefox. The regex looks like:

```
MSIE*|Firefox*
```

That's it. With this configuration, all requests from Internet Explorer and Firefox, the only 2 browsers that can be installed on their computers, are subject to user authentication. All other requests, most particularly web applications, bypass authentication. To further customize the approach, Best Corp could create other authentication rules and/or add proxy filtering rules (filter.config) to deny or bypass specific applications by User-Agent value.

Authentication using Captive Portal

Beginning with 7.8.3, Content Gateway provides a Captive Portal option when adding an authentication rule. Captive Portal may be especially helpful in handling mobile and other personal devices brought in to your Web Security Gateway networks.

This feature:

- ◆ Redirects users to a web portal page for authentication.
- ◆ Supports captive, interactive (prompted) user authentication of IP addresses (users) that match the Captive Portal rule.
- ◆ Can be used with LDAP and Legacy NTLM; IWA and RADIUS are not supported.
- ◆ Handles credential caching and expiration per the global configuration; cookie authentication and caching are not supported for 7.8.3 but are supported beginning with 7.8.4.

Note that most applications on mobile devices do not share cookies. For those applications, IP-based identification will be required. See the Credential Caching section of [Global authentication options](#) for more information.

Also, for web applications that use Ajax, where Ajax is configured to prevent cookies, cookie-mode cannot support sites that include cross-origin requests (CORS) that rely on Ajax.

- ◆ Allows the authentication form (web portal page) to be customized to suit your needs.
- ◆ Supports only basic authentication.
- ◆ Beginning with 7.8.4, provides the option to display the authentication page using either HTTP or HTTPS.

When adding an authentication rule (see [Creating an authentication rule](#)), a new option is provided. Navigate to **Configure > Security > Access Control > Authentication Rules** and click **Enable** (for 7.8.3) or **Enabled for HTTPS/HTTP Authentication page** (beginning with 7.8.4) next to Captive Portal to select the feature. Users who match the rule are redirected to the new web portal authentication page.

- ◆ This option is disabled if an IWA domain is included in the Auth Sequence list.
- ◆ When this option is enabled, an error message will display if an IWA domain is selected for inclusion in the Auth Sequence list.

Note that when Content Gateway receives an unauthenticated POST request from a user who matches a Captive Portal rule, it redirects the user to the web portal authentication page and does not record the POST data. After successful authentication, the original POST data must be input again.

**Note**

If the requested URL is configured for tunneling or bypass, no user authentication is performed.

When a rule is added with the Captive Portal option enabled, users are reminded that they can customize the pre-defined web portal page. Go to the new Captive Portal Page Customization tab of **Configure > Security > Access Control**. Edit the text and HTML to suit your needs. For example, you may want to include your company logo in place of the Websense logo.

Customizing the web portal page

The web portal page is an HTML form that is presented to the user for interactive authentication.

Default contents are provided on the Captive Portal Page Customization tab of **Configure > Security > Access Control**. It is recommended that you customize the form to convey to users who see it that this logon portal is part of your network and organization. For example, you might:

- ◆ Replace the Websense logo with your organization's logo. To do that:
 - Edit the src tag and replace the png file name with your company logo file.
 - Copy your png file to /opt/WCG/config/ui_files/images.
- ◆ Include text to explain why the user is seeing this page

The form must be a valid HTML document, defined with valid HTML syntax.

The following variables are used in the document to ensure that it is delivered to the users properly. It is recommended that you do not change their placement or usage.

- %P is replaced with the protocol of the current transaction
- %h is replaced with “redirect_host:8080”
- %u is replaced with the URL request for the portal page
- \$\$DOMAIN is replaced with the basic authentication domain defined in the configuration variable proxy.config.proxy.authenticate.basic.realm. (See [Authentication basic realm](#) for more information.)

When you have entered all of the syntax, click **Preview** to preview the page you have created. When you are happy with the way the portal page looks, click **Apply** to save the content to a file. If you want to return to the default, pre-defined portal page syntax, click **Restore to Default Page**.

The customized Captive Portal page is saved to **auth_form.html**, which is stored in **/opt/WCG/config**. In addition, css and image files can be used to define the portal page. CSS files must be stored in **/opt/WCG/config/ui_files** and image files must be stored in **/opt/WCG/config/ui_files/images**, by default.



Note

The css and image files also reside in **/opt/WCG/ui/configure/auth_form** and **/opt/WCG/ui/configure/auth_form/images**, respectively, for use by the **Preview** feature. Copy any new files to those directories to use **Preview**.

Add a variable to records.config to use a different name for the saved Captive Portal page or store the css and image files in a different directory.

| Configuration Variable | Data Type | Default Value | Description |
|---------------------------------|-----------|------------------|--|
| proxy.config.auth.form_filename | STRING | auth_form.html | Specifies the file that defines the Captive Portal authentication page. Changing this filename is not recommended. |
| proxy.config.internal.file.path | STRING | /config/ui_files | Specifies the location of any css and image files used to define the Captive Portal authentication page. The full default path is /opt/WCG/config/ui_files. Image files are located in an /images sub-directory. |

Troubleshooting authentication rules

In rule-based authentication, problems often present as:

- ◆ Users are *not* challenged for credentials when a challenge is expected
- ◆ Users *are* challenged for credentials when no challenge is expected
- ◆ User authentication is performed against the wrong domain

These problems occur in one of the following phases of user authentication processing:

- ◆ General user authentication logic (outlined below)
- ◆ Rule definition and matching
- ◆ User authentication protocol processing (IWA, NTLM, LDAP; for IWA troubleshooting, see [Troubleshooting Integrated Windows Authentication](#).)

Rule-based authentication logic

Rule-based authentication applies the following logic:

1. The rules in **filter.config** are checked and applied. This action occurs first in every type of Content Gateway user authentication. If a filtering rule is matched, the rule is applied and user authentication processing stops. See [Filtering Rules, page 181](#).
2. If no filtering rule matches, user authentication rule matching is performed.
 - a. The requestor's IP address is checked, top-down, against the rule set.
 - b. If the IP address matches a rule, the source port is checked.
 - c. If the IP address matches a rule, the User-Agent value is checked.
 - d. The first rule matched is applied. **If no rule matches, no authentication is attempted.**
3. If a rule is matched, the specified authentication protocol is applied against the specified domain. All rule configuration details are applied.
4. If the user is authenticated, the request proceeds or is denied per Web Security policy.
5. The transaction is logged.

To see how the logic is applied in a running environment, you can temporarily enable user authentication debug output. Among other details, the debug output shows the parsing of rules and matching. See [Enabling and disabling user authentication debug output](#).

Troubleshooting

When rule-based authentication doesn't produce the expected results, it is recommended that you troubleshoot the problem in the following order:

1. **Check Network Address Translation (NAT)**

Confirm that there is no unexpected IP address NAT. Network address translation has the result that the original source IP address is changed to another address before user authentication is performed. In the Content Gateway manager, go to **Configure > Networking > ARM > General** and examine the rules in

ipnat.config. Addresses can also be NATed by other devices in the network, such as downstream proxies or firewalls.

2. Check the rules in filter.config

Confirm that there is no unexpected matching of a **filter.config** rule. Among other purposes, filter.config rules can be used to bypass user authentication. See [Filtering Rules](#).

3. Check rule matching

Using the IP address of a user who is or is not being challenged as expected, walk through each rule, top to bottom, examining the settings to find the first match. Be meticulous in your analysis. A common problem is that the IP address falls within a too-broad IP address range.

If the rule uses an alias, confirm that the alias is present in the User Service of the primary domain controller.

For explicit clients configured to send traffic to a specific port, check both the rule and the configuration of the client's browser.

4. Check the domain

If you are getting the match you expect, verify that the domain is reachable and that the user is a member of the domain. If yes, troubleshoot the problem at the authentication protocol level. For IWA, see [Troubleshooting Integrated Windows Authentication](#).

5. When Content Gateway is in a proxy chain

If Content Gateway is a member of a proxy chain, verify that X-Forwarded-For headers are sent by the downstream proxy and read by Content Gateway.

- Use a packet sniffer to inspect inbound packets from the downstream proxy. Look for properly formed X-Forwarded-For headers.
- In the Content Gateway manager, go to **Configure > My Proxy > Basic**, scroll to the bottom of the page and verify that **Read authentication from child proxy** is enabled. If it's not, select **On**, click **Apply**, and then restart Content Gateway.

Enabling and disabling user authentication debug output



Warning

Debug output should not be left enabled. Debug output slows proxy performance and can fill the file system with log output.

Debug log information is written to: **/opt/WCG/logs/content_gateway.out**

To enable user authentication debug information, edit: **/opt/WCG/config/records.config**

```
(root)# vi /opt/WCG/config/records.config
```

Find and modify the following parameters and assign values as shown:

```
CONFIG proxy.config.diags.debug.enabled INT 1
CONFIG proxy.config.diags.debug.tags STRING
    http_xauth.* | auth_* | winauth.* | ldap.* | ntlm.*
```

Save and close the file. Force Content Gateway to reread the file with the command:

```
(root)# /opt/WCG/bin/content_line -x
```

Follow the flow of debug information with the **tail -f** command:

```
(root)# tail -f /opt/WCG/logs/content_gateway.out
```

Use **Ctrl+C** to terminate the command.

When you have collected the debug output you want (after one or several user authentication processes is complete), disable debug output by editing **records.config** and modifying the parameter value as shown.

```
(root)# CONFIG proxy.config.diags.debug.enabled INT 0
```

Save and close the file. Force Content Gateway to reread the file with the command:

```
(root)# /opt/WCG/bin/content_line -x
```

Mac and iPhone/iPad authentication

Websense Web Security solutions can be used to authenticate or identify Mac and iPhone/iPad users for user- or group-based filtering.

For Mac computers, see:

- ◆ [Authentication for Mac computers](#)
 - [Enabling transparent identification of Mac users with DC Agent](#)
 - [Authenticating Mac users with Content Gateway](#)
 - [Typical steps for joining a Mac to an Active Directory domain](#)

For iPhones/iPads, see:

- ◆ [Authentication for iPhones and iPads](#)

For a list of Frequently Asked Questions regarding Mac and iPhone/iPad authentication, see [How do I use Websense Web Security solutions to authenticate or identify Mac users for user- or group-based filtering?](#)

Authentication for Mac computers

Web Security solutions can be used to authenticate or identify Mac users for user- or group-based filtering. These restriction apply:

- ◆ Authentication and identification require that users belong to an Active Directory.
- ◆ Protocol block messages cannot be displayed on Macs.

- ◆ Websense Remote Filtering Client and Web Endpoint are not supported on Mac OS X systems.

If your organization uses DC Agent for transparent user identification, see [Enabling transparent identification of Mac users with DC Agent](#).

If your organization uses Content Gateway to authenticate users, see [Authenticating Mac users with Content Gateway](#).

Manual (prompted) authentication can also be used to enable user and group-based filtering of Mac users.

Enabling transparent identification of Mac users with DC Agent

In order for DC Agent to identify the user on a Mac workstation, the Mac must mount a file share on the domain controller. This can be done by configuring the Mac to use a file share on the domain controller machine as the user's home directory, or by mounting another share with the domain controller.



Note

If the Mac only logs to the domain without mounting a file share, it will not be visible to DC Agent.

Configuration summary:

- ◆ Ensure that each participating Mac user is a member of a common Active Directory. See your Active Directory documentation.
- ◆ Create a home folder for each Mac user, and make sure that it is accessible to the user. See the first paragraph of this section.

When the user logs on to the properly configured Mac OS X system, the Mac mounts a network directory as the user's home directory, the DC Agent user map is populated, and user and group-based policies can be applied to user requests. When requests are blocked, browser-based block pages are displayed normally.

Authenticating Mac users with Content Gateway

Using the Integrated Windows Authentication (IWA) feature of Content Gateway, Mac users can be transparently authenticated when the user is a member of an Active Directory domain and the Mac computer is joined to the Active Directory domain. For more information see [Integrated Windows Authentication](#).

Configuration summary:

- ◆ Ensure that each Mac computer is joined to the Active Directory domain. See [Typical steps for joining a Mac to an Active Directory domain](#).
- ◆ Ensure that each participating Mac user is a member of a common Active Directory. See your Active Directory documentation.
- ◆ Ensure that Content Gateway is joined to the Active Directory domain.
 - If Content Gateway is not configured for IWA, see [Integrated Windows Authentication](#) and apply the configuration instructions.

- If Content Gateway is already configured for IWA and your Mac users belong to the currently joined domain, there is nothing to do.
- If Content Gateway is already configured for IWA and your Mac users belong to a different Active Directory domain, use the Rule-Based Authentication feature. See [Rule-Based Authentication](#) and follow the configuration instructions.
- ◆ When Content Gateway is an explicit proxy, configure participating Mac systems and browsers to send HTTP, HTTPS, and FTP requests to the Fully Qualified Domain Name (FQDN) of Content Gateway. Alternatively, specify the IP address of Content Gateway if NTLM is adequate.

If Content Gateway is a transparent proxy, no additional Mac system or browser configuration is required.



Important

Safari users may be prompted for credentials the first time they open a browser. The user should enter their credentials and check the “Remember password in keychain” check box.

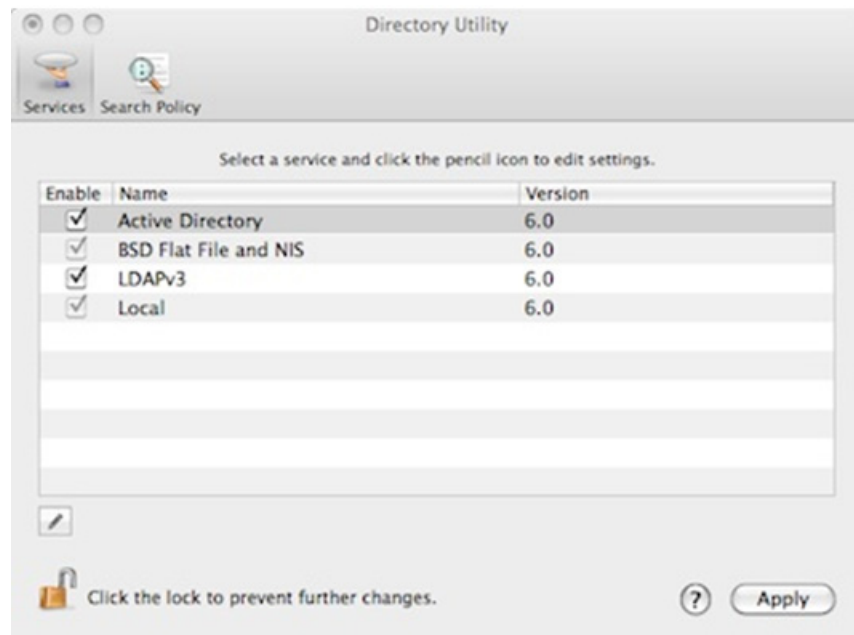
Firefox users may receive an “Proxy Authentication Required” error message. This is a known issue in Firefox (<http://support.mozilla.org/en-US/questions/926378>) and is easily corrected by changing the browser configuration. In **About:Config** set the following options to **false**:

- ◆ `network.automatic-ntlm-auth.allow-proxies`
 - ◆ `network.negotiate-auth.allow-proxies`
-

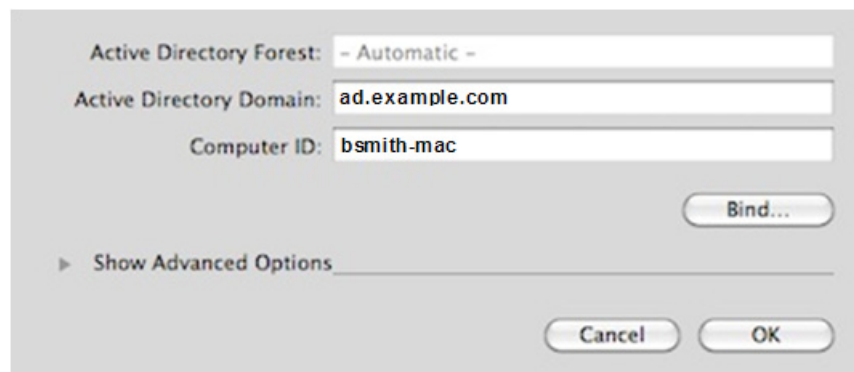
Typical steps for joining a Mac to an Active Directory domain

1. Using an account with Administrator privileges, log on to the Mac computer that you want to join to an Active Directory domain.
2. Open the **Directory Utility**. On OS X 10.6 (Snow Leopard), go to:
`/System/Library/CoreServices`
3. If necessary, click the padlock icon and enter your password to unlock the Directory Utility.

- Select the box next to **Active Directory** to enable Active Directory support.



- Highlight Active Directory and click on the Pencil icon to configure the Active Directory connection.
- Under **Domain**, enter the Fully Qualified Domain Name (FQDN).
- Under **Computer ID**, enter the computer name.



- Click Bind. You are prompted for network credentials and a computer OU. Enter your OU admin account and password, and the computer OU location. For example:

```
ou=computers,ou=orgunits,dc=ad,dc=example,dc=com
```

Your machine will be bound to the specified Active Directory.
- Click **Apply** in the Directory Utility to save your changes and restart the machine.

Authentication for iPhones and iPads

Proxy-based user authentication is supported by the Content Gateway (proxy) component of Web Security Gateway and Web Security Gateway Anywhere, resulting in user- or group-based filtering.

User identification via DC Agent is not supported and, therefore, there is no user- or group-based filtering solution with Web Filter or Web Security. Filtering can be provided to those devices based on IP address or network range.

Content Gateway user authentication has the following features and restrictions:

- ◆ Works with the authentication method configured in Content Gateway. Users must belong to the associated user directory.
- ◆ Supports the Safari browser. Other browsers may not work as expected.
- ◆ Transparent authentication is not supported. The user is always prompted for credentials.
- ◆ Works in transparent and explicit Content Gateway deployments.
- ◆ Many iPhone and iPad apps do not work well with Content Gateway (or any Web proxy) because they are not well programmed to handle proxy user authentication.

Explicit proxy settings can be configured in the iOS Network settings area.



16

Working With Log Files

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Event log files, page 242](#)
- ◆ [Managing event log files, page 243](#)
- ◆ [Event log file formats, page 245](#)
- ◆ [Rolling event log files, page 251](#)
- ◆ [Splitting event log files, page 254](#)
- ◆ [Collating event log files, page 256](#)
- ◆ [Viewing logging statistics, page 259](#)
- ◆ [Viewing log files, page 260](#)
- ◆ [Example event log file entries, page 261](#)

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 for error conditions in the Content Gateway manager; see [Working with alarms, page 121](#), for details.)

All system information messages are logged with the system-wide logging facility **syslog** under the daemon facility. The **syslog.conf** configuration file (stored in the **/etc** directory) specifies where these messages are logged. A typical location is **/var/log/messages**.

The **syslog** process works on a system-wide basis, so it is the single repository for messages from all Content Gateway processes, including **content_gateway**, **content_manager**, and **content_cop**.

Each log entry in the log contains information about the date and time the error was logged, the hostname of the proxy server that reported the error, and a description of the error or warning.

See [Websense Content Gateway error messages, page 495](#), 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

Help | Content Gateway | Version 7.8.x

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, 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 243.
- ◆ 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 245.
- ◆ 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 251.
- ◆ Separate log files according to hosts
You can configure the proxy to create separate log files for different protocols based on the host. See [Splitting event log files](#), page 254.

- ◆ 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 256](#).
- ◆ View statistics about the logging system
Content Gateway provides statistics about the logging system. You can access the statistics through the Content Gateway manager or through the command line interface. See [Viewing logging statistics, page 259](#).
- ◆ 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 261](#).

Managing event log files

Help | Content Gateway | Version 7.8.x

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 244](#).

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 244](#)), Content Gateway enters a low space state and takes the following actions:

- ◆ If the autodelete option (discussed in [Rolling event log files, page 251](#)) 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.
When Content Gateway is on a V-series appliance, the size is set to 5120 (5 GB) and cannot be changed.
When Content Gateway is installed on a stand-alone server, the default size is 20480 (20 GB) and the size is configurable.



Note

All files in the logging directory contribute to the space used, even if they are not log files.

4. In the **Headroom** field, enter the tolerance for the log space limit. The default value is 100 MB.
If the **Auto-Delete Rolled Files** option is enabled in the **Log Rolling** section, autodeletion is triggered when the amount of free space available in the logging directory is less than the headroom. For information about log file rolling, see [Rolling event log files, page 251](#).
5. Click **Apply**.

Event log file formats

Help | Content Gateway | Version 7.8.x

Websense Content Gateway supports the following log file formats:

- ◆ *Standard formats*, such as Squid or Netscape (see [Using standard formats, page 245](#))
- ◆ the Content Gateway *custom format* (see [Custom format, page 246](#))

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 249](#).



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.



Important

When IPv6 is enabled, Event log entries are normalized to IPv6 format.

For example, “10.10.41.200” is logged as “::ffff:10.10.41.200”.

To filter on a client at “10.10.41.200” in a custom log, use:

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

Using standard formats

Help | Content Gateway | Version 7.8.x

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 246](#).

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

Help | Content Gateway | Version 7.8.x

The XML-based custom log format is more flexible than the standard log file formats, giving you more control over the type of information in your log files. Create a custom log format if you need data for analysis that is not available in the standard formats. You can decide what information to record for each Content Gateway transaction and create filters to define which transactions to log.

The heart of the custom logging feature is an XML-based logging configuration file (**logs_xml.config**) that enables you to create modular descriptions of logging objects. The **logs_xml.config** file uses three types of objects to create custom log files:

- ◆ The **LogFormat** defines the content of the log file using printf-style format strings.
- ◆ The **LogFilter** defines a filter so that you include or exclude certain information from the log file.
- ◆ The **LogObject** specifies all the information needed to produce a log file. For example:
 - The name of the log file (required).
 - The format to be used (required). This can be a standard format (Squid or Netscape) or a previously defined custom format (a previously defined **LogFormat** object).
 - The file mode (ASCII, Binary, or ASCII_PIPE). The default is ASCII.

The ASCII_PIPE mode writes log entries to a UNIX named pipe (a buffer in memory). Other processes can then read the data using standard I/O functions. The advantage of using this option is that Content Gateway does not have to write to disk, freeing disk space and bandwidth for other tasks.



Note

When the buffer is full, Content Gateway drops log entries and issues an error message indicating how many entries were dropped. Content Gateway writes only complete log entries to the pipe; therefore, only full records are dropped.

- Any filters you want to use (previously defined **LogFilter** objects).
- The collation servers that are to receive the log files.
- The protocols you want to log (if the protocols tag is used, Content Gateway logs only transactions from the protocols listed; otherwise, all transactions for all protocols are logged).
- The origin servers you want to log (if the servers tag is used, Content Gateway logs only transactions for the origin servers listed; otherwise, transactions for all origin servers are logged).
- The header text you want the log files to contain. The header text appears at the beginning of the log file, just before the first record.
- The log file rolling options.

**Note**

To generate a custom log format, you must specify at least one **LogObject** definition. One log file is produced for each **LogObject** definition. You can create a custom log format in the Content Gateway manager or by editing a configuration file.

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 408.
3. Click **Apply**.

Creating summary log files

Help | Content Gateway | Version 7.8.x

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 408](#).

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**.

Applying logs_xml.config file changes to all nodes in a cluster

Help | Content Gateway | Version 7.8.x

After modifying the **logs_xml.config** file on one Content Gateway node, enter the following command from the Content Gateway **bin** directory (**/opt/WCG/bin**):

```
content_line -x
```

Content Gateway applies the changes to all nodes in the cluster. The changes take effect immediately.

Choosing binary or ASCII

Help | Content Gateway | Version 7.8.x

You can configure Content Gateway to create event log files in either of the following:

- ◆ **ASCII**: these files can be processed using standard, off-the-shelf log-analysis tools. However, Content Gateway must perform additional processing to create the files in ASCII, resulting in an increase in overhead. Also, ASCII files tend to be larger than the equivalent binary files. ASCII log files have a **.log** filename extension by default.
- ◆ **Binary**: these files generate lower system overhead, as well as generally occupying less space on the disk, depending on the type of information being logged. You must, however, use a converter application before you can read or analyze these files using standard tools. Binary log files use a **.blog** filename extension by default.

While binary log files typically require less disk space, this is not always the case. For example, the value 0 (zero) requires only one byte to store in ASCII but requires four bytes when stored as a binary integer. If you define a custom format that logs IP addresses, a binary log file would require only four bytes of storage per 32-bit address. However, the same IP address stored in dot notation would require around 15 characters (bytes) in an ASCII log file.

For standard log formats, you select **Binary** or **ASCII** on the **Configure > Subsystems > Logging > Formats** tab in the Content Gateway manager. See [Setting standard log file format options, page 246](#). For the custom log format, you specify

ASCII or Binary mode in the **LogObject**. Refer to [Custom format](#), page 246.

**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 408, 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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 hostname of the Content Gateway server that generated the log file.
- ◆ Two timestamps separated by a hyphen (-). The first time stamp is a lower bound for the time stamp of the first record in the log file. The lower bound is the time when the new buffer for log records is created. Under low load, the first time stamp in the filename can be different from the timestamp of the first entry. Under normal load, the first time stamp in the filename and the time stamp of the first entry are similar.

The second time stamp is an upper bound for the time stamp of the last record in the log file (this is normally the rolling time).

- ◆ The suffix **.old**, which makes it easy for automated scripts to find rolled log files.

The timestamps have the following format:

```
%Y%M%D.%Hh%Mm%SS-%Y%M%D.%Hh%Mm%SS
```

The following table describes the format:

| Code | Definition | Example |
|------|--|---------|
| %Y | The year in four-digit format | 2000 |
| %M | The month in two-digit format, from 01-12 | 07 |
| %D | The day in two-digit format, from 01-31 | 19 |
| %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

Help | Content Gateway | Version 7.8.x

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 the Content Gateway manager or the **records.config** file described above.

Splitting event log files

Help | Content Gateway | Version 7.8.x

By default, Websense Content Gateway uses standard log formats and generates log files that contain HTTP and FTP transactions in the same file. However, you can enable host log splitting if you prefer to log transactions for different origin servers in separate log files.

HTTP host log splitting

Help | Content Gateway | Version 7.8.x

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 255](#)).

When HTTP host log splitting is enabled, Content Gateway generates separate log files for HTTP/FTP transactions, based on the origin server.

For example, if the `log_hosts.config` file contains the two origin servers **uni.edu** and **company.com**, and the Squid format is enabled, Content Gateway generates the following log files:

| Log Filename | Description |
|------------------------------------|--|
| <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.log</code> | All HTTP and FTP transactions for other hosts |

Content Gateway also enables you to create XML-based custom log formats that offer even greater control over log file generation based on protocol and host name. See [Custom format, page 246](#).

Setting log splitting options

Help | Content Gateway | Version 7.8.x

1. Navigate to **Configure > Subsystems > Logging > Splitting**.
2. Enable the **Split Host Logs** option to record all HTTP and FTP transactions for each origin server listed in the `log_hosts.config` file in a separate log file. Disable the **Split Host Logs** option to record all HTTP and FTP transactions for each origin server listed in the `log_hosts.config` file in the same log file.
3. Click **Apply**.

Editing the log_hosts.config file

Help | Content Gateway | Version 7.8.x

The default **log_hosts.config** file is located in **/opt/WCG/config**. To record HTTP and FTP transactions for different origin servers in separate log files, you must specify each origin server's hostname on a separate line in the file.

**Note**

You can specify keywords in the **log_hosts.config** file to record in a separate log file all transactions from origin servers that contain the specified keyword in their names. For example, if you specify the keyword **sports**, Content Gateway records all HTTP and FTP transactions from **sports.yahoo.com** and **www.foxsports.com** in a log file called **squid-sports.log** (if the Squid format is enabled).

**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 **/opt/WCG/config**.
2. Enter the hostname 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 (**/opt/WCG/bin**):

```
./content_line -x
```

Collating event log files

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 246](#).

3. In the **To Collation Server** field, enter the hostname of the collation server. This could be the Content Gateway collation server or a stand-alone collation server.
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

Help | Content Gateway | Version 7.8.x

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 258.
2. Copy the **sac** binary from the Content Gateway **bin** directory (**/opt/WCG/bin**) to the machine serving as the stand-alone collator.
3. Create a directory called **config** in the directory that contains the **sac** binary.
4. Create a directory called **internal** in the **config** directory you created in [Step 3](#). This directory will be used internally by the stand-alone collator to store lock files.
5. Copy the **records.config** file (**/opt/WCG/config**) from a Content Gateway node configured to be a log collation client to the **config** directory you created in [Step 3](#) on the stand-alone collator.

The **records.config** file contains the log collation secret and port you specified when configuring nodes to be collation clients. The collation port and secret must be the same for all collation clients and servers.

6. Open the **records.config** file on the stand-alone collator and edit the following variable:

| Variable | Description |
|--------------------------------------|--|
| <i>proxy.config.log2.logfile_dir</i> | 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. |

7. Save and close the file.
8. Enter the following command:

```
sac -c config
```

Viewing logging statistics

Help | Content Gateway | Version 7.8.x

Content Gateway generates statistics about the logging system that help you see the following information:

- ◆ How many log files (formats) are currently being written.
- ◆ The current amount of space being used by the logging directory, which contains all of the event and error logs.
- ◆ The number of access events that have been written to log files since Content Gateway installation. This counter represents one entry in one file. If multiple formats are being written, a single event will create multiple event log entries.

- ◆ The number of access events skipped (because they were filtered out) since Content Gateway installation.
- ◆ The number of access events that have been written to the event error log since Content Gateway installation.

You can view the statistics from the Monitor tab in the Content Gateway manager or retrieve them through the command-line interface. See [Monitoring Traffic](#), page 117.

Viewing log files

Help | Content Gateway | Version 7.8.x

Related topics:

- ◆ [Squid format](#), page 261
- ◆ [Netscape examples](#), page 262

You can view the system, event, and error log files that Content Gateway creates from the Content Gateway manager. You can view an entire log file, a specified last number of lines in the log file, or all lines that contain a specified string.

You can also delete a log file or copy it to your local system.



Note

You must have the correct user permissions to copy and delete log files.



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 the 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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

This section shows examples of a log file entry in each of the standard log formats supported by Content Gateway:

- ◆ [Squid format, page 261](#)
- ◆ [Netscape examples, page 262](#)
- ◆ [Netscape Extended format, page 263](#)
- ◆ [Netscape Extended-2 format, page 263](#)

Squid format

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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 Cache result codes in Squid- and Netscape-format log files , page 264. The proxy response status code (the HTTP response status code from Content Gateway to client). |
| 5 | The length of the Content Gateway response to the client in bytes, including headers and content. |
| 6 | The client request method: GET, POST, and so on. |
| 7 | The client request canonical URL; blanks and other characters that might not be parsed by log analysis tools are replaced by escape sequences. The escape sequence is a percentage sign followed by the ASCII code number of the replaced character in hex. |
| 8 | The authenticated client's user name. A hyphen (-) means that no authentication was required. |
| 9 | The proxy hierarchy route; the route Content Gateway used to retrieve the object. The proxy request server name; the name of the server that fulfilled the request. If the request was a cache hit, this field contains a hyphen (-). |
| 10 | The proxy response content type; the object content type taken from the Content Gateway response header. |

Netscape examples

Help | Content Gateway | Version 7.8.x

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. |

| Field | Description |
|-------|---|
| | Netscape Extended |
| 8 | The origin server's response status code. |
| 9 | The server response transfer length; the body length in the origin server's response to the proxy, in bytes. |
| 10 | The client request transfer length; the body length in the client's request to the proxy, in bytes. |
| 11 | The proxy request transfer length; the body length in the proxy request to the origin server. |
| 12 | The client request header length; the header length in the client's request to the proxy. |
| 13 | The proxy response header length; the header length in the proxy response to the client. |
| 14 | The proxy request header length; the header length in the proxy request to the origin server. |
| 15 | The server response header length; the header length in the origin server's response to the proxy. |
| 16 | The time Content Gateway spent processing the client request; the number of seconds between the time that the client established the connection with the proxy and the time that the proxy sent the last byte of the response back to the client. |
| | Netscape Extended-2 |
| 17 | The proxy hierarchy route; the route Content Gateway used to retrieve the object. |
| 18 | The client finish status code: FIN if the client request completed successfully or INTR if the client request was interrupted. |
| 19 | The proxy finish status code: FIN if the Content Gateway request to the origin server completed successfully or INTR if the request was interrupted. |
| 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 Cache result codes in Squid- and Netscape-format log files, page 264 . |

Cache result codes in Squid- and Netscape-format log files

Help | Content Gateway | Version 7.8.x

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. |
| 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. |

| Cache Result Code | Description |
|--------------------------|--|
| 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. |

A

Statistics

Help | Content Gateway | Version 7.8.x

This section describes the following statistics accessed on the Content Gateway manager **Monitor** tab:

- ◆ [My Proxy](#), page 267
- ◆ [Protocols](#), page 273
- ◆ [Security](#), page 275
- ◆ [Subsystems](#), page 280
- ◆ [Networking](#), page 282
- ◆ [Performance](#), page 287
- ◆ [SSL](#), page 290

My Proxy

Help | Content Gateway | Version 7.8.x

My Proxy statistics are divided into the following categories:

- ◆ [Summary](#), page 267
- ◆ [Node](#), page 269
- ◆ [Graphs](#), page 270
- ◆ [Alarms](#), page 270
- ◆ [Diagnostics](#), page 271

Summary

Help | Content Gateway | Version 7.8.x

| Statistic/Field | Description |
|-----------------|---|
| | Subscription Details |
| Feature | Lists available features, such as scanning options, Data Security, and ThreatScope. |

| Statistic/Field | Description |
|---|--|
| Purchased Status | Indicates if a feature has been purchased. |
| Expiration Date | If a feature has been purchased, displays the expiration date of the subscription. |
| | More Detail |
| Subscription key | Displays the subscription key. See Entering your subscription key , page 16. |
| Last successful subscription download time | Displays the time of the last successful validation of the subscription key. The check is made once a day. |
| Connection status | Displays the Content Gateway connection status to Policy Server, Policy Broker, and Filtering Service. |
| Registration status | Displays the Content Gateway registration status with the Forensics Repository. |
| | Scanning Data Files |
| Engine Name | Displays the name of each scanning engine. |
| Engine Version | Displays the version number of the scanning engine. |
| Data File Version | Displays the version number of the data file currently in use by the scanning engine. |
| Last update | Displays the time and date when Content Gateway last successfully loaded that analytics data files, settings, and policies. |
| Last time Content Gateway loaded data | Displays the time and date when Content Gateway last successfully loaded databases, settings, and policies. |
| Last time Content Gateway checked for updates | Displays the time and date when Content Gateway last successfully communicated with the Websense download server to check for data file updates. |
| | Node Details |
| Node | Name of the Content Gateway node or cluster. |
| On/Off | Indicates if the proxy and manager services are running. |
| Objects Served | The total number of objects served by the node. |
| Ops/Sec | The number of operations per second processed by the node. |
| Hit Rate | The percentage of HTTP requests served from the cache, averaged over the past 10 seconds. |

| Statistic/Field | Description |
|-----------------------|---|
| Throughput (Mbit/sec) | <p>The number of megabits per second passing through the node (and cluster).</p> <p>The proxy updates the throughput statistic after it transfers an entire object. For larger files, the byte count increases sharply at the end of a transfer. The complete number of bytes transferred is attributed to the last 10-second interval, although it can take several minutes to transfer the object.</p> <p>This transient inaccuracy is more noticeable with a light load.</p> |
| HTTP Hit (ms) | The amount of time it takes for an HTTP object that is fresh in the cache to be served to the client. |
| HTTP Miss (ms) | The 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 | The percentage of HTTP requests served from the cache, averaged over the past 10 seconds. This value is refreshed every 10 seconds. |
| errors | The percentage of requests that end in early hangups. |
| aborts | The percentage of aborted requests. |
| active clients | The current number of open client connections. |
| active servers | The current number of open origin server connections. |
| node IP address | The IP address assigned to the node. If virtual IP addressing is enabled, several virtual IP addresses could be assigned to this node. |
| cache free space | The amount of free space in the cache. |
| HostDB hit rate | The ratio of host database hits to total host database lookups, averaged over a 10-second period. |

Node

Help | Content Gateway | Version 7.8.x

| 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 |

| Statistic | Description |
|------------------------------|--|
| 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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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 117](#).



Important

The graph is displayed in your browser using a Java applet. You should have the latest version of Java installed on your PC (at least version 1.7). To validate your access to Content Gateway statistics, you will be prompted for Content Gateway logon credentials.

Alarms

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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 121](#).

Diagnostics

Help | Content Gateway | Version 7.8.x

Beginning with 7.8.4, use the tools provided to help diagnose communication or connection issues, trace network packets, or capture network packets.

Diagnostic tools are provided for:

- ◆ [Automatic diagnostic tests, page 271](#)
- ◆ [Manual diagnostic tests, page 272](#)

Automatic diagnostic tests

By default, the page opens to the **Automatic** tab. Click **Run Diagnostics** to execute all of the tests listed in the table. Connectivity is tested from the Content Gateway host machine to each of the servers listed under **Test**. In addition, the availability of the DNS servers is confirmed.

- ◆ The IPv4 default gateway
- ◆ The IPv6 default gateway
- ◆ Your primary DNS server
- ◆ Your secondary DNS server
- ◆ download.websense.com (a Websense download server)
- ◆ ddsdom.websense.com (a Websense download server)
- ◆ ddsint.websense.com (a Websense download server)
- ◆ my.websense.com (customer account portal)

Once the diagnostics are run, additional information is provided:

- ◆ **Result** indicates whether the test is running, passed, failed, or could not complete.
- ◆ **Latency** provides the round-trip latency of the Ping command used to test the connection. The value, reported in milliseconds, is the amount of time between the command being sent and the response being received from the server.

An empty latency value does not necessarily indicate a problem. Rather, it indicates either (a) that the test passed, but the packet that holds the value was

banned by something in the network, or (b) that the test failed, and thus no latency value could be obtained.

If the value seems high (a full 10 seconds, for example) when compared to other latency values, it may indicate a problem in the network.

- ◆ **Details** offers additional information for any test that failed or could not complete.

Below the table, the Last update information reflects the date and time the connections were last tested. Each time you access the page, the results of the last test will display.

Manual diagnostic tests

The **Manual** tab offers 4 commands typically run from the Linux command line.

- ◆ **Ping**, used to determine if a remote device can be reached across the network.
- ◆ **Traceroute**, used to determine the path network packets take and measure delays across the network.
- ◆ **NSlookup**, used to obtain domain name or IP address mapping.
- ◆ **TCPDump**, used to analyze network packets.

Click the radio button next to the command you want to execute and enter parameters for the command in the entry field provided.

- ◆ Enter a server name or IP address for **Ping** or **Traceroute**.
- ◆ Enter a server name for **NSlookup**.
- ◆ Enter valid parameters for **TCPDump**. Click the link provided for additional information on using TCPDump with Content Gateway. View the same technical article using [this link](#).

Click the **Run** button next to your selected command to execute the test. The results for Ping, Traceroute, and NSlookup display in the **Test Results** section at the bottom of the pane.

Test results for TCPDump are typically too long to easily display and review in the Test Results window. When TCPDump is run, the Test Results window simply indicates the success or failure of the command.

As TCPDump runs, output is written to /opt/WCG/logs/tcpdump.pcap. This file is overwritten each time TCPDump is executed. When a test is successful, a link is provided so that you can download and view or save a copy of the most recent file.

To avoid disk space problems, tcpdump.pcap is limited to 10,000 packets. Once that limit is reached, no additional output is written to the file.



Important

TCPDump uses a lot of system resources. Try to avoid using it during peak hours when the system is busy.

As each command executes, the **Run** button becomes a **Stop** button. Click **Stop** to abort the command.

Protocols

Help | Content Gateway | Version 7.8.x

Protocol statistics are divided into the following categories:

- ◆ [HTTP, page 273](#)
- ◆ [FTP, page 275](#)

For [SSL](#) statistics, click the SSL button at the bottom of the Monitor tab.

HTTP

Help | Content Gateway | Version 7.8.x

| Statistic | Description |
|--------------------------|---|
| | General |
| Client | |
| Total Document Bytes | Total amount of HTTP data served to clients since installation. |
| Total Header Bytes | Total amount of HTTP header data served to clients since installation. |
| Total Connections | Total number of HTTP client connections since installation. |
| Current Connections | Current number of HTTP client connections |
| Transactions in Progress | Total number of HTTP client transactions in progress. |
| Server | |
| Total Document Bytes | Total amount of HTTP data received from origin servers since installation. |
| Total Header Bytes | Total amount of HTTP header data received from origin servers since installation. |
| Total Connections | Total number of HTTP server connections since installation. |
| Current Connections | Current number of HTTP server connections |
| Transactions in Progress | Total number of HTTP server connections currently in progress. |
| | Transaction |
| Hits | |
| Fresh | Percentage of hits that are fresh and their average transaction times. |

| Statistic | Description |
|-----------------------------|---|
| 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. |
| 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. |

| Statistic | Description |
|-----------------------------|---|
| 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

Help | Content Gateway | Version 7.8.x

| 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. |

Security

Help | Content Gateway | Version 7.8.x

Security statistics are divided into the following categories:

- ◆ [Integrated Windows Authentication](#), page 276
- ◆ [LDAP](#), page 278
- ◆ [Legacy NTLM](#), page 278
- ◆ [SOCKS](#), page 279

- ◆ [Data Security, page 279](#)



Note

Even when multiple authentication rules are used, Content Gateway reports authentication statistics discreetly for each authentication method (IWA, LDAP, Legacy NTLM).

Integrated Windows Authentication

Help | Content Gateway | Version 7.8.x

| Statistic | Description |
|--------------------------------|---|
| | <p>Diagnostic Test</p> <p>This function runs diagnostic tests on the Kerberos connection to the selected domain. Results are displayed on screen and written to /opt/WCG/logs/content_gateway.out and /opt/WCG/logs/smbadmin.log.</p> |
| Domain drop down box | Select a joined domain. Unless Rule-Based Authentication is configured, there will only be 1 joined domain. |
| Run Test button | Click to initiate a test. |
| | <p>Active Directory Joined Domains list</p> <p>Lists all joined AD domains.</p> <p>The Content Gateway Hostname DNS is the name that clients must specify in their browser proxy settings for Kerberos authentication to occur.</p> |
| | <p>Kerberos request counters</p> |
| Total Kerberos requests | The total number of Kerberos authentication requests. |
| Authentication succeeded | The number of Kerberos authentication requests that resulted in successful authentication. |
| Authentication failed | The number of Kerberos authentication requests that resulted in authentication failure. |
| Kerberos errors | The number of Kerberos process errors. |
| | <p>NTLM request counters</p> |
| Total NTLM requests | The total number of NTLM authentication requests. |
| Authentication succeeded | The number of NTLM authentication requests that resulted in successful authentication. |
| Authentication failed | The number of NTLM authentication requests that resulted in authentication failure. |
| NTLM request errors | The number of NTLM process errors. |
| NTLM within negotiate requests | The number of NTLM requests encapsulated in Negotiate requests. |

| Statistic | Description |
|---|---|
| | Basic authentication request counters |
| Total basic authentication requests | The total number of basic authentication requests. |
| Authentication succeeded | The number of basic authentication requests that resulted in successful authentication. |
| Authentication failed. | The number of basic authentication requests that resulted in authentication failure. |
| Basic authentication request errors | The number of basic authentication process errors. |
| | Performance counters |
| Kerberos - Average time per transaction | The average time, in milliseconds, to complete a Kerberos transaction. |
| NTLM - Average time per transaction | The average time, in milliseconds, to complete a NTLM transaction. |
| Basic - Average time per transaction | The average time, in milliseconds, to complete a basic transaction. |
| Average helper latency per transaction | The average time for Samba to process an authentication request. |
| Time authentication spent offline | <p>The time, in seconds, that Content Gateway was unable to perform NTLM authentication due to service or connectivity failures. (This measure does not apply to Kerberos because no communication with the DC is needed.)</p> <p>If the Fail Open option is enabled (<i>Global authentication options</i>), proxy requests may proceed without authentication.</p> <p>The counter is incremented when connectivity is reestablished after a failure.</p> |
| Number of times authentication servers or services went offline | The number of times that connectivity with authentication servers or services has been lost. |
| | Top lists counters |
| | <p>These user authentication lists provide a view into which User-Agent values and client IP addresses are most active. Four counters tally the top 20 User-Agent and client IP addresses that are passing or failing user authentication.</p> |
| Button: Reset Top Lists to Zero | Resets all Top Lists counters to zero. |
| Top User-Agents passing authentication | Lists the top 20 User-Agent matches by number of authentication attempts that pass authentication. |
| Top User-Agents failing authentication | Lists the top 20 User-Agent matches by number of authentication attempts that fail authentication. |

| Statistic | Description |
|--|---|
| Top Client IP addresses passing authentication | Lists the top 20 client IP addresses by number of authentication attempts that pass authentication. |
| Top Client IP addresses failing authentication | Lists the top 20 client IP addresses by number of authentication attempts that fail authentication. |

LDAP

Help | Content Gateway | Version 7.8.x

| 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. |
| | Successful Authentications (7.8.2) |
| Authentication Succeeded | Number of times authentication was successful. |
| | Unsuccessful Authentications |
| Authentication Denied | Number of times the LDAP Server denied authentication. |
| Authentication Timeouts | Number of times authentication timed out. |
| Authentication Cancelled | Number of times authentication was terminated after LDAP authentication was started and before it was completed. Note: This does not count the number of times that an authentication request was cancelled by the client by clicking “Cancel” in the dialog box that prompts for credentials. |

Legacy NTLM

Help | Content Gateway | Version 7.8.x

| 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. |

| Statistic | Description |
|--------------------------|--|
| | Successful Authentications (7.8.2) |
| Authentication Succeeded | Number of times authentication was successful. |
| | Unsuccessful Authentications |
| Authentication Denied | Number of times the NTLM server denied authentication. |
| Authentication Cancelled | Number of times authentication was cancelled. |
| Authentication Rejected | Number of times authentication failed because the queue was full. |
| | Queue Size |
| Authentication Queued | Number of requests that are currently queued because all of the domain controllers are busy. |

SOCKS

[Help](#) | [Content Gateway](#) | Version 7.8.x

| Statistic | Description |
|--|---|
| On-Appliance SOCKS Server (when Content Gateway is on a V-Series appliance) | Indicates whether the on-appliance SOCKS server is on (enabled) or off (disabled). |
| Unsuccessful Connections | Number of unsuccessful connections to the SOCKS server since Content Gateway was started. |
| Successful Connections | Number of successful connections to the SOCKS server since Content Gateway was started. |
| Connections in Progress | Number of connections to the SOCKS server currently in progress. |

Data Security

[Help](#) | [Content Gateway](#) | Version 7.8.x

| 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. |

| Statistic | Description |
|-----------------------|---|
| 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. |
| Total Bytes Scanned | Total number of bytes scanned by Data Security. |
| Average Response Time | Average time needed to by Data Security to complete a scan since the last time Content Gateway was started. |

Subsystems

Help | Content Gateway | Version 7.8.x

Subsystems statistics are divided into the following categories:

- ◆ [Cache](#), page 280
- ◆ [Clustering](#), page 282
- ◆ [Logging](#), page 282

Cache

Help | Content Gateway | Version 7.8.x



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. |

| Statistic | Description |
|------------------|--|
| 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. |
| | 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

Help | Content Gateway | Version 7.8.x

| Statistic | Description |
|------------------|-----------------------------|
| Clustering Nodes | Number of clustering nodes. |

Logging

Help | Content Gateway | Version 7.8.x

| Statistic | Description |
|---------------------------------|---|
| Currently Open Log Files | Number of event log files (formats) that are currently being written. |
| Space Used for Log Files | Current amount of space being used by the logging directory, which contains all of the event and error logs. |
| Number of Access Events Logged | Number of access events that have been written to log files since Content Gateway installation. This counter represents one entry in one file. If multiple formats are being written, a single access creates multiple event log entries. |
| Number of Access Events Skipped | Number of access events skipped (because they were filtered out) since Content Gateway installation. |
| Number of Error Events Logged | Number of access events that have been written to the event error log since Content Gateway installation. |

Networking

Help | Content Gateway | Version 7.8.x

Networking statistics are divided into the following categories:

- ◆ [System](#), page 283
- ◆ [ARM](#), page 283
- ◆ [ICAP](#), page 284
- ◆ [WCCP](#), page 285
- ◆ [DNS Resolver](#), page 286
- ◆ [Virtual IP](#), page 287

System

Help | Content Gateway | Version 7.8.x

| Statistic/Field | Description |
|-----------------|---|
| | General |
| Hostname | The hostname assigned to this Content Gateway machine. |
| Search Domain | Search domain that this Content Gateway machine uses. |
| IPv4 or IPv6 | |
| Default Gateway | IP address of the default gateway used to forward packets from this Content Gateway machine to other networks or subnets. |
| Primary DNS | IP address of the primary DNS server that this Content Gateway machine uses to resolve host names. |
| Secondary DNS | Secondary DNS server that this Content Gateway machine uses to resolve host names. |
| Tertiary DNS | Third DNS server that this Content Gateway machine uses to resolve host names. |
| | NIC <interface_name> |
| Status | Indicates whether the NIC is up or down. |
| Start on Boot | Indicates whether the NIC is configured to start on boot. |
| IPv4 or IPv6 | |
| IP address | The assigned IP address of the NIC. |
| Netmask | The netmask that goes with the IP address. |
| Gateway | The configured default gateway IP address for the NIC. |

ARM

Help | Content Gateway | Version 7.8.x

| 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. |

| Statistic | Description |
|----------------------------------|---|
| | Bypass Statistics |
| Total Connections Bypassed | Total number of connections bypassed by the ARM. |
| Connections Dynamically Bypassed | Total number of connections dynamically bypassed. See Dynamic bypass rules, page 71 . |
| DNS Packets Bypassed | Number of DNS packets bypassed by the ARM. |
| Connections Shed | Total number of connections shed. See Connection load shedding, page 73 . |
| | 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

Help | Content Gateway | Version 7.8.x

| 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. |

| Statistic | Description |
|--------------------|---|
| 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

Help | Content Gateway | Version 7.8.x

WCCP v2 statistics are displayed only if WCCP version v2 is enabled.

| Statistic/Field | Description |
|-----------------------------|--|
| | WCCP v2.0 Statistics |
| WCCP Fragmentation | |
| Total Fragments | Total number of WCCP fragments. |
| Fragmentation Table Entries | Number of entries in the fragmentation table. |
| Out of Order Fragments | Number of fragments out of order. |
| Matches | Number of fragments that match a fragment in the fragmentation table. |
| Service group name | |
| Service Group ID | Service Group ID for the protocol being serviced. |
| Configured mode | The forward, return and assignment settings. |
| IP Address | IP address to which the router is sending traffic. |
| Leader's IP Address | IP address of the leader in the WCCP cache farm. |
| Number of Buckets Assigned | Number of buckets assigned to this Content Gateway node. Determined by the value of Weight and the current active nodes. |
| 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. |

| Statistic/Field | Description |
|------------------------|---|
| Router IP Address | <p>IP address of the WCCP router sending traffic to Content Gateway.</p> <p>Note: If the WCCP router is configured with multiple IP addresses, as for example when the router is configured to support multiple VLANs, the IP address reported in Monitor > Networking > WCCP statistics, and in packet captures, may differ from the IP address configured here. This is because the router always reports traffic on the highest active IP address.</p> <p>One way to get the router to always report the same IP address is to set the router's loopback address to a value higher than the router's highest IP address, then the loopback address is always reported as the router's IP address. This is the recommended configuration.</p> |
| Router ID Received | The number of times that Content Gateway has received WCCP protocol messages from the router(s). |
| Router Negotiated mode | The return, forward, and assignment modes negotiated with the router. |

DNS Proxy

Help | Content Gateway | Version 7.8.x

| Statistic | Description |
|----------------|---|
| Total Requests | Total number of DNS requests received from clients. |
| Hits | Number of DNS cache hits. |
| Misses | Number of DNS cache misses. |

DNS Resolver

Help | Content Gateway | Version 7.8.x

| 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 |

| Statistic | Description |
|-------------------|--|
| 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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The Virtual IP table displays the virtual IP addresses that are managed by the proxies in the cluster.

Client Connection Status

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

| Statistic | Description |
|---|---|
| | Clients Connections |
| Current Unique Clients Connected | |
| Total Unique Clients that have Connected | Total since Content Gateway last started. |
| Total Clients that have Exceeded the Limits | Total clients that exceeded the connection limits since Content Gateway last started. See Configure > Connection Management > Client Connection Control . |
| Total Clients for which Connections were Closed | Total since Content Gateway last started. |

Performance

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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.

A description is given adjacent to each graph. Click on a graph to get the daily, weekly, monthly, and yearly on a single screen.

These graphs are available (sorted alphabetically):

- Active Client Connections
- Active Native FTP Client Connections
- Active Origin Server Connections
- Active Parent Proxy Connections
- Bandwidth Savings
- Cache Read
- Cache Reads Per Second
- Cache Writes
- Cache Writes Per Second
- Completed Client Transactions Per Second
- Content Gateway Manager Memory Usage
- Content Gateway Uptime
- CPU Available
- CPU Busy
- Data Security Module Memory Usage
- Disk Cache Usage
- DNS Cache Usage

- DNS Lookup Latency
- HTTP Abort Latency
- HTTP and HTTPS Transactions Per Second
- HTTP Cache Hit Latency
- HTTP Cache Miss Latency
- HTTP Connection Errors & Aborts (Count)
- HTTP Connection Errors & Aborts (Percentage)
- HTTP Document Hit Rate
- HTTP Error Latency
- HTTP Hits & Misses (Count)
- HTTP Hits & Misses (Percentage)
- HTTP POST and FTP PUT Transactions Per Second
- Microsoft Internet Explorer Browser Requests (Percentage)
- MRTG Runtime
- Network Reads
- Network Writes
- RAM Cache Read I/O Hit Rate
- RAM Cache Usage
- System Memory
- TCP CLOSE_WAIT Connections
- TCP Connect Rate
- TCP ESTABLISHED Connections
- TCP FIN_WAIT_1 Connections
- TCP FIN_WAIT_2 Connections
- TCP LAST_ACK Connections
- TCP Segments Transmitted
- TCP Throughput
- TCP TIME_WAIT Connections
- Throughput in Bytes
- Throughput in Error and Dropped Packets
- Throughput in Packets
- Transaction Buffer Memory Usage
- WCCP Exceptional Input Fragments
- WCCP Fragment Table Size
- WCCP Input Fragments
- Web Security Scanned Transactions (Percentage)
- Web Security Slow Scanned Transactions
- Web Security Slow Transactions

- Websense Content Gateway Memory Usage

SSL

Help | Content Gateway | Version 7.8.x

The following tabs monitor and report on SSL traffic.

[SSL Key Data](#), page 290

[CRL Statistics](#), page 291

[Reports](#), page 291

SSL Key Data

Help | Content Gateway | Version 7.8.x

These fields provide information about SSL connections and activity.

| Statistic/Field | Description |
|---|---|
| | SSL Inbound Key Data |
| Is alive | Online indicates that SSL support is enabled. |
| Current SSL connections | The number of active inbound SSL requests (browser to Content Gateway). |
| Total SSL server connections | The number of browser requests. |
| Total finished SSL server connections | The number of browser requests that resulted in decryption. |
| Total SSL server renegotiation requests | The number of browser requests renegotiated due to handshake failures or invalid certificates between the browser and Content Gateway. |
| | SSL Outbound Key Data |
| Is alive | Online indicates that SSL support is enabled. |
| Current SSL connections | The number of active outbound SSL requests (Content Gateway to origin server). |
| Total SSL client connections | The number of Content Gateway requests to origin servers. |
| Total finished SSL client connections | The number of requests where data went from Content Gateway to the origin server. |
| Total SSL client renegotiation requests | The number of requests that were renegotiated due to handshake failures or invalid certificates between Content Gateway and the origin server |

| Statistic/Field | Description |
|----------------------------------|---|
| Total SSL session cache hits | The number of times that a request was validated by a key in the session cache. |
| Total SSL session cache misses | The number of times that a request could not be validated by a key in the session cache. |
| Total SSL session cache timeouts | The number of times that keys were removed from the session cache because the timeout period expired. |

CRL Statistics

Help | Content Gateway | Version 7.8.x

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 163. |
| | 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. |

Reports

Help | Content Gateway | Version 7.8.x

See [Creating SSL-related reports](#), page 124 for information on creating reports on certificate authorities or incidents.

B

Commands and Variables

Help | Content Gateway | Version 7.8.x

WebSense Content Gateway commands

Use the command line to execute individual commands and when scripting multiple commands in a shell.

Run commands as 'root'.

Execute Content Gateway commands from the Content Gateway **bin** directory.



Note

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

`./`

For example:

```
./content_line -p
```

| Command | Description |
|--|--|
| <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 Cop, Content Gateway, Content Gateway Manager, and Analytics Server. |
| <code>WCGAdmin help</code> | Displays a list of the WCGAdmin commands |
| <code>content_line -h</code> | Displays the list of Content Gateway commands. |
| <code>content_line -p</code> <code>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 294 . |
| <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 420 , for a list of the configuration variables you can specify. |
| <code>content_line -x</code> | Initiates a Content Gateway configuration file reread. Executing this command is similar to clicking Apply in the Content Gateway manager. |
| <code>content_line -y</code> | Clears Websense dynamically signed certificates from the cache and the SSL sqlite database. |
| <code>content_line db_clear -y</code> | Clears Websense dynamically signed certificates from the SSL sqlite database. |
| <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 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 on the local node. |
| <code>content_line -W</code> | Enables WCCP router communication. |
| <code>content_line -w</code> | Disables WCCP router communication. After changing the Content Gateway WCCP configuration, or the router WCCP configuration, force WCCP communication down for 60 seconds to force WCCP to negotiate a new connection. |
| <code>content_line -N snapshot_name</code> | Perform a Content Gateway snapshot (backup). See Taking configuration snapshots, page 115 . |
| <code>content_line -n snapshot_name</code> | Restore a Content Gateway snapshot. See Restoring configuration snapshots, page 115 . |

Websense Content Gateway variables

Help | Content Gateway | Version 7.8.x

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 420.

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 121, and [Command-line interface](#), page 112.

Statistics

Help | Content Gateway | Version 7.8.x

The following table lists the variables you can specify on the command line to view individual statistics. See [Statistics](#), page 267 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> |

| Statistic | Variable |
|--|---|
| 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 response header bytes to client from cache | <i>proxy.process.http.user_agent_response_from_cache_header_total_size</i> |
| Total response document bytes to client from cache | <i>proxy.process.http.user_agent_response_from_cache_document_total_size</i> |
| Total connections to client | <i>proxy.process.http.current_client_connections</i> |
| Current unique clients connected | <i>proxy.process.http.client.unique_clients.active</i> |
| Total unique clients that have connected | <i>proxy.process.http.client.unique_clients.total</i> |
| Total clients that exceeded limit | <i>proxy.process.http.client.exceeding_limit</i> |
| Total clients for which connections were closed | <i>proxy.process.http.client.closed_connections</i> |
| Open HTTP client connections | <i>proxy.process.http.current_active_http_client_connections</i> |
| Open HTTPS client connections | <i>proxy.node.process.http.current_active_https_client_connections</i> |
| Client Requests (IPv4 +IPv6) | <i>proxy.process.http.real_client_requests</i> |
| Client IPv6 Requests | <i>proxy.process.http.real_client_ipv6_requests</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> |
| 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> |

| Statistic | Variable |
|--------------------------------|--|
| | 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> |
| | WCCP |
| Enabled | <i>proxy.config.wccp.enabled</i> |
| WCCP interface | <i>proxy.local.wccp2.ethernet_interface</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> |
| 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> |

| Statistic | Variable |
|---------------------------------|---|
| 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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

Options are grouped as follows on the left side of the Configure pane:

[My Proxy](#), page 299

[Protocols](#), page 311

[Content Routing](#), page 327

[Security](#), page 332

[Subsystems](#), page 354

[Networking](#), page 360

My Proxy

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The My Proxy options are:

[Basic](#), page 300

[Subscription](#), page 304

[UI Setup](#), page 305

[Snapshots](#), page 308

[Logs](#), page 310

Basic

Help | Content Gateway | Version 7.8.x

Configure > My Proxy > Basic > General

| | |
|------------------------|--|
| Restart | <p>Restarts the proxy and manager services (the content_gateway and content_manager processes). You must restart the proxy and manager services after modifying certain configuration options. A message is displayed in the manager when a restart is required.</p> <p>IMPORTANT: In a cluster configuration, the Restart button restarts the proxy and manager services on all nodes in the cluster.</p> |
| Proxy Name | <p>Specifies the name of your Content Gateway node. By default, this is the hostname of the machine running Content Gateway.</p> <p>If this node is part of a cluster, this option specifies the name of the Content Gateway cluster. In a cluster, all nodes must share the same name.</p> |
| Alarm email | Specifies the email address to which Content Gateway sends alarm notifications. |
| Features | |
| Protocols: FTP | <p>When this option is enabled, Content Gateway accepts FTP requests from FTP clients.</p> <p>If this option is changed you must restart Content Gateway.</p> |
| Protocols: HTTPS | <p>Enables/disables Content Gateway HTTPS traffic management and security analysis. After selecting HTTPS On, you must provide additional information on the Configure > Protocols > HTTPS page and on the Configure > SSL pages. See <i>Working With Encrypted Data</i>, page 139.</p> |
| Networking: WCCP | <p>Enable this option to use a WCCP v2-enabled router for transparent redirection to Content Gateway. WCCP v1 is not supported.</p> <p>See <i>Transparent interception with WCCP v2 devices</i>, page 50.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Networking: DNS Proxy | <p>When this option is enabled, Content Gateway resolves DNS requests on behalf of clients. This option offloads remote DNS servers and reduces response time for DNS lookups. See <i>DNS Proxy Caching</i>, page 107.</p> |
| Networking: Virtual IP | <p>When this option is enabled, Content Gateway maintains a pool of virtual IP addresses that it assigns to the nodes in a cluster as necessary. See <i>Virtual IP failover</i>, page 91.</p> |

| | |
|---|---|
| Networking: IPv6 (Prior to 7.8.4, explicit proxy only) | <p>When this option is enabled, Content Gateway provides support for IPv6.</p> <p>Prior to 7.8.3, support is provided for explicit proxy only.</p> <p>IPv6 addresses can be used on any dual stack Ethernet interface that services client and/or Internet traffic.</p> <p>IPv4 addresses must be used to communicate with all TRITON components.</p> <p>To see a complete description of the feature and an important list of restrictions, see Support for IPv6, page 82.</p> |
| Networking: Data Security | <p>Enables a connection to Websense Data Security. There are 2 options:</p> <ul style="list-style-type: none"> • Automatic registration with the Data Security Management Server (must be version 7.8.x) • ICAP communication to a remote Data Security Suite deployment (may be version 7.1, or earlier) <p>See Working With Websense Data Security, page 129.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Networking: Data Security: Integrated on-box | <p>Enables registration with the on-box Data Security components and the Data Security Management Server. See Registering and configuring Data Security, page 132.</p> |
| Networking: Data Security: ICAP | <p>Enables ICAP for use with Data Security Suite. See Configuring the ICAP client, page 135.</p> |
| Security: SOCKS | <p>When SOCKS is enabled, Content Gateway communicates with your SOCKS servers. See Configuring SOCKS firewall integration, page 185.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Authentication: None | <p>Content Gateway supports several types of user authentication.</p> <p>When this option is selected, the proxy does not perform user authentication. This is the default setting.</p> |
| Authentication: Integrated Windows Authentication | <p>When Integrated Windows Authentication (IWA) is enabled, users are authenticated by IWA before they are allowed access to content.</p> <p>See Integrated Windows Authentication, page 197.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Authentication: LDAP | <p>When LDAP is enabled, users are authenticated by an LDAP server before they are allowed access to content. See LDAP authentication, page 204.</p> <p>If you change this option, you must restart Content Gateway.</p> |

| | |
|--|--|
| Authentication: Radius | <p>When RADIUS is enabled, users are authenticated by a RADIUS server before they are allowed access to content. See RADIUS authentication, page 208.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Authentication: Legacy NTLM | <p>When legacy NTLM (NTLMSSP) is enabled, users in a Windows network are authenticated by a Domain Controller before they are allowed access to content. See Legacy NTLM authentication, page 202.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Authentication: Rule-Based Authentication | <p>When Rule-Based Authentication is enabled, users are authenticated based on the parameters of the rule that they match. Rule-based authentication supports multiple realm, multiple domain, and other user authentication scenarios. See Rule-Based Authentication, page 211.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| 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> |
| Authentication: Send authentication to parent proxy | <p>Enables or disables the insertion of X-Authenticated-User header values in outgoing requests. This option is disabled by default.</p> <p>Enable this option when Content Gateway is the child (downstream) proxy in a chain and the parent (upstream) proxy wants X-Authenticated-User values to facilitate authentication.</p> |

Configure > My Proxy > Basic > 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 activate management clustering mode. The nodes in the cluster share configuration information and you can administer all the nodes at the same time.</p> <p>For complete information about clustering, see Clusters, page 85.</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, eth1.</p> <p>It is recommended that you use a dedicated secondary interface.</p> <p>Node configuration information is multicast, in plain text, to other Content Gateway nodes on the same subnet. Therefore, Websense recommends that clients be located on a separate subnet from Content Gateway nodes (multicast communications for clustering are not routed).</p> <p>On V-Series appliances, P1 (eth0) is the recommended interface. However, you may also use P2 (eth1) if you want to isolate cluster management traffic.</p> <p>See Changing clustering configuration, page 86.</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>See Changing clustering configuration, page 86.</p> |

Subscription

Help | Content Gateway | Version 7.8.x

Configure > My Proxy > Subscription > Subscription Management

| | |
|------------------|--|
| Subscription Key | <p>Displays the subscription key you received from Websense, Inc.</p> <p>If Content Gateway is used with Web Security Gateway or Web Security Gateway Anywhere, this is the subscription key you entered in the Web Security manager.</p> <p>If Content Gateway is deployed with only Websense Data Security Suite, you must enter your Content Gateway subscription key in this field.</p> |
|------------------|--|

Configure > My Proxy > Subscription > Scanning

| | |
|--|--|
| Policy Server | |
| IP address | The IP address of the Websense Web Security Policy Server. This value is specified when Content Gateway is installed. |
| Port | The port used by the Websense Web Security Policy Server. The default port is 55806. |
| Filtering Service | |
| IP address | Specify the IP address of the Websense Web Security Filtering Service. This value is specified when Content Gateway is installed. |
| Port | Specify the port used by the Websense Web Security Filtering Service. The default port is 15868. |
| Communication Timeout | <p>Specifies the timeout, in milliseconds, in which Policy Server and Filtering Service must respond before a communication timeout condition occurs and the Action for Communication Errors setting is applied.</p> <p>The default value is 5000 ms (5 seconds).</p> |
| Action for Communication Errors | |
| Permit traffic | Permits all traffic if communication with Policy Server or Filtering Service fails. |
| Block traffic | Blocks all traffic if communication with Policy Server or Filtering Service fails. |

| Scanning Data Files Update (7.8.3 and later) | |
|---|--|
| Delay time | Specifies the length of time scanning data file downloads are delayed. The default value is No delay. See the Scanning Data Files Update section of Providing system information . |

UI Setup

Help | Content Gateway | Version 7.8.x

Configure > My Proxy > UI Setup > General

| | |
|-------------------------|---|
| UI Port | Specifies the port on which browsers can connect to the Content Gateway manager. The default port is 8081. If you change this setting, you must restart Content Gateway. |
| HTTPS: Enable/Disable | Enables or disables support for SSL connections to the Content Gateway manager (enabled by default). SSL provides protection for remote administrative monitoring and configuration. To use SSL for Content Gateway manager connections, you must install an SSL certificate on the Content Gateway server machine. For more information, see Using SSL for secure administration , page 178. |
| HTTPS: Certificate File | Specifies the name of the SSL certificate file used to authenticate users who want to access the Content Gateway manager. |
| Monitor Refresh Rate | Specifies how often Content Gateway manager refreshes the statistics on the Monitor pane. The default value is 30 seconds. |
| Default Help Language | Specifies the language that Content Gateway Manager Help displays by default. If a page is not available in the default language, another language may be substituted. |

Configure > My Proxy > UI Setup > Login

| | |
|----------------------|--|
| Basic Authentication | <p>Enables/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 the Content Gateway manager.</p> |
| Administrator: Login | <p>Specifies the administrator login. The default is 'admin'.</p> <p>The administrator login is the master login that has access to both Configure and Monitor mode in the Content Gateway manager.</p> <p>Note: Content Gateway checks the administrator login only if the Basic Authentication option is enabled.</p> |

| | |
|-------------------------|---|
| Administrator: Password | <p>Lets you change the administrator password that controls access to the Content Gateway manager.</p> <p>Enter the current password in the Old Password field. Enter the new password in the New Password field, re-enter it in the New Password (Retype) field, and then click Apply.</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 the Content Gateway manager, Content Gateway updates the records.config file. If you forget the administrator password and cannot access the Content Gateway manager, see Accessing the Content Gateway manager if you forget the master administrator password, page 15.</p> |
| Additional Users | <p>Lists the current user accounts and lets you add new user accounts. User accounts determine who has access the Content Gateway manager and which activities they can perform. You can create a list of user accounts if a single administrator login and password is not sufficient security for your needs.</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 177.</p> <p>Note: Content Gateway checks the user login and password only if the Basic Authentication option is enabled.</p> |

Configure > My Proxy > UI Setup > Access

| | |
|----------------|---|
| Access Control | <p>Displays a table listing the rules in the mgmt_allow.config file. Rules specify the remote hosts allowed to access the Content Gateway manager. The entries in this file ensure that only authenticated users can change configuration options and view performance and network traffic statistics.</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> |

| | |
|------------------|---|
| Edit File | Opens the configuration file editor so that you can edit and add rules to the mgmt_allow.config file. |
| | mgmt_allow.config Configuration File Editor |
| rule display box | Lists the mgmt_allow.config file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. Content Gateway applies the rules in the order listed, starting from the top. |
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| IP Action | Lists the type of rules you can add. An ip_allow rule allows the remote hosts specified in the Source IP field to access the Content Gateway manager. An ip_deny rule denies the remote hosts specified in the Source IP field access to the Content Gateway manager. |
| Source IP | Specifies the IP addresses that are allowed or denied access to the Content Gateway manager. You can enter a single IP address (111.111.11.1) or a range of IP addresses (0.0.0.0-255.255.255.255). |
| 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

Help | Content Gateway | Version 7.8.x

Configure > My Proxy > Snapshots > 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. |

| | |
|---|--|
| Snapshots: Restore/Delete Snapshot | Lists the snapshots that are stored on this node. Select the snapshot that you want to restore or delete from the drop-down list. |
| Snapshots: Restore Snapshot from "directory_name" Directory | Restores the snapshot selected in the Restore/Delete Snapshot drop-down box. In a cluster configuration, snapshots are restored on all nodes in the cluster. |
| Snapshots: Delete Snapshot from "directory_name" Directory | Deletes the snapshot selected in the Restore/Delete Snapshot drop-down box. |

Configure > My Proxy > Snapshots > FTP server

| | |
|-----------------------------|--|
| FTP Server | Specifies the name of the FTP server from which you want to restore a configuration snapshot or to which you want to save a configuration snapshot. |
| Login | Specifies the login needed to access the FTP server. |
| Password | Specifies the password needed to access the FTP server. |
| Remote Directory | Specifies the directory on the FTP server from which you want restore, or in which you want to save a configuration snapshot. |
| Restore Snapshot | Lists the configuration snapshots on the FTP server that you can restore. This field appears after you have logged on to the FTP server successfully. |
| Save Snapshot to FTP Server | Specifies the name of the configuration snapshot you want to take and save on the FTP server. This field appears after you have logged on to the FTP server successfully. |

Logs

Help | Content Gateway | Version 7.8.x

Configure > My Proxy > Logs > System

| | |
|---|--|
| Log File | Lists the system log files you can view, delete or copy to your local system. Content Gateway lists the system log files logged with the system-wide logging facility syslog under the daemon facility. |
| Action: Display the selected log file | When this option is enabled, Content Gateway displays the first MB of the system log file selected in the Log File drop-down list. To view the entire file, select “Save the selected log file in local filesystem” and view the file with a local viewer. |
| Action: Display last lines of the selected file | When this option is enabled, Content Gateway displays the last specified number of lines in the selected system log file. |
| Action: Display lines that match in the selected log file | When this option is enabled, Content Gateway displays all the lines in the selected system log file that match the specified string. |
| Action: Remove the selected log file | When this option is enabled, Content Gateway deletes the selected log file. |
| Action: Save the selected log file in local filesystem | When this option is enabled, Content Gateway saves the selected log file on the local system in a location you specify. |

Configure > My Proxy > Logs > Access

| | |
|---------------------------------------|---|
| Log File | Lists the event or error log files you can view, delete, or copy to your local system. Content Gateway lists the event log files located in the directory specified in the Logging Directory field under Subsystems/Logging and by the configuration variable proxy.config.log2.logfile_dir in the records.config file. The default directory is logs in the Content Gateway installation directory. |
| Action: Display the selected log file | When this option is enabled, Content Gateway displays the first MB of the event or error log file selected in the Log File drop-down list. To view the entire file, select “Save the selected log file in local filesystem” and view the file with a local viewer. |

| | |
|---|---|
| Action: Display last lines of the selected file | When this option is enabled, Content Gateway displays the last specified number of lines in the event or error log file selected from the Log File drop-down list. |
| Action: Display lines that match in the selected log file | When this option is enabled, Content Gateway displays all the lines in the selected event or error log file that match the specified string. |
| Remove the selected log file | When this option is enabled, Content Gateway deletes the selected log file. |
| Action: Save the selected log file in local filesystem | When this option is enabled, Content Gateway saves the selected log file on the local system in a location you specify. |

Protocols

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The Protocol configuration options are divided into the following categories:

[HTTP](#), page 311

[HTTP Responses](#), page 322

[HTTP Scheduled Update](#), page 323

[HTTPS](#), page 325

[FTP](#), page 326

HTTP

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

Configure > Protocols > HTTP > General

| | |
|-----------------------------------|--|
| HTTP Proxy Server Port | Specifies the port that Content Gateway uses when acting as a Web proxy server for HTTP traffic or when serving HTTP requests transparently. The default port is 8080. If you change this option, you must restart Content Gateway. |
| Secondary HTTP Proxy Server Ports | For explicit proxy configurations only , specifies additional ports on which Content Gateway listens for HTTP traffic. Transparent proxy configurations always send all HTTP traffic to port 8080. |

| | |
|-----------------------------------|--|
| Unqualified Domain Name Expansion | <p>Enables or disables .com name expansion. When this option is enabled, Content Gateway attempts to resolve unqualified hostnames by redirecting them to the expanded address, prepended with www. and appended with .com. For example, if a client makes a request to <i>company</i>, Content Gateway redirects the request to www.company.com.</p> <p>If local domain expansion is enabled (see DNS Resolver, page 373), 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 on which Content Gateway allows tunneling. This is a space separated list that also accepts port ranges (e.g. 1-65535).</p> <p>When SSL is not enabled, all traffic destined for the specified ports is allowed to tunnel to an origin server.</p> <p>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> |
| HTTPS ports | <p>When SSL support is enabled, specifies ports on which HTTPS traffic is decrypted and policy is applied. Note that Content Gateway receives HTTPS traffic on the port specified in Configure > Protocols > HTTPS > HTTPS Proxy: Server Port.</p> <p>When SSL support is disabled, traffic to these ports is not decrypted. However, filtering policy is applied based on:</p> <ul style="list-style-type: none"> • Explicit proxy: the server hostname in the CONNECT request. • Transparent proxy: the SNI hostname or the server hostname in the server's certificate. If the hostname in the server's certificate includes a wildcard (*), the lookup is performed on the destination IP address. |

| | |
|-----------------------------------|--|
| Unqualified Domain Name Expansion | <p>Enables or disables .com name expansion. When this option is enabled, Content Gateway attempts to resolve unqualified hostnames by redirecting them to the expanded address, prepended with www. and appended with .com. For example, if a client makes a request to <i>company</i>, Content Gateway redirects the request to www.company.com.</p> <p>If local domain expansion is enabled (see DNS Resolver, page 373), 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 on which Content Gateway allows tunneling. This is a space separated list that also accepts port ranges (e.g. 1-65535).</p> <p>When SSL is not enabled, all traffic destined for the specified ports is allowed to tunnel to an origin server.</p> <p>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> |
| HTTPS ports | <p>When SSL support is enabled, specifies ports on which HTTPS traffic is decrypted and policy is applied. Note that Content Gateway receives HTTPS traffic on the port specified in Configure > Protocols > HTTPS > HTTPS Proxy: Server Port.</p> <p>When SSL support is disabled, traffic to these ports is not decrypted. However, filtering policy is applied based on:</p> <ul style="list-style-type: none"> • Explicit proxy: the server hostname in the CONNECT request. • Transparent proxy: the SNI hostname or the server hostname in the server's certificate. If the hostname in the server's certificate includes a wildcard (*), the lookup is performed on the destination IP address. |

| | |
|--|--|
| FTP over HTTP: Anonymous Password | Specifies the anonymous password Content Gateway must use for FTP server connections that require a password. This option affects FTP requests from HTTP clients. |
| FTP over HTTP: Data Connection Mode | <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> |

Configure > Protocols > HTTP > Cacheability

| | |
|-----------------------------------|---|
| Caching: HTTP Caching | <p>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.</p> <p>Note: HTTPS content is never cached.</p> |
| Caching: FTP over HTTP Caching | <p>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.</p> |

| | |
|------------------------------|---|
| 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 Expires, max-age, or last-modified 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 max-age 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 Expires or Cache-Control headers; Content Gateway considers all HTTP objects without Expires or Cache-Control 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 cache.config file. This is the default option.</p> <p>For more information about revalidation, see Revalidating HTTP objects, page 25.</p> |

| | |
|--|--|
| Behavior: Add “no-cache” to MSIE Requests | <p>Specifies when Content Gateway adds no-cache 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 no-cache headers to all requests from Microsoft Internet Explorer.</p> <p>Select To IMS MSIE Requests to add no-cache headers to IMS (If Modified Since) Microsoft Internet Explorer requests.</p> <p>Select Not to Any MSIE Requests to never add no-cache headers to requests from Microsoft Internet Explorer.</p> |
| Behavior: Ignore “no-cache” in Client Requests | <p>When this option is enabled, Content Gateway ignores no-cache 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 no-cache 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. |

| | |
|--|---|
| 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. |

| | |
|-----------------------------------|--|
| 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> |
| Secondary Specifiers: Port | <p>Specifies the port in a requested URL.</p> |

| | |
|-------------------------------------|---|
| 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. |

Configure > Protocols > HTTP > Privacy

| | |
|------------------------------------|--|
| Insert Headers: Client-IP | <p>When enabled, Content Gateway inserts the Client-IP header into outgoing requests to retain the client's IP address.</p> <p>This option is mutually exclusive with the Remove Headers: Client-IP option. When Insert Headers: Client-IP is enabled the Remove Headers: Client-IP option is automatically disabled.</p> <p>Insert Headers: Client-IP and Remove Headers: Client-IP can both be disabled.</p> |
| Insert Headers: Via | <p>When enabled, Content Gateway inserts a Via header into the outgoing request. The Via header informs the destination server of proxies through which the request was sent.</p> |
| Insert Headers: X-Forwarded-For | <p>When enabled, Content Gateway inserts an X-Forwarded-For header into the outgoing request. The X-Forwarded-For value contains the originating IP address.</p> |
| Remove Headers: Client-IP | <p>When this option is enabled, Content Gateway removes the Client-IP header from outgoing requests to protect the privacy of your users.</p> <p>This option is mutually exclusive with the Insert Headers: Client-IP option. When Remove Headers: Client-IP is enabled the Insert Headers: Client-IP option is automatically disabled.</p> <p>Remove Headers: Client-IP and Insert Headers: Client-IP can both be disabled.</p> |
| Remove Headers: Cookie | <p>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.</p> |
| Remove Headers: From | <p>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.</p> |

| | |
|-------------------------------|---|
| 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. |

Configure > Protocols > HTTP > Timeouts

See [this knowledge base article](#) for a discussion of HTTP timeout options.

| | |
|------------------------------------|---|
| 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> |
| Active Timeouts: Origin Server Request | <p>Specifies how long Content Gateway waits for fulfillment of a connection request to an origin server. If Content Gateway does not establish connection to an origin server before the timeout expires, Content Gateway terminates the connection request.</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

Help | Content Gateway | Version 7.8.x

Configure > Protocols > HTTP Responses > 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> |
|---------------------------|--|

Configure > Protocols > HTTP Responses > 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> |

Incorporating images, animated gifs, and Java applets on the response page

Content Gateway can respond to clients with only a single text or HTML document.

However, you can provide references on your custom response pages to images, animated gifs, Java applets, or objects other than text that are located on a Web server.

Add links in the **body_factory** template files in the same way you would for any image in an HTML document, with the full URL in the SRC attribute.

It is recommended that you do not run the Web server and Content Gateway on the same system, to prevent both programs from trying to serve documents on the same port number.

HTTP Scheduled Update

Help | Content Gateway | Version 7.8.x

Configure > Protocols > HTTP Scheduled Updates > General

| | |
|---------------------------------|---|
| Scheduled Update | 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. |
| Maximum Concurrent Updates | 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. |
| Retry on Update Error: Count | 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. |
| Retry on Update Error: Interval | Specifies the delay in seconds between each scheduled update retry for a URL in the event of failure. The default value is 2 seconds. |

Configure > Protocols > HTTP Scheduled Updates > 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 GET 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

Help | Content Gateway | Version 7.8.x

Configure > Protocols > HTTPS

This page is displayed only when HTTPS is enabled on **Configure > My Proxy > Basic > General**

| | |
|-------------------------|--|
| HTTPS Proxy Server Port | Specifies the port that Content Gateway uses when acting as a Web proxy server for HTTPS traffic. The default value is 8080. See also, Configure > Protocols > HTTP > General: HTTPS Ports . |
|-------------------------|--|

| | |
|--------------------------|--|
| Tunnel Skype | <p>Enables/disables the tunneling of Skype traffic when HTTPS is enabled and Content Gateway is an explicit proxy.</p> <p>To complete the configuration, you must ensure that all users who are allowed to use Skype have a Web Security Filtering policy that permits internet telephony. This is required regardless of whether Skype is used with HTTPS enabled or not.</p> <p>Also, if Skype is not prevented, after the handshake it will route traffic over a non-HTTP port. To force Skype traffic to go through Content Gateway, a GPO should be used, as described in the Skype IT Administrators Guide.</p> <p>Note: This option is not necessary if HTTPS is not enabled.</p> <p>Note: This option is not valid when Content Gateway is a transparent proxy.</p> |
| Tunnel Unknown Protocols | <p>Enables and disables tunneling of HTTPS requests when the SSL handshake results in an unknown protocol error.</p> <p>Tunneled connections are not decrypted or inspected.</p> <p>When Content Gateway is an explicit proxy, a URL lookup is performed and policy is applied before the SSL connection request is made with the server. Therefore, tunneled transactions appear in the Web Security transaction log.</p> <p>When Content Gateway is a transparent proxy, if there is an SNI a URL lookup is done on the hostname in the SNI. Otherwise no URL lookup is possible and tunneled transactions are not logged. This is because an initial connection with the server is required to get the Common Name from the SSL certificate. It is used for the URL lookup. If the connection handshake fails and this option is enabled, the connection is tunneled without the proxy being aware of it.</p> <p>Important: This setting persists after the HTTPS feature is disabled (on Configure > My Proxy > Basic > General). Therefore, disable this option before disabling HTTPS support.</p> |

FTP

Help | Content Gateway | Version 7.8.x



Note

The FTP configuration options appear on the Configure pane only if you have enabled FTP processing in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Protocols > FTP > General

| | |
|--------------------------------|--|
| FTP Proxy Server Port | Specifies the port that Content Gateway uses to accept FTP requests. The default port is 2121. |
| Listening Port Configuration | <p>Specifies how FTP opens a listening port for a data transfer.</p> <p>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.</p> <p>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.</p> |
| Default Data Connection Method | <p>Specifies the default method used to set up data connections with the FTP server.</p> <p>Select Proxy Sends PASV to send a PASV to the FTP server and let the FTP server open a listening port.</p> <p>Select Proxy Sends PORT to set up a listening port on the Content Gateway side of the connection first.</p> |
| Shared Server Connections | When enabled, server control connections can be shared between multiple anonymous FTP clients. |

Configure > Protocols > FTP > Timeouts

| | |
|-------------------------------------|---|
| Keep-Alive Timeout: Server Control | Specifies the timeout value when the FTP server control connection is not used by any FTP clients. The default value is 90 seconds. |
| Inactivity Timeouts: Client Control | Specifies how long FTP client control connections can remain idle. The default value is 900 seconds. |
| Inactivity Timeouts: Server Control | Specifies how long the FTP server control connection can remain idle. The default value is 120 seconds. |
| Active Timeouts: Client Control | Specifies the how long FTP client control connections can remain open. The default value is 14400 seconds. |
| Active Timeouts: Server Control | Specifies how long the FTP server control connection can remain open. The default value is 14400 seconds. |

Content Routing

Help | Content Gateway | Version 7.8.x

The Content Routing configuration options are divided into the following categories:

[Hierarchies](#), page 328

[Mapping and Redirection](#), page 330

[Browser Auto-Config](#), page 332

Hierarchies

Help | Content Gateway | Version 7.8.x

Configure > Content Routing > Hierarchies

| | |
|------------------------------------|--|
| Parent Proxy | Enables or disables the HTTP parent caching option. When this option is enabled, Content Gateway can participate in an HTTP cache hierarchy. You can point your Content Gateway server at a parent network cache (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 95. |
| No DNS and Just Forward to Parent | When enabled, and if HTTP parent caching is enabled, Content Gateway does no DNS lookups on requested hostnames. If rules in the parent.config file are set so that only selected requests are sent to a parent proxy, Content Gateway skips name resolution only for requests that are going to the parent proxy. Name resolution is performed as usual for requests that are not sent to a parent proxy. If the parent proxy is down and the child proxy can go directly to origin servers, the child performs DNS resolution. |
| Uncacheable Requests Bypass Parent | When enabled, and if parent caching is enabled, Content Gateway bypasses the parent proxy for uncacheable requests. |
| HTTPS Requests Bypass Parent | When enabled, Content Gateway bypasses the parent proxy for HTTPS requests. |
| Tunnel Requests Bypass Parent | When enabled, Content Gateway bypasses parent proxy for non-HTTPS tunnel requests. |
| Parent Proxy Cache Rules | Displays a table listing the rules in the parent.config file that identify the HTTP parent proxies used in an HTTP cache hierarchy and configure selected URL requests to bypass parent proxies. Rules are applied from the list top-down; the first match is applied. |
| Refresh | Updates the table to display the most up-to-date rules in the parent.config file. |
| Edit File | Opens the configuration file editor so that you can edit and add rules to the parent.config file. |
| | parent.config Configuration File Editor |
| rule display box | Lists the 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. |

| | |
|---------------------------------|---|
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| Primary Destination Type | Lists the primary destination types: dest_domain is a requested domain name. dest_host is a requested hostname. dest_ip is a requested IP address. url_regex is a regular expression to be found in a URL. |
| Primary Destination Value | Specifies the value of the primary destination type. For example: If the primary destination is dest_domain , a value for this field can be yahoo.com If the primary destination type is dest_ip , the value for this field can be 123.456.78.9. If the primary destination is url_regex , a value for this field can be politics. |
| Parent Proxies | Specifies the IP addresses or hostnames of the parent proxies and the port numbers used for communication. Parent proxies are queried in the order specified in the list. If the request cannot be handled by the last parent server in the list, it is routed to the origin server. Separate each entry with a semicolon; for example: parent1:8080; parent2:8080 |
| Round Robin | Select true for the proxy to go through the parent cache list in a round-robin based on client IP address. Select strict for the proxy to serve requests strictly in turn. For example, machine proxy1 serves the first request, proxy2 serves the second request, and so on. Select false if you do not want round-robin selection to occur. |
| Go direct | Select true for requests to bypass parent hierarchies and go directly to the origin server. Select false if you do not want requests to bypass parent hierarchies. |
| Secondary Specifiers: Time | Specifies a time range, using a 24-hour clock, such as 08:00-14:00. If the range crosses midnight, enter this as two comma-separated ranges. For example, if a range extends from 6:00 in the evening until 8:00 in the morning, enter the following: 18:00 - 23:59, 0:00 - 8:00 |
| Secondary Specifiers: Prefix | Specifies a prefix in the path part of a URL. |
| Secondary Specifiers: Suffix | Specifies a file suffix in the URL, such as .htm or .gif. |
| Secondary Specifiers: Source IP | Specifies the IP address or range of IP addresses of the clients. |

| | |
|-------------------------------------|---|
| Secondary Specifiers: Port | Specifies the port in a requested URL. |
| Secondary Specifiers: Method | Specifies a request URL method. For example: get post put trace |
| Secondary Specifiers: Scheme | Specifies the protocol of a requested URL. This must be either HTTP or FTP. |
| Secondary Specifiers: User-Agent | Specifies a request header User-Agent value. |

Mapping and Redirection

Help | Content Gateway | Version 7.8.x

Configure > Content Routing > Mapping and Redirection

| | |
|--------------------------------|---|
| Serve Mapped Hosts Only | Select Required if you want the proxy to serve requests only to origin servers listed in the mapping rules of the remap.config file. If a request does not match a rule in the remap.config file, the browser receives an error. This option provides added security for your Content Gateway system. |
| Retain Client Host Header | When this option is enabled, Content Gateway retains the client host header in a request (it does not include the client host header in the mapping translation). |
| Redirect No-Host Header to URL | Specifies the alternate URL to which to direct incoming requests from older clients that do not provide a <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 Host: headers to a particular server. |
| URL Remapping Rules | Displays a table listing the mapping rules in the remap.config file so that you can redirect HTTP requests permanently or temporarily without the proxy having to contact any origin servers. Note: Mapping a URL to another URL in the same domain requires that a “/” be specified in From Path Prefix field. See the example following this table. |
| Refresh | Updates the table to display the most up-to-date rules in the remap.config file. |
| Edit File | Opens the configuration file editor so that you can edit and add rules to the remap.config file. |

| | remap.config Configuration File Editor |
|-----------------------------|---|
| rule display box | Lists the remap.config file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a 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 remap.config 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). map and reverse_map are not supported. |
| From Scheme | Specifies the protocol of the mapping rule. rtsp and mms are not supported. Note: Mapping a URL of one protocol (scheme) to a different protocol (scheme) 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. Sometimes it is desirable to redirect a URL to a sub-page in the same domain. For example, to redirect “www.cnn.com” to “www.cnn.com/tech”. To make this rule work you must specify “/” in the From Path Prefix field. If it is not specified, the redirection results in a URL that recursively adds the page specifier to the URL. For example, “www.example.com/tech” becomes “www.example.com/tech/tech/tech/tech/tech/tech/tech/tech/...” |
| From Query (Optional) | Specifies the query of the URL to map from. |
| To Scheme | Must match From Scheme . |
| 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. |
| To Query (Optional) | Specifies the query of the URL to map to. |
| {undefined} | Specifies the media protocol type of the mapping rule. Not supported. |

Browser Auto-Config

Help | Content Gateway | Version 7.8.x

Configure > Content Routing > Browser Auto-Config > PAC

| | |
|-------------------------|--|
| Auto-Configuration Port | Specifies the port Content Gateway uses to download the auto-configuration file to browsers. The port cannot be assigned to any other process. The default port is 8083. If you change this option, you must restart Content Gateway. |
| PAC Settings | Lets you edit the PAC file (proxy.pac). See Using a PAC file, page 40 . |

Configure > Content Routing > Browser Auto-Config > WPAD

| | |
|---------------|---|
| WPAD Settings | Lets you edit the wpad.dat file. See Using WPAD, page 42 . |
|---------------|---|

Security

Help | Content Gateway | Version 7.8.x

The Security configuration options are divided into the following categories:

[Connection Control, page 333](#)

[FIPS Security, page 333](#)

[Data Security, page 334](#)

[Access Control, page 335](#)

[SOCKS, page 351](#)

Connection Control

Help | Content Gateway | Version 7.8.x

Configure > Security > Connection Control

| Option | Description |
|------------------|---|
| | Proxy Access |
| Access Control | Displays the rules in the <i>ip_allow.config</i> file that control which clients can access Content Gateway. By default, all remote hosts are allowed to access the proxy. |
| Refresh | Updates the table to display the most up-to-date rules in the ip_allow.config file. |
| Edit File | Opens the configuration file editor for to the ip_allow.config file. |
| | ip_allow.config Configuration File Editor |
| rule display box | Lists the <i>ip_allow.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| IP Action | Lists the type of rules you can add. An ip_allow rule allows the clients listed in the Source IP field to access the proxy. An ip_deny rule denies the clients listed in the Source IP field access to the proxy. |
| Source IP | Specifies the IP address or range of IP addresses of the clients. |
| Apply | Applies the configuration changes. |
| Close | Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost. |

FIPS Security

Help | Content Gateway | Version 7.8.x

Configure > Security > FIPS

This option pertains to HTTPS traffic and the FIPS 140-2 cryptography standard.

By default, when handling HTTPS traffic, Content Gateway accepts connections from sites using SSLv3. SSLv3 is not FIPS 140-2 compliant.

When FIPS mode is enabled:

- ◆ HTTPS connections use only TLSv1 or higher
- ◆ HTTPS connections use FIPS 140-2 approved algorithms
- ◆ Content Gateway generates SHA-256 certificates in response to origin server certificate requests



Warning

Once enabled, FIPS 140-2 mode cannot be disabled without reinstalling Content Gateway. If Content Gateway is on an appliance, the appliance must be reimaged.



Important

Due to a system limitation, FIPS 140-2 mode cannot be used with NTLM user authentication (IWA fallback to NTLM or Legacy NTLM).

For complete information, see [FIPS 140-2 Mode, page 179](#).

| Option | Description |
|-----------------------------------|---|
| FIPS Enable/Disable radio buttons | <p>By default, Content Gateway is installed in non-FIPS 140-2 mode.</p> <p>To switch to FIPS 140-2 mode, select the Enabled radio button, click Apply, and restart Content Gateway.</p> <p>Warning: Once enabled, FIPS 140-2 mode cannot be disabled without reinstalling Content Gateway. For Websense appliance installations, reinstallation requires reimaging the system.</p> |

Data Security

Help | Content Gateway | Version 7.8.x



Note

The Data Security configuration options appear on the Configure menu only if you have enabled Data Security on the **Configure > My Proxy > Basic > General** tab and selected **Integrated on-box** in the **Features** table.

Configure > Security > Data Security

| 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 | Specifies whether decrypted traffic should be sent to Websense Data Security for analysis or sent directly to the destination. |
| Analyze FTP Uploads | Specifies whether to send FTP upload requests to Websense Data Security for analysis. The FTP proxy feature must be enabled. See FTP , page 326. |

Registration screen fields:

| Option | Description |
|---------------------------------------|--|
| Data Security Management server IP | Specifies the IP address of the Data Security Management Server. This is where data security policy configuration and management is performed. |
| Data Security administrator user name | Specifies the account name of a Websense Data Security administrator. The administrator must have Deploy Settings privileges. |
| Data Security administrator password | 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 of the fields. |

Access Control

Help | Content Gateway | Version 7.8.x

Use the **Access Control** tabs to:

- ◆ Create custom filtering rules
- ◆ Configure proxy user authentication

The [Filtering](#) tab is always available on the **Access Control** page.

Other tabs are dynamic based on the authentication method selected in the **Authentication** section of **Configure > My Proxy > Basic**.

If an authentication method is enabled, the [Global Configuration Options](#) tab is always displayed.

If **Integrated Windows Authentication** is selected, these tabs display:

- [Integrated Windows Authentication](#)

- [Global Configuration Options](#)

If **LDAP** is selected, these tabs display:

- [LDAP](#)
- [Global Configuration Options](#)

If **Radius** is selected, these tabs display:

- [Radius](#)
- [Global Configuration Options](#)

If **NTLM** is selected, these tabs display:

- [NTLM](#)
- [Global Configuration Options](#)

If **Rule-Based Authentication** is selected, these tabs display:

- [Domains](#)
- [Authentication Rules](#)
- [Global Configuration Options](#)

The tables below describe the purpose of each field on each tab. Use your browser's Search feature to find the field that you're looking for.

For a complete description of Content Gateway user authentication features, see [Content Gateway user authentication](#), page 189.

Configure > Security > Access Control > Filtering

Filtering rules can be used to:

- ◆ Deny or allow URL requests
- ◆ Insert custom headers
- ◆ Allow specified applications, or requests to specified websites to bypass user authentication
- ◆ Keep or strip header information from client requests
- ◆ Prevent specified applications from transiting the proxy

Rules are ordered checked prior to user authentication (if configured). Rules are applied based on first match in a top-down traversal of the list. If no rule matches, the request is allowed to proceed.

Rules are stored in [filter.config](#).

After adding, deleting, or modifying a rule, restart Content Gateway.

For complete information about filtering rules, see [Filtering Rules](#), page 181.

| | |
|---------------------------|--|
| Filtering | <p>Displays an order list of filtering rules.</p> <p>Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 ThreatScope destinations.</p> |
| 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 | <p>Specifies the rule type:</p> <p>Select allow to allow particular URL requests to bypass authentication.</p> <p>Select deny to deny requests for objects from specific destinations. When a request is denied, the client receives an access denied message.</p> <p>Select keep_hdr to specify which client request header information you want to keep.</p> <p>Select strip_hdr to specify which client request header information you want to strip.</p> <p>Select add_hdr to cause a custom header to be added to the request. This rule type requires that values be defined for Custom Header and Header Value. Add custom headers to satisfy specific requirements of a destination domain. See Filtering Rules, page 181.</p> <p>The radius rule type is not supported.</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 host name.</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 | Specifies the value of the Primary Destination Type. For example, if the Primary Destination Type is dest_ip , the value for this field might be 123.456.78.9. |

| | |
|--|--|
| Additional Specifiers: Header Type | Specifies the client request header information that you want to keep or strip. This option applies to only keep_hdr or strip_hdr rule types. |
| Additional Specifiers: Realm (optional) | Not supported. |
| Additional Specifiers: Proxy Port (optional) | Specifies the proxy port to match for this rule. |
| Additional Specifiers: Custom Header (optional) | For use when the rule type is add_hdr . Specifies the custom header name that the destination domain expects to find in the request. |
| Additional Specifiers: Header Value (optional) | For use when the rule type is add_hdr . Specifies the custom header value that the destination domain expects to be paired with the custom header. |
| Secondary Specifiers: Time | Specifies a time range, such as 08:00-14:00. |
| Secondary Specifiers: Prefix | Specifies a prefix in the path part of a URL. |
| Secondary Specifiers: Suffix | Specifies a file suffix in the URL. |
| Secondary Specifiers: Source IP | Specifies the IP address of the client. |
| Secondary Specifiers: Port | Specifies the port in a requested URL. |
| Secondary Specifiers: Method | Specifies a request URL method: <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) rtsp and mms are not supported. |
| Secondary Specifiers: User-Agent | Specifies the Request header User-Agent value. Use this field to create application filtering rules that: <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. |

Configure > Security > Access Control > Global Configuration Options

Use this page to specify global options for:

- ◆ The fail open/fail closed action to take when user authentication fails
- ◆ Credential caching
- ◆ For transparent proxy, an alternate hostname for the proxy that all clients on the network can resolve. Required.

For more information, see [Global authentication options](#), page 193.



Note

The user interface setting to disable the NTLM cache for explicit proxy has been removed. Although not recommended, the cache can be disabled for explicit proxy traffic in records.config by setting the value of **proxy.config.ntlm.cache.enabled** to **0** (zero).

Global Configuration Options

| | |
|-----------|--|
| Fail Open | <p>Disabled – Prevents requests from proceeding to the Internet when an authentication failure occurs.</p> <p>Enabled only for critical service failures (default) – Allows requests to proceed if authentication fails because there is no response from the domain controller or because the client is sending badly formatted messages.</p> <p>Enabled for all authentication failures – Allows requests to proceed for all authentication failures, including password failures.</p> <p>When a fail open setting is enabled, if a Web Security XID agent is configured an attempt is made to identify the requestor and apply user-based policy. Otherwise, if a policy has been assigned to the client’s IP address, that policy is applied. Otherwise, the Default policy is applied.</p> <p>Important: When user authentication is rule-based with a domain list:</p> <ul style="list-style-type: none"> • If Enabled only for critical service failures is selected, when a critical service failure occurs fail open is not applied. An error always results in fail closed. • If Enabled for all authentication failures, including incorrect password is selected, after trying basic credentials with every domain in the list, fail open is applied. <p>Important: The Fail Open setting does not apply when IWA is the authentication method and the client fails to retrieve a kerberos ticket from the domain controller (DC) because the DC is down. The Fail Open setting does apply with IWA when IWA falls back to NTLM and authentication fails.</p> |
|-----------|--|

| | |
|---|---|
| <p>Credential Caching: Caching Method</p> | <p>Cache using IP address only – specifies that all credentials are cached with IP address surrogates. This is the recommended method when all clients have unique IP addresses.</p> <p>Cache using Cookies only – specifies that all credentials are cached with cookie surrogates. This is recommended when all clients share IP addresses, as with multi-host servers such as Citrix servers, or when traffic is NATed by a device that is forwarding traffic to Content Gateway.</p> <p>Cache using both IP addresses and Cookies – specifies to use cookie surrogates for the IP addresses listed in the cookie caching list, and to use IP address surrogates for all other IP addresses. This is recommended when the network has a mix of clients, some with unique IP addresses and some using multi-user hosts or that are subject to NATing.</p> <p>For a description of surrogate credentials, see Surrogate credentials.</p> <p>Important:</p> <ul style="list-style-type: none"> • Cookie mode caching does not work with applications that do not support cookies, or with browsers in which cookie support has been disabled. • When the browser is Internet Explorer, the full proxy hostname in the form “http://host.domain.com” must be added to the Local intranet zone. • When the browser is Chrome, it must be configured to allow third-party cookies or configured for an exception to allow cookies from the proxy hostname in the form “host.domain.com”. • When the IP address is set for cookie mode and the request method is CONNECT, no caching is performed. • Cookie mode caching is not performed for FTP requests. • Cookie mode caching is not supported by the 7.8.3 version of Captive Portal. Support for cookie mode caching was added for Captive Portal in 7.8.4. |
| <p>Credential Caching: Time-To-Live</p> | <p>Specifies the duration, in minutes, that an entry in the cache is retained. When the TTL expires, the entry is removed and the next time that that user submits a request, the user is authenticated. If the authentication succeeds, an entry is placed in the cache.</p> |

| | |
|--|---|
| Purge LDAP cache on authentication failure | Specifies that when an LDAP user authentication failure occurs, Content Gateway will delete the authorization record for that client from the LDAP cache. |
| Redirect Hostname | For transparent proxy, specifies an alternate hostname for the proxy that all clients on the network can resolve. Required. For complete information see Redirect Hostname , page 196. |

Configure > Security > Access Control > IWA

The Integrated Windows Authentication (IWA) page appears only if you have enabled IWA in the Features table on the **Configure > My Proxy > Basic > General** tab.

Use this page to join or unjoin the Windows domain. When a domain has been joined, the page provides a summary of the domain attributes and an Unjoin button.

For a complete description, see [Integrated Windows Authentication](#), page 197.

Integrated Windows Authentication

| | |
|--------------------------|---|
| Domain Name | Specifies the fully qualified Windows domain name. |
| Administrator Name | Specifies the Windows Administrator user name. |
| Administrator Password | Specifies the Windows Administrator password. Note: The name and password are used only during the join and are not stored. |
| Domain Controller | Specifies how to locate the domain controller: <ul style="list-style-type: none"> • Auto-detect using DNS • DC name or IP address If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list. |
| Content Gateway Hostname | Specifies the Content Gateway hostname. Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on V-Series appliances (V-Series adds 4 characters to the hostname to ensure that the hostname is unique across modules (Doms)). IMPORTANT: Once the domain is joined the hostname cannot be changed. If it is, IWA will immediately stop working until the domain is unjoined and then rejoined with the new hostname. |
| Join Domain | Click Join Domain to join the domain. |

Configure > Security > Access Control > LDAP

The LDAP configuration options appear on the Configure pane only if you have enabled LDAP in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information on configuring LDAP see [LDAP authentication](#), page 204.

| LDAP | |
|--------------------------------------|---|
| LDAP Server: Hostname | Specifies the hostname of the LDAP server. If you change this option, you must restart Content Gateway. |
| LDAP Server: Port | Specifies the port used for LDAP communication. The default port number is 389. To use the default Global Catalog server port, specify port 3268. If Secure LDAP is enabled, set the port to 636 or 3269 (the secure LDAP ports). If you change this option, you must restart Content Gateway. |
| LDAP Server: Secure LDAP | Specifies whether Content Gateway will use secure communication with the LDAP server. If enabled, set the LDAP Port field (above) to 636 or 3269 (the secure LDAP ports). |
| LDAP Server: Server Type | Specifies the search filter. Select either Microsoft Active Directory or other directory services. |
| LDAP Server: Bind Distinguished Name | Specifies the Full Distinguished Name (fully qualified name) of a user in the LDAP-based directory service. For example: <code>CN=John Smith, CN=USERS, DC=MYCOMPANY, DC=COM</code> Enter a maximum of 128 characters in this field. If you do not specify a value for this field, the proxy attempts to bind anonymously. |
| LDAP Server: Password | Specifies a password for the user identified in the Bind_DN field. |
| LDAP Server: Base Distinguished Name | Specifies the base Distinguished Name (DN). You can obtain this value from your LDAP administrator. You must specify a correct base DN; otherwise LDAP authentication will fail to operate. If you change this option, you must restart Content Gateway. |

Configure > Security > Access Control > Radius

The Radius configuration options appear on the Configure pane only if you have enabled Radius in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information on configuring Radius, see [RADIUS authentication](#), page 208.

| Radius | |
|--|--|
| Primary Radius Server: Hostname | Specifies the hostname or IP address of the primary RADIUS authentication server. If you change this option, you must restart Content Gateway. |
| Primary Radius Server: Port | Specifies the port that Content Gateway uses to communicate with the primary RADIUS authentication server. The default port is 1812. If you change this option, you must restart Content Gateway. |
| Primary Radius Server: Shared Key | Specifies the key to use for encoding. If you change this option, you must restart Content Gateway. |
| Secondary Radius Server (optional): Hostname | Specifies the hostname or IP address of the secondary RADIUS authentication server. If you change this option, you must restart Content Gateway. |
| Secondary Radius Server (optional): Port | Specifies the port that Content Gateway uses to communicate with the secondary RADIUS authentication server. The default port is 1812. If you change this option, you must restart Content Gateway. |
| Secondary Radius Server (optional): Shared Key | Specifies the key to use for encoding. If you change this option, you must restart Content Gateway. |

Configure > Security > Access Control > NTLM

The NTLM configuration options appear on the Configure pane only if you have enabled NTLM in the Features table on the **Configure > My Proxy > Basic > General** tab.

For more information on configuring NTLM, see [Legacy NTLM authentication](#), page 202.

NTLM

| | |
|-----------------------------|--|
| Domain Controller Hostnames | <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> |
| Load Balancing | <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> |

Configure > Security > Access Control > Domains

The Domains tab appears in the Access Control list only if you have enabled **Rule-Based Authentication** in the Features table on **Configure > My Proxy > Basic > General**.

Use this tab to create and maintain a list of domains that can be specified in authentication rules. Use the Authentication Rules tab to define authentication rules. Be sure to set the [Global authentication options](#), page 193.



Important

You must configure the Domains list before you configure authentication rules.

If you have never configured rule-based authentication, see [Rule-Based Authentication, page 211](#), for complete information.

Domains

| | |
|---------------------------------------|--|
| Domain List | <p>An unordered list of domains that have been identified for use in authentication rules.</p> <p>Use the Edit button to change some attributes associated with the domain.</p> <p>Use the Delete or Unjoin button to remove a domain from the list.</p> <p>The domain list is stored in <code>auth_domains.config</code>.</p> |
| Domain list: New Domain button | <p>Use the New Domain button to add a domain to the Domains list. The screen is expanded to allow for specification of the domain.</p> |
| | <p>New Domain action</p> |
| Domain Details: Domain Identifier | <p>Specify a unique name for the domain. The name is used only by Content Gateway; it does not change any attribute of the actual domain or directory.</p> <p>Important: You cannot change the domain identifier after it has been added to the list. To change the name, delete the entry from the list and re-add it with the new name.</p> |
| Domain Details: Authentication Method | <p>Specify the authentication method: IWA, Legacy NTLM, or LDAP. Radius is not supported.</p> <p>When you select an authentication method, configuration options specific to that method are added to the page.</p> <p>Important: You cannot change the authentication method after you add the domain to the list. To change the authentication method, delete the entry from the list and re-add the domain specifying the new authentication method.</p> |
| Domain Details: Aliasing | <p>Specify an alias to send to the filtering service for all users who match this rule (optional). The alias must be static. It can be empty (blank). The alias must exist in the primary domain controller (the DC visible to the filtering service). See Unknown users and the 'alias' option, page 216.</p> |
| IWA Domain Details | <p>These options are presented when IWA is specified as the authentication method.</p> |

| | |
|-----------------------------------|--|
| Domain Name | Specify the fully qualified domain name. For example: corp-domain.example.com |
| Administrator Name | Specify a Windows Active Directory domain administrator user name. |
| Administrator Password | Specify the corresponding domain administrator password. Note: The name and password are used only during the join and are not stored. |
| Domain Controller | Specify how to locate the domain controller: <ul style="list-style-type: none"> • Auto-detect using DNS • DC name or IP address If the domain controller is specified by name or IP address, you can also specify backup domain controllers in a comma separated list. |
| Content Gateway Hostname | Specify the Content Gateway hostname. Because IWA uses the hostname as a NetBIOS name when registering with Kerberos, the hostname cannot exceed 15 characters in length (a NetBIOS restriction), or 11 characters on V-Series appliances (V-Series adds 4 characters to the hostname to ensure that the hostname is unique across modules (Doms). Warning: Once the domain is joined the hostname cannot be changed. If it is, IWA will immediately stop working until the domain is unjoined and then rejoined with the new hostname. |
| Join Domain | Click Join Domain to join the domain. |
| Legacy NTLM Domain Details | |
| Domain Controller | Specify the IP address and port number of the primary domain controller (if no port is specified, Content Gateway uses port 139), followed by a comma separated list of secondary domain controllers to be used for load balancing and failover. |
| Load Balance | Select the check box to balance the load across multiple NTLM DCs. Note: When multiple domain controllers are specified, even if load balancing is disabled, when the load on the primary domain controller reaches the maximum number of connections allowed, new requests are sent to a secondary domain controller as a short-term failover provision, until such time that the primary domain controller can accept new connections. |
| LDAP Domain Details | |
| LDAP Server Name | Specify the LDAP server name. |
| LDAP Server Port | Specify the LDAP Server Port (optional) The default is 389. |
| LDAP Base Distinguished Name | Specify the LDAP Base Distinguished Name. |

| | |
|------------------|--|
| LDAP Server Type | Set the search filter to “sAMAccountName” for Active Directory, or “uid” for other directory services. |
| Bind Domain Name | Specify the LDAP bind account distinguished name. For example: CN=John Smith,CN=USERS,DC=MYCOMPANY,DC=COM The field length is limited to 128 characters. If no value is specified, Content Gateway attempts to bind anonymously. |
| Bind Password | Specify the LDAP bind account password. |
| Secure LDAP | Specify whether Content Gateway will use secure communication with the LDAP server. If enabled, you must set the LDAP port to one of the secure ports: 636 or 3269. |

Configure > Security > Access Control > Authentication Rules

The Authentication Rules tab appears in the Access Control list only if you have enabled **Rule-Based Authentication** in the Features table on the **Configure > My Proxy > Basic > General** tab.

Use this tab to create and maintain authentication rules. Use the **Domains** tab to build and maintain a list of domains that can be used in authentication rules. You must configure the Domains list before you define authentication rules.

Be sure to set the [Global authentication options, page 193](#).



Important

If you have never configured rule-based authentication, see [Rule-Based Authentication, page 211](#), for complete information.

Authentication Rules

| | |
|--------------------------|--|
| Authentication Rule List | Displays a table of the ordered list of rules defined for user authentication. Rules are defined for sets of clients to be authenticated against one or more IWA, LDAP and NTLM domains. See Rule-Based Authentication, page 211 . |
| Refresh | Updates the table to display the current rules in the auth_rules.config file. |
| Edit File | Opens the authentication rule editor. Warning: Do not edit rules directly in the configuration file. |

| | auth_rules.config Configuration File Editor |
|------------------|---|
| rule display box | <p>Lists, in order, the current rule set. When user authentication is performed, the list is traversed, top-down and the first match is applied.</p> <p>Select a rule to edit it.</p> <p>The arrows to the left of the box allow you to move the selected rule up or down in the list.</p> <p>The “X” button deletes the selected rule.</p> <p>Rules cannot be more than 512 characters.</p> |
| Add | Adds a new rule. |
| Set | Updates the selected rule with the current values. |
| Status | <p>Specifies whether the rule is enabled (active) or disabled after the rule is saved and Content Gateway is restarted.</p> <p>You can create a rule and not enable it until other elements of your network are ready to support it.</p> |
| Rule Name | Specifies a unique, descriptive name for the rule. It is recommended that the name not exceed 50 characters. |
| Source IP | <p>Specifies IP addresses or IP address ranges for this rule (must be entered without any spaces).</p> <p>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</p> |
| Proxy Port | <p>Specifies the inbound port for traffic when Content Gateway is deployed as an explicit proxy. If undefined, all ports match, as configured on Configure > Protocols > HTTP > General.</p> <p>Transparent proxy deployment should leave this field undefined.</p> |
| User-Agent | <p>Specifies 1 or more regular expressions used to match text in the User-Agent string, for example to match common browsers.</p> <p>Regexes must be POSIX-compliant.</p> <p>The “^” operator is not supported.</p> <p>When the field is empty, all User-Agent values match.</p> <p>You can edit the field directly.</p> <p>To insert a predefined regex for a common browser, select it from the drop down list and click Add.</p> <p>Multiple regexes can be specified. Use the “ ” character to separate entries (logical ‘or’).</p> <p>For more information, including regex examples, see Authentication based on User-Agent, page 227.</p> |

| | |
|----------------|---|
| Auth Sequence | <p>Specifies 1 or more domains to use for authentication. Select a domain from the Domains drop down list (populated from the Domains List), and click Include to add it to the list.</p> <p>If you add more than one domain, you can set the order by selecting an entry and using the up and down arrows. You can delete a selected domain with the “X” button.</p> <p>Best practice: If you know what domain a set of users belongs to, create a rule just for that group.</p> <p>Best practice: Place the rule with the largest number of users authenticating with known domain membership at the top of the list. These are the fastest authentications.</p> <p>Best practice: If you don’t know what domain a set of users belongs to, specify the fewest number of domains needed to authenticate the users in the set.</p> <p>Best practice: It is always better to create targeted rules because attempting to authenticate against a large set of domains can introduce noticeable latency.</p> <p>Important: When user authentication is rule-based with a domain list:</p> <ul style="list-style-type: none"> • For each user, the first successful authentication is cached and used in subsequent authentications. If IP address caching is configured, an IP address surrogate is cached. If Cookie Mode is configured, a cookie surrogate is cached. <p>For Fail Open:</p> <ul style="list-style-type: none"> • If Enabled only for critical service failures is selected, the fail open setting is not applied. The user continues to be prompted for credentials until there is a timeout. • If Enabled for all authentication failures, including incorrect password is selected, after trying basic credentials with every domain in the list, fail open is applied. |
| Captive Portal | <p>Click Enabled (7.8.3) or Enabled for HTTPS/HTTP Authentication page (beginning with 7.8.4) to redirect users to a customizable web portal page for authentication.</p> <p>See <i>Authentication using Captive Portal</i> for details.</p> |
| Apply | <p>Applies the configuration changes.</p> <p>Important: If the rule specifies a regex for User-Agent, the regex is validated when Apply is clicked. If the regex is not valid, the rule is deleted and must be recreated.</p> |
| Close | <p>Exits the configuration file editor.</p> <p>Click Apply before you click Close; otherwise, all configuration changes will be lost.</p> |

SOCKS

Help | Content Gateway | Version 7.8.x

For more information about Content Gateway support for SOCKS, see [Configuring SOCKS firewall integration](#), page 185.



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.

Configure > Security > Access Control > SOCKS > 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> |
|---------------|--|

Configure > Security > Access Control > SOCKS > Proxy

| | |
|------------------|---|
| SOCKS Proxy | <p>Enables or disables the SOCKS Proxy option. As a SOCKS proxy, Content Gateway can receive SOCKS packets (usually on port 1080) from the client, and forward requests directly to the SOCKS server.</p> <p>For more information about the SOCKS Proxy option, see Configuring SOCKS firewall integration, page 185.</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> |

Configure > Security > Access Control > SOCKS > Server

| | |
|----------------------------------|---|
| On-Appliance SOCKS server | <p>Displays only when Content Gateway is on a Websense appliance.</p> <p>Enables or disables the on-appliance SOCKS server.</p> <p>The SOCKS proxy option must be enabled to route client requests through the SOCKS server.</p> <p>You can configure Content Gateway to use other SOCKS servers in your network by editing socks_server.config. See the next entry.</p> |
| Socks Servers table | Displays a table of configured SOCKS servers. For information about adding and configuring SOCKS servers, see Configuring SOCKS servers , page 186. |
| Refresh | Updates the table to display the current entries in socks_server.config . |
| Edit File | Opens the configuration file editor for socks_server.config . |
| | socks_server.config Configuration File Editor |
| entry display box | Lists the SOCKS servers that have been configured for use with Content Gateway. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected entry up or down in the list. |
| Add | Adds an entry to the server list. |
| Set | Updates the selected entry. Select a server from the list; modify the settings; click Set to update the entry. |
| Clear Fields | Clears all fields for the selected server. |
| SOCKS Server Name | Specify a name that helps distinguish this SOCKS server from other SOCKS servers. |
| SOCKS Server Host | Specify the SOCKS server IP address, or a hostname that is resolvable by your internal DNS service. |
| SOCKS Port | Specify the port on which the SOCKS server listens. |
| Default SOCKS Server | Select this option to make this SOCKS server the default SOCKS server. |
| SOCKS User Name | When SOCKS authentication is used, specify the SOCKS user name with which to authenticate. |
| SOCKS Password | When SOCKS authentication is used, specify the password that goes with the specified user. |
| Apply | Applies the configuration changes. |
| Close | <p>Exits the configuration file editor.</p> <p>Click Apply before you click Close; otherwise, all configuration changes are lost.</p> |

| | |
|---------------------------|--|
| Socks Server Rules | <p>Displays a table listing the rules in the socks.config file that specify the SOCKS servers that Content Gateway must go through to access specific origin servers, and the order in which Content Gateway goes through the SOCKS server list.</p> <p>You can also specify the origin servers that you want the proxy to access directly, without going through a SOCKS server.</p> |
| Refresh | Updates the table to display the current rules in the socks.config file. |
| Edit File | Opens the configuration file editor for the socks.config file. |
| | socks.config Configuration File Editor |
| rule display box | Lists the <i>socks.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| Rule Type | <p>Select Route through SOCKS server to specify the origin servers that you want the proxy to route through a SOCKS server.</p> <p>Select Do not route through SOCKS server to specify the origin servers that you want the proxy to access directly, bypassing the SOCKS server(s).</p> |
| Destination IP | <p>For Route through SOCKS server, specify either a single IP address <i>or</i> a range of IP addresses of origin servers for which Content Gateway must use the SOCKS servers specified in the SOCKS Servers field below.</p> <p>For Do not route through SOCKS server, specify the IP addresses of the origin servers that you want the proxy to access directly (without going through the SOCKS server). You can enter a single IP address, a range of IP addresses, or a list of IP addresses. Separate each entry in the list with a comma. Do not specify the all networks broadcast address: 255.255.255.255.</p> |
| SOCKS Server | For a Route through SOCKS server rule , select the SOCKS server(s) through which to route requests. |
| Round Robin | Specifies how strictly Content Gateway will follow round robin. You can select strict , or false . |
| Apply | Applies the configuration changes. |
| Close | <p>Exits the configuration file editor.</p> <p>Click Apply before you click Close; otherwise, all configuration changes will be lost.</p> |

Configure > Security > Access Control > SOCKS > Options

| | |
|---------------------------------|--|
| Server Connection Timeout | Specifies how many seconds Content Gateway waits attempting to connect to a SOCKS server before timing out. |
| Connection Attempts Per Server | Specifies how many times Content Gateway attempts to connect to a given SOCKS server before marking the server as unavailable. |
| Server Pool Connection Attempts | Specifies how many times Content Gateway attempts to connect to a given SOCKS server in the pool before giving up. |

Subsystems

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The Subsystems configuration options are divided into the following categories:

[Cache](#), page 354

[Logging](#), page 356

[Networking](#), page 360

Cache

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

Configure > Subsystems > Cache > 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 cache.config file. |
| Ram Cache Size | Specifies the size of the RAM cache, in bytes. The default size is 104857600 (100 MB). A value of “-1” directs Content Gateway to automatically size the RAM cache to 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. |

Configure > Subsystems > Cache > Partition

Configure > Subsystems > Cache > Hosting

| | |
|---|---|
| Cache Hosting | Displays a table listing the rules in the hosting.config 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 hosting.config file. |
| Edit File | Opens the configuration file editor for the hosting.config file. The configuration file editor page is described below. |
| 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

Help | Content Gateway | Version 7.8.x

Configure > Subsystems > Logging > General

| | |
|-----------------------------|---|
| Logging | <p>Enables or disables event logging so that transactions are recorded into event log files and/or error log files.</p> <p>Select Log Transactions and Errors to log transactions into your selected event log files and errors in the error log files.</p> <p>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.</p> <p>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.</p> <p>Select Disabled to turn off logging.</p> |
| Log Directory | <p>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</p> |
| Log Space: Limit | <p>Specifies the maximum amount of space (in megabytes) allocated to the logging directory for the log files.</p> <p>When Content Gateway is on a Websense appliance, the size is set to 5120 (5 GB) and cannot be changed.</p> <p>When Content Gateway is installed on a stand-alone server, the default size is 20480 (20 GB) and the size is configurable.</p> <p>Note: Transaction logs can consume a lot of space. Make sure that this limit is smaller than the actual space available on the partition that contains the logging directory.</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 251.</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 | Specifies the amount of time Content Gateway enters data in log files before rolling them to .old files. The minimum value is 300 seconds (five minutes). The default value is 21600 seconds (6 hours). The maximum value is 86400 (1 day). |
| Log Rolling: Auto-Delete Rolled Files | Enables 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 . |
| Reverse DNS lookup for Threat Tracking | Enables or disables reverse DNS lookups to facilitate inclusion of the client host name in the Web Security Threat dashboard, and in logs and reports. Caution: To achieve the expected results and avoid unexpected network behaviors, before enabling this option be sure that reverse DNS is configured in your network. |

Configure > Subsystems > Logging > Formats

| | |
|--|--|
| Squid Format: Enable/Disable | Enables or disables the Squid log format. |
| Squid Format: ASCII/Binary | Select ASCII or Binary as the type of log files to be created. |
| Squid Format: Filename | Specifies the name used for Squid log files. The default filename is squid.log . |
| Squid Format: Header | Specifies the text header you want Squid log files to contain. |
| Netscape Common Format: Enable/Disable | Enables or disables the Netscape Common log format. |
| Netscape Common Format: ASCII/Binary | Select ASCII or Binary as the type of log file to be created. |
| Netscape Common Format: Filename | Specifies the name used for Netscape Common log files. The default filename is common.log . |
| Netscape Common Format: Header | Specifies the text header you want Netscape Common log files to contain. |
| Netscape Extended Format: Enable/Disable | Enables or disables the Netscape Extended log format. |
| Netscape Extended Format: ASCII/Binary | Select ASCII or Binary as the type of log file to be created. |
| Netscape Extended Format: Filename | Specifies the name used for Netscape Extended log files. The default filename is extended.log . |
| Netscape Extended Format: Header | Specifies the text header you want Netscape Extended log files to contain. |
| Netscape Extended 2 Format: Enable/Disable | Enables or disables the Netscape Extended-2 log format. |

| | |
|---|---|
| Netscape Extended 2 Format: ASCII/Binary | Select ASCII or Binary as the type of log file to be created. |
| Netscape Extended 2 Format: Filename | Specifies the name used for Netscape Extended-2 log files. The default filename is extended2.log . |
| Netscape Extended 2 Format: Header | Specifies the text header you want Netscape Extended-2 log files to contain. |

Configure > Subsystems > Logging > Splitting

| | |
|-----------------|--|
| Split ICP Logs | <p>When enabled, Content Gateway records ICP transactions in a separate log file.</p> <p>When disabled, Content Gateway records ICP transactions in the same log file with HTTP and FTP entries.</p> |
| Split Host Logs | <p>When enabled, Content Gateway creates a separate log file for each of the hosts listed in the log_hosts.config file.</p> <p>When disabled, Content Gateway records transactions for all hosts in the same log file.</p> |

Configure > Subsystems > Logging > 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 256.</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 LogObject 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. |

Configure > Subsystems > Logging > Custom

| | |
|-----------------------------|--|
| Custom Logging | Enables or disables custom logging. |
| Custom Log File Definitions | Displays the <i>logs_xml.config</i> file so that you can configure custom (XML-based) logging options. |

Networking

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The Networking configuration options are divided into the following categories:

[Connection Management](#), page 360

[ARM](#), page 362

[WCCP](#), page 368

[DNS Proxy](#), page 372

[DNS Resolver](#), page 373

[ICAP](#), page 376

[Virtual IP](#), page 377

[Health Check URLs](#), page 378

Connection Management

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The options on the Connection Management pages allow you to tune several important properties of proxy behavior, including connection throttling and load shedding, individual client connection limits and rates, and how to respond to low memory conditions.

By default, Content Gateway accepts 45,000 connections. A connection throttle event occurs when client or origin server connections reach 90% of half the configured limit (20,250 by default). When a connection throttle event occurs, Content Gateway continues processing all existing connections and queues new client connection requests until the connection count falls below the limit.

If you think that Content Gateway is hitting the connection limits, you should monitor the Performance graphs to get an accurate reading of connection activity. In particular, check the **Active Client Connections** and **TCP ESTABLISHED Connections** graphs. You can also check error messages in the system log file, error log file, or event log files.

Configure > Networking > Connection Management > Throttling

| | |
|----------------------------|--|
| Throttling Net Connections | <p>Specifies the maximum number of network connections that Content Gateway accepts. The default value is 45,000.</p> <p>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.</p> <p>Do not set this variable below the minimum value of 100.</p> |
|----------------------------|--|

Configure > Networking > Connection Management > Load Shedding

| | |
|---------------------|---|
| Maximum Connections | <p>Specifies the maximum number of client connections allowed before the ARM starts forwarding incoming requests directly to the origin server. The default value is 1 million connections.</p> <p>If you change this option, you must restart Content Gateway.</p> |
|---------------------|---|

Configure > Networking > Connection Management > Client Connection Control

Specifies:

- ◆ Client concurrent connection limits
- ◆ Client connection rate limits
- ◆ Proxy response when a limit is exceeded
- ◆ A list of clients excepted from the limits

| | |
|--|--|
| Concurrent Connection Limit: Maximum concurrent connections | Specifies the maximum number of concurrent HTTP/HTTPS connections a client is allowed. The default is 1000. The supported range is: 1 - 45000 |
| Concurrent Connection Limit: Alert when limit exceeded | <p>When enabled, causes Content Gateway to generate an alert when a client exceeds the maximum concurrent connection limit.</p> <p>In addition to displaying the alert in the Content Gateway manager, it is also logged in <code>/var/log/messages</code> and <code>content_gateway.out</code>.</p> |
| Concurrent Connection Limit: Close excessive connections when limit exceeded | When enabled, causes Content Gateway to close excessive connections when the limit is exceeded. |

| | |
|---|---|
| Connection Rate Limit: Maximum connection rate | Specifies the maximum connections per second, averaged over a minute, that a client can make. The default is 100. The supported range is: 1 - 1000 |
| Connection Rate Limit: Alert when limit exceeded | When enabled, causes Content Gateway to generate an alert when a client exceeds the maximum connection rate limit. In addition to displaying the alert in the Content Gateway manager, it is also logged in /var/log/messages and content_gateway.out. |
| Connection Rate Limit: Close excessive connections when limit exceeded | When enabled, causes Content Gateway to close excessive connections when the limit is exceeded. |
| Exceptions | Specifies IP addresses and/or IP address ranges to which connection limits are not applied. IP addresses can be IPv4 or IPv6 (IPv6 support must be enabled). Multiple addresses or ranges can be specified in a comma-separated list. |

Configure > Networking > Connection Management > Low Memory Mode

Specifies whether Content Gateway suspends analysis of web traffic when the host system experiences a low-memory condition. In this state, Web Security URL filtering is applied as usual.

| | |
|--------------------------------------|--|
| Low Memory Mode: Enabled/Disabled | Select Enabled to suspend content analysis when there is a low memory condition. |
| Low Memory Mode Duration | Specifies the length of time, in minutes, that content analysis is suspended. If the low memory condition resolves itself before the timer expires, analysis resumes and the low memory mode trigger resets. If the timer expires, analysis resumes and the low memory mode trigger is not reset. |

ARM

Help | Content Gateway | Version 7.8.x

The Adaptive Redirection Module (ARM) performs several essential functions including sending device notifications for cluster communication interface failover and inspection of incoming packets before the IP layer sees them, readdressing them for Content Gateway processing.

The ARM is always active. For more information, see [The ARM](#), page 48.

Configure > Networking > ARM > General

| | |
|--|--|
| 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 ipnat.conf file. |
| Edit File | Opens the configuration file editor for the ipnat.conf file. |
| | ipnat.conf Configuration File Editor |
| rule display box | Lists the <i>ipnat.conf</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| Ethernet Interface | Specifies the Ethernet interface that traffic will use to access the Content Gateway machine: for example, <code>eth0</code> on Linux. |
| Connection Type | Specifies the connection type that applies for the rule: TCP or UDP. |
| Destination IP | Specifies the IP address from which traffic is sent. 0.0.0.0 matches all IP addresses. |
| Destination CIDR | Specifies the IP address in CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24. Entering a value in this field is optional. |
| Destination Port | Specifies the traffic destination port: for example, 80 for HTTP traffic. |
| Redirected Destination IP | Specifies the IP address of your Content Gateway server. |
| Redirected Destination Port | Specifies the proxy port: for example, 8080 for HTTP traffic. |
| User Protocol (Optional) | When dns is selected, the ARM redirects DNS traffic to Content Gateway; otherwise, DNS traffic is bypassed. |
| Apply | Applies the configuration changes. |
| Close | Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes are discarded. |

| | |
|---|---|
| <p>IP Spoofing: Enabled/ Disabled</p> | <p>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 77.</p> <p>WARNING: IP spoofing requires precise control of the routing paths on your network, overriding the normal routing process for traffic running on TCP port 80 and 443.</p> |
| <p>Range Based IP Spoofing: Enabled/ Disabled</p> | <p>Enables or disables the range-based IP spoofing extension. This extension supports the specification of IP addresses and ranges of addresses that are mapped to specified IP addresses for spoofing.</p> <p>Many groups can be specified. However, use this feature judiciously because list traversal adds overhead to every connection request. The larger the list, the more overhead.</p> <p>The list is traversed in order (as displayed). The first match is applied.</p> <p>Clients that don't match a grouping are spoofed with their own IP address (basic IP spoofing).</p> <p>For more information, see IP spoofing, page 77.</p> |
| <p>Range Based IP Spoofing: Address table</p> | <p>In the Client IP Addresses field, enter a comma separated list of individual IP addresses and/or IP address ranges. Do not use spaces.</p> <p>You can use:</p> <ul style="list-style-type: none"> • A simple IP address, such as 123.45.67.8 • CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24. • A range separated by a dash, such as 1.1.1.1-2.2.2.2 • Any combination of the above, separated by commas, such as: 1.1.1.0/24,25.25.25.25,123.1.23.1-123.1.23.123 <p>In the Specified IP Address field, enter the IP address to use with matching clients. This is the spoofed IP address.</p> <p>To add a row to the table, click Add Row.</p> <p>To remove a row from the table, delete the contents of the cells. When you click Apply the empty row(s) is removed</p> <p>The table always has a minimum of 5 rows.</p> <p>Restart Content Gateway to put changes into effect.</p> |

Configure > Networking > ARM > Static Bypass

Static bypass rules route requests around the proxy (bypass). Rules can be defined for clients (sources), origin servers (destinations), or both (pairs). See [Static bypass rules](#), page 72.



Important

This feature is for transparent proxy deployments only.

| | |
|---------------------|--|
| Static Bypass table | Lists the configured static bypass rules. When Content Gateway is serving transparent traffic, the proxy uses these rules to determine whether to bypass incoming client requests or attempt to serve them transparently. Rules are stored in bypass.config |
| 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 bypass.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 | Specifies the rule type: A bypass rule bypasses specified incoming requests. A deny_dyn_bypass rule prevents the proxy from bypassing specified incoming client requests dynamically (a deny bypass rule can prevent Content Gateway from bypassing itself). |
| Source IP | Specifies the source IP address in incoming requests that the proxy must bypass or deny bypass. The IP address can be one of the following: A simple IP address, such as 123.45.67.8 In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24. A range separated by a dash, such as 1.1.1.1-2.2.2.2 Any combination of the above, separated by commas, such as 1.1.1.0/24, 25.25.25.25, 123.1.23.1-123.1.23.123 |

| | |
|----------------|---|
| Destination IP | <p>Specifies the destination IP address of incoming requests that the proxy must bypass or deny bypass. The IP address can be one of the following:</p> <p>A simple IP address, such as 123.45.67.8</p> <p>In CIDR (Classless Inter-Domain Routing) format, such as 1.1.1.0/24</p> <p>A range separated by a dash, such as 1.1.1.1-2.2.2.2</p> <p>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</p> |
| 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> |

Configure > Networking > ARM > Dynamic Bypass

| | |
|-----------------------------|--|
| Dynamic Bypass | <p>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.</p> |
| 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> |

| | |
|--------------------|--|
| 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

Help | Content Gateway | Version 7.8.x



Note

The WCCP configuration options appear on the Configure pane only if you have enabled WCCP in the Features table on the **Configure > My Proxy > Basic > General** tab.

The options defined in the **wccp.config** configuration file control the use of WCCP with Content Gateway. Entries should be defined and maintained using the editor provided on **Configure > Networking > WCCP**.

Administrators should have a good working knowledge of WCCP.

Only WCCP v2 is supported.

It is recommended that you consult the documentation and the manufacturer's support site for information regarding optimal configuration and performance of your WCCP v2 device. Most devices should be configured to take best advantage of hardware-based redirection. With Cisco devices, the most recent version of IOS is usually best.

For every active WCCP service group, there must be a corresponding ARM NAT rule. See [ARM](#), page 362.

For a complete description of Content Gateway support for WCCP v2, see [Transparent interception with WCCP v2 devices](#), page 50.

| Option | Description |
|---------------------|---|
| WCCP Service Groups | Displays a table of the service groups defined in the wccp.config file. WCCP service group configuration defines WCCP behavior. Column fields are explained in the Configuration Editor entries below. |
| Refresh | Refreshes the table to display the current definitions in the wccp.config file. |
| Edit File | Opens wccp.config in the configuration file editor. |

| Option | Description |
|--|--|
| Synchronize in the Cluster | <p>When there are several Content Gateway nodes in a cluster:</p> <p>Enable this option to cause the WCCP configuration (wccp.config) to be synchronized in the cluster. This allows configuration changes to be made on any node in the cluster.</p> <p>Disable this option to cause the WCCP configuration to not be synchronized in the cluster. This requires that changes to the WCCP configuration be made individually on each node. A common use case for this is to control which service groups are enabled/disabled on each node, and to use proportional load distribution with weight.</p> <p>If after being disabled this option is enabled, the configuration on the node on which the option is enabled is used to initially synchronize the cluster.</p> |
| wccp.config Configuration File Editor | |
| Service group display box | <p>Lists the WCCP service group definitions.</p> <p>Select an entry in the list to edit it.</p> <p>Use the “X” button to delete the selection.</p> <p>List order has no meaning; therefore, the up and down arrows can be ignored.</p> |
| Add | <p>Adds a new service group definition. After Add is clicked, the new definition is displayed in the box at the top of the page.</p> |
| Set | <p>Accepts modifications to the selected service group definition, displaying the new values in the box at the top of the page.</p> |
| Service Group Information | |
| Service Group Status | <p>Enables or disables the service group.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Service Group Name | <p>Specifies a unique service group name. This is as an aid to administration.</p> |
| Service Group ID | <p>Specifies a service group ID between 0-255. This ID must also be configured on the router(s).</p> <p>If the specified number is already in use, an error is displayed when Add or Set is clicked.</p> |
| Protocol | <p>Specifies the protocol, TCP or UDP, that applies to this service group.</p> |
| Ports | <p>Specifies up to 8 ports in a comma separated list.</p> |
| Network Interface | <p>Specifies the Ethernet interface on this Content Gateway host system to use with this service group. On a V10000 appliance, eth0 is bound to P1 and eth1 is bound to P2.</p> |
| Mode Negotiation | |

| Option | Description |
|---------------------------|--|
| Special Device Profile | <p>Select ASA Firewall to specify that traffic is routed to the proxy by a Cisco ASA firewall. When this option is selected, GRE is automatically selected as the Packet Forward Method and Packet Return Method. These settings are required and cannot be changed.</p> |
| Packet Forward Method | <p>Specifies the preferred encapsulation method used by the WCCP router to transmit intercepted traffic to the proxy. If the router supports GRE and L2, the method specified here is used.</p> <p>Important: GRE and Multicast are incompatible.</p> <p>Important: If you change the forward or return method configuration while there is an active connection with the WCCP device, in order to re-negotiated the method you must force the current connection to terminate. Typically, this means turning off the service group on the WCCP device for 60 seconds. See the documentation for your WCCP device.</p> |
| Packet Return Method | <p>Specifies the preferred packet encapsulation method used to return intercepted traffic to the WCCP router.</p> <p>Note: If Content Gateway is configured with a Forward/Return method that the router does not support, the proxy attempts to negotiate a method supported by the router.</p> <p>Note: Selecting L2 requires that the router or switch be Layer 2-adjacent (in the same subnet) as Content Gateway.</p> |
| Advanced Settings | |
| Assignment Method | <p>Specifies the method that the router will use to distribute intercepted traffic across multiple proxy servers. Choices are HASH and MASK.</p> <p>The MASK value is applied up to 6 significant bits (in a cluster, a total of 64 buckets are created).</p> <p>See your WCCP documentation for more information about assignment method. Use the value recommended in the manufacturer's documentation for your device.</p> |
| Distribution attribute(s) | <p>Specifies the attribute that the assignment method uses to determine which requests are distributed to which proxy servers.</p> <p>If the assignment method is HASH, select one or more distribution attributes.</p> <p>If the assignment method is MASK, select one distribution attribute.</p> |

| Option | Description |
|---------------------------------|---|
| Weight | <p>This option is only useful when Synchronize in the Cluster is disabled.</p> <p>Specifies the distribution of requests to servers in a cluster by proportional weighting. Set weight to a value that is the desired proportion of the total flow of traffic.</p> <p>When all cluster members have a value of 0 (the default), distribution is equal. If any member has a non-zero value, distribution is proportional, relative to the weight values of other members. Members that continue to have a value of zero, receive no traffic.</p> <p>See <i>WCCP load distribution</i>, page 53.</p> |
| Reverse Service Group ID | <p>For use when IP spoofing is enabled.</p> <p>When IP spoofing is enabled, the proxy advertises a reverse service group for each enabled WCCP forward service group. The reverse service group must be applied along the return path of origin server responses to the proxy.</p> |
| Router Information | |
| Security (optional) | <p>Enables or disables security so that the router and Content Gateway can authenticate each other.</p> <p>If you enable security in Content Gateway, you must also enable security on the router. See your router documentation.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Security:Password | <p>Specifies the password used for authentication. The password must be the same password as that configured on the router and can be a maximum of eight characters long.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Multicast (optional) | <p>Enables or disables WCCP multicast mode.</p> <p>Important: Cannot be used with GRE packet Forward/Return method.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| Multicast: IP Address | <p>Specifies the multicast IP address.</p> <p>If you change this option, you must restart Content Gateway.</p> |
| WCCP Routers: Router IP Address | <p>Specifies the IP addresses of up to 10 WCCP v2-enabled routers.</p> <p>If you change this option, you must restart Content Gateway.</p> |

| Option | Description |
|---|---|
| WCCP Routers: Local GRE Tunnel Endpoint IP Address | <p>If GRE is selected for Packet Return Method, also specify Local GRE Tunnel Endpoint IP Addresses, except when the device is an ASA firewall.</p> <p>These are Content Gateway tunnel endpoints for the associated Router IP Addresses.</p> <p>A Local GRE Tunnel Endpoint IP Address:</p> <ul style="list-style-type: none"> • Must be IPv4 • Must be unique for every router in the table • Must not be assigned to any other device • Must be a routable IP address • Should reside on the same subnet as the proxy. If it is not, you must define a route for it. • Is not intended to be a client-facing proxy IP address • Is bound to the physical interface specified for the service group (on a V-Series appliance, eth0 = P1; eth1 = P2) |
| WCCP Routers: GRE Tunnel Next Hop Router IP Address | Specify a GRE Tunnel Next Hop Router IP Address (must be in IPv4 format) when GRE Packet Return Method is configured and Content Gateway does not have a route back to the WCCP router. You can use “ping” to test connectivity to the router. |

DNS Proxy

Help | Content Gateway | Version 7.8.x



Note

The DNS Proxy configuration options appear on the Configure pane only if you have enabled DNS Proxy in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Networking > DNS Proxy

| | |
|----------------|---|
| DNS Proxy Port | Specifies the port that Content Gateway uses for DNS traffic. The default port is 5353. |
|----------------|---|

DNS Resolver

Help | Content Gateway | Version 7.8.x

Configure > Networking > DNS Resolver > Resolver

| | |
|---------------------------|--|
| Local Domain Expansion | Enables or disables local domain expansion so that Content Gateway can attempt to resolve unqualified hostnames by expanding to the local domain. For example, if a client makes a request to an unqualified host named <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> . |
| DNS Preference | Specifies the IP version preference when IPv6 support is enabled in Content Gateway and a web server supports both IPv4 and IPv6. Select IPv4 to cause the proxy to prefer IPv4. Select IPv6 to cause the proxy to prefer IPv6. |
| DNS Preference Exceptions | List IPv4 / IPv6 preference rules for specific origin servers. |
| Refresh | Updates the table to display the most up-to-date rules. Click this button after you have added or modified rules with the configuration file editor. |
| Edit File | Opens the configuration file editor. dns_prefer_exception.config File Editor |
| rule display box | Displays an ordered list of the dns_prefer_exception.config file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a new rule to the rule display box. Enter information in the fields provided before clicking this button. |
| Set | Updates the selected rule with the values in the entry fields. |
| Name | Specify a unique name to aid in administering rules. |
| Destination Host | Specify the destination hostname. |
| Preferred Format | Specify the preferred IP version, IPv4 or IPv6. |
| Apply | Applies the configuration changes. |
| Close | Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes are lost. |

Configure > Networking > DNS Resolver > Host Database

These settings pertain to all DNS name resolution performed by Content Gateway, including DNS Proxy.

| | |
|--------------------|--|
| DNS Lookup Timeout | <p>Specifies the maximum number of seconds the proxy can wait for a lookup response from the DNS server.</p> <p>Specifies how long, in seconds, the proxy will wait before making a second DNS request if there is no response to the first request. The value is stored in “proxy.config.hostdb.lookup_timeout”. The default value is 120 seconds.</p> <p>Important: This setting is not used. Instead the records.config entry “proxy.config.dns.lookup_timeout” is used. The default value is 20 seconds.</p> <p>proxy.config.dns.lookup_timeout specifies how long the proxy will wait for the DNS response after sending the request.</p> |
| Foreground Timeout | <p>Specifies how long DNS entries remain in the host database before they are flagged as stale. This setting is used only when “proxy.config.hostdb.ttl_mode” is not zero (the default value is 0, which means use the time-to-live (ttl) value set by the DNS server. See HostDB, page 466.</p> <p>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.</p> <p>The default is 86400 seconds (144 minutes).</p> <p>Caution: Setting the foreground timeout too low might slow response time. Setting it too high risks accumulation of incorrect information.</p> |
| Failed DNS Timeout | <p>Specifies how long, in seconds, that a hostname is retained in the failed DNS lookup cache (default = 60). When the timeout expires, the hostname is removed from the cache and the next request for that hostname is sent to the DNS server.</p> <p>A DNS lookup failure is considered to have occurred when:</p> <ul style="list-style-type: none"> • There is no DNS response • There is a DNS response error code, including NXDOMAIN • There is an error parsing the DNS response code (there is a malformed response). <p>Zero (0) is not a legal value.</p> |

Configure > Networking > DNS Resolver > Split DNS

| | |
|---------------------------|---|
| Split DNS | Enables or disables the Split DNS option. When enabled, Content Gateway can use multiple DNS servers, depending on your security requirements. For example, you can configure the proxy to look to one set of DNS servers to resolve hostnames on your internal network, while allowing DNS servers outside the firewall to resolve hosts on the Internet. For information about using Split DNS, see <i>Using the Split DNS option</i> , page 188. |
| Default Domain | Specifies the default domain used for split DNS requests. If a hostname does not include a domain, Content Gateway appends the default domain name to the hostname before choosing which DNS server to use. |
| DNS Servers Specification | Displays a table listing the rules in the <i>splitdns.config</i> file that control which DNS server the proxy uses for resolving hosts under specific conditions. |
| Refresh | Updates the table to display the most up-to-date rules in the splitdns.config file. Click this button after you have added or modified rules with the configuration file editor. |
| Edit File | Opens the configuration file editor so that you can edit and add rules to the splitdns.config file. The configuration file editor page is described below. |
| | splitdns.config Configuration File Editor |
| rule display box | Lists the <i>splitdns.config</i> file rules. Select a rule to edit it. The buttons on the left of the box allow you to delete or move the selected rule up or down in the list. |
| Add | Adds a new rule to the rule display box at the top of the configuration file editor page. Enter information in the fields provided before you click this button. |
| Set | Updates the rule display box at the top of the configuration file editor page. Select a rule and change its properties before you click this button. |
| Primary Destination Type | Specifies that DNS server selection is based on the destination domain (dest_domain), destination host (dest_host), or on a regular expression (url_regex). |
| Primary Destination Value | Specifies the value of the primary destination. Place the symbol “!” at the beginning of the value to specify the NOT logical operator. |
| DNS Server IP | Specifies the DNS server to use with the primary destination specifier. You can specify a port using a colon (:). If you do not specify a port, 53 is used. You can specify multiple DNS servers separated by spaces or by semicolons (;). |

| | |
|--------------------------------|---|
| Default Domain Name (Optional) | Specifies the default domain name to use for resolving hosts. Only one entry is allowed. If you do not provide the default domain, the system determines its value from /etc/resolv.conf . |
| Domain Search List (Optional) | Specifies the domain search order. You can specify multiple domains separated by spaces or by semicolons (;). If you do not provide the search list, the system determines the value from /etc/resolv.conf . |
| Apply | Applies the configuration changes. |
| Close | Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes are lost. |

ICAP

Help | Content Gateway | Version 7.8.x



Note

The ICAP configuration option appears on the Configure pane only if you have enabled **ICAP** in the **Features** table on the **Configure > My Proxy > Basic > General** tab.

ICAP provides an alternate interface to Websense Data Security and other data security services that are ICAP-conversant. A primary and backup URI can be specified, and failover and load balancing can be configured. See [Configuring the ICAP client, page 135](#) and the subsection for [ICAP failover and load balancing, page 136](#).

Configure > Networking > ICAP

| | |
|-----------------------|---|
| 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. An optional secondary URI service can be specified immediately after the first by adding a comma and the URI of the second service, no spaces. |
| Analyze HTTPS Content | Select whether decrypted traffic should be sent to 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 326 . |

| | |
|---------------------------------|---|
| 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. |
| Action for Large files | Select whether to allow traffic or send a block page if a file larger than the size limit specified in Data Security is sent. The default size limit in Data Security is 12 MB. |

Virtual IP

Help | Content Gateway | Version 7.8.x



Note

The Virtual IP configuration options appear on the Configure pane only if you have enabled Virtual IP in the Features table on the **Configure > My Proxy > Basic > General** tab.

Configure > Networking > Virtual IP

| | |
|----------------------|---|
| Virtual IP Addresses | Displays a table listing the virtual IP addresses managed by Content Gateway. |
| Refresh | Updates the table to display the most up-to-date list of virtual IP addresses. Click this button after you have added to or modified the list of virtual IP addresses with the configuration file editor. |
| Edit File | Opens the configuration file editor so that you can edit and add to the list of virtual IP addresses. |
| | vaddrs.config Configuration File Editor |
| rule display box | Lists the virtual IP addresses. Select a virtual IP address to edit it. The buttons on the left of the box allow you to delete or move the selected virtual IP address up or down in the list. |
| Add | Adds a new virtual IP address to the rule display box at the top of the configuration file editor page. |
| Set | Updates the rule display box at the top of the configuration file editor page. |
| Virtual IP Address | Specifies the virtual IP address managed by Content Gateway. |
| Ethernet Interface | Specifies the network interface assigned to the virtual IP address. |
| Sub-Interface | Specifies the subinterface ID. This is a number between 1 and 255 that the interface uses for the address. |

| | |
|-------|---|
| Apply | Applies the configuration changes. |
| Close | Exits the configuration file editor. Click Apply before you click Close ; otherwise, all configuration changes will be lost. |

Health Check URLs

Help | Content Gateway | Version 7.8.x

Content Gateway includes 3 URLs that return proxy health and performance information in the HTTP response. These URLs are designed to help load balancers optimize performance by acquiring and adjusting for real-time state information of each proxy node.

The default port for health check URLs is 8083. The value can be changed in records.config by assigning the desired value to **proxy.config.admin.autoconf_port**

Configure > Networking > Health Check URLs

Force Health Checks to Report Proxy Down

| | |
|---|--|
| Enable/Disable | <p>When enabled, all health check URLs sent to this proxy report WSDOWN.</p> <p>The URL response will be similar to:</p> <pre>HTTP/1.0 503 Service Unavailable Server: Content Gateway Manager 7.7.0 Date: Thu, 26 Jul 2012 20:26:14 GMT Cache-Control: no-store Pragma: no-cache Content-type: text/plain Content-length: 6 WSDOWN</pre> |
| Health Check URLs | <p>The load balancer should consider the service down if the URL request fails for the following reasons:</p> <ul style="list-style-type: none"> • No TCP connection -- proxy down • Response too slow -- proxy deadlocked or not responsive • Invalid response |
| http://[Content Manager IP address]:8083/health.basic | Checks connectivity with Content Gateway and responds with WSUP or WSDOWN. |

| | |
|---|---|
| http://[Content Manager IP address]:8083/health.app.filtering | Checks the health of Filtering Service responses to Content Gateway requests and reports WSUP or WSDOWN. |
| http://[Content Manager IP address]:8083/health.load | <p>If the health.basic URL reports WSDOWN, this URL also reports WSDOWN.</p> <p>Otherwise, health.load returns:</p> <ul style="list-style-type: none"> • CPU usage (operating system load average) • Connection usage (number of open connections) • Bandwidth usage <p>How these values are calculated and how they can be customized is described below.</p> <p>The default response will look similar to:</p> <pre> HTTP/1.0 200 OK Server: Content Gateway Manager 7.7.0 Date: Thu, 26 Jul 2012 20:26:14 GMT Cache-Control: no-store Pragma: no-cache Content-type: text/plain Content-length: xx Load=2253 Conns=5150 Mbps=6.42 </pre> |

A format file, `/opt/WCG/config/health.load.template`, allows for customization of the response format.

Format specifiers are:

- %L = Load (integer)
- %C = Connections integer
- %B = Bandwidth in Mbps (double)
- %% = %

The default **health.load.template** file is:

```

Load=%L
Conns=%C
Mbps=%B

```

Here is **health.load.template** modified to respond with an xml-like format:

```

<load>
<item name="Load" value="%L" />
<item name="Conns" value="%C" />

```

```
<item name="Mbps" value="%B" />
```

```
</load>
```

How the values are calculated:

The **Load** value, **%L**, is derived from the LINUX system load average. To make the value comparable across machines with varying numbers of cores, the number is divided by the number of cores on the system.

The calculation is:

```
// load avg values are 0.00 precision
double avgs[3];
// get load averages for 1, 5, and 15 minutes
getloadavg(avgs, 3);
// 5 minute_load_average * 10000 / number_of_cores
Load = avgs[1] * 10000 / get_nprocs();
```

The **Connection** value, **%C**, is the sum of `proxy.process.http.current_server_connections` and `proxy.process.http.current_client_connections`.

The **Bandwidth** value, **%B**, is the value of `proxy.node.client_throughput_out`.

**Note**

HTTP connection and bandwidth information can be viewed in the Content Gateway manager on the **Monitor > Protocols > HTTP** page.

SSL

Help | Content Gateway | Version 7.8.x

The SSL configuration options are divided into the following categories:

- ◆ Certificates (see [Managing certificates](#), page 154)
- ◆ Decryption/Encryption (see [SSL configuration settings for inbound traffic](#), page 156 and [SSL configuration settings for outbound traffic](#), page 158)
- ◆ Validation (see [Validating certificates](#), page 159)
- ◆ Incidents (see [Managing HTTPS website access](#), page 164)
- ◆ Client certificates (see [Client certificates](#), page 169)
- ◆ Customization (see [Customizing SSL connection failure messages](#), page 171)
- ◆ Internal Root CA (see [Internal Root CA](#), page 145)

D

Event Logging Formats

Help | Content Gateway | Version 7.8.x

Custom logging fields

Related topic:

- ◆ [Logging format cross-reference, page 384](#)

| %<field symbol> | Description |
|---------------------------------------|---|
| <i>{HTTP header field name}</i> cqh | 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. This field cannot be used in custom log filters. |
| <i>{HTTP header field name}</i> cqhua | Logs the information in the requested field of the client request HTTP header; for example, %<{User-Agent}cqhua> logs the User-Agent: field in client request headers. |
| <i>{HTTP header field name}</i> pqh | 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. This field cannot be used in custom log filters. |
| <i>{HTTP header field name}</i> psh | 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. This field cannot be used in custom log filters. |
| <i>{HTTP header field name}</i> ssh | 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. This field cannot be used in custom log filters. |

| %<field symbol> | Description |
|------------------------------|---|
| caun | The client authenticated user name; result of the RFC931/ <i>ident</i> 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. |
| cqhm | The HTTP method in the client request to Content Gateway: GET, POST, and so on (subset of <i>cqtx</i>). |
| 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 <code>http://www.company.com</code> HTTP/1.0 |
| cqu | The client request URI; universal resource identifier (URI) of the request from client to Content Gateway (subset of <i>cqtx</i>). |
| cquc | The client request canonical URL; differs from <i>cqu</i> 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.). |

| %<field symbol> | Description |
|------------------------------|--|
| 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. |
| 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 | <p>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.</p> <p>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.</p> |
| shn | The host name of the origin server. |
| sscl | The server response transfer length; response length, in bytes, from origin server to Content Gateway. |
| sshl | The server response header length; the header length in the origin server's response to Content Gateway in bytes. |
| sshv | The server response HTTP version (1.0, 1.1, and so on). |

| %<field symbol> | Description |
|------------------------------|--|
| <code>sssc</code> | The server response status code; the HTTP response status code from origin server to Content Gateway. |
| <code>ttms</code> | 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. |
| <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>tts</code> field displays 1. |
| <code>tts</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. |
| <code>wc</code> | The predefined or custom category of the URL for the data being scanned. For example, "News and Media". |
| <code>wct</code> | The content type of the web page. For example, "text/html; charset=UTF-8". |
| <code>wds</code> | The scan disposition string such as CATEGORY_BLOCKED, PERMIT_ALL, FILTERED_AND_PASSED, etc. |
| <code>wsr</code> | The scan recommended bit ("true" or "false"). The URL database identifies and recommends data that should be analyzed further. Depending on the policy used, the data may or may not be analyzed further. |
| <code>wstms</code> | The scan time in milliseconds that it took to scan a downloaded file or page. |
| <code>wui</code> | The authenticated user's ID used to select the policy for scanning data of the client request. |

Logging format cross-reference

Help | Content Gateway | Version 7.8.x

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 |
|----------------|-------------------------------|
| time | cqts |
| elapsed | ttms |
| client | chi |
| action/code | cr/pssc |
| size | psql |
| method | cqhm |
| url | cquc |
| ident | caun |
| hierarchy/from | phr/pqsn |
| content | psct |

For example, if you want to create a custom format called `short_sq` based on the first three Squid fields, enter a line in the **logs.config** file as follows:

```
format:enabled:1:short_sq:%<cqts> %<ttms>
%<chi>:short_sq:ASCII:none
```

See [Custom format, page 246](#), for more information about defining custom log files.

Netscape Common logging formats

| Netscape Common | Content Gateway Field Symbols |
|-----------------|-------------------------------|
| host | chi |
| usr | caun |
| [time] | [cqt] |
| “req” | “cqt” |
| s1 | pssc |
| c1 | pscl |

Netscape Extended logging formats

| Netscape Extended | 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 |

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 |

| Netscape Extended-2 | Content Gateway Field Symbols |
|----------------------------|--------------------------------------|
| h1 | cqhl |
| h2 | pshl |
| h3 | pqhl |
| h4 | sshl |
| xt | tts |
| route | phr |
| pfs | cfsc |
| ss | pfsc |
| crc | crc |

E

Configuration Files

Help | Content Gateway | Version 7.8.x

Websense Content Gateway contains the following configuration files that you can edit to customize the proxy.

- ◆ [auth_domains.config](#), page 391
- ◆ [auth_rules.config](#), page 393
- ◆ [bypass.config](#), page 395
- ◆ [cache.config](#), page 397
- ◆ [filter.config](#), page 400
- ◆ [hosting.config](#), page 403
- ◆ [ip_allow.config](#), page 405
- ◆ [ipnat.conf](#), page 406
- ◆ [log_hosts.config](#), page 407
- ◆ [logs_xml.config](#), page 408
- ◆ [mgmt_allow.config](#), page 415
- ◆ [parent.config](#), page 416
- ◆ [partition.config](#), page 419
- ◆ [records.config](#), page 420
- ◆ [remap.config](#), page 484
- ◆ [socks.config](#), page 486
- ◆ [socks_server.config](#), page 488
- ◆ [splitdns.config](#), page 489
- ◆ [storage.config](#), page 491
- ◆ [update.config](#), page 491
- ◆ [wccp.config](#), page 493

Specifying URL regular expressions (`url_regex`)

Help | Content Gateway | Version 7.8.x

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.

auth_domains.config

Help | Content Gateway | Version 7.8.x

The **auth_domains.config** file stores the list of domains that have been identified for use with [Rule-Based Authentication](#), page 211.

Domains must be identified (added to this file) using the interface in the Content Gateway manager on the **Configure > Security > Access Control > Domains** tab. Do not edit this configuration file.

Format

Each line in **auth_domains.config** consists of a set of tags; each tag is followed by its value. For example:

```
type=<auth_method> name=<unique_name> use_alias=<0 or 1> <additional tags>
```

The set of tags varies depending on the selected authentication method.

The following table lists all of the tags.

| Tag | Allowed value |
|-----------|---|
| type | Specifies the authentication method: IWA, NTLM, LDAP |
| name | Specifies a unique name for the domain. This is not the actual domain name, but rather a name that is unique to the proxy and rule-based authentication. |
| use_alias | Specifies the user name sent to filtering service if authentication is successful. <ul style="list-style-type: none"> • 0 = send actual authenticated user name (default). • 1 = send a blank username • 2 = send the string specified in auth_name_string |
| alias | Only active if use_alias=2. Specifies the static string to send as the user name for all successful authentications using this rule. |

The following table lists the additional tags used with IWA domains.

| IWA Tag | Allowed Value |
|---------------|---|
| winauth_realm | Specifies the joined Windows domain to use with the rule. Content Gateway must be joined and active in that domain. |

The following table lists the additional tags used with NTLM domains.

| NTLM Tag | Allowed Value |
|----------------------------|---|
| dc_list | Takes the IP address and port number of the primary domain controller (if no port is specified, Content Gateway uses port 139), followed by a comma separated list of secondary domain controllers to be used for load balancing and failover. |
| 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 additional tags used with LDAP domains.

| LDAP Tag | Allowed Value |
|------------------------|---|
| server_name | Specifies the fully qualified domain name of the LDAP server. |
| server_port (optional) | Specifies the LDAP server port. The default is 389. To use the default Global Catalog server port, specify port 3268. If Secure LDAP is enabled, set the port to 636 or 3269 (the secure LDAP ports). |
| base_dn (optional) | Specifies the LDAP base distinguished name. |
| uid_filter (optional) | Specifies the type of service, if different from that configured on the LDAP tab. Enter sAMAccountName 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). |

auth_rules.config

Help | Content Gateway | Version 7.8.x

The **auth_rules.config** file stores rules that direct specified IP addresses and IP address ranges, and/or traffic on specified inbound ports (explicit proxy only), and/or matching Request header User-Agent values to authenticate with distinct domain controllers. One or more domain controllers can be specified in an ordered list. This feature is called *Rule-Based Authentication*, page 211.

Rule-based authentication rules must be defined in the Content Gateway manager on the **Configure > Security > Access Control > Authentication Rules** tab. Do not edit this configuration file.

- ◆ Rule-based authentication is supported for Integrated Windows Authentication (IWA), legacy NTLM, and LDAP authentication only.

- ◆ Each authentication rule can specify source IP addresses, inbound port (explicit proxy only), and/or a User-Agent regex
- ◆ Each authentication rule can specify one or more domains in an ordered list. Domains are identified on the **Configure > Security > Access Control > Authentication Rules** tab. That process includes specifying the authentication method (IWA, Legacy NTLM, LDAP).
- ◆ When a rule matches, authentication is performed against one or more domains in the ordered list. The first successful authentication ends domain list traversal and the authenticating domain is cached for later use.
- ◆ Authentication rules are applied from the list top-down; only the first match is applied. If no rule matches, no user authentication is performed.



Note

If all the users in your network can be authenticated by domain controllers that share trust relationships, you probably don't need rule-based authentication.

However, rule-based authentication can be useful in any deployment that needs to perform special authentication handling based on IP address, inbound proxy port (explicit proxy), and/or User-Agent values.

Format

Each line in **auth_rules.config** contains an authentication rule that consists of a set of tags, each followed by its value. Authentication rules have the format:

rule_name=<name> src_ip=<IP addresses> user_agent=<regex> <additional tags>

The following table lists all of the tags.

| Tags | Allowed value |
|-----------------------|--|
| rule_name | A short, unique name. |
| enabled | Specifies whether the rule will be active: <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled |
| src_ip | Takes a comma separated list of IP addresses and IP address ranges. No spaces. If this field is empty, all IP addresses match. |
| user_agent (optional) | Takes a regular expression that is applied to the user-agent string. |
| proxy_port (optional) | Takes a port number. Valid with explicit proxy only. Client applications must be configured to send requests to the correct port. |

| Tags | Allowed value |
|--------------------|---|
| domain_list | An ordered, comma separated list of domains the Content Gateway will attempt to authenticate a matching user with. |
| use_captive_portal | Specifies whether Captive Portal (available beginning with 7.8.3) is used. <ul style="list-style-type: none"> • 0 = disabled • 1 = enabled using HTTP • 2 = enabled using HTTPS (available beginning with 7.8.4) |

bypass.config

[Help](#) | [Content Gateway](#) | Version 7.8.x

The **bypass.config** file contains *static* bypass rules that Content Gateway uses in transparent proxy mode. Static bypass rules instruct Content Gateway to bypass certain incoming client requests so that they are served by the origin server.

The **bypass.config** file also accepts *dynamic* deny bypass rules. See [Dynamic deny bypass rules](#), page 396.

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 |
|--|--|
| <i>src ipaddress</i> | <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 |
| <i>dst ipaddress</i> | <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 |
| <i>src ipaddress</i> AND <i>dst ipaddress</i> | <p>Specifies the source and destination IP address pair that the proxy must bypass.</p> <p><i>ipaddress</i> can be a single IP address, an IP address range, or a combination of both separated by commas</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 396](#).



Note

For the dynamic deny bypass rules to work, you must enable the **Dynamic Bypass** option in the 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

Help | Content Gateway | Version 7.8.x

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 (`/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> . |
| tll-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

[Help](#) | [Content Gateway](#) | Version 7.8.x

Filtering rules stored in **filter.config** allow you to:

- ◆ Deny or allow URL requests
- ◆ Keep or strip header information from client requests
- ◆ Insert custom headers
- ◆ Allow specified applications or requests to specified web sites to bypass authentication
- ◆ Prevent specified applications from transiting the proxy

Filtering rules should be defined in the Content Gateway manager on the **Configure > Security > Access Control > Filtering** tab. See [Creating filtering rules, page 181](#).



Important

After you modify the file, run `content_line -x` from the Content Gateway **bin** directory (`/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.

Three filtering rules are configured by default. The first denies traffic on port 25 to all destinations. The second and third bypass user authentication for connections to 2 ThreatScope destinations.

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 destination types.

| Primary Destination Type | Allowed Value |
|--------------------------|---|
| dest_domain | A requested domain name |
| dest_host | A requested hostname |
| dest_ip | A requested IP address |
| url_regex | A regular expression to be found in a URL |

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 . |
| add_hdr | The custom header value you want to add. Requires specification of the custom header and a header value. For example: <pre>add_hdr="header_name:header_value"</pre> |

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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

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 102.



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 102. For a description of the **partition.config** file, see [partition.config](#), page 419.

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 419).



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
```

ip_allow.config

Help | Content Gateway | Version 7.8.x

The **ip_allow.config** file controls client access to the proxy. You can specify ranges of IP addresses that are allowed to use Content Gateway.



Important

After you modify the file, run `content_line -x` from the Content Gateway **bin** directory (`/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.

The action `ip_allow` allows the specified clients to access the proxy.

The action `ip_deny` denies the specified clients to access the proxy.

By default, the **ip_allow.config** file contains the following line, which allows all clients to access the proxy. Comment out or delete this line before adding rules to restrict access.

```
src_ip=0.0.0.0-255.255.255.255 action=ip_allow
```

Examples

The following example allows all clients to access the proxy:

```
src_ip=0.0.0.0-255.255.255.255 action=ip_allow
```

The following example allows all clients on a specific subnet to access the proxy:

```
src_ip=123.12.3.000-123.12.3.123 action=ip_allow
```

The following example denies all clients on a specific subnet to access the proxy:

```
src_ip=123.45.6.0-123.45.6.123 action=ip_deny
```

ipnat.conf

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 254](#)).



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 (`/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

Help | Content Gateway | Version 7.8.x

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 409](#).
- ◆ `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 410](#).
- ◆ `LogObject` specifies an object that contains a particular format, a local filename, filters, and collation servers. See [LogObject, page 411](#).



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 predefined 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 381 . The specified field can be of two types: Simple: for example, %<cpu> 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 384.</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>.</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 the 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 the 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 the 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

[Help](#) | [Content Gateway](#) | Version 7.8.x

The `mgmt_allow.config` file specifies the IP addresses of remote hosts allowed access or denied access to the Content Gateway manager.



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (`/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 the Content Gateway manager.

action must specify either `ip_allow` to allow access to the 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 the Content Gateway manager. Comment out or delete this line before adding rules to restrict access.

```
src_ip=0.0.0.0-255.255.255.255 action=ip_allow
```

Examples

The following example configures Content Gateway to allow only one user to access the Content Gateway manager:

```
src_ip=123.12.3.123 action=ip_allow
```

The following example configures Content Gateway to allow a range of IP addresses to access the Content Gateway manager:

```
src_ip=123.12.3.000-123.12.3.123 action=ip_allow
```

The following example configures Content Gateway to deny the IP address 123.45.67.8 access to the Content Gateway manager:

```
src_ip=123.45.67.8 action=ip_deny
```

parent.config

Help | Content Gateway | Version 7.8.x

The **parent.config** file identifies the HTTP parent proxies used in an HTTP cache hierarchy. Use this file to perform the following configuration:

- ◆ Set up parent cache hierarchies, with multiple parents and parent failover
- ◆ Configure selected URL requests to bypass parent proxies

Rules are applied from the list top-down; the first match is applied. Bypass rules are usually placed above parent proxy designation rule(s).

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 96.



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (`/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
```

The following rule is a typical destination bypass rule:

```
dest_domain=example.com go_direct=true
```



Important

Every line in the `parent.config` file must contain *either* a `parent=` or `go_direct=` directive.

A bypass rule that includes `parent=` *and* `go_direct=true`, causes the specified `dest_domain` to be sent to the parent while all other domains are bypassed (the opposite of the usual intended action).

partition.config

Help | Content Gateway | Version 7.8.x

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

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The **records.config** file is a list of configurable variables used by Content Gateway.

Most values are set using controls in the Content Gateway manager. Some options can be set only by editing variables in the **records.config** file.



Warning

Do not change the **records.config** variables unless you are certain of the effect. Many variables are coupled, meaning that they interact with other variables. Changing a single variable in isolation can cause Content Gateway to fail.

Whenever possible, use the Content Gateway manager to configure Content Gateway.



Important

After you modify this file, to apply the changes run `content_line -x` from the Content Gateway **bin** directory (`/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.winauth.enabled** is a yes/no flag. A value of 0 (zero) disables the option. A value of 1 enables the option.

```
CONFIG proxy.config.winauth.enabled INT 0
```

In the following example, the variable sets the cluster startup timeout to 10 seconds.

```
CONFIG proxy.config.cluster.startup_timeout INT 10
```

Configuration variables

[Help](#) | [Content Gateway](#) | [Version 7.8.x](#)

The following tables describe the configuration variables listed in the **records.config** file.

System variables

Local manager

Virtual IP manager

Alarm configuration

ARM

Load shedding configuration (ARM)

Authentication basic realm

LDAP

RADIUS authentication

NTLM

Integrated Windows Authentication

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

DNS proxy

HostDB

Logging configuration

URL remap rules

Scheduled update configuration

WCCP configuration

SSL Decryption

ICAP

Connectivity, analysis, and boundary conditions

System variables

Help | Content Gateway | Version 7.8.x

| 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 | 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. |

| Configuration Variable Data Type | Data Type | Default Value | Description |
|---|--------------|------------------------------|--|
| <code>proxy.config.cli_binary</code> | STRING | <code>content_line</code> | Specifies the name of the executable that runs the content_line interface. |
| <code>proxy.config.watch_script</code> | STRING | <code>content_cop</code> | Specifies the name of the executable that runs the content_cop process. |
| <code>proxy.config.env_prep</code> | STRING | <code>example_prep.sh</code> | Specifies the script that is executed before the content_manager process spawns the content_gateway process. |
| <code>proxy.config.config_dir</code> | STRING | <code>config</code> | Specifies the directory, relative to <code>bin_path</code> (above), that contains the Content Gateway configuration files. |
| <code>proxy.config.temp_dir</code> | STRING | <code>/tmp</code> | Specifies the directory used for Content Gateway temporary files |
| <code>proxy.config.alarm_email</code> | STRING | <code>websense</code> | 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. |
| <code>proxy.config.syslog_facility</code> | STRING | <code>LOG_DAEMON</code> | Specifies the facility used to record system log files. See <i>Working With Log Files</i> , page 241. |
| <code>proxy.config.cop.core_signal</code> | 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. |
| <code>proxy.config.cop.sleep_time</code> | 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. |
| <code>proxy.config.cop.linux_min_swapfree_kb</code> | INT | 10240 | This variable is not used. |

| Configuration Variable Data Type | Data Type | Default Value | Description |
|---|--------------|---------------------|---|
| 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. |
| proxy.config.output.logfile.log_dir_usage_percent | INT | 35 | Specifies the percentage of space allocated by proxy.config.log2.max_space_mb_for_logs , that can be used for logs in /opt/WCG/logs/ except for content_gateway.out . Content_gateway.out can use up to the log directory limit. |
| proxy.config.snapshot_dir | 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. |
| proxy.config.attach_debugger_script | 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. |
| proxy.config.healthcheck_force_offline | INT | 0 | When enabled (1), forces URL health checks to report proxy down. See, Health Check URLs , page 378. |

Local manager

Help | Content Gateway | Version 7.8.x

| 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"> • 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. |
| <code>proxy.config.cluster.log_bogus_mc_msgs</code> | INT | 1 | Enables (1) or disables (0) logging of invalid multicast messages. |
| <code>proxy.config.admin.html_doc_root</code> | STRING | ui | Specifies the document root for the 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 | -1 | Specifies the port used for retrieving and setting statistics and configuration variables. This port is disabled by default. |
| <code>proxy.config.admin.admin_user</code> | STRING | admin | Specifies the administrator ID that controls access to the Content Gateway manager. |

| Configuration Variable | Data Type | Default Value | Description |
|--------------------------------------|-----------|---------------|---|
| proxy.config.admin.admin_password | STRING | | Specifies the encrypted administrator password that controls access to the Content Gateway manager. You cannot edit the password; however, you can specify a value of NULL to clear the password. <i>See Accessing the Content Gateway manager if you forget the master administrator password, page 15.</i> |
| proxy.config.admin.basic_auth | INT | 1 | Enables (1) or disables (0) basic user authentication to control access to the Content Gateway manager. <i>Note: If basic authentication is not enabled, any user can access the Content Gateway manager to monitor and configure Content Gateway.</i> |
| proxy.config.admin.use_ssl | INT | 1 | Enables the Content Gateway manager SSL option for secure communication between a remote host and the Content Gateway manager. |
| proxy.config.admin.ssl_cert_file | STRING | server.pem | Specifies the filename of the SSL certificate installed on the Content Gateway system for secure communication between a remote host and the Content Gateway manager. |
| 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 | root | 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 the 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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|--|---|
| <code>proxy.config.admin.log_resolve_hostname</code> | INT | 1 | When enabled (1), the hostname of the client connecting to the Content Gateway manager is recorded in the lm.log file. When disabled (0), the IP address of the client connecting to the Content Gateway manager is recorded in the lm.log file. |
| <code>proxy.config.admin.subscription</code> | STRING | NULL | Not used. |
| <code>proxy.config.admin.supported_cipher_list</code> | STRING | AES128-SHA, DHE-RSA-AES128-SHA, DHE-DSS-AES128-SHA, DES-CBC3-SHA, EDH-RSA-DES-CBC3-SHA, EDH-DSS-DES-CBC3-SHA | A comma-separated list, no spaces, of ciphers supported by Content Gateway. No validation is performed on the string. |
| <code>proxy.config.lm.display_reset_alarm</code> | INT | 0 | When enabled (1), email is sent to the administrator (proxy.config.alarm_email) whenever Content Gateway resets. Default is 0. |

Process manager

[Help](#) | [Content Gateway](#) | Version 7.8.x

| 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

[Help](#) | [Content Gateway](#) | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|---|
| <code>proxy.config.arm.enabled</code> | INT | 1 | Specifies that the ARM is enabled or disabled. Warning: Do not disable the ARM. In all deployments, it must be running to support proper proxy function. |
| <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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|---|
| <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 in version 7.8.1, IP addresses are logged, instead of domain names.</p> <p>When enabled in version 7.8.2 and later, domain names are logged, instead of IP addresses, unless <code>proxy.config.arm.use_hostname_for_wisp_and_reporting</code> (see below) is disabled.</p> <p>When disabled, domain names are logged. See Reducing DNS lookups, page 74 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.arm.use_hostname_for_wisp_and_reporting</code> | INT | 1 | <p>Enables (1) or disables (0) the ability to capture hostname (instead of IP address) when Always Query Destination is enabled for transparent proxy deployments. See preceding entry.</p> <p>Note: This variable must be manually added to the config file and is available for version 7.8.2 and later.</p> |
| <code>proxy.config.http.outgoing_ip_spoofing_enabled</code> | INT | 0 | <p>Enables (1) or disables (0) the IP spoofing option that allows Content Gateway to establish connections to origin servers with the client IP address instead of the Content Gateway IP address.</p> <p>See IP spoofing, page 77.</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 71 . |
| <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)

Help | Content Gateway | Version 7.8.x

| 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. |
| <code>proxy.config.http.client.connection_control.enabled</code> | INT | 1 | Disables (0) or enables (1) the ability to limit the number of connections from a single computer. |
| <code>proxy.config.http.client.concurrent_connection_control.close.enabled</code> | INT | 1 | Disables (0) or enables (1) closing connections on reaching the concurrent connection limit. |
| <code>proxy.config.http.client.concurrent_connection_control.alert.enabled</code> | INT | 0 | Disables (0) or enables (1) alerting on violation of the concurrent connection limit. |
| <code>proxy.config.http.client.concurrent_connection_control.max_connections</code> | INT | 1000 | Configures the maximum number of concurrent connections allowed from one client IP address. |
| <code>proxy.config.http.client.connection_rate_control.close.enabled</code> | INT | 0 | Disables (0) or enables (1) closing connections on reaching the connection rate limit. |
| <code>proxy.config.http.client.connection_rate_control.alert.enabled</code> | INT | 1 | Disables (0) or enables (1) alerting on exceeding the connection rate limit. |
| <code>proxy.config.http.client.connection_rate_control.second</code> | INT | 100 | Configures the maximum connections per second allowed from one client IP. |
| <code>proxy.config.http.client.connection_control.exceptions</code> | STRING | NULL | Specifies a comma separated list of IP addresses for which the connection limits do not apply. |

Authentication basic realm

Help | Content Gateway | Version 7.8.x

| 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 | 0 | Specifies the type of client authentication. <ul style="list-style-type: none"> • 0 = None • 1 = LDAP • 2 = RADIUS • 3 = Legacy NTLM • 4 = Integrated Window Authentication • 5 = Rule-Based Authentication |
| <code>proxy.config.multiauth.enabled</code> | INT | 0 | Enables (1) or disables (0) rule-based authentication. Tells Content Gateway to use the <code>auth_rules.config</code> file. |
| <code>proxy.config.multiauth.domain.max</code> | INT | 50 | Beginning with 7.8.3, specifies the maximum number of domains that can be added or joined on Configure > Security > Access Control > Domains |
| <code>proxy.config.auth.form_filename</code> | STRING | <code>auth_form.html</code> | Beginning with 7.8.3, specifies the file that defines the Captive Portal authentication page. This variable must be added manually. Changing this filename is not recommended. |
| <code>proxy.config.internal.file.path</code> | STRING | <code>/opt/WCG/config/ui_files</code> | Beginning with 7.8.3, specifies the location of any css and image files used to define the Captive Portal authentication page. The full default path is <code>/opt/WCG/config/ui_files</code> . Image files are located in an <code>/images</code> sub-directory. This variable must be added manually. |
| <code>proxy.config.ssl.auth_server_port</code> | INT | 4443 | Beginning with 7.8.4, specifies the local port used for the HTTPS Captive Portal page. |

LDAP

Help | Content Gateway | Version 7.8.x

| 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 LDAP authentication , page 204. |
| <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. |
| <code>proxy.config.ldap.proc.encode_convert</code> | INT | 0 | Enables (1) or disables (0) the support of passwords with special characters. The variable <code>proxy.config.ldap.proc.encode_name</code> is required when this variable is enabled. This variable must be added manually. See this page for additional information. |
| <code>proxy.config.ldap.proc.encode_name</code> | STRING | NULL | Specifies the encoding name to be used when <code>proxy.config.ldap.proc.encode_convert</code> is enabled. This variable must be added manually. |

RADIUS authentication

Help | Content Gateway | Version 7.8.x

| 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. |
| <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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|---|
| <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

Help | Content Gateway | Version 7.8.x

| 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. |

| Configuration Variable | Data Type | Default Value | Description |
|--------------------------------------|-----------|---------------|--|
| proxy.config.ntlm.dc.load_balance | INT | 0 | <p>Enables (1) or disables (0) 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> |
| proxy.config.ntlm.dc.max_connections | INT | 10 | Specifies the maximum number of connections Content Gateway can have open to the domain controller. |
| proxy.config.ntlm.cache.enabled | INT | 1 | <p>Enables (1) or disables (0) the NTLM cache. Applies only when Content Gateway is an explicit proxy.</p> <p>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.</p> |
| 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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <code>proxy.config.ntlm.cache.storage_size</code> | 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. |
| <code>proxy.config.ntlm.cache_exception.list</code> | 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. |
| <code>proxy.config.ntlm.fail_open</code> | 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. |
| <code>proxy.config.ntlm.check_account_passwd</code> | INT | 0 | Beginning with 7.8.3, enables (1) or disables (0) whether Content Gateway will create a log file entry when users are locked out after multiple failed password errors. Filter.config can be edited for user agents causing the lockout. NOTE: This variable must be added to the config file and should only be used for debugging purposes and then disabled. |

Integrated Windows Authentication

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|---|
| <code>proxy.config.winauth.enabled</code> | INT | 0 | Enables (1) or disables (0) Integrated Windows Authentication (Kerberos). |
| <code>proxy.config.winauth.realm</code> | STRING | NULL | Specifies the name of the Windows Active Directory domain. By entering "*", all domain controllers found in the DNS SRV records will be used. |
| <code>proxy.config.winauth.dc.list</code> | STRING | NULL | Specifies a comma separated list of domain controllers. |
| <code>proxy.config.winauth.log_denied_requests</code> | INT | 1 | Enables (1) or disables (0) logging of denied authentication requests. |

Transparent authentication

Help | Content Gateway | Version 7.8.x

| 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. |
| <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. <p>In either mode, the length of time before a client must re-authenticate is determined by the value of <code>proxy.config.http.transparent_auth_session_time</code>.</p> |
| <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

Help | Content Gateway | Version 7.8.x

| 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 ports used for tunneling. This is a space-separated list that can also include ranges of ports, e.g. 1-65535. Content Gateway allows tunnels only to the specified ports. |
| <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"> • 0 = Never use HTTP 1.1 when communicating with origin servers. • 1 = Always use HTTP 1.1 when communicating with origin servers. • 2 = Use HTTP 1.1 if the origin server has previously used HTTP 1.1. • 3 = Use HTTP 1.1 if the client request is HTTP 1.1 and the origin server has previously used HTTP 1.1. <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.share_server_sessions_max</code> | INT | 2500 | <p>Specifies the maximum number of server sessions that can be reused.</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> |

| Configuration Variable | Data Type | Default Value | Description |
|--------------------------------------|-----------|---------------|---|
| proxy.config.http.record_heartbeat | INT | 0 | Enables (1) or disables (0) content_cop heartbeat logging. |
| proxy.config.http.large_file_support | INT | 1 | When enabled (1), Content Gateway supports downloading of files larger than 2 GB. |

Parent proxy configuration

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| proxy.config.http.parent_proxy_routing_enable | INT | 0 | Enables (1) or disables (0) the HTTP parent caching option. See Hierarchical Caching , page 95. |
| proxy.config.http.parent_proxy.retry_time | INT | 300 | Specifies the amount of time allowed between connection retries to a parent cache that is unavailable. |
| proxy.config.http.parent_proxy.fail_threshold | INT | 10 | Specifies the number of times the connection to the parent cache can fail before Content Gateway considers the parent unavailable. |
| proxy.config.http.parent_proxy.total_connect_attempts | INT | 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). |
| proxy.config.http.parent_proxy.per_parent_connect_attempts | INT | 2 | Specifies the total number of connection attempts allowed per parent if multiple parents are used. |
| proxy.config.http.parent_proxy.connect_attempts_timeout | INT | 30 | Specifies the timeout value, in seconds, for parent cache connection attempts. |
| proxy.config.http.forward.proxy_auth_to_parent | 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. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| <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 |
| <code>proxy.local.http.parent_proxy.disable_ssl_connect_tunneling</code> | INT | 0 | When enabled (1), HTTPS requests bypass the parent proxy. |
| <code>proxy.local.http.parent_proxy.disable_unknown_connect_tunneling</code> | INT | 0 | When enabled (1), non-HTTPS tunnel requests bypass the parent proxy. |

HTTP connection timeouts (secs)

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|---|
| <code>proxy.config.http.keep_alive_no_activity_timeout_in</code> | INT | 60 | Specifies how long Content Gateway keeps connections to clients open for a subsequent request after a transaction ends. |
| <code>proxy.config.http.keep_alive_no_activity_timeout_out</code> | INT | 60 | Specifies how long Content Gateway keeps connections to origin servers open for a subsequent transfer of data after a transaction ends. |
| <code>proxy.config.http.transaction_no_activity_timeout_in</code> | INT | 120 | Specifies how long Content Gateway keeps connections to clients open if a transaction stalls. |
| <code>proxy.config.http.transaction_no_activity_timeout_out</code> | INT | 120 | Specifies how long Content Gateway keeps connections to origin servers open if the transaction stalls. |
| <code>proxy.config.http.transaction_active_timeout_in</code> | 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 | 120 | 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

Help | Content Gateway | Version 7.8.x

| 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.streaming_connect_attempts_timeout</code> | INT | 1800 | Specifies the timeout value in seconds for a streaming content 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. |

Negative response caching

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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 |
|--|-----------|---------------|--|
| <code>proxy.config.http.anonymize_remove_user_agent</code> | INT | 0 | When enabled (1), Content Gateway removes the User-Agent header that accompanies transactions to protect the privacy of your site and users. |
| <code>proxy.config.http.anonymize_remove_cookie</code> | INT | 0 | When enabled (1), Content Gateway removes the Cookie header that accompanies transactions to protect the privacy of your site and users. |
| <code>proxy.config.http.anonymize_remove_client_ip</code> | INT | 1 | When enabled (1), Content Gateway removes Client-IP headers for more privacy. |
| <code>proxy.config.http.anonymize_insert_client_ip</code> | INT | 0 | When enabled (1), Content Gateway inserts Client-IP headers to retain the client's IP address. |
| <code>proxy.config.http.append_xforwards_header</code> | INT | 0 | When enabled (1), Content Gateway appends X-Forwards headers to outgoing requests. |
| <code>proxy.config.http.anonymize_other_header_list</code> | STRING | NULL | Specifies the headers that Content Gateway will remove from outgoing requests. |
| <code>proxy.config.http.snarf_username_from_authorization</code> | 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> . |
| <code>proxy.config.http.insert_squid_x_forwarded_for</code> | INT | 0 | When enabled (1), Content Gateway adds the client IP address to the X-Forwarded-For header. |
| <code>proxy.config.http.insert_x_authenticated_user</code> | INT | 0 | When enabled (1), Content Gateway inserts the X-Authenticated-User header to advertise the proxy authenticated user. |

Security

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <code>proxy.config.http.cache.cache_responses_to_cookies</code> | 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 |
| <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. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| <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. |
| <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. |
| <code>proxy.config.http.cache.cache_301_responses</code> | INT | 0 | Enables (1) or disables (0) caching of “301” response pages. |

Heuristic expiration

Help | Content Gateway | Version 7.8.x

| 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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|----------------------------------|--|
| <code>proxy.config.body_factory.enable_customizations</code> | INT | 0 | Specifies whether customizable response pages are enabled or disabled and which response pages are used: <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 | 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. |
| <code>proxy.config.body_factory.template_sets_dir</code> | STRING | <code>config/body_factory</code> | Specifies the customizable response page default directory. |
| <code>proxy.config.body_factory.response_suppression_mode</code> | INT | 0 | Specifies when Content Gateway suppresses generated response pages: <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

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| FTP over HTTP | | | |
| <code>proxy.config.ftp.data_connection_mode</code> | INT | 1 | Specifies the FTP connection mode: <ul style="list-style-type: none"> • 1 = PASV then PORT • 2 = PORT only • 3 = PASV only |
| <code>proxy.config.ftp.control_connection_timeout</code> | INT | 300 | Specifies how long Content Gateway waits for a response from the FTP server. |
| <code>proxy.config.ftp.rc_to_switch_to_PORT</code> | 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 | | | |
| <code>proxy.config.ftp.ftp_enabled</code> | INT | 0 | Enables (1) or disables (0) processing of FTP requests from FTP clients. |
| <code>proxy.config.ftp.cache_enabled</code> | INT | 0 | Enables (1) or disables (0) caching of FTP objects. When this option is disabled, Content Gateway always serves FTP objects from the FTP server. |
| <code>proxy.config.ftp.file_fresh_mdtm_checking_enabled</code> | INT | 0 | Only applies when FTP caching is enabled. When Enabled (1), Content Gateway sends an 'MDTM' command before the 'RETR' command to get the last modified time of file(s). If the file is in cache and the last_contact time is the same as 'MDTM' response, the proxy serves the cache file to the client. |
| <code>proxy.config.ftp.logging_enabled</code> | INT | 1 | Enables (1) or disables (0) logging of FTP transactions. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| <code>proxy.config.ftp.proxy_server_port</code> | INT | 2121 | Specifies the port used for FTP connections. |
| <code>proxy.config.ftp.open_lisn_port_mode</code> | INT | 1 | Specifies how FTP opens a listening port for a data transfer: <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 | 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. |
| <code>proxy.config.ftp.max_lisn_port</code> | INT | 65535 | Specifies the highest port in the range of listening ports used by Content Gateway for data connections when the FTP client sends a PASV or Content Gateway sends a PORT to the FTP server. |
| <code>proxy.config.ftp.server_data_default_pasv</code> | INT | 1 | Specifies the default method used to set up server side data connections: <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 | Specifies how an FTP server control connection is shared between different FTP client sessions: <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.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. |
| <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' or '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, 'ls -al' or 'ls *.txt'). |

SOCKS processor

Help | Content Gateway | Version 7.8.x

| 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 185. |
| <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. |
| <code>proxy.config.socks.socks_server_enabled</code> | INT | 0 | Note: Configure only if Content Gateway is installed on an appliance. |
| <code>proxy.config.socks.socks_server_port</code> | INT | 61080 | Note: Configure only if Content Gateway is installed on an appliance. |

Net subsystem

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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

Help | Content Gateway | Version 7.8.x

| 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 397). |
| <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 41 MB per GB of disk. |
| <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

Help | Content Gateway | Version 7.8.x

| 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 <i>Using the Split DNS option</i>, page 188 |
| <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. |
| <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. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|--|
| <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. The default value is lower than <code>proxy.config.hostdb.lookup_timeout</code> and, therefore, takes precedence. |
| <code>proxy.config.dns.retries</code> | INT | 5 | Specifies the number of times a DNS lookup is retried before giving up. |
| <code>proxy.config.dns.prefer_ipv4</code> | INT | 1 | When a name resolves to both IPv4 and IPv6 addresses, specifies the preferred address type. |
| <code>proxy.config.ipv6.ipv6_enabled</code> | INT | 0 | Specifies to enable (1) or disable (0) support for IPv6. |

DNS proxy

Help | Content Gateway | Version 7.8.x

| Configuration Variable Data Type | Data Type | Default Value | Description |
|---|-----------|---------------|---|
| <code>proxy.config.dns.proxy.enabled</code> | INT | 0 | Enables (1) or disables (0) the DNS proxy caching option that lets you resolve DNS requests on behalf of clients. This option off-loads remote DNS servers and reduces response time for DNS lookups. See DNS Proxy Caching, page 107 . |
| <code>proxy.config.dns.proxy_port</code> | INT | 5353 | Specifies the port that Content Gateway uses for DNS traffic. |

HostDB

Help | Content Gateway | Version 7.8.x

| 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 | <p>Specifies the host database time to live (ttl) mode.</p> <p>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.</p> <p>You can specify:</p> <p>0 = obey the ttl values set by the name servers (default)</p> <p>1 = ignore the ttl values set by name servers and use the value set by the Content Gateway configuration variable <code>proxy.config.hostdb.timeout</code>. Set this variable to a value appropriate for your environment.</p> <p>2 = use the lower of the two values (the one set by the name server or the one set by Content Gateway)</p> <p>3 = use the higher of the two values (the one set by the name server or the one set by Content Gateway)</p> |
| <code>proxy.config.hostdb.timeout</code> | INT | 86400 | Specifies the foreground timeout, in seconds. |
| <code>proxy.config.hostdb.fail.timeout</code> | INT | 60 | Specifies the time for which a failed DNS will be cached in seconds. |
| <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

Help | Content Gateway | Version 7.8.x

| 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 241. |
| <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 | 5120 or 20480 | Specifies the amount of space allocated to the logging directory, in megabytes. When Content Gateway is on a V-series appliance, the size is 5120 (5 GB) and cannot be changed. When Content Gateway is installed on a stand-alone server, the default size is 20480 (20 GB) and the size is configurable. |
| <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 | /opt/WCG/logs | 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 408.</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 251.</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 |
|---|-----------|---------------|---|
| <code>proxy.config.log2.common_log_name</code> | STRING | common | Specifies the Netscape common log filename. |
| <code>proxy.config.log2.common_log_header</code> | STRING | NULL | Specifies the Netscape common log file header text. |
| <code>proxy.config.log2.extended_log_enabled</code> | INT | 1 | Enables (1) or disables (0) the Netscape extended log file format. |
| <code>proxy.config.log2.extended_log_is_ascii</code> | INT | 1 | Specifies the Netscape extended log file type: <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary |
| <code>proxy.config.log2.extended_log_name</code> | STRING | extended | Specifies the Netscape extended log filename. |
| <code>proxy.config.log2.extended_log_header</code> | STRING | NULL | Specifies the Netscape extended log file header text. |
| <code>proxy.config.log2.extended2_log_enabled</code> | INT | 0 | Enables (1) or disables (0) the Netscape Extended-2 log file format. |
| <code>proxy.config.log2.extended2_log_is_ascii</code> | INT | 1 | Specifies the Netscape Extended-2 log file type: <ul style="list-style-type: none"> • 1 = ASCII • 0 = binary |
| <code>proxy.config.log2.extended2_log_name</code> | STRING | extended2 | Specifies the Netscape Extended-2 log filename. |
| <code>proxy.config.log2.extended2_log_header</code> | STRING | NULL | Specifies the Netscape Extended-2 log file header text. |
| <code>proxy.config.log2.separate_host_logs</code> | 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 254). |
| <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. <p>For information on sending XML-based custom formats to the collation server, see logs_xml.config, page 408.</p> |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <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 251 . |
| <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. |
| <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

Help | Content Gateway | Version 7.8.x

| 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> |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <code>proxy.config.url_remap.remap_required</code> | INT | 0 | Set this variable to 1 if you want Content Gateway to serve requests only from origin servers listed in the mapping rules of the <code>remap.config</code> file. If a request does not match, the browser will receive an error. |
| <code>proxy.config.url_remap.pristine_host_hdr</code> | INT | 0 | Set this variable to 1 if you want to retain the client host header in a request during remapping. |

Scheduled update configuration

[Help](#) | [Content Gateway](#) | Version 7.8.x

| 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. |
| <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. |

SNMP configuration

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|-------------|
| proxy.config.snmp.master_agent_enabled | INT | 0 | |
| proxy.config.snmp_encap_enabled | INT | 0 | |

Plug-in configuration

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|--------------------------------|-----------|----------------|---|
| proxy.config.plugin.plugin_dir | STRING | config/plugins | Specifies the directory in which plugins are located. |

WCCP configuration

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|---------------------------|-----------|---------------|-----------------------------------|
| proxy.config.wccp.enabled | INT | 0 | Enables (1) or disables (0) WCCP. |

FIPS (Security Configuration)

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|---------------------------------------|-----------|---------------|---|
| proxy.config.fips.security_enabled | INT | 0 | Preserves FIPS settings for customers upgrading from v7.5.3 FIPS to v7.7 only. |
| proxy.config.fips.security_enabled_ui | INT | 0 | Preserves FIPS UI settings for customers upgrading from v7.5.3 FIPS to v7.7 only. |

SSL Decryption

Help | Content Gateway | Version 7.8.x



Important

All SSL decryption settings 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.enabled | INT | 1 | When enabled (1), Content Gateway accepts SSL connections and performs URL filtering before establishing a connection with the origin server. See proxy.config.ssl_decryption.use_decryption to enable SSL decryption. |
| proxy.config. ssl_decryption. use_decryption | INT | 0 | When enabled (1), Content Gateway accepts and decrypts SSL traffic. See Working With Encrypted Data , page 139. |
| 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_server_port | INT | 8080 | The port on which Content Gateway listens for client SSL traffic. |
| proxy.config. administrator_id | STRING | NULL | Do not change. Holds the encrypted administrator ID. |
| proxy.config. ssl_decryption. tunnel_skype | INT | 0 | When enabled (1), Content Gateway identifies and tunnels Skype traffic (explicit proxy deployments only). User policies must be adjusted accordingly. See the configuration information in Enabling SSL support , page 143. |
| proxy.config. ssl_decryption. tunnel_unknown_protocols | INT | 0 | Enables (1) or disables the tunneling of unrecognized protocols using SSL ports. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| <code>proxy.config.ssl_decryption.tunnel_unknown_protocols_timeout</code> | INT | 10 | Specifies the time in seconds that Content Gateway waits for the “client hello” response before tunneling the request as an unknown protocol. |
| <code>proxy.config.ssl.server.SSLv2</code> | INT | 0 | When enabled (1), Content Gateway accepts SSLv2 connections from clients. (In this case, “server” refers to Content Gateway’s role as server to the client.) |
| <code>proxy.config.ssl.server.SSLv3</code> | INT | 1 | When enabled (1), Content Gateway accepts SSLv3 connections from clients. (In this case, “server” refers to Content Gateway’s role as server to the client.) |
| <code>proxy.config.ssl.server.TLSv1</code> | INT | 1 | When enabled (1), Content Gateway accepts TLSv1 connections from clients. (In this case, “server” refers to Content Gateway’s role as server to the client.) |
| <code>proxy.config.ssl.server.TLSv1.1</code> | INT | 1 | When enabled (1), Content Gateway accepts TLSv1.1 connections from clients. (In this case, “server” refers to Content Gateway’s role as server to the client.) |
| <code>proxy.config.ssl.server.TLSv1.2</code> | INT | 1 | When enabled (1), Content Gateway accepts TLSv1.2 connections from clients. (In this case, “server” refers to Content Gateway’s role as server to the client.) |
| <code>proxy.config.ssl.client.SSLv2</code> | INT | 0 | When enabled (1), Content Gateway accepts SSLv2 connections from origin servers. (In this case, “client” refers to Content Gateway’s role as client to the origin server.) |
| <code>proxy.config.ssl.client.SSLv3</code> | INT | 1 | When enabled (1), Content Gateway accepts SSLv3 connections from origin servers. (In this case, “client” refers to Content Gateway’s role as client to the origin server.) |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|-------------------------|---|
| proxy.config.ssl.client.TLSv1 | INT | 1 | When enabled (1), Content Gateway accepts TLSv1 connections from origin servers. (In this case, “client” refers to Content Gateway’s role as client to the origin server.) |
| proxy.config.ssl.client.TLSv1.1 | INT | 1 | When enabled (1), Content Gateway accepts TLSv1.1 connections from origin servers. (In this case, “client” refers to Content Gateway’s role as client to the origin server.) |
| proxy.config.ssl.client.TLSv1.2 | INT | 1 | When enabled (1), Content Gateway accepts TLSv1.2 connections from origin servers. (In this case, “client” refers to Content Gateway’s role as client to the origin server.) |
| proxy.config.ssl.server.cipherlist_option | STRING | Default | Specifies the client-to-proxy cipher setting. Values are: Default High Medium Low See SSL configuration settings for inbound traffic, page 156 . |
| proxy.config.ssl.server.cipherlist | STRING | ALL:!ADH:@STRENGTH:!EXP | The cipher list that corresponds to the cipher_list option setting. |
| proxy.config.ssl.client.cipherlist_option | STRING | Default | Specifies the cipher list setting. Values are: Default High Medium Low See SSL configuration settings for outbound traffic, page 158 . |
| proxy.config.ssl.client.cipherlist | STRING | ALL:!ADH:@STRENGTH:!EXP | The cipher list that corresponds to the cipher_list option setting. See the preceding entry. |
| proxy.config.ssl.server.session_cache | INT | 1 | Enables (1) or disables the SSL server session cache. |
| proxy.config.ssl.server.session_cache_timeout | INT | 300 | Specifies the SSL server session cache timeout period. The default is 300 seconds (5 minutes). |
| proxy.config.ssl.client.session_cache | INT | 1 | Enables (1) or disables the SSL client session cache. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|-----------------------------|--|
| <code>proxy.config.ssl.client.session_cache_timeout</code> | INT | 300 | Specifies the SSL client session cache timeout period. The default is 300 seconds (5 minutes). |
| <code>proxy.config.ssl.client.certification_level</code> | INT | 0 | Specifies whether client certificates are not needed, optional, or required. certification level should be: 0 = no client certificates 1 = client certificates optional 2 = client certificates required |
| <code>proxy.config.ssl.server.cert.filename</code> | STRING | <code>server.crt.pem</code> | Specifies the server certificate filename. |
| <code>proxy.config.ssl.server.private_key.filename</code> | STRING | <code>Domainkey.pem</code> | Specifies the private key for the server certificate. |
| <code>proxy.config.ssl.server.private_key.path</code> | STRING | <code>/config</code> | Specifies the private key path for the server certificate. |
| <code>proxy.config.ssl.CA.cert.filename</code> | STRING | NULL | Specifies the name of the file containing the list of CAs that Content Gateway will accept from a client. When the connection is from the client to Content Gateway and the value of <code>proxy.config.ssl.client.certification_level</code> is 1 or 2, Content Gateway sends the CA list to client. |
| <code>proxy.config.ssl.CA.cert.path</code> | STRING | NULL | Specifies the path to the CA list files. See the preceding entry. |
| <code>proxy.config.ssl.client.cert.policy</code> | INT | 1 | For SSL certificate incidents, specifies whether to tunnel an incident (0), or block the request and create an entry in the incident list (1). |
| <code>proxy.config.ssl.client.verify.server</code> | INT | 0 | Enables (1) or disables the Certificate Verification Engine (CVE). See Validating certificates , page 159. |
| <code>proxy.config.ssl.cert.verify.denycnmismatch</code> | INT | 1 | Enables (1) or disables the CVE check: “Deny certificates where the common name does not match the URL” The setting applies only when the CVE is enabled. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| proxy.config.ssl.cert.verify.allowcnwild | INT | 1 | Enables (1) or disables the CVE check: "Allow wildcard certificates" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.denyexpired | INT | 1 | Enables (1) or disables the CVE check: "No expired or not yet valid certificates" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.certchain | INT | 1 | Enables (1) or disables the CVE check: "Verify entire certificate chain" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.checkcrl | INT | 1 | Enables (1) or disables the CVE check: "Check certificate revocation by CRL" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.checkocsp | INT | 0 | Enables (1) or disables the CVE check: "Check certificate revocation by OCSP" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.blockunknownocsp | INT | 0 | Enables (1) or disables the CVE check: "Block certificates with Unknown OCSP state" The setting applies only when the CVE is enabled. |
| proxy.config.ssl.cert.verify.denymd5cert | INT | 0 | Enables (1) denial of certificates that use an MD5 signature. |
| proxy.config.ssl.cert.verify.revprefer | INT | 1 | Specifies the preferred method for the certificate revocation check. 1 = CRL 2 = OCSP |
| proxy.config.ssl.cert.verify.blocknouri | INT | 0 | Enables (1) or disables the CVE check: "Block certificates with no CRL URI and with no OCSP URI" |
| proxy.config.ssl.cert.verify.bypassfail INT 0 | INT | 0 | Enables (1) the certificate check failure bypass option that allows users to proceed to a site after the certificate check has failed. |
| proxy.config.ssl.cert.verify.bypasscache | INT | 1 | Enables (1) the verification timeout cache. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|--|
| proxy.config.ssl.cert.verify.bypasscachetimeout | INT | 6 | Specifies the time, in seconds, that an entry in verification bypass cache times out and is purged. |
| proxy.config.ssl_decryption_bypass.tunnel_non-ssl_traffic | INT | 0 | Beginning with 7.8.3, enables (1) or disables (0) tunneling of non-ssl traffic. With 7.8.3, available for explicit proxy only. Available for transparent proxy beginning with 7.8.4. This variable must be added manually. |

ICAP

Help | Content Gateway | Version 7.8.x

| 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 129. |
| proxy.config.icap.ICAPUri | STRING | NULL | Specifies the Uniform Resource Identifier for the ICAP service. A backup server can be specified in a comma-separated list. 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. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---------------|---|
| proxy.config.icap.FailOpen | INT | 1 | Set to: <ul style="list-style-type: none"> • 1 to allow traffic when the ICAP server(s) is down • 0 to send a block page if the ICAP server(s) is down |
| proxy.config.icap.BlockHugeContent | 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 by Data Security Suite is sent. The default size limit in DSS is 12 MB. • 1 to allow traffic |
| proxy.config.icap.AnalyzeSecureContent | 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. |
| proxy.config.icap.AnalyzeFTP | INT | 1 | When enabled (1), send native FTP upload file transfers to ICAP server for analysis. |
| proxy.config.icap.ActiveTimeout | INT | 5 | The read/response timeout in seconds. The activity is considered a failure if the timeout is exceeded. |
| proxy.config.icap.RetryTime | INT | 5 | The recovery interval, in seconds, to test whether a down server is back up. |
| proxy.config.icap.LoadBalance | INT | 1 | When to ICAP servers are specified, set to: <ul style="list-style-type: none"> • 1 to distribute requests to all available servers • 0 to distribute requests to only the primary server. |

Data Security

Help | Content Gateway | Version 7.8.x

| Configuration Variable | Data Type | Default Value | Description |
|---------------------------------------|-----------|---------------|--|
| proxy.config.dss.enabled | INT | 0 | Enables (1) or disables (0) support for on-box Data Security. See Working With Websense Data Security , page 129. |
| proxy.config.dss.AnalyzeFTP | INT | 1 | When enabled (1), send native FTP upload file transfers to the on-box Data Security policy engine for analysis. |
| proxy.config.dss.AnalyzeSecureContent | 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. |
| proxy.config.dss.analysis_timeout | INT | 10000 | Specifies the maximum length of time, in milliseconds, that a single file analysis can take before analysis is aborted. |
| proxy.config.dss.UsingLoginID | INT | 0 | Beginning with 7.8.3, enables (1) or disables (0) sending Login ID rather than full user name to Data Security. This variable must be added manually. |

Connectivity, analysis, and boundary conditions

Help | Content Gateway | Version 7.8.x

| 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. |

| Configuration Variable | Data Type | Default Value | Description |
|---|-----------|---------------|---|
| 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.wse_server_timeout | INT | 5000 | Specifies the maximum length of time, in milliseconds, for communication with Filtering Service. |
| wtg.config.ssl_bypassed_categories | STRING | NULL | <p>This variable takes a list of category identifiers that will bypass SSL decryption.</p> <p>Do not change the value of this variable. It is included strictly as a troubleshooting aid.</p> <p>Use the Web Security Manager to specify categories to bypass SSL decryption.</p> |
| wtg.config.ssl_decryption_bypass_ip_based | INT | 0 | <p>Specifies that the SSL category bypass process use only the IP address (not the hostname) when performing a category lookup.</p> <p>0 = disabled 1 = enabled</p> |
| wtg.config.ssl_fail_open | INT | 1 | <p>Specifies whether SSL sites are decrypted if Filtering Service becomes unreachable.</p> <p>0 = disable – causes all SSL sites to be decrypted when Filtering Service is unreachable.</p> <p>1 = enable – causes all SSL sites not to be decrypted when Filtering Service is unreachable</p> |
| wtg.config.fail_open | INT | 1 | <p>Specifies whether Content Gateway will permit or block the request when Websense Web filtering (Filtering Service) is unavailable.</p> <p>Set to:</p> <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_analytic_scan</code> | INT | 1 | Specifies how Content Gateway behaves should analytic scanning become non-functional. Set to: <ul style="list-style-type: none"> • 0 to block traffic • 1 to perform a lookup in the URL master database and apply policy Note: An alarm is raised whenever analytics scanning becomes non-functional. |
| <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.max_mem_allowed</code> | INT | 1500 | Specifies in megabytes, the maximum amount of memory, which when consumed, causes Content Gateway to perform more extensive memory monitoring. |
| <code>wtg.config.lowmem_behavior</code> | INT | 0 | Enables (1) or disables (0) bypass of scanning, but still filters. |
| <code>wtg.config.lowmem_timeout</code> | INT | 120 | Timeout value (in minutes) for the low-memory management. After that time, resets to “no management”. |
| <code>wtg.config.rdnsclients</code> | INT | 0 | Enables (1) or disables (0) logging of clients’ hostnames in the log records by doing reverse DNS on each. |

| Configuration Variable | Data Type | Default Value | Description |
|--|-----------|---|--|
| <code>wtg.config. ip_ranges_not_to_scan</code> | STRING | 10.0.0.0- 10.255.255.255, 172.16.0.0- 172.31.255.255, 192.168.0.0- 192.168.255.255 | Specifies internal IP address ranges not to scan. By default, the list is the standard private non-routable IP addresses. Address ranges are hyphenated with each range separated by a comma. This is especially helpful in explicit proxy deployments in which a PAC file is not used and you want to exclude the standard internal IP addresses from being scanned. |
| <code>wtg.config. scan_ip_ranges</code> | INT | 1 | Enables (1) or disables (0) bypass of the internal IP address ranges specified in <code>wtg.config.ip_ranges_not_to_scan</code> . See above. |

remap.config

Help | Content Gateway | Version 7.8.x

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, restart the proxy or run `content_line -x` from the Content Gateway **bin** directory (`/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> | <p>Enter one of the following:</p> <ul style="list-style-type: none"> • 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). <p>Note: map and reverse_map are not supported.</p> |
| <i>target</i> | <p>Enter the origin or <i>from</i> URL. You can enter up to four components:</p> <p><i>scheme://host:port/path_prefix</i> <i>scheme</i> can be http, https, or ftp.</p> |
| <i>strict URL matching flag</i> | <p>Enable Match URL Exactly to force matching to be exact against the entire requested URL.</p> <p>Without this option, the URL is compared up to the end of the target (From Path Prefix). If there is a match, the redirect is applied. This can cause unwanted matching, when the redirect URL includes the base URL. See Mapping and Redirection, page 330.</p> |
| <i>replacement</i> | <p>Enter the destination or <i>to</i> URL. You can enter up to four components:</p> <p><i>scheme://host:port/path_prefix</i> <i>scheme</i> can be http, https, or ftp.</p> |



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

Help | Content Gateway | Version 7.8.x

The **socks.config** file specifies:

- ◆ SOCKS servers that the proxy must use to access specific origin servers, and the order in which the proxy goes through the SOCKS server list.
- ◆ Origin servers that Content Gateway accesses directly, *without* going through a SOCKS server.



Note

It is recommended that all SOCKS configuration be performed in the Content Gateway manager.



Important

After you modify this file, you must restart the proxy.

Traffic that does not match a manually configured rule is handled via a default rule. A default rule is constructed for each SOCKS server with the **default** option enabled in the **Socks Servers** table. Default rules are created automatically and displayed on the SOCKS Server page. Default rules are not written in the **socks.config** file. The destination IP address is 'All'.

Format

To specify SOCKS servers that the proxy must use to reach specific origin servers, add rules to the **socks.config** file in the following format:

```
dest_ip=ipaddress socksparent="alias1" [round_robin=value]
```

where:

ipaddress is the origin server IP address or range of IP addresses separated by - or /.

alias1 is the alias name of the SOCKS server named in the **SOCKS Servers** list.

value is either *strict* if you want Content Gateway to try the SOCKS servers one by one, or *false* if you do not want round-robin selection to occur.

To specify origin servers that you want Content Gateway to access directly, *without* going through the SOCKS server(s), enter a rule in **socks.config** 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 that you want Content Gateway to access directly. Do not specify the all networks broadcast address: 255.255.255.255.



Note

Each rule in **socks.config** can consist of a maximum of 400 characters. The order of the rules in the **socks.config** file is not significant.

Examples

The following example configures the proxy to send requests to the origin servers associated with the range of IP addresses 123.15.17.1 - 123.14.17.4 through the SOCKS server aliases 'alias1' and 'alias2'. Because the optional specifier **round_robin** is set to **strict**, the proxy sends the first request to alias1, the second request to alias2, the third request to alias1, and so on.

```
dest_ip=123.14.15.1 - 123.14.17.4
socksparent="alias; alias2" round_robin=strict
```

The following example configures the proxy to access the origin server associated with the IP address 11.11.11.1 directly, *without* going through the SOCKS server:

```
no_socks 11.11.11.1
```

The following example configures Content Gateway to access the origin servers associated with the range of IP addresses 123.14.15.1 - 123.14.17.4 and the IP address 113.14.18.2 directly, *without* going through the SOCKS server:

```
no_socks 123.14.15.1 - 123.14.17.4, 113.14.18.2
```

socks_server.config

[Help](#) | [Content Gateway](#) | Version 7.8.x

The **socks_server.config** file specifies the SOCKS servers available to Content Gateway.

Format

To specify SOCKS servers use the following format:

```
alias=name host=IP_address|domain_name port=port_number  
[username=user_name password=password] default=true|false
```

where:

name is the name of a SOCKS server.

IP_address or *domain_name* is an IP address or a domain name that can be resolved by your DNS service.

port_number is the port on which the SOCKS server is listening.

username and *password* are the username/password pair for SOCKS 5 authentication. The password is encrypted.

Set default to *true* to make the specified server a default SOCKS server. When the default server option is on, the SOCKS server is used when no SOCKS rule matches.

If no SOCKS server is designated a default server, traffic that doesn't match a rule is not routed through a SOCKS server.

Examples:

This example adds the SOCKS server 'default1' at 127.0.0.1 on port 61080. It is designated a default SOCKS server.

```
alias=default1 host=127.0.0.1 port=61080 default=true
```

This example adds a SOCKS server that uses authentication. Note that the password, "465751475058" is not the real password. It is encrypted.

```
alias=test1 host=socks5.example.com port=1080 username=test  
password=465751475058 default=false
```

If this file is modified, you must restart Content Gateway.



Note

Each rule in **socks_server.config** cannot exceed 400 characters.

splitdns.config

Help | Content Gateway | Version 7.8.x

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 188](#).



Important

After you modify this file, run `content_line -x` from the Content Gateway **bin** directory (`/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

Help | Content Gateway | Version 7.8.x

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

Help | Content Gateway | Version 7.8.x

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 (`/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.

wccp.config

Help | Content Gateway | Version 7.8.x

The `wccp.config` file stores the WCCP configuration information and service group settings. When WCCP is enabled on the **Configure > MyProxy > Basic** page, WCCP service group settings can be configured on the **Configure > Networking > WCCP** page. Service groups must be defined if WCCP is to be used for transparent redirection to Content Gateway. For more information, see [Transparent interception with WCCP v2 devices](#), page 50.

F

Error Messages

Help | Content Gateway | Version 7.8.x

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 errors

| Message | Description |
|---|---|
| Accept port is not between 1 and 65535. Please check configuration. | The port specified in the records.config file that accepts incoming HTTP requests is not valid. |
| Ftp accept port is not between 1 and 65535. | The port specified in the records.config file that accepts incoming FTP requests is not valid. |
| Self loop is detected in parent proxy configuration. | The name and port of the parent proxy are the same as that of Content Gateway. This creates a loop when Content Gateway attempts to send requests to the parent proxy. |
| Could not open the ARM device | The ARM failed to load. The most common reason for this is that the host system has an incompatible system kernel. To see if the ARM is loaded, run: <code>/sbin/lsmmod grep arm</code> |

| Message | Description |
|--|---|
| content_manager failed to set cluster IP address | The content_manager process could not set the cluster IP address. Check the cluster IP address. Make sure that it is not already used by another device in the network. |
| Unable to initialize storage. (Re)Configuration required. | Cache initialization failed during startup. The cache configuration should be checked and configured or reconfigured. |

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 connection | The specified client is not allowed to connect to Content Gateway. The client IP address is not listed in the ip_allow.config file. |
| 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, 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 381 . |
| 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. |

| Message | Description |
|---|--|
| 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 491). |
| All disks are bad, cache disabled | There is a problem with the cache disk(s) and caching has been disabled. Please verify that the cache disks are working and have been properly formatted for caching. See Configuring the Cache , page 99. |
| Missing DC parameter <missing_param> on auth.profile line | A required parameter was not specified. Please provide a value for the missing parameter. |
| Bad DC parameter <bad_param> - <dc_name> | A specified Domain Controller parameter is invalid. Please enter a valid value for the cited parameter. |
| [ParentSelection] <error_description> for default parent proxy | Proxy chaining is not working due to misconfiguration of the parent proxy in the child proxy. Please check the chaining configuration of parent proxy values in the child proxy. |
| WCCP2: Cannot find Interface name. Please check that the variable proxy.local.wccp2.ethernet_interface is set correctly | No value is specified for the WCCP interface. In the Content Gateway manager, check Configure > Networking > WCCP > General , or assign a value to proxy.local.wccp2.ethernet_interface in records.config . |
| ARMManager: Unable to read network interface configuration | There is a format or configuration error in ipnat.conf . In the Content Gateway manager, go to Configure > Networking > ARM > General and click Edit File to view and correct ipnat.conf . |

Alarm messages

Help | Content Gateway | Version 7.8.x

The following table describes alarm messages that you may see in the Content Gateway manager.

| Message | Description/Solution |
|--|---|
| The Content Gateway subscription has expired. | Please contact your Websense customer service representative or Technical Support for assistance. |
| Content Gateway subscription download failed. | Content Gateway was unable to connect to the download server to verify the subscription information. Please check your connection to the download server. |
| After several attempts, Content Gateway failed to connect to the Websense Database Download Service. Please troubleshoot the connection. | Verify that Content Gateway is able to access the Internet. Check firewall and upstream proxy server settings that might prevent Content Gateway from connecting to the download server. |
| After several attempts, Content Gateway failed to connect to the Policy Server. Please troubleshoot the connection. | Verify that there is network connectivity between Content Gateway and Web Security. Sometimes firewall settings block connectivity. Also confirm that the Policy Server service is running on the Web Security host. |
| After several attempts, Content Gateway failed to connect to the Policy Broker. Please troubleshoot the connection. | Verify that there is network connectivity between Content Gateway and Web Security. Sometimes firewall settings block connectivity. Also confirm that the Policy Broker service is running on the Web Security host. |
| After several attempts, Content Gateway failed to connect to the Filter service. Please troubleshoot the connection. | Verify that there is network connectivity between Content Gateway and Web Security. Sometimes firewall settings block connectivity. Also confirm that the Filter Service process is running on the Web Security host. |
| Communication with the analytics engine has failed. Please restart Content Gateway. | Restart Content Gateway. |
| SSL decryption has been disabled due to an internal error, please restart Content Gateway. | There was a fatal error in SSL Support. Please restart Content Gateway. |
| [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. |

| Message | Description/Solution |
|--|---|
| [Rollback::Rollback] Unable to read or write config file <i>filename</i> | Go to the Content Gateway config directory and make sure the indicated file exists. Check its permissions and change them if necessary. |
| [Content Gateway Manager] Configuration File Update Failed <i>error_number</i> | Go to the Content Gateway config directory and check the indicated file permissions; change them if necessary. |
| Access logging suspended - configured space allocation exhausted. | The space allocated to the event log files is full. You must either increase the space or delete some log files to enable access logging to continue. To prevent this from happening, consider rolling log files more frequently and enabling the autodelete feature. See Rolling event log files, page 251 . |
| 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 251 . |
| 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 381 , 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: <pre>icap://hostname:port/path</pre> <p>See Working With Websense Data Security, page 129 for additional details on the format of the URI.</p> |

| 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 the Content Gateway manager. See Working With Websense Data Security, page 129 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. |
| Windows domain [domain name] unreachable or bad membership status | This alarm can indicate any of the following: 1. The Active Directory is unreachable. The AD server is either down or there is a network connectivity problem. 2. The AD is reachable, but there is a configuration problem that prevents it from communicating with Content Gateway. For example, the alarm is generated if the AD has multiple Sites and the subnet that Content Gateway resides on has not been added to one of them. |
| The Scanning Data Files Update option (My Proxy > Subscription) is set to 'suspend updates'. To get the best Websense protection, set it to 'no delay', or, on a backup system, use a time-based option. | Included in 7.8.3 and later, this alarm is a reminder that downloads of the security scanning data files used by Content Gateway analysis has been suspended. It is recommended that you not clear this alarm until the delay time has been reset. |

HTML messages sent to clients

Help | Content Gateway | Version 7.8.x

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 504](#). 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 |
| 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 |

| Title | HTTP Code | Description | Customizable Filename |
|-----------------------------------|-----------|--|-----------------------|
| 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 |
| 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 |

| Title | HTTP Code | Description | Customizable Filename |
|-------------------------------|-----------|--|----------------------------|
| 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 |
| 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

Help | Content Gateway | Version 7.8.x

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 |
| 501 | Not Implemented |
| 502 | Bad Gateway |
| 504 | Gateway Timeout |

Index

A

- absolute freshness limit, setting, 24
- access web site, 164
- accessing Content Gateway Manager, 12, 177
- Adaptive Redirection Module. See ARM.
- adding certificate authorities, 154, 155
- adding nodes to a cluster, 89
- adding virtual IP addresses, 92
- administration tools, 7
- administrator ID, 12
- administrator ID, changing, 177
- administrator ID, setting, 176
- administrator password, 176
- administrator password default
 - administrator ID, 12
- aging factor
 - modifying, 23
- alarm messages, 498
- alarm script file, 123
- alarms, 7, 121
 - clearing, 122
 - email notification, 123
 - viewing, 122
- Alarms button, 271
- alerts, 7
- allow certificates, 155
- allowing requests, 400
- alternates for caching, 36
- analytic_server process, 6
- applets, caching, 22
- ARM, 5, 47, 72
 - bypass and WCCP, 49
 - enabling, 48
 - redirection rules, 48
 - static bypass rules, 70
- ASCII log files, 249
- ASCII_PIPE mode, 246, 412
- auth_domains.config file, 391
- auth_rules.config file, 393
- autodeleting log files, 243

B

- backing up certificates, 156
- backup domain controllers, 191
- binary log files, 249
- browsers to use with Content Gateway Manager, 11
- bypass rules
 - deny, 396
 - dynamic, 71
 - static, 72

- viewing, 73
- bypass.config file, 395
 - examples, 397
 - format, 396
- bypassing certificate validation, 162
- bypassing parent proxies, 96, 416

C

- cache
 - changing capacity, 100
 - child, 95
 - clearing, 104
 - contents, 100, 491
 - hit, 22
 - miss, 22
 - partitioning, 102
 - scheduling updates to, 27, 491
 - statistics, 118
- cache affinity, 50, 52
- cache capacity, 491
- cache fault tolerance, 99
- cache pinning, 29
- cache request overview, 21
- cache space
 - managing, 419
- cache statistics, 270
- cache.config file, 23, 397
- cache-control headers, 25
- cached objects
 - expiration date, 22
 - freshness, 23
 - FTP, 22
 - HTTP, 22
- caching, 23
 - applets and scripts, 22
- caching alternates, 36
- caching cookied content, 35
- caching dynamic content, 35
- caching FTP objects, 38
- certificate authorities
 - adding, 154, 155
- certificate error, 12
- certificate revocation lists
 - updating, 163
- certificate status, 154
- certificate validation, bypassing, 162
- certificates, 145, 155
 - allow, 155
 - backing up, 156
 - bypassing validation, 162
 - deleting, 154
 - deny, 155
 - generating, 145
 - importing, 146

- managing, 154
 - revocation status, 163
 - sub-certificate authority, 148
 - changing, 177
 - changing cache capacity, 100
 - changing certificate status, 154, 155
 - changing records.config variables, 113
 - changing size
 - RAM cache, 104
 - changing status of, 155
 - changing virtual IP addresses, 92
 - child cache, 95
 - clearing alarms, 122
 - clearing the cache, 104
 - client access control lists, 70
 - client access to proxy cache, 70, 175, 405
 - client no-cache directives, 31
 - clustering
 - adding nodes, 89
 - management, 86
 - management-only, 4
 - modes, 4
 - collating event log files, 256
 - command line interface
 - variables, 294
 - command-line interface, 19
 - commands, 293
 - commands
 - content_line -h, 19
 - WCGAdmin start, 19
 - configuration
 - remote, 178
 - configuration files, 113
 - filter.config, 400
 - configuration information, sharing, 86
 - configuration options, 186
 - changing in the records.config file, 113
 - configuration snapshots
 - deleting, 116
 - restoring, 115
 - taking, 115
 - configuration variables (records.config), 420
 - Configure mode
 - Content Gateway Manager, 111
 - configuring Content Gateway, 111
 - configuring user authentication
 - NTLM, 203
 - RADIUS, 209
 - configuring Websense Content Gateway, 19, 113
 - via the command line, 112
 - Content Gateway Manager, 118, 178
 - accessing, 12
 - alarms, 121
 - Alarms button (Monitor), 271
 - Configure mode, 12, 111
 - controlling access, 176
 - logging on, 12
 - Monitor mode, 117
 - Performance button (Monitor), 120
 - starting, 12
 - starting Monitor mode, 117
 - supported browsers, 11
 - user accounts, 177
 - viewing statistics, 117
 - content_cop process, 6
 - content_gateway process, 6
 - content_line -h command, 19
 - content_manager process, 6
 - controlling
 - access to Content Gateway Manager, 176, 415
 - client access to proxy cache, 175
 - controlling host access to Content Gateway Manager, 178
 - controlling information leakage, 131
 - cookies. See caching cookied content
 - custom logging, 246
 - fields, 381
 - customer support, 8
- ## D
- Date header, 23
 - default, 12
 - deleting certificates, 154
 - deleting configuration snapshots, 116
 - deny bypass rules, 71, 396
 - deny certificates, 155
 - denying requests, 400
 - deployment options, 3
 - directory services, user authentication, 191
 - disabling
 - FTP over HTTP caching, 38
 - HTTP caching, 35
 - logging, 242
 - disk usage
 - restricting, 419
 - displaying cache statistics, 118, 270
 - DNS
 - proxy caching, 107
 - resolver, 6
 - DNS servers
 - specifying, 188, 489
 - dynamic bypass rules, 71
 - deny bypass, 71, 396
 - setting, 72
 - dynamic bypass statistics, viewing, 72
 - dynamic content
 - caching, 35

E

- editing virtual IP addresses, 92
- emailing alarms, 123
- encryption, 15
- error log files, 242
- error messages, 495
 - HTML, 500
- event log entries, examples, 261
- event log files
 - collating, 256
 - converting binary to ASCII, 250
 - managing, 243
 - splitting, 254
 - statistics, 259
 - summary logs, 247
- expiration date, 22
- Expires header, 23
- explicit proxy, 21
 - HTTPS PAC file, 142
 - SSL, 141
- explicit proxy caching, 4

F

- Fault tolerance
 - cache, 99
- files
 - auth_domains.config, 391
 - auth_rules.config, 393
 - bypass.config, 395
 - cache.config, 23, 397
 - hosting.config, 403
 - ip_allow.config, 176, 405
 - ipnat.conf, 406
 - log_hosts.config, 256
 - logs_xml.config, 246
 - mgmt_allow.config, 178, 415
 - parent.config, 96, 416
 - partition.config, 419
 - records.config, 23, 420
 - socks_server.config, 488
 - socks.config, 486
 - splitdns.config, 188, 489
 - storage.config, 100, 491
 - update.config, 491
 - wccp.config, 493
- filter.config file, 400
 - examples, 402
 - format, 401
- filtering rules, 400
- FIPS 140-2, 179
- force immediate update option, 29
- forcing object caching, 36
- freshness computations, 23

- FTP client application, 43
- FTP objects
 - caching, 38
 - freshness, 26

G

- getting started, 11
- Graphs button
 - Content Gateway Manager, 118
 - statistics, 270
- Graphs button (Monitor), 270

H

- header requirements, 24
- headers
 - cache-control, 25
 - Expires, 23
 - Last-Modified, 23
 - max-age, 23
 - WWW-Authenticate, 34
- headroom limit (logging), 243
- health alerts, 7
- hierarchical caching, 4, 95
 - HTTP hierarchies, 95
 - parent failover, 96
- host access, 178
- host access to Content Gateway Manager, 178
- host database, 5
- host log splitting, 254
- hosting.config file, 403
- hostname length restriction with IWA, 198
- hostname, changing when using IWA user
 - authentication, 198, 218
- HTML error messages, 500
- HTTP
 - alternates, 36
 - cache hierarchies, 95, 96, 416
 - host, separate logs, 254
- HTTP object freshness, 23
- HTTP response messages, 504
- HTTPS support
 - enabling, 143

I

- ICAP, 131
- ICAP (Internet Content Adaptation Protocol)
 - protocols supported, 131
- ICAP Service URI, 136
- immediate update, 29
- incidents, 164
- increasing cache capacity, 100
- information leakage, controlling, 131
- integrated windows authentication, 197

- interception strategies, 49
- internal root CA, 145
 - backup, 153
- IP spoofing, 77
- ip_allow.config file, 176, 405
 - examples, 405
 - format, 405
- ipnat.conf file, 406
- IWA, 197
 - changing the domain, 199
 - configuration, 198
 - configuration summary, 197
 - finding the domain controller, 199
 - hostname length restriction, 198
 - hostname, changing, 198, 218
 - troubleshooting, 200

J

- Java, 11
- JavaScript, 11

K

- keeping header information, 400
- Kerberos, 197

L

- Last-Modified header, 23
- LDAP authentication rules, 400
- LDAP proxy authentication, 205
- listing cache contents, 100, 491
- listing commands, 19
- log collation, 256
- log collation server, 257
- log files
 - autodeleting, 243
- log formats, 245
- log_hosts.config file, 256
- logcat application, 250
- LogFilter specification, 410
- LogFormat specification, 410
- logging
 - aggregate summaries, 247
 - ASCII_PIPE, 246, 412
 - choosing log file formats, 245
 - collating log files, 256
 - converting binary files to ASCII, 250
 - custom logging fields, 381
 - disabling, 242
 - example log entries, 261
 - file splitting, 254
 - headroom limit, 243
 - managing log files, 243
 - Netscape Common formats, 385

- Netscape Extended formats, 386
- Netscape Extended-2 formats, 386
- offset hour, 253
- rolling intervals, 253
- Squid formats, 385
- stand-alone collator (SAC), 258
- statistics, 259
- time stamps, 253
- WELF, 415

- logging on
 - Windows 7, 13
- logging on to Content Gateway Manager, 12
- LogObject specification, 411
- logs_xml.config file, 246

M

- management clustering, 86
- management-only clustering, 4
- manager alarms, 7
- managing certificates, 154
- max-age header, 23
- messages
 - certificate validation failure, 172
 - connection failure, 172
- mgmt_allow.config file, 178, 415
- modifying aging factor, 23
- monitor
 - remote, 178
- Monitor mode, 117
- multiple realm user authentication
 - authentication logic, 232
 - changing a rule, 224
 - troubleshooting, 232
 - use cases, 225
- My Proxy
 - statistics, 267
- My Proxy button
 - Monitor tab, 117

N

- naming rolled log files, 252
- Netscape Common logging formats, 385
- Netscape Extended logging formats, 386
- Netscape Extended-2 logging formats, 386
- Networking
 - statistics, 282
- Networking button
 - Content Gateway Manager Monitor tab, 120
- nodes
 - adding to a cluster, 89
- NTLM authentication rules, 400
- NTLM proxy authentication, 202, 203
- NTLMv2, 197

O

- object caching, forcing, 36
- object freshness
 - aging factor, 23
- object store, 99
- offset hour, 253
- Online certification status protocol, 163
- origin server, 21
- orphan log files, 257
- outbound traffic
 - SSL, 158
- outgoing content, examining, 131
- overwriting dynamic bypass rules, 397

P

- PAC file
 - HTTPS, 142
 - with HTTPS, 141
- parent cache, 95
- parent failover, 96
- parent proxies
 - bypassing, 96, 416
- parent.config file, 96, 416
- partition.config file, 419
- partitioning the cache, 102
- partitions, 491
- passphrase, 148
- password, 12, 177
- password encryption, 15
- password, setting administrator, 176
- Performance graphs, 7
- pin-in-cache, 399
- print_bypass utility, 73
- processes (Websense Content Gateway), 6
- Protocol
 - statistics, 273
- Protocols button
 - Content Gateway Manager Monitor tab, 119
- proxy
 - explicit, 21
 - transparent, 21
- proxy cache
 - client access to, 175
 - controlling client access to, 405
- proxy caching
 - cache-control headers, 25
 - client no-cache directives, 31
 - cooked content, 35
 - disabling HTTP caching, 35
 - dynamic content, 35
 - FTP object freshness, 26
 - header requirements, 24
 - HTTP alternates, 36

- revalidating HTTP objects, 25
- scheduling cache updates, 27
- server no-cache directives, 33
- whether to cache, 30
- WWW-Authenticate headers, 34
- proxy user authentication, 189
- PUSH, 401

R

- RADIUS proxy authentication, 208, 209
- RAM cache, 99, 104
- raw disk, 491
- Read authentication from child proxy, 302
- records.config file, 23
- records.config file, 420
- Redirect Hostname, 196
- redirecting requests (ARM), 47
- reducing cache capacity, 101
- regular expressions, 390
- remote monitoring and configuration, 178
- restoring configuration snapshots, 115
- restoring Websense Content Gateway configurations, 114, 115
- restricting access to Content Gateway Manager, 176
- revalidation, 25
- revocation status, 163
- rolled log files, 252
- rolling intervals, 253
- root CA
 - internal, 145
 - backup, 153
- routers
 - configuring, 55
- rules
 - filtering, 400

S

- SAC (stand-alone collator), 258
- sample records.config file, 114
- saving configurations, 115
- scheduling cache updates, 27
- scheduling updates, 491
- script file for alarms, 123
- scripts, caching, 22
- Secure Sockets Layer, 178
- Security
 - statistics, 275
- security, 175
 - Content Gateway Manager
 - access, 176
 - options, 1, 175
 - proxy user authentication, 189

- SOCKS, 185
 - split DNS, 188
 - SSL for secure administration, 178
 - Security button
 - Content Gateway Manager Monitor tab, 119
 - security certificate alerts, 12
 - server no-cache directives, 33
 - service group ID number, 55
 - service groups, 59
 - disabling WCCP processing, 59
 - enabling WCCP processing, 56
 - guidelines for configuring, 56
 - setting absolute freshness limit, 24
 - setting administrator ID and password, 176
 - setting administrator password, 176
 - sharing configuration information, 86
 - snapshots
 - deleting, 116
 - restoring, 115
 - taking, 115
 - SOCKS, 186
 - proxy options, 187
 - SOCKS servers
 - specifying, 486, 488
 - socks_server.config file, 488
 - socks.config file, 486
 - split DNS, 188
 - splittedns.config file, 188, 489
 - splitting event log files, 254
 - spoofing, 77
 - Squid logging formats, 385
 - SSL, 178
 - certificates, 178
 - enabling (Content Gateway Manager), 178
 - outbound traffic, 158
 - SSL support
 - enabling, 143
 - stand-alone collators, 258
 - starting, 19
 - Content Gateway Manager Configure mode, 111
 - Content Gateway Manager Monitor mode, 117
 - starting Content Gateway Manager, 12
 - starting Websense Content Gateway, 19
 - static bypass rules, 72, 397
 - statistics
 - My Proxy, 267
 - Networking, 282
 - Protocol, 273
 - Subsystems, 280
 - viewing from Content Gateway Manager, 117
 - viewing from the command line, 121
 - viewing in Content Gateway Manager, 117
 - status
 - changing certificate, 155
 - status, certificate, 154
 - storage.config file, 100, 491
 - format, 491
 - stripping header information, 400
 - Subsystems
 - statistics, 280
 - Subsystems button
 - Content Gateway Manager Monitor tab, 119
 - summary log files, 247
 - system status, 7
- ## T
- technical support, 8
 - time in cache, 23
 - time stamps (log files), 252
 - traffic analysis options, 7
 - traffic graphs, see Performance graphs, 7
 - transaction logging, 8
 - transparent proxy, 21
 - interception strategies, 49
 - transparent proxy caching, 47
 - L4 switch, 50
 - policy-based routing, 67
 - software solutions, 68
 - WCCP, 51
 - transparent proxy user authentication
 - Redirect Hostname, 196
 - troubleshooting
 - Integrated Windows Authentication, 200
- ## U
- update.config file, 28, 491
 - updates
 - scheduling, 491
 - URL regular expressions, 390
 - url_regix, 390
 - user accounts, 177
 - user authentication, 189
 - backup domain controllers, 191
 - browser limitations, 192
 - integrated windows, 197
 - integrated windows, configuration
 - summary, 197
 - Kerberos, 197
 - LDAP, 205
 - Multiple realm authentication
 - authentication logic, 232
 - changing a rule, 224
 - troubleshooting, 232
 - use cases, 225
 - NTLM, 202
 - NTLMv2, 197
 - RADIUS, 208

supported directories, 191
transparent, 191

V

validation, bypassing certificate, 162
variables
 records.config file, 113, 420
verifying that Websense Content Gateway is
 running, 18
verifying URLs, 29
viewing alarms, 122
viewing bypass rules, 73
viewing certificates
 certificates
 viewing, 154
viewing dynamic bypass statistics, 72
viewing logging statistics, 259
viewing statistics
 from Content Gateway Manager, 117
 from the command line, 121
virtual IP addresses, 92
 adding, 92
 editing, 92
virtual IP failover, 4, 91

W

WCCP, 51
 enabling, 60
 load balancing, 53
 service groups, 59

wccp
 wccp.config file, 493
WCCP 2.0
 security, 59
WCCP processing
 disabling, 59
 enabling, 56
WCCP2 routers
 configuring, 55
WCGAdmin start command, 19
Web browser authentication support limits, 192
web proxy caching, 3, 21
Web Security user identification, 189
web site access, 164
Websense Content Gateway, 19
 verifying, 18
Websense Content Gateway components, 4
Websense Content Gateway configurations
 saving, 114
Websense Content Gateway Manager, 270
 Monitor mode, 12
Websense Content Gateway processes, 6
WELF, 415
Windows 7, 13
WWW-Authenticate headers, 34

X

X-Authenticated-User, 302
X-Forwarded-For, 302
XML custom log formats, 246, 408