



TRITON - Web Security Help

Websense® Web Security
Websense Web Filter

v7.6

©1996–2015, Websense Inc.

All rights reserved.

10240 Sorrento Valley Rd., San Diego, CA 92121, USA

Published 2015

Printed in the United States and Ireland

The products and/or methods of use described in this document are covered by U.S. Patent Numbers 5,983,270; 6,606,659; 6,947,985; 7,185,015; 7,194,464 and RE40,187 and other patents pending.

This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent in writing from Websense Inc.

Every effort has been made to ensure the accuracy of this manual. However, Websense Inc., makes no warranties with respect to this documentation and disclaims 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.

Trademarks

Websense is a registered trademark and TRITON is a trademark of Websense, Inc., in the United States and certain international markets. Websense has numerous other unregistered trademarks in the United States and internationally. All other trademarks are the property of their respective owners.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation in the United States and/or other countries.

eDirectory and Novell Directory Services are a registered trademarks of Novell, Inc., in the U.S and other countries.

Adobe, Acrobat, and Acrobat Reader are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Red Hat is a registered trademark of Red Hat, Inc., in the United States and other countries. Linux is a trademark of Linus Torvalds, in the United States and other countries.

This product includes software distributed by the Apache Software Foundation (<http://www.apache.org>).

Copyright (c) 2000. The Apache Software Foundation. All rights reserved.

Other product names mentioned in this manual may be trademarks or registered trademarks of their respective companies and are the sole property of their respective manufacturers.

Contents

Topic 1	Getting Started	17
	Overview	18
	Working in TRITON - Web Security	18
	Reviewing, saving, and discarding changes	23
	Customizing the Today page	26
	Monitoring Web Security status	27
	History: Last 30 Days	27
	Time and bandwidth saved	29
	Customize the History page	30
	Your subscription	31
	Managing your account through the MyWebsense Portal	31
	Configuring your account information	32
	The Websense Master Database	34
	Real-time database updates	35
	Real-Time Security Updates™	35
	Configuring database downloads	35
	Websense Technical Support	37
Topic 2	Internet Usage Filters	39
	Filtering categories and protocols	40
	New Master Database categories and protocols	42
	Special categories	42
	Risk classes	43
	Security protocol groups	46
	Filtering actions	46
	Using quota time to limit Internet access	48
	Search filtering	48
	Working with filters	49
	Creating a category filter	50
	Editing a category filter	51
	Creating a protocol filter	52
	Editing a protocol filter	53
	Websense-defined category and protocol filters	55
	Category and protocol filter templates	56
	Configuring Websense filtering settings	57
Topic 3	Clients	61

	Working with clients	62
	Working with computers and networks	63
	Working with users and groups	64
	Directory services	65
	Windows Active Directory (Mixed Mode)	66
	Windows Active Directory (Native Mode)	66
	Novell eDirectory and Oracle (Sun Java) Directory Server	67
	Advanced directory settings	68
	Working with custom LDAP groups	69
	Adding or editing a custom LDAP group	70
	Adding a client	71
	Searching the directory service	72
	Changing client settings	72
	Password override	73
	Account override	74
	Moving clients to roles	75
	Working with hybrid filtering clients	76
Topic 4	Internet Filtering Policies	79
	The Default policy	80
	Working with policies	81
	Creating a policy	82
	Editing a policy	83
	Assigning a policy to clients	85
	Filtering order	85
	Filtering a site	86
Topic 5	Block Pages	91
	Blocking graphical advertisements	92
	Blocking embedded pages	93
	Protocol block messages	94
	Working with block pages	95
	Customizing the block message	96
	Changing the size of the message frame	98
	Changing the logo that displays on the block page	98
	Using block page content variables	99
	Reverting to the default block pages	100
	Creating alternate block messages	100
	Using an alternate block page on another machine	101
	Determining why a request was blocked	102
	Request blocked by Filtering Service	102
	Request blocked by hybrid filtering	103
Topic 6	Use Reports to Evaluate Filtering	105

What is Internet browse time?	106
Presentation reports	107
Copying a presentation report	109
Defining the report filter	109
Selecting clients for a report	110
Selecting categories for a report	111
Selecting protocols for a report	112
Selecting actions for a report	113
Setting report options	113
Customizing the report logo	114
Confirming report filter definition	115
Working with Favorites	116
Running a presentation report	116
Scheduling presentation reports	118
Setting the schedule	120
Selecting reports to schedule	121
Setting the date range	122
Selecting output options	123
Viewing the scheduled jobs list	123
Viewing job history	125
Reviewing scheduled presentation reports	125
Investigative reports	127
Summary reports	129
Anonymizing investigative reports	132
The Anonymous option	132
Multi-level summary reports	133
Flexible detail reports	134
Columns for flexible detail reports	136
User Activity Detail reports	138
User activity detail by day	138
User activity detail by month	140
Category mapping	140
Standard reports	142
Favorite investigative reports	143
Saving a report as a Favorite	144
Generating or deleting a Favorite report	144
Modifying a Favorite report	145
Scheduling investigative reports	146
Managing scheduled investigative reports jobs	148
Outliers reports	149
Output to file	150
Printing investigative reports	150
Accessing self-reporting	151
Real-Time Monitor	152

	Real-Time Monitor in Multiple Policy Server Deployments	154
Topic 7	Scanning and SSL Bypass Options	155
	Scanning options	157
	Content categorization.	158
	Tunneled protocol detection	160
	Security threats: Content scanning	161
	Security threats: File scanning	162
	Advanced options	164
	Scanning exceptions	166
	Data files used with scanning	168
	Reporting on scanning activity	168
	How scanning is logged	170
	SSL decryption bypass	171
Topic 8	Configure Hybrid Filtering	173
	Activate your hybrid filtering account	174
	Define the locations filtered by hybrid filtering	175
	Adding filtered locations	176
	Editing filtered locations	177
	Managing explicit proxies	179
	Adding an explicit proxy	179
	Editing an explicit proxy	180
	Specify sites not filtered by hybrid filtering	180
	Adding unfiltered destinations	181
	Editing unfiltered destinations	182
	Configure user access to hybrid filtering	183
	Adding domains	184
	Editing domains	185
	Customizing hybrid block pages	185
	Enabling HTTPS notification pages	186
	What is a PAC file?	187
	Send user and group data to the hybrid service	189
	Configure Directory Agent settings for hybrid filtering	189
	Configure how data is gathered for hybrid filtering	191
	Oracle (Sun Java) Directory Server and hybrid filtering	192
	Novell eDirectory and hybrid filtering	193
	Adding and editing root contexts	193
	Optimizing search results	195
	Schedule communication with hybrid filtering	196
	Monitor communication with the hybrid service	197
	View hybrid service authentication reports	198
Topic 9	Filter Users Off Site	201

Using remote filtering software	202
When server communication fails.	203
Configuring Remote Filtering settings	204
Configure remote filtering to ignore FTP or HTTPS traffic	205
Configure the Remote Filtering Client heartbeat interval	206
Applying hybrid filtering to off-site users	206
Configuring hybrid filtering for off-site users.	207
Off-site user self-registration	207
Topic 10 Protect Vital Information.	209
Topic 11 Refine Filtering Policies	211
Restricting users to a defined list of Internet sites	212
Limited access filters and filtering precedence.	212
Creating a limited access filter	213
Editing a limited access filter	214
Adding sites from the Edit Policy page.	215
Copying filters and policies to roles.	216
Building filter components.	217
Working with categories.	218
Editing categories and their attributes.	218
Reviewing all customized category attributes	219
Making global category filtering changes	220
Renaming a custom category.	220
Creating a custom category.	221
Filtering based on keyword.	222
Defining keywords.	223
Redefining filtering for specific sites	224
Defining unfiltered URLs	225
Recategorizing URLs.	226
Prioritizing Security Risk categorization	227
Blocking posts to sites in some categories.	228
Working with protocols	229
Filtering protocols.	230
Editing custom protocols.	231
Adding or editing protocol identifiers	232
Renaming a custom protocol	233
Making global protocol filtering changes	233
Creating a custom protocol	234
Adding to a Websense-defined protocol	235
Using Bandwidth Optimizer to manage bandwidth	236
Configuring the default Bandwidth Optimizer limits	237
Managing traffic based on file type	238
Working with file types.	239

Adding custom file types	240
Adding file extensions to a file type	240
Using regular expressions.	241
Using the Toolbox to verify filtering behavior.	242
URL Category	242
Check Policy	242
Test Filtering	243
URL Access.	243
Investigate User.	243
Identifying a user to check policy or test filtering	244
Topic 12 User Identification	245
Transparent identification.	245
Transparent identification of remote users	246
Manual authentication	247
Configuring user identification methods	248
Setting authentication rules for specific machines	249
Defining exceptions to user identification settings	250
Revising exceptions to user identification settings	251
Secure manual authentication	252
Generating keys and certificates	252
Activating secure manual authentication	253
Accepting the certificate within the client browser	254
DC Agent	255
Configuring DC Agent	256
Logon Agent.	258
Configuring Logon Agent	259
RADIUS Agent	261
Processing RADIUS traffic.	262
Configuring the RADIUS environment	262
Configuring RADIUS Agent.	263
Configuring the RADIUS client	264
Configuring the RADIUS server.	265
eDirectory Agent	265
Special configuration considerations	266
Configuring eDirectory Agent	267
Adding an eDirectory server replica	269
Configuring eDirectory Agent to use LDAP	269
Enabling full eDirectory Server queries	270
Configuring multiple agents.	271
Configuring different settings for an agent instance.	272
INI file parameters	274
Configuring an agent to ignore certain user names	275

	Identification of hybrid filtering users	276
	Deploying Web Endpoint	277
	Distributing Web Endpoint via GPO	279
	Deploying Websense Authentication Service	280
	Websense Directory Agent	281
	Directory Agent and User Service	281
	When users are not identified	283
Topic 13	Delegated Administration and Reporting	285
	The fundamentals of delegated administration	285
	Delegated administration roles	286
	Delegated administrators	287
	Delegated administration and reporting permissions	288
	Administrators in multiple roles	290
	Multiple administrators accessing TRITON - Web Security	290
	Preparing for delegated administration	291
	Creating a Filter Lock	293
	Locking categories	293
	Locking protocols	294
	Preparing delegated administrators	295
	Managing delegated administration roles	296
	Adding roles	297
	Editing roles	298
	Adding Administrators	301
	Adding managed clients	302
	Managing role conflicts	304
	Updating delegated administration roles	304
	Delete roles	305
	Delete managed clients	306
	Managing Super Administrator clients	307
	Performing delegated administrator tasks	307
	View your user account	307
	View your role definition	308
	Add clients to the Clients page	309
	Create policies and filters	310
	Apply policies to clients	311
	Generate reports	311
	Reviewing administrator accounts	312
	Enabling network accounts	312
Topic 14	Web Security Server Administration	313
	Websense Web Security product components	314
	Filtering components	315

Reporting components	317
User identification components	318
Interoperability components	319
Understanding the Policy Database	319
Working with Policy Server	320
Reviewing Policy Server connections	320
Adding or editing Policy Server instances	321
Working in a multiple Policy Server environment	322
Changing the Policy Server IP address	322
Working with Filtering Service	324
Review Filtering Service details	325
Review Master Database download status	325
Resuming Master Database downloads	326
Working with Content Gateway	326
Managing Content Gateway connections	327
Viewing and exporting the audit log	328
Stopping and starting Websense services	329
Websense Web Security installation directories	331
Alerting	332
Flood control	333
Configuring general alert options	333
Configuring system alerts	335
Configuring category usage alerts	336
Adding category usage alerts	337
Configuring protocol usage alerts	337
Adding protocol usage alerts	338
Reviewing current system status	339
Backing up and restoring your Websense data	340
Scheduling backups	342
Running immediate backups	343
Maintaining the backup files	344
Restoring your Websense data	345
Discontinuing scheduled backups	346
Command reference	346
Topic 15 Reporting Administration	349
Planning your configuration	349
Configuring distributed logging	350
Managing access to reporting tools	352
Basic configuration	352
Assigning categories to risk classes	353
Configuring reporting preferences	354

	Configuring how filtered requests are logged	355
	Log Server Configuration utility	357
	Configuring WebCatcher	358
	Stopping and starting Log Server	359
	Introducing the Log Database	360
	Database jobs.	361
	Log Database administration settings.	361
	Configuring database partition options	362
	Configuring Log Database maintenance options	364
	Configuring how URLs are logged.	365
	Configuring Internet browse time options	366
	Configuring investigative reports	368
	Database connection and report defaults.	368
	Display and output options	370
	Self-reporting	372
Topic 16	Network Configuration	375
	Hardware configuration	376
	Filtering IPv6 traffic	378
	Network Agent configuration.	379
	Configuring global settings	379
	Configuring local settings	381
	Configuring NIC settings	382
	Configuring monitoring settings for a NIC	384
	Adding or editing IP addresses	385
	Verifying Network Agent configuration	385
Topic 17	Troubleshooting	387
	Installation and subscription issues	387
	There is a subscription problem	388
	Unable to verify the subscription key	388
	After upgrade, users are missing from TRITON - Web Security	388
	Master Database issues.	389
	The initial filtering database is being used	389
	The Master Database is more than 1 week old	389
	The Master Database does not download	390
	Subscription key.	390
	Internet access	391
	Verify firewall or proxy server settings	392
	Insufficient disk space on the Filtering Service machine	392
	Insufficient memory on the Filtering Service machine	393
	Restriction applications	394
	Master Database download does not occur at the correct time.	394
	Contacting Technical Support for database download issues	395

Filtering issues	395
Filtering service is not running	395
User Service is not available	396
High CPU usage on the Filtering Service machine.	397
Sites are incorrectly categorized as Information Technology.	397
Keywords are not being blocked.	398
Custom or limited access filter URLs are not filtered as expected.	398
A user cannot access a protocol or application as expected	398
An FTP request is not blocked as expected.	399
Websense software is not applying user or group policies	399
Remote users are not filtered by correct policy.	399
Network Agent issues.	400
Network Agent is not installed	400
Network Agent is not running.	400
Network Agent is not monitoring any NICs	401
Network Agent can't communicate with Filtering Service.	401
Update Filtering Service IP address or UID information	401
Insufficient memory on the Network Agent machine.	402
High CPU usage on the Network Agent machine.	402
User configuration and identification issues	402
User and group-based policies are not applied	403
DC Agent has insufficient permissions.	404
DC Agent unable to access required file.	404
General DC Agent troubleshooting.	405
Users are incorrectly filtered by the Default policy.	406
Troubleshooting Logon Agent	406
Group Policy Objects.	407
Domain controller visibility.	407
NetBIOS.	407
User profile issues	408
Troubleshooting eDirectory Agent	409
Enabling eDirectory Agent diagnostics	410
eDirectory Agent miscounts eDirectory Server connections.	410
Running eDirectory Agent in console mode	411
Troubleshooting RADIUS Agent	411
Running RADIUS Agent in console mode	412
I cannot add users and groups to TRITON - Web Security	412
Directory service connectivity and configuration	413
Directory service configuration in TRITON - Web Security	413
User identification and Windows Server 2008	414
Turning on the Computer Browser service	414
Changing DC Agent, Logon Agent, and User Service permissions	415
User Service on Linux	415
Remote users are not prompted for manual authentication.	417

Remote users are not being filtered correctly	417
Block message issues	417
No block page appears for a blocked file type	418
Users receive a browser error instead of a block page	418
A blank white page appears instead of a block page.	419
Protocol block messages don't appear as expected.	419
A protocol block message appears instead of a block page	419
Log, status message, and alert issues	420
Where do I find error messages for Websense components?	420
Websense Health alerts	421
Two log records are generated for a single request.	422
Usage Monitor is not available	423
Usage Monitor is not running	423
Policy Server and Policy Database issues	423
I forgot my password.	424
I cannot log on to Policy Server	424
The Websense Policy Database service fails to start	424
Policy Server stops unexpectedly	425
Delegated administration issues	425
Managed clients cannot be deleted from role	425
Logon error says someone else is logged on at my machine	426
Some users cannot access a site in the Unfiltered URLs list	426
Recategorized sites are filtered according to the wrong category.	426
I cannot create a custom protocol	426
Log Server and Log Database issues	426
Log Server is not running	427
Low disk space on the Log Server machine	427
No Log Server is installed for a Policy Server	428
Log Database was not created.	429
Log Database is not available	430
Log Database size	431
Configure Log Server to use a database account	431
Log Server is not recording data in the Log Database	431
Updating the Log Server connection account or password.	432
Configuring user permissions for Microsoft SQL Server	432
Log Server cannot connect to the directory service	433
Wrong reporting page displayed.	434
Investigative report and presentation report issues	434
Presentation Reports Scheduler not connected to Log Database	434
Inadequate disk space to generate presentation reports	435
Scheduled jobs in presentation reports failed	435
Data on Internet browse time reports is skewed	435

Bandwidth is larger than expected	436
Some protocol requests are not being logged	436
All reports are empty.	436
Database partitions.	436
SQL Server Agent job	437
Log Server configuration	437
Microsoft Excel output is missing some report data.	438
Saving presentation reports output to HTML	438
Error generating presentation report, or report does not display.	439
Investigative reports search issues	439
General investigative reports issues	439
Other reporting issues.	440
Low memory on the Real-Time Monitor machine	440
Real-Time Monitor is not running	440
Real-Time Monitor is not responding.	441
Cannot access certain reporting features.	442
No charts appear on Today or History pages	442
Interoperability issues.	442
Content Gateway is not running	443
Content Gateway is not available	443
Content Gateway non-critical alerts	443
Administrator unable to access other TRITON modules	446
Sync Service is not available.	446
Sync Service has been unable to download log files	447
Sync Service has been unable to send data to Log Server	448
Hybrid filtering data does not appear in reports	448
Disk space is low on the Sync Service machine.	449
The Sync Service configuration file	449
Directory Agent is not running	450
Directory Agent cannot connect to the domain controller	450
Directory Agent communication issues	451
Directory Agent does not support this directory service.	452
The Directory Agent configuration file.	452
Directory Agent command-line parameters	454
Alerts were received from the hybrid service	454
Unable to connect to the hybrid service	455
Hybrid service unable to authenticate connection.	455
Missing key hybrid configuration information	456
User access no longer filtered following upgrade.	456
Troubleshooting tips and tools	457
Where is the Websense “bin” directory?.	457
The Windows Services dialog box	458
The Windows Event Viewer	458

The Websense log file.....458

1

Getting Started

To learn to use Websense Web Security and Websense Web Filter and find answers to your questions, you can browse this guide, or use select one of the following topics as a launch point.

First steps <ul style="list-style-type: none">• <i>Working in TRITON - Web Security</i>• <i>Your subscription</i>• <i>The Websense Master Database</i>• <i>Websense Web Security product components</i>	Initial solutions <ul style="list-style-type: none">• <i>Installation and subscription issues</i>• <i>Master Database issues</i>• <i>Troubleshooting tips and tools</i>
Start filtering <ul style="list-style-type: none">• <i>Filtering categories and protocols</i>• <i>Adding a client</i>• <i>Working with policies</i>• <i>Assigning a policy to clients</i>	Filtering solutions <ul style="list-style-type: none">• <i>Filtering issues</i>• <i>Network Agent issues</i>• <i>User configuration and identification issues</i>• <i>Block message issues</i>
Using reports <ul style="list-style-type: none">• <i>Presentation reports</i>• <i>Investigative reports</i>• <i>Real-Time Monitor</i>• <i>Using the Toolbox to verify filtering behavior</i>	Reporting solutions <ul style="list-style-type: none">• <i>Log, status message, and alert issues</i>• <i>Log Server and Log Database issues</i>• <i>Investigative report and presentation report issues</i>
Advanced tools <ul style="list-style-type: none">• <i>Redefining filtering for specific sites</i>• <i>Restricting users to a defined list of Internet sites</i>• <i>Delegated Administration and Reporting</i>	Other solutions <ul style="list-style-type: none">• <i>Delegated administration issues</i>• <i>Policy Server and Policy Database issues</i>

Overview

Working in conjunction with integration devices—including proxy servers, firewalls, routers, and caching appliances—Websense Web Security provides the engine and configuration tools to develop, monitor, and enforce Internet access policies.

Together, a series of Websense components (described in *Websense Web Security product components*, page 314) provide Internet filtering, user identification, alerting, reporting, and troubleshooting capabilities.

An overview of the new features included in this Websense software version can be found in the [Release Notes](#), available from the [Websense Technical Library](#).

After installation, Websense software applies the **Default** policy to monitor Internet usage without blocking requests. This policy governs Internet access for all clients in the network until you define your own policies and assign them to clients. Even after you have created your custom filtering settings, the Default policy is applied any time a client is not governed by any other policy. See *The Default policy*, page 80, for more information.

The process of creating filters, adding clients, defining policies, and applying policies to clients is described in:

- ◆ *Internet Usage Filters*, page 39
- ◆ *Clients*, page 61
- ◆ *Internet Filtering Policies*, page 79

A single, browser-based tool—the TRITON™ Unified Security Center—provides a central, graphical interface to the general configuration, policy management, and reporting functions of your Websense Web Security, Data Security, and Email Security software. See *Working in TRITON - Web Security*, page 18, for more information.

You can define levels of access to the TRITON Unified Security Center to allow certain administrators to manage one or more TRITON modules. Within the Web Security module, you can further refine access permissions to allow administrators to configure Internet filtering behavior, manage filtering policies, perform reporting tasks, and more. See *Delegated Administration and Reporting*, page 285, for more information.

Working in TRITON - Web Security

Related topics:

- ◆ *Navigating in TRITON - Web Security*, page 21
- ◆ *Today: Health, Security, and Value Since Midnight*, page 24
- ◆ *History: Last 30 Days*, page 27

The TRITON Unified Security Center is the central configuration interface used to manage Websense Web Security, Email Security, and Data Security software. It includes a Web Security module (called TRITON - Web Security) used to customize filtering behavior, monitor Internet usage, generate Internet usage reports, and manage configuration and settings for Websense Web Security.



Important

Internet Explorer 8 Compatibility View is **not** supported with the TRITON console. If you experience odd behavior or page layouts in Internet Explorer 8, make sure that Compatibility View button (between the URL and the Refresh button in the browser address bar) is not selected.

At installation, the TRITON Unified Security Center is set up to give full access to all modules to a single administrator account: **admin**. The password for this account is set during installation.

When you log on using the admin account, or any other administrator account with Web Security permissions, Status > Today page for the Web Security module is displayed.

- ◆ If this is your first time logging on to TRITON - Web Security, you are offered the option of launching a Quick Start tutorial. If you are new to Websense software, or new to this version of Websense software, completing a Quick Start tutorial is highly recommended.
- ◆ If you are using an account with permissions to access multiple TRITON modules, use the TRITON toolbar to switch between modules. See [Navigating in TRITON - Web Security, page 21](#).
- ◆ If you are using delegated administration, and have created administrative roles, you may be prompted to select a role to manage. See [Delegated Administration and Reporting, page 285](#).

At logon, TRITON - Web Security connects to the default (base) Policy Server specified during installation. To manage another Policy Server, select its IP address from the Policy Server drop-down list in the Web Security toolbar.

A TRITON session ends 30 minutes after the last action taken in the user interface (clicking from page to page, entering information, caching changes, or saving changes). A warning message is displayed 5 minutes before the session ends.

- ◆ If there are uncached changes on the page or cached changes pending, the changes are lost when the session ends. Remember to click **OK** to cache and **Save All** to save and implement any changes.
- ◆ If TRITON - Web Security is open in multiple tabs of the same browser window, all instances share the same session. If the session times out in one tab, it times out in all tabs.

- ◆ If TRITON - Web Security is open in multiple browser windows on the same computer, the instances share the same session **if**:
 - You are using Microsoft Internet Explorer and use the Ctrl-N shortcut to open a new instance of TRITON - Web Security.
 - You are using Mozilla Firefox.

If the session times out in one window, it times out in all windows.

- ◆ In the following instances, you can open multiple TRITON - Web Security instances that do not share a session. In these situations, if one window times out, the others are not affected.
 - Launch multiple Internet Explorer windows independently of one another.
 - Use the File > New Session command to open a new Internet Explorer 8 window.
 - Use Internet Explorer to open one connection to TRITON - Web Security, and then use Firefox to open another connection.

If you close the browser without logging off of the TRITON Unified Security Center, or if the remote machine from which you are accessing a TRITON module shuts down unexpectedly, you may be temporarily locked out. Websense software typically detects this issue within about 2 minutes and end the interrupted session, allowing you to log on again.

Navigating in TRITON - Web Security

The TRITON - Web Security interface can be divided into 6 main areas:

1. Banner
2. TRITON toolbar
3. Web Security toolbar
4. Left navigation pane
5. Right shortcut pane
6. Content pane



This guide describes the options available to the **admin** account. Delegated administrators may see a subset of the features described. See [Delegated Administration and Reporting, page 285](#), for more information.

The banner

The banner, located at the top of the browser page, shows:

- ◆ The **user name** associated with your administrative logon account
- ◆ A **Log Off** button, for when you're ready to end your administrative session

The TRITON toolbar

The TRITON toolbar, located under the banner, allows you to:

- ◆ Move between modules of the TRITON Unified Security Center. Once you have opened multiple TRITON modules, use the mouse or the Alt-Tab (Windows) keyboard shortcut to move back and forth between modules.
- ◆ Connect to Appliance Manager for any V-Series **Appliances** deployed in your network.

- ◆ Configure global **TRITON Settings** that affect all installed modules.
- ◆ Access **Help**, tutorials, and Websense Technical Support resources.

The Web Security toolbar

The Web Security toolbar, located under the TRITON toolbar, is used to:

- ◆ Switch between the **Main** and **Settings** tabs of the left navigation pane.
- ◆ See which Policy Server you are currently connected to, and, in multiple Policy Server deployments, switch between Policy Server instances (see [Working with Policy Server](#), page 320).
- ◆ View your administrative **Role**, switch between roles, or release policy permissions for the current role.



Tip

If you have policy management and reporting permissions, but only reporting features are displayed, another administrator may be logged on to the role. Only one administrator at a time can access policy management features for each role.

- ◆ **View Pending Changes** (via the small magnifying glass icon) and **Save All** pending changes. If there are no cached changes waiting to be saved, these buttons are disabled.

See [Reviewing, saving, and discarding changes](#), page 23, for more information.

The left and right navigation panes

The left navigation pane has two tabs: **Main** and **Settings**. Use the **Main** tab to access status, reporting, and policy management features and functions. Use the **Settings** tab to manage your Websense account and perform global system administration tasks.

The right shortcut pane contains links to useful tools and common administrative tasks.

- ◆ **Common Tasks** provides shortcuts to frequently-performed administrative tasks. Click an item in the list to jump to the page where the task is performed.
- ◆ The **Toolbox** contains quick lookup tools that you can use to verify your filtering setup. See [Using the Toolbox to verify filtering behavior](#), page 242, for more information.

Both the left and right navigation panes can be minimized by clicking the double arrow (<< or >>) icon at the top of the pane. Click the reverse icon (>> or <<) to view the pane.

Mouse over a shortcut icon on the minimized left navigation pane to see a menu of related features without maximizing the pane.

Reviewing, saving, and discarding changes

When you make a change in TRITON - Web Security, you must typically first cache the change by clicking **OK** at the bottom of the page, then save the change to the Policy Database, which causes the change to take effect, by clicking **Save All**.

- ◆ Some fields or sections in TRITON - Web Security have their own **Save** or **Save Now** buttons. Changes to these features are saved and implemented immediately, rather than first being cached and later being saved.
- ◆ Some types of changes require you to click OK on both a subordinate page and a main page to cache changes.



Important

Avoid double- or triple-clicking the OK button. Multiple, rapid clicks to the same button can cause display problems in Mozilla Firefox that can be solved only by exiting and reopening the browser.

Use the **View Pending Changes** page to review cached changes. Changes to a single area of functionality are typically grouped into a single entry in the cache list. For example, if you add 6 clients and delete 2 clients, the cache list indicates only that changes were made to Clients. Changes to a single Settings page, on the other hand, may result in multiple entries in the cache list. This occurs when a single Settings page is used to configure multiple Websense software functions.

- ◆ To save all of the cached changes, click **Save All Changes**.
- ◆ To abandon all of the cached changes, click **Cancel All Changes**.

After choosing Save All or Cancel All, the View Pending Changes and Save All buttons in the toolbar are updated, and you are returned to the last page you selected. There is no undo for either the Save All or Cancel All functions.

Use the Audit Log to review the details of changes made in TRITON - Web Security. See [Viewing and exporting the audit log, page 328](#), for more information.

Today: Health, Security, and Value Since Midnight

Related topics:

- ◆ [Navigating in TRITON - Web Security, page 21](#)
- ◆ [History: Last 30 Days, page 27](#)
- ◆ [Customizing the Today page, page 26](#)

The **Status > Today: Health, Security and Value Since Midnight** page appears first when you log on to TRITON - Web Security. It displays alert messages and graphical charts that show the current state of your filtering software, focusing on Internet filtering activity in your network. The charts on this page cover the 24-hour period beginning at 12:01 a.m. according to the time set on the Log Database machine.



Note

On Linux machines, for best results, use Firefox 3.5 or later to view the Today page.

At the top of the page, 2 summary sections provide a quick overview of current status:

- ◆ The **Health Alert Summary** shows the status of your Websense software. If an error or warning appears in the summary, click the alert message to open the Alerts page, where more detailed information is available (see [Reviewing current system status, page 339](#)).

Information in the Health Alert Summary is updated every 30 seconds.

- ◆ Under **Today's Value**, see examples of how Websense filtering has protected your network today. Depending on your subscription type, this may show information about Malicious, Adult, and Spyware sites blocked, and about sites recategorized as a result of Content Gateway scanning and analysis.

This section also shows the total number of Internet requests handled so far today, the total number of requests blocked, and the number of real-time database updates processed.

Find longer term aggregate and trend information for many of these counters on the Status > History page. See [History: Last 30 Days, page 27](#).

Below the summary information, up to 4 Flash charts provide information about filtering activities.

- ◆ The default set of displayed charts varies based on subscription type.
- ◆ Charts are visible to Super Administrators, and to those delegated administrators with permission to view reports on the Today page. See [Editing roles, page 298](#).

Information in the charts is updated every 2 minutes. You may need to scroll down to see all of the charts.

Chart Name	Description
Current Filtering Load	Shows the number of filtered Internet requests processed into the Log Database per 10-minute interval.
Top Security Risks by Requests	Shows which Security Risk categories have received the most requests today to help you determine whether filtering policies are providing the right protection for your network.
Top Categories by Requests	Shows the categories that are being accessed most today to provide a high-level overview of potential security, bandwidth, or productivity concerns.
Policy Enforcement by Risk Class	Shows how many requests to each risk class have been permitted and blocked today (see Risk classes, page 43) to help you evaluate whether the current policies are effective.
Top Protocols by Bandwidth	Shows which protocols are using the most bandwidth in your network today. Use this information to evaluate bandwidth needs, and the potential need for policy changes.
Computers Requesting Security Risk Sites	Shows which computers have accessed Security Risk sites today. You may want to check these machines to make sure they are not infected with any viruses or spyware.
Top Blocked Users	Shows which users have requested the most blocked sites today to provide insight into compliance with your organization's Internet use standards.
Top Uncategorized Sites	Shows which sites not categorized by the Websense Master Database have been accessed most today. Go to Common Tasks > Recategorize URL to assign a site to a category for filtering.

If your subscription includes Websense Web Security Gateway or Websense Web Security Gateway Anywhere, these charts are also available.

Chart Name	Description
Top Categories by Scanned Requests	Shows the top categories to which requested sites were assigned after scanning determined that they no longer fit their original category.
Security Classification of Scanned Requests	Shows how many scanned requests were assigned to new categories because the content had been changed or the site was compromised.

Click any bar chart to open an investigative report with more details.

Up to 4 buttons appear above the page:

- ◆ **Database Download**, available to Super Administrators only, shows Master Database download status and provides the option to initiate or interrupt a download (see [Review Master Database download status, page 325](#)).

- ◆ *(Web Security Gateway Anywhere)* **Hybrid Service Status**, available to Super Administrators only, shows the results of recent attempts by on-premises components to communicate with hybrid filtering (see [Monitor communication with the hybrid service](#), page 197).
- ◆ **Customize**, available to Super Administrators only, allows you to change which charts appear on the page (see [Customizing the Today page](#), page 26).
- ◆ **Print** opens a secondary window with a printer-friendly version of the charts displayed on the Today page. Use browser options to print the page.
- ◆ **Status Monitor** releases the current administrator's policy permissions and enters a monitoring mode that allows access to the Status > Today, History, and Alerts pages, and to Real-Time Monitor, without timing out (see [Monitoring Web Security status](#), page 27).

Below the Internet activity and filtering charts, the **Filtering Service Summary** shows the status of each Filtering Service associated with the current Policy Server. Click the Filtering Service IP address to see more information about that Filtering Service instance, including its Network Agent and Content Gateway connections. See [Review Filtering Service details](#), page 325.

Customizing the Today page

Related topics:

- ◆ [Today: Health, Security, and Value Since Midnight](#), page 24
- ◆ [Customize the History page](#), page 30

Use the **Today > Customize** page to select up to 4 charts for the Status > Today page. Only Super Administrators with unconditional policy permissions (including admin) can customize the Today page.

The charts that you select appear on the Today page for all Super Administrators, and for delegated administrators who have permission to view charts on the Today page. See [Editing roles](#), page 298.

Some charts show potentially sensitive information, such as user names or IP addresses. Be sure that the charts you select are appropriate for all of the administrators who may view them.

To select charts, mark or clear the check box next to the chart name. When you are finished making selections, click **OK** to return to the Today page and view the charts. To return to the Today page without making changes, click **Cancel**.

For a short description of the information displayed in each chart, see [Today: Health, Security, and Value Since Midnight](#), page 24.

Monitoring Web Security status

For security purposes, a TRITON - Web Security session ends after 30 minutes of inactivity. You can, however, enter a Status Monitor mode that lets you monitor filtering and alerting data without timing out.

- ◆ You must log off of other TRITON modules to enter Status Monitor mode in TRITON - Web Security.
- ◆ In Status Monitor mode, information on the Status > Today, History, and Alerts pages and the Reporting > Real-Time Monitor page continues to update normally until you close the browser or log off.

To initiate Status Monitor mode, first save or discard any pending changes, then:

- ◆ Select **Status Monitor** mode from the **Role** drop-down list in the Web Security toolbar.
- ◆ Click the **Status Monitor** button in the toolbar at the top of the Status > Today, History, or Alerts. page.

To stop monitoring Web Security status, log off of TRITON - Web Security or close the browser.

History: Last 30 Days

Related topics:

- ◆ [Today: Health, Security, and Value Since Midnight, page 24](#)
- ◆ [Navigating in TRITON - Web Security, page 21](#)
- ◆ [Customize the History page, page 30](#)

Use the **Status > History: Last 30 Days** page to get an overview of filtering behavior for up to the past 30 days. The charts on the page are updated daily at 12:01 a.m. to incorporate data from the previous day, as determined by the time on the Log Database machine.



Note

On Linux machines, for best results, use Firefox 3.5 or later to view the History page.

The exact time period covered by the charts and summary tables depends on how long Websense software has been filtering. During the first month that Websense software is installed, the page shows data for the number of days since installation. After that, the reports cover the 30 days prior to today.

The **Value Estimates** at the top of the page provide an estimate of time and bandwidth savings afforded by Websense software, as well as a summary of blocked requests for categories that are of importance to many organizations.

Mouse over the **Time** or **Bandwidth** item (under Saved) for an explanation of how the estimate was calculated (see [Time and bandwidth saved](#), page 29). You can click **Customize** to change the way the values are calculated.

The **Blocked Requests** area further illustrates how Websense software has protected your network by listing several categories of interest to many organizations, and showing the total number of blocked requests to each during the time period.

Depending on the reporting permissions granted to the role, delegated administrators may not see the charts described below. See [Editing roles](#), page 298.

The page also includes up to 4 Flash charts with filtering highlights. You may need to scroll down to see all the charts. Information in the charts is updated once each day. Click a chart to launch an investigative report with more details.

Chart Name	Description
Internet Activity by Requests	Review the number of filtered Internet requests processed into the Log Database each day.
Top Security Risks by Requests	See which Security Risk categories have been accessed recently, and determine whether filtering policies are providing the right protection for your network.
Top Categories by Requests	See which categories have been accessed most. Get a high level overview of potential security, bandwidth, or productivity concerns.
Top Uncategorized Sites	Shows which sites not categorized by the Websense Master Database have been accessed most. Go to Common Tasks > Recategorize URL to assign a site to a category for filtering.
Top Protocols by Bandwidth	Shows which protocols have been using the most bandwidth in your network. Use this information to evaluate bandwidth needs and potential policy changes.
Policy Enforcement by Risk Class	Shows how many requests to each risk class have been permitted and blocked (see Risk classes , page 43). Evaluate whether the current policies are effective or whether changes are needed.
Top Blocked Users	Shows which users' Internet requests have been blocked the most. Gain insight into compliance with your organization's Internet use standards.
Policy Enforcement Summary	Provides an overview of recently permitted requests, blocked requests to sites in the Security Risk class, and blocked requests to other sites. Consider which aspects of filtering need a more detailed evaluation.

These additional reports are available if your subscription includes Websense Web Security Gateway or Websense Web Security Gateway Anywhere.

Chart Name	Description
Top Categories by Scanned Requests	Shows the top categories to which requested sites were assigned after scanning determined that they no longer fit their original category.
Security Classification of Scanned Requests	Shows how many scanned requests were assigned to new categories because the content had been changed or the site was compromised.
Top Web 2.0 Categories by Requests	Shows the top most frequently requested categories within the set of Web 2.0 sites. Use this information to learn more about Internet usage patterns and to discover potential productivity issues.
Top Web 2.0 Sites by Bandwidth	Shows the top Web 2.0 sites that consume the most bandwidth. Use this information to evaluate whether policy changes are needed to manage bandwidth.

Additional reports are available with a Websense Web Security Gateway Anywhere subscription only.

Chart Name	Description
Hybrid Service: Requests Processed	Shows how many requests by users from your organization were permitted and blocked by the hybrid service.
Hybrid Service: Bandwidth Usage	Shows the bandwidth consumed by Internet requests from users filtered by the hybrid service in your organization.

Two buttons appear above the page:

- ◆ **Customize**, available to Super Administrators only, opens a page where you can change which charts appear on the page, and to change how estimated savings are calculated (see [Customize the History page, page 30](#)).
- ◆ **Print**, available to all administrators, opens a secondary window with a printable version of the charts displayed on the History page. Use browser options to print this page, which omits all the navigation options found in the main TRITON - Web Security window.

Time and bandwidth saved

In addition to the improved security that Websense Web filtering offers, it also helps minimize the time and bandwidth lost to unproductive Internet activity.

The Saved section of the Value Estimates area presents an estimate of these time and bandwidth savings. These values are calculated as follows:

- ◆ Time saved: multiply the **typical time taken per visit** by the **sites blocked**. Initially, Websense software uses a default value as the average number of seconds that a user spends viewing a requested Web site. The sites blocked value

represents the total number of requests blocked during the time frame covered in the History page.

- ◆ **Bandwidth saved:** multiply the **typical bandwidth per visit** by the number of **sites blocked**. Initially, Websense software uses a default value as the average number of bytes consumed by the average Web site. The sites blocked value represents the total number of requests blocked during the time frame covered in the History page.

See [Customize the History page, page 30](#), for information on how to change the values used in these calculations to reflect usage at your organization.

Customize the History page

Related topics:

- ◆ [History: Last 30 Days, page 27](#)
- ◆ [Customizing the Today page, page 26](#)

Use the **History > Customize** page to determine which charts appear on the Status > History page, and to determine how time and bandwidth savings are calculated.

Mark the check box next to each chart name, up to 4, that you want to include on the History page. For a short description of each chart, see [History: Last 30 Days, page 27](#). Only Super Administrators with unconditional policy permissions (including admin) can customize the charts on the History page.

Some charts show potentially sensitive information, such as user names. Be sure that the charts you select are appropriate for all of the administrators who may view them.

Both Super Administrators and delegated administrators can customize the way that time and bandwidth savings are calculated. Delegated administrators access these fields by clicking the **Customize** link in the popup that describes the time and bandwidth saved calculations.

Enter new average time and bandwidth measurements to use as the basis for the calculation:

Option	Description
Average seconds saved per blocked page	Enter the average number of seconds that your organization estimates a user spends viewing individual pages. Websense software multiplies this value by the number of pages blocked to determine the time savings shown on the History page.
Average bandwidth [KB] saved per blocked page	Enter an average size, in kilobytes (KB), for pages viewed. Websense software multiplies this value by the number of pages blocked to determine the bandwidth savings shown on the History page.

When you are finished making changes, click **OK** to return the History page and view the new charts or time and bandwidth estimates. To return to the History page without making changes, click **Cancel**.

Your subscription

Websense subscriptions are issued on a per-client basis. A client is a user or IP address in your network.

When you purchase a subscription, a subscription key is provided via email.

Before you can begin filtering, you must enter a valid subscription key (see [Configuring your account information, page 32](#)). This lets you download the Master Database (see [The Websense Master Database, page 34](#)), which enables Websense software to filter clients.

After the first successful database download, TRITON - Web Security displays the number of clients your subscription includes and your subscription type (Web Filter, Web Security, Web Security Gateway, or Web Security Gateway Anywhere).

Websense software maintains a subscription table of clients filtered each day. The subscription table is cleared each night. The first time a client makes an Internet request after the table has been cleared, its IP address is entered in the table.

When the number of clients listed in the table reaches the subscribed maximum, any previously-unlisted client that requests Internet access exceeds the subscription. If this occurs, the client exceeding the subscription level is either blocked entirely from the Internet or given unfiltered Internet access, depending on a setting that you configure. Likewise, when a subscription expires, all clients are either entirely blocked or unfiltered, depending on this setting.

To configure filtering behavior when a subscription is exceeded or expires, see [Configuring your account information, page 32](#).

To configure Websense software to send email warnings when the subscription approaches or exceeds its limit, see [Configuring system alerts, page 335](#).

The number of categories filtered depends on your Websense subscription. Websense software filters all sites in all categories activated by your purchase.

Managing your account through the MyWebsense Portal

Websense, Inc., maintains a customer portal at mywebsense.com that you can use to access product updates, patches and hotfixes, product news, evaluations, and technical support resources for your Websense software.

When you create an account, the account is associated with your Websense subscription key or keys. This helps to ensure your access to information, alerts, and patches relevant to your Websense product and version.

Multiple members of your organization can create MyWebsense accounts associated with the same subscription key.

Configuring your account information

Related topics:

- ◆ [Your subscription, page 31](#)
- ◆ [Configuring database downloads, page 35](#)
- ◆ [Working with protocols, page 229](#)

Use the **Settings > General > Account** page to enter or review subscription information, and to determine how Websense software responds when the subscription expires, or the subscription level is exceeded.

Also use the page to direct Websense software to send category and protocol usage data to Websense, Inc., anonymously. This information may be used to help optimize the Websense Master Database to make filtering more efficient (see [The Websense Master Database, page 34](#)) and to contribute to the Websense Security Labs ThreatSeeker[®] Network (see websense.com/content/Threatseeker.aspx).

After installing Websense software, or any time you receive a new subscription key, use the **Subscription key** field to enter the key, and then click **Apply**. A check is done to verify the key syntax, and then Filtering Service attempts to download the Master Database.

- ◆ If a key is displayed, but the Subscription key field is disabled, you are connected to a secondary Policy Server. This means that the Policy Server instance gets its key information from the primary Policy Server whose IP address appears below the number of subscribed users.
- ◆ Use the Settings > General > Policy Servers page to manage subscription keys in multiple Policy Server environments (see [Working in a multiple Policy Server environment, page 322](#)).
- ◆ If the key syntax is correct, but the Master Database download fails because the key is invalid or expired, an health alert message is displayed on the Status > Today and Status > Alerts pages.

After the first successful Master Database download, the Account page displays the following information:

Key expires	End date for your current subscription. After this date, you must renew the subscription to continue downloading the Master Database and filtering your network.
Subscribed users	<i>Websense Web Security Gateway</i> : Sum of users filtered by on-premises components, hybrid filter, and remote filtering software

Subscribed network users	Number of in-network users that can be filtered.
Subscribed remote users	Number of users that can be filtered outside the network (requires optional remote filtering components).
Primary Policy Server	IP address of the Policy Server instance from which this Policy Server receives subscription key information. Appears only when viewing information for a secondary Policy Server.

1. Select **Block users when subscription expires or is exceeded** to:
 - Block all Internet access for all users when the subscription expires.
 - Block all Internet access for users who exceed the number of subscribed users.

If this option is not selected, users have unfiltered Internet access in these situations.
2. Mark **Send category and protocol data to Websense, Inc.** to have Websense software collect usage data about Websense-defined categories and protocols, and submit it anonymously to Websense, Inc.

This usage data helps Websense, Inc., to continually enhance the filtering capabilities of Websense software.
3. (*Websense Web Security Gateway Anywhere*) To activate or update the connection between the on-premises and hybrid portions of your Web security software:
 - Enter the **Contact email address** for your Web security administrators. This is typically a group email alias that is monitored frequently. Alerts about hybrid filtering issues are sent to this address. Failing to respond appropriately to an alert could lead to temporary disconnection of your hybrid service.
 - Enter the **Country** and **Time zone** in which the administrators are located.

Users are not filtered by the hybrid service until this information has been provided and validated. For more information, see [Configure Hybrid Filtering](#), page 173.
4. When you are finished making changes, click **OK**. Changes are not implemented until you click Save All.

The Websense Master Database

Related topics:

- ◆ [Real-time database updates](#), page 35
- ◆ [Real-Time Security Updates™](#), page 35
- ◆ [Configuring database downloads](#), page 35
- ◆ [Review Master Database download status](#), page 325
- ◆ [Resuming Master Database downloads](#), page 326

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content (see [Filtering categories and protocols](#), page 40).

- ◆ **Categories** are used to group Web sites (identified by URL and IP address) with similar content.
- ◆ **Protocol** definitions group Internet communications protocols used for similar purposes, like transferring files, or sending instant messages.

A limited version of the filtering database is installed during Websense software installation, but it is a good idea to download the full Master Database as soon as possible to enable comprehensive Internet filtering capabilities. To download the Master Database for the first time, enter your subscription key on the **Settings > General > Account** page (see [Configuring your account information](#), page 32).

If Websense software must go through a proxy to perform the download, also use the **Settings > General > Database Download** page to configure proxy settings (see [Configuring database downloads](#), page 35).

The process of downloading the full database may take a few minutes or more than 60 minutes, depending on factors such as Internet connection speed, bandwidth, available memory, and free disk space.

After the initial download, Websense software downloads database changes on a schedule that you establish (see [Configuring database downloads](#), page 35). Because the Master Database is updated frequently, by default, database downloads are scheduled to happen daily.

If the Master Database is more than 14 days old, Websense software does not filter Internet requests.

To initiate a database download at any time, or to view the status of the last database download, the date of the last download, or the current database version number, go to **Status > Today** and click **Database Download**.

Real-time database updates

In addition to scheduled downloads, Websense software performs emergency updates to the database as needed. A real-time update might be used, for example, to recategorize a site that was temporarily miscategorized. These updates ensure that sites and protocols are filtered appropriately.

Websense software checks for database updates every hour.

The most recent updates are listed on the **Status > Alerts** page (see [Reviewing current system status](#), page 339).

Real-Time Security Updates™

In addition to receiving the standard real-time database updates, users of Websense Web Security can enable Real-Time Security Updates to receive security-related updates to the Master Database as soon as they are published by Websense, Inc.

Real-Time Security Updates provide an added layer of protection against Internet-based security threats. Installing these updates as soon as they are published reduces vulnerability to new phishing (identify fraud) scams, rogue applications, and malicious code infecting mainstream Web sites or applications.

Filtering Service checks for security updates every 5 minutes, but because the updates are much smaller than full database updates, they tend not to disrupt normal network activity.

Use the **Settings > Database Download** page to enable Real-Time Security Updates (see [Configuring database downloads](#), page 35).

Configuring database downloads

Related topics:

- ◆ [Configuring your account information](#), page 32
- ◆ [The Websense Master Database](#), page 34
- ◆ [Review Master Database download status](#), page 325

Use the **Settings > General > Database Download** page to establish the schedule for automatic Master Database downloads. Also, provide important information about any proxy server or firewall Websense software must pass through to download the database.

1. Select the **Download days** for automatic downloads.

You must download the Master Database at least once every 14 days for Websense software to continue filtering uninterrupted. If you deselect all download days,

Websense software automatically attempts a download when the database is 7 days old.

**Note**

Download days are disabled when Real-Time Security Updates are enabled (see Step 3). Downloads are automatically performed every day to assure that the most up-to-date standard database is available for the security updates.

2. Select the starting time (**From**) and the ending time (**To**) for the **Download timeframe**. If no times are selected, the database download occurs between 21:00 (9 p.m.) and 06:00 (6 a.m.).

Websense software selects a random time during this period to contact the Master Database server. To configure alerts for download failures, see [Configuring system alerts](#), page 335.

**Note**

After downloading the Master Database, or updates to it, CPU usage can reach 90% while the database is loaded into local memory.

3. (*Websense Web Security*) Select **Enable real-time security updates** to have Websense software check for security updates to the Master Database every 5 minutes. When a security update is detected, it is downloaded immediately. Real-time security updates rapidly protect your network from vulnerability to threats like new phishing (identity fraud) scams, rogue applications, and malicious code infecting a mainstream Web site or application.
4. Select **Use proxy server or firewall** if Websense software must access the Internet through a proxy server or a proxying firewall to download the Master Database. Then, configure the following.

Server IP or name	Enter the IP address or name of the machine hosting the proxy server or firewall.
Port	Enter the port number through which the database download must pass (default is 8080).

- If the proxy server or firewall configured in step 4 requires authentication to reach the Internet, select **Use authentication**, and then enter the **User name** and **Password** that Websense software should use to gain Internet access.



Note

If Use authentication is selected, the proxy server or firewall must be configured to accept clear text or basic authentication to enable Master Database downloads.

By default, the user name and password are encoded to match the character set for the Policy Server machine's locale. This encoding can be configured manually via the **Settings > General > Directory Services** page (see [Advanced directory settings](#), page 68).

Websense Technical Support

Technical information about Websense software and services is available 24 hours a day at www.websense.com/support/, including:

- ◆ the searchable Websense Knowledge Base (made up of a Solution Center and Technical Library)
- ◆ forums, Webinars, and show-me tutorials
- ◆ product documents and in-depth technical papers
- ◆ answers to frequently asked questions

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

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

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

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

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

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

For telephone requests, please have ready:

- ◆ Websense subscription key
- ◆ Access to the Websense management console.
- ◆ Access to the machine running reporting tools and the database server (Microsoft SQL Server or SQL Server Express)
- ◆ Familiarity with your network's architecture, or access to a specialist

2

Internet Usage Filters

Related topics:

- ◆ [Filtering categories and protocols, page 40](#)
- ◆ [Working with filters, page 49](#)
- ◆ [Configuring Websense filtering settings, page 57](#)
- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Refine Filtering Policies, page 211](#)

Policies govern user Internet access. A policy is a schedule that tells Websense software how and when to filter access to Web sites and Internet applications. At their simplest, policies consist of:

- ◆ **Category filters**, used to apply actions (permit, block) to Web site categories
- ◆ **Protocol filters**, used to apply actions to Internet applications and non-HTTP protocols



Note

In Websense Web Security Gateway Anywhere environments, hybrid filtering does not enforce protocol filters.

- ◆ A schedule that determines when each filter is enforced

Policy-based filtering lets you assign varying levels of Internet access to clients (for example, users, groups, or IP addresses in your network). First, create filters to define precise Internet access restrictions, and then use the filters to construct a policy.

In a first-time installation, Websense software creates a **Default** policy and uses it to begin monitoring Internet requests as soon as a subscription key is entered (see [The Default policy](#), page 80). Initially, the Default policy permits all requests.

**Note**

When you upgrade from an earlier Websense software version, existing policy settings are preserved. After upgrading, review your policies to ensure that they are still appropriate.

To apply different filtering restrictions to different clients, start by defining category filters. You might define:

- ◆ One category filter that blocks access to all Web sites except those in the Business and Economy, Education, and News and Media categories
- ◆ A second category filter that permits all Web sites except those that represent a security risk and those containing adult material
- ◆ A third category filter that monitors access to Web sites without blocking them (see [Creating a category filter](#), page 50)

To accompany these category filters, you might define:

- ◆ One protocol filter that blocks access to Instant Messaging and Chat, P2P File Sharing, Proxy Avoidance, and Streaming Media protocol groups.
- ◆ A second protocol filter that permits all non-HTTP protocols except those associated with security risks and proxy avoidance
- ◆ A third protocol filter that permits all non-HTTP protocols (see [Creating a protocol filter](#), page 52)

Once you have defined a set of filters that correspond to your organization's Internet access regulations, you can add them to policies and apply them to clients (see [Internet Filtering Policies](#), page 79).

Filtering categories and protocols

The Websense Master Database organizes similar Web sites (identified by URLs and IP addresses) into **categories**. Each category has a descriptive name, like Adult Material, Gambling, or Peer-to-Peer File Sharing. You can also create your own, custom categories to group sites of particular interest to your organization (see [Creating a custom category](#), page 221). Together, the Master Database categories and user-defined categories form the basis for Internet filtering.

Websense, Inc., does not make value judgments about categories or sites in the Master Database. Categories are designed to create useful groupings of the sites of concern to subscribing customers. They are not intended to characterize any site or group of sites or the persons or interests who publish them, and they should not be construed as

such. Likewise, the labels attached to Websense categories are convenient shorthand and are not intended to convey, nor should they be construed as conveying, any opinion or attitude, approving or otherwise, toward the subject matter or the sites so classified.

The up-to-date list of Master Database categories is available at:

websense.com/global/en/ProductsServices/MasterDatabase/URLCategories.php

To suggest that a site be added to the Master Database, or that a site be moved from one category to another, click **Suggest New Category** in the right shortcut pane of TRITON - Web Security. You are prompted to log on to MyWebsense, then taken to the Site Lookup Tool.

When you create a **category filter** in TRITON - Web Security, you choose which categories to block and which to permit.

In addition to housing URL categories, the Websense Master Database includes protocol groups used to manage non-HTTP Internet traffic. Each protocol group defines similar types of Internet protocols (like FTP or IRC) and applications (like MSN Messenger or BitTorrent). The definitions are verified and updated as frequently as nightly.

As with categories, you can define custom protocols for use in Internet filtering.

The up-to-date list of Master Database protocols is available at:

websense.com/global/en/ProductsServices/MasterDatabase/ProtocolCategories.php

When you create a **protocol filter**, you choose which protocols to block and which to permit.

**Note**

In Websense Web Filter and Web Security deployments, Network Agent must be installed to enable protocol-based filtering.

With Websense Web Security Gateway, it is possible to filter non-HTTP protocols that tunnel over HTTP ports without using Network Agent. See [Tunneled protocol detection](#), page 160, for more information.

In Websense Web Security Gateway Anywhere environments, hybrid filtering does not enforce protocol filters.

Some Websense-defined protocols allow blocking of outbound Internet traffic destined for an external server—for example, a specific instant messaging server. Only Websense-defined protocols with dynamically-assigned port numbers can be blocked as outbound traffic.

New Master Database categories and protocols

When new categories and protocols are added to the Master Database, each is assigned a default filtering action, like **Permit** or **Block** (see [Filtering actions](#), page 46).

- ◆ The default action is applied in all active category and protocol filters (see [Working with filters](#), page 49). To change the way the category or protocol is filtered, you can:
 - Edit each active filter individually. Use this option if you want to give different groups of clients different levels of access to the category or protocol.
 - Edit the attributes of the category or protocol to apply the same action in all filters. See [Making global category filtering changes](#), page 220, and [Making global protocol filtering changes](#), page 233.
- ◆ The default action is based on feedback regarding whether or not the sites or protocols in question are generally considered business-appropriate.

You can configure Websense software to generate a system alert and notify you whenever new categories or protocols are added to the Master Database. See [Alerting](#), page 332, for more information.

Special categories

The Master Database contains special categories to help you manage specific types of Internet usage. The following categories are available in all editions of Websense software:

- ◆ The **Special Events** category is used to classify sites considered hot topics to help you manage event-related surges in Internet traffic. For example, the official World Cup site might generally appear in the Sports category, but be moved to the Special Events category during the World Cup Finals.

Updates to the Special Events category are added to the Master Database during scheduled downloads. Sites are added to this category for a short period of time, after which they are either moved to another category or deleted from the Master Database.

- ◆ The **Productivity** category focuses on preventing time-wasting behavior.
 - Advertisements
 - Freeware and Software Download
 - Instant Messaging
 - Online Brokerage and Trading
 - Pay-to-Surf
- ◆ The **Bandwidth** category focuses on saving network bandwidth.
 - Internet Radio and TV
 - Internet Telephony
 - Peer-to-Peer File Sharing

- Personal Network Storage and Backup
- Streaming Media

Websense Web Security includes additional security categories:

- ◆ **Websense Security Filtering** (also known simply as **Security**) focuses on Internet sites containing malicious code, which can bypass virus-detection software programs.
 - Bot Networks
 - Keyloggers
 - Malicious Web Sites
 - Malicious Embedded Link
 - Malicious Embedded iFrame
 - Phishing and Other Frauds
 - Potentially Unwanted Software
 - Spyware
 - Suspicious Embedded Link
- ◆ **Extended Protection** focuses on potentially malicious Web sites.
 - **Elevated Exposure** contains sites that camouflage their true nature or identity, or that include elements suggesting latent malign intent.
 - **Emerging Exploits** holds sites found to be hosting known and potential exploit code.
 - **Potentially Damaging Content** includes sites likely to contain little or no useful content.

The Extended Protection group filters potentially malicious Web sites based on *reputation*. Site reputation is based on early signs of potential malicious activity. An attacker might target a URL containing a common misspelling, for example, or otherwise similar to a legitimate URL. Such a site could be used to distribute malware to users before traditional filters can be updated to reflect these sites as malicious.

When Websense security researchers detect that a site includes a potential threat, the site is added to the Extended Protection category until researchers are 100% confident of the site's final categorization.

Risk classes

Related topics:

- ◆ [Assigning categories to risk classes, page 353](#)
- ◆ [Presentation reports, page 107](#)
- ◆ [Investigative reports, page 127](#)

The Websense Master Database groups categories into **risk classes**. Risk classes suggest possible types or levels of vulnerability posed by sites in the group of categories.

Risk classes are used primarily in reporting. The Today and History pages include graphs where Internet activity is displayed by risk class, and you can generate presentation or investigative reports organized by risk class.

Risk classes may also be helpful in creating category filters. Initially, for example, the Basic Security category filter blocks all of the default categories in the Security Risk class. You might use the risk class groupings as a guideline when you create your own category filters, to help decide whether a category should be permitted, blocked, or restricted in some way.

Websense software includes 5 risk classes, listed below. By default, Websense software groups the following categories into each risk class.

- ◆ A category can appear in multiple risk classes, or not be assigned to any risk class.
- ◆ The groupings may be changed periodically in the Master Database.

Legal Liability

Adult Material (including Adult Content, Lingerie and Swimsuit, Nudity, and Sex)

Bandwidth > Peer-to-Peer File Sharing

Gambling

Illegal or Questionable

Information Technology > Hacking and Proxy Avoidance

Militancy and Extremist

Racism and Hate

Tasteless

Violence

Weapons

Network Bandwidth Loss

Bandwidth (including Internet Radio and TV, Internet Telephony, Peer-to-Peer File Sharing, Personal Network Storage and Backup, and Streaming Media)

Entertainment > MP3 and Audio Download Services

Productivity > Advertisements and Freeware, as well as Software Download

Business Usage

Business and Economy (including Financial Data and Services, as well as Hosted Business Applications)

Education > Educational Materials, as well as Reference Materials

Government (including Military)

Information Technology and (including Computer Security, Search Engines and Portals, URL Translation Sites, and Web Collaboration)

Business Usage

Travel
Vehicles

Security Risk

Bandwidth > Peer-to-Peer File Sharing
Extended Protection (including Elevated Exposure, Emerging Exploits, and Potentially Damaging Content) [*Websense Web Security*]
Information Technology > Hacking, Proxy Avoidance, and Web and Email Spam
Parked Domain
Productivity > Freeware and Software Download
Security (including Bot Networks, Keyloggers, Malicious Embedded iFrame, Malicious Embedded Link, Malicious Web Sites, Phishing and Other Frauds, Potentially Unwanted Software, Spyware, and Suspicious Embedded Link) [*Websense Web Security*]

Productivity Loss

Abortion (including both Pro-Choice and Pro-Life)
Adult Material > Sex Education
Advocacy Groups
Bandwidth > Internet Radio and TV, Peer-to-Peer File Sharing, and Streaming Media
Drugs (including Abused Drugs, Marijuana, Prescribed Medications, and Supplements and Unregulated Compounds)
Education (including Cultural Institutions as well as Educational Institutions)
Entertainment (including MP3 and Audio Download Services)
Gambling
Games
Government > Political Organizations
Health
Information Technology > Web and Email Spam, as well as Web Hosting
Internet Communication (including General Email, Organizational Email, Text and Media Messaging, and Web Chat)
Job Search
News and Media (including Alternative Journals)
Parked Domain
Productivity (including Freeware and Software Download, Instant Messaging, Message Boards and Forums, Online Brokerage and Trading, and Pay-to-Surf)
Religion (including Non-Traditional Religions and Occult and Folklore as well as Traditional Religions)
Shopping (including both Internet Auctions and Real Estate)

Productivity Loss

Social Organizations (including Professional and Worker Organizations, Service and Philanthropic Organizations, and Social and Affiliation Organizations)

Society and Lifestyles (including Alcohol and Tobacco, Gay or Lesbian or Bisexual Interest, Hobbies, Personals and Dating, Restaurants and Dining, and Social Networking)

Special Events

Sports (including Sport Hunting and Gun Clubs)

Travel

Vehicles

Super Administrators can change the categories assigned to each risk class on the **Settings > General > Risk Class** page (see [Assigning categories to risk classes](#), page 353).

Security protocol groups

In addition to the Security and Extended Protection categories, Websense Web Security includes two protocol groups intended to help detect and protect against spyware and malicious code or content transmitted over the Internet.

- ◆ The **Malicious Traffic** protocol group includes the **Bot Networks** protocol, aimed at blocking command-and-control traffic generated by a bot attempting to connect with a botnet for malicious purposes.
- ◆ The **Malicious Traffic (Cannot block)** protocol group is used to identify traffic that may be associated with malicious software.
 - **Email-Borne Worms** tracks outbound SMTP traffic that may be generated by an email-based worm attack.
 - **Other** tracks inbound and outbound traffic suspected of connection with malicious applications.

The Malicious Traffic protocol group is blocked by default, and can be configured within your protocol filters (see [Editing a protocol filter](#), page 53). The Malicious Traffic (Cannot block) protocols can be logged for reporting, but no other filtering action can be applied.

Filtering actions

Category and protocol filters assign an **action** to each category or protocol. This is the action that Websense filtering software takes in response to a client's Internet request. The actions that apply to both categories and protocols are:

- ◆ **Block** the request. Users receive a block page or block message, and are not able to view the site or use the Internet application.
- ◆ **Permit** the request. Users can view the site or use the Internet application.

- ◆ Evaluate current **Bandwidth** usage before blocking or permitting the request. When this action is enabled, and bandwidth usage reaches a specified threshold, further Internet requests for a specific category or protocol are blocked. See [Using Bandwidth Optimizer to manage bandwidth, page 236](#).

Additional actions can be applied only to categories.

**Note**

The Confirm and Quota options should not be used when individual clients (users or computers) are managed by multiple Policy Servers.

The timing information associated with these features is not shared among Policy Servers, and affected clients could be granted more or less Internet access than you intend.

- ◆ **Confirm**—Users receive a block page, asking them to confirm that the site is being accessed for business purposes. If a user clicks **Continue**, she can view the site.
Clicking Continue starts a timer. During the configured time period (60 seconds by default), the user can visit other sites in Confirm categories without receiving another block page. Once the time period ends, browsing to any other Confirm site results in another block page.
The default time can be changed on the **Settings > General > Filtering** page.
- ◆ **Quota**—Users receive a block page, asking them whether to use quota time to view the site. If a user clicks **Use Quota Time**, he can view the site.
Clicking Use Quota Time starts two timers: a quota session timer and a total quota allocation timer.
 - If the user requests additional quota sites during a default **session** period (10 minutes by default), he can visit those sites without receiving another block page.
 - **Total** quota time is allocated on a daily basis. Once it is used up, each client must wait until the next day to access sites in quota categories. The default daily quota allocation (60 minutes by default) is set on the **Settings > General > Filtering** page. Daily quota allocations can also be granted to clients on an individual basis. See [Using quota time to limit Internet access, page 48](#), for more information.
- ◆ **Block Keywords**—When you define keywords and enable keyword blocking, users requesting a site whose URL contains a blocked keyword are not allowed to access the site. See [Filtering based on keyword, page 222](#).
- ◆ **Block File Types**—When file type blocking is enabled, users attempting to download a file whose type is blocked receive a block page, and the file is not downloaded. See [Managing traffic based on file type, page 238](#).

Using quota time to limit Internet access

When a user clicks Use Quota Time, she can view sites in any quota category until the quota session ends. The default quota session time (configured via the **Settings > General > Filtering** page) is 10 minutes.



Note

The Quota option should not be used when individual clients are managed by multiple Policy Servers.

The timing information associated with this feature is not shared among Policy Servers, and affected clients could be granted more or less Internet access than you intend.

Once the quota session ends, a request for a quota site results in another quota block message. Users who have not depleted their daily quota allocation can start a new quota session.

Once quota time is configured, Websense software uses a priority list to determine how to respond when a user requests a site in a quota category. The software looks for quota time configured for:

1. The user
2. The computer or network client
3. Groups to which the user belongs

If a user is a member of multiple groups, Websense software grants quota time according to the **Use more restrictive blocking** setting on the **Settings > General > Filtering** page (see [Configuring Websense filtering settings, page 57](#)).

4. Default quota time

Internet applets, such as Java or Flash applets, may not respond as expected to quota time restrictions. Even if it is accessed from a quota-restricted site, an applet that runs within the browser can continue running beyond the configured quota session time.

This is because such applets are downloaded completely to a client machine and run just like applications, without communicating back to the original host server. If the user clicks the browser's Refresh button, however, Websense software detects the communication to the host server, and then blocks the request according to applicable quota restrictions.

Search filtering

Search filtering is a feature offered by some search engines that helps to limit the number of inappropriate search results displayed to users.

Ordinarily, Internet search engine results may include thumbnail images associated with sites matching the search criteria. If those thumbnails are associated with blocked

sites, Websense software prevents users from accessing the full site, but does not prevent the search engine from displaying the image.

When you enable search filtering, Websense software activates a search engine feature that stops thumbnail images associated with blocked sites from being displayed in search results. Enabling search filtering affects both local and remote filtering clients.

Websense, Inc., maintains a database of search engines with search filtering capabilities. When a search engine is added to or removed from the database, an alert is generated (see [Alerting](#), page 332).

Search filtering is activated via the **Settings > General > Filtering** page. See [Configuring Websense filtering settings](#), page 57, for more information.

Working with filters

Related topics:

- ◆ [Filtering categories and protocols](#), page 40
- ◆ [Internet Filtering Policies](#), page 79
- ◆ [Creating a category filter](#), page 50
- ◆ [Creating a protocol filter](#), page 52
- ◆ [Creating a limited access filter](#), page 213

Use the **Policy Management > Filters** page in TRITON - Web Security to view, create, and modify category and protocol filters, and to work with other filtering tools.

The Filters page is divided into 3 main sections:

- ◆ **Category Filters** determine which categories to block and permit.
- ◆ **Protocol Filters** determine which non-HTTP protocols to block and permit.

Network Agent must be installed to enable full protocol-based filtering.

With Websense Web Security Gateway, it is possible to filter non-HTTP protocols that tunnel over HTTP ports without using Network Agent. See [Tunneled protocol detection](#), page 160, for more information.

In Websense Web Security Gateway Anywhere environments, the hybrid service does not provide protocol filtering.

- ◆ **Limited Access Filters** define a restrictive list of permitted Web sites (see [Restricting users to a defined list of Internet sites](#), page 212).

Category, protocol, and limited access filters form the building blocks of **policies**. Each policy is made up of at least one category or limited access filter, and one protocol filter, applied to selected clients on a specific schedule.

- ◆ To review or edit an existing category, protocol, or limited access filter, click the filter name. For more information, see:
 - [Editing a category filter](#), page 51
 - [Editing a protocol filter](#), page 53
 - [Editing a limited access filter](#), page 214
- ◆ To create a new category, protocol, or limited access filter, click **Add**. For more information, see:
 - [Creating a category filter](#), page 50
 - [Creating a protocol filter](#), page 52
 - [Creating a limited access filter](#), page 213

To duplicate an existing filter, mark the check box next to the filter name, and then click **Copy**. The copy is given the name of the original filter with a number appended for uniqueness, and then added to the list of filters. Edit the copy just as you would any other filter.

If you have created delegated administration roles (see [Delegated Administration and Reporting](#), page 285), Super Administrators can copy filters that they have created to other roles for use by delegated administrators.

To copy filters to another role, first mark the check box next to the filter name, and then click **Copy to Role**. See [Copying filters and policies to roles](#), page 216, for more information.

Creating a category filter

Related topics:

- ◆ [Working with filters](#), page 49
- ◆ [Editing a category filter](#), page 51

Use the **Policy Management > Filters > Add Category Filter** page to create a new category filter. You can work from a predefined template, or make a copy of an existing category filter to use as the basis for the new filter.

1. Enter a unique **Filter name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
Filter names can include spaces, dashes, and apostrophes.
2. Enter a short **Description** of the filter. This description appears next to the filter name in the Category Filters section of the Filters page, and should explain the filter's purpose.

The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. Select an entry from the drop-down list to determine whether to use a template or make a copy of an existing filter. For more information about templates, see [Category and protocol filter templates, page 56](#).
4. To see and edit the new filter, click **OK**. The filter is added to **Category Filters** list on the Filters page.

To customize the filter, click the filter name, and then continue with [Editing a category filter](#).

Editing a category filter

Related topics:

- ◆ [Filtering categories and protocols, page 40](#)
- ◆ [Filtering actions, page 46](#)
- ◆ [Using quota time to limit Internet access, page 48](#)
- ◆ [Working with filters, page 49](#)
- ◆ [Working with categories, page 218](#)

Use the **Policy Management > Filters > Edit Category Filter** page to make changes to existing category filters.



Important

When you edit a category filter, the changes affect every policy that enforces the filter.

Policies that enforce a category filter with the same name in another delegated administration role are not affected.

The filter name and description appear at the top of the page.

- ◆ Click **Rename** to change the filter name.
- ◆ Simply type in the **Description** field to change the filter description.

The number next to **Policies using this filter** shows how many policies currently use the selected filter. If the category filter is active, click **View Policies** for a list of policies that enforce the filter.

The bottom portion of the page shows a list of categories and the actions currently applied to each.

1. Select an entry in the **Categories** list to view category information or to change the filtering action associated with the selected category.
2. Before making changes to the action applied to a category, use the **Category Details** section to review any special attributes associated with the category.

- To review recategorized or unfiltered URLs assigned to the category, if any, click **See custom URLs in this category**. See [Redefining filtering for specific sites](#), page 224.
 - To review keywords assigned to the category, click **See keywords in this category**. See [Filtering based on keyword](#), page 222.
 - To review regular expressions used to define custom URLs or keywords for the category, click **See regular expressions in this category**.
3. Use the buttons at the bottom of the category list to change the action applied to the selected category. For more information about the available actions, see [Filtering actions](#), page 46.
- Delegated administrators cannot change the action assigned to categories that have been locked by a Super Administrator.
4. Use the check boxes to the right of the Categories list to apply advanced filtering actions to the selected category:
- To change the way that keywords are used in filtering the selected category, mark or clear **Block keywords**. [Filtering based on keyword](#), page 222
 - To determine whether users can access certain types of files from sites in the selected category, mark or clear **Block file types**. See [Managing traffic based on file type](#), page 238.
If you have chosen to block file types, select one or more file types to block.
 - To specify whether access to sites in the category is limited based on certain bandwidth thresholds, mark or clear **Block with Bandwidth Optimizer**. See [Using Bandwidth Optimizer to manage bandwidth](#), page 236.
If you have chosen to block based on bandwidth, specify which threshold limits to use.
5. Repeat steps 1 through 3 to make changes to the filtering actions applied to other categories.
6. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

To activate a new category filter, add it to a policy and assign the policy to clients. See [Internet Filtering Policies](#), page 79.

Creating a protocol filter

Related topics:

- ◆ [Filtering categories and protocols](#), page 40
- ◆ [Filtering actions](#), page 46
- ◆ [Editing a protocol filter](#), page 53
- ◆ [Working with protocols](#), page 229

Use the **Policy Management > Filters > Add Protocol Filter** page to define a new protocol filter. You can work from a predefined template or make a copy of an existing protocol filter to use as the basis for the new filter.

1. Enter a unique **Filter name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Filter names can include spaces, dashes, and apostrophes.

2. Enter a short **Description** of the filter. This description appears next to the filter name in the Protocol Filters section of the Filters page, and should explain the filter's purpose.

The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. Select an entry from the drop-down list to determine whether to use a template (see [Category and protocol filter templates](#), page 56) or make a copy of an existing filter as a basis for the new filter.
4. To see and edit the new filter, click **OK**. The filter is added to **Protocol Filters** list on the Filters page.

To finish customizing the new filter, continue with [Editing a protocol filter](#).

Editing a protocol filter

Related topics:

- ◆ [Filtering categories and protocols](#), page 40
- ◆ [Creating a protocol filter](#), page 52
- ◆ [Filtering actions](#), page 46
- ◆ [Working with protocols](#), page 229
- ◆ [Using Bandwidth Optimizer to manage bandwidth](#), page 236

Use the **Policy Management > Filters > Edit Protocol Filter** page to make changes to existing protocol filters.



Important

Changes that you make here affect all policies that enforce this filter.

Policies that enforce a protocol filter with the same name in another delegated administration role are not affected.

The filter name and description appear at the top of the page.

- ◆ Click **Rename** to change the filter name.

- ◆ Simply type in the **Description** field to change the filter description.

The number next to **Policies using this filter** shows how many policies currently use the selected filter. If the protocol filter is active, click **View Policies** for a list of policies that enforce the filter.

The bottom portion of the page shows a list of protocols and the actions currently applied to each.

To change the way that protocols are filtered and logged:

1. Select a protocol in the **Protocols** list. Advanced filtering actions for the selected protocol appear to the right of the list.
2. Use the **Permit** and **Block** buttons at the bottom of the Protocols list to change the action applied to the selected protocol.



Note

Websense software can block TCP-based protocol requests, but not UDP-based protocol requests.

Some applications use both TCP- and UDP-based messages. If an application's original network request is made via TCP, and then subsequent data is sent using UDP, Websense software blocks the initial TCP request and thus blocks subsequent UDP traffic.

UDP requests may be logged as blocked, even when they are permitted.

To apply the same action to the other protocols in the selected protocol group, click **Apply to Group**.

3. If you want information about use of the selected protocol available for alerting or reporting, mark the **Log protocol data** check box.
4. To impose bandwidth limits on the use of this protocol, click **Block with Bandwidth Optimizer**, and then supply the bandwidth thresholds to use. See [Using Bandwidth Optimizer to manage bandwidth, page 236](#), for more information.
5. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

To activate a new protocol filter, add it to a policy and apply the policy to clients (see [Internet Filtering Policies](#), page 79).

**Note**

You can create policies that start enforcing a protocol filter at a specific time. If users initiate a protocol session before that filter goes into effect, they can continue to access the protocol, even if the filter blocks it, for as long as the session continues. Once a user terminates the session, additional requests for the protocol are blocked.

Websense-defined category and protocol filters

Websense software includes several sample category and protocol filters. You can use these filters as they are, or modify them to suit your filtering needs. If you do not need the predefined filters, many of them can also be deleted.

The predefined category filters are:

- ◆ Basic
- ◆ Basic Security
- ◆ Block All
- ◆ Default
- ◆ Monitor Only
- ◆ Permit All

The Block All and Permit All category filters are not listed on the Filters page, though they can be added to policies. These filters play a special role in filtering, and cannot be deleted or edited. When an Internet request is filtered, Websense software first checks to see if the Block All or Permit All filter applies, before performing any additional filtering checks (see [Filtering a site](#), page 86).

The predefined protocol filters are:

- ◆ Basic Security
- ◆ Default
- ◆ Monitor Only
- ◆ Permit All

The Permit All protocol filter, like its equivalent category filter, is not listed on the Filters page and cannot be edited or deleted. It is also prioritized when filtering is performed.

The Default category and protocol filters can be edited, but cannot be deleted. In upgrade environments, if there are gaps in the Default policy, the Default filters are used to filter requests during periods when no other filter applies.

Category and protocol filter templates

When you create a new category or protocol filter, you can begin by making a copy of an existing filter on the Filters page, selecting an existing filter as a model on the Add Filter page, or using a filter **template**.

Websense software includes 5 category filter templates:

- ◆ **Monitor Only** and **Permit All** permits all categories.
- ◆ **Block All** blocks all categories.
- ◆ **Basic** blocks the most frequently blocked categories and permits the rest.
- ◆ **Default** applies the Block, Permit, Continue, and Quota actions to categories.
- ◆ **Basic Security** blocks only the default categories in the Security Risk class (see [Risk classes, page 43](#)).

Websense software also includes 3 protocol filter templates:

- ◆ **Monitor Only** and **Permit All** permit all protocols.
- ◆ **Basic Security** blocks the P2P File Sharing and Proxy Avoidance protocols, as well as Instant Messaging File Attachments (if subscribed) and Malicious Traffic (Websense Web Security).
- ◆ **Default** blocks the Instant Messaging / Chat protocols, as well as the P2P File Sharing, Proxy Avoidance, Instant Messaging File Attachments (if subscribed), and Malicious Traffic (Websense Web Security).

Although you can modify or delete most Websense-defined category and protocol filters, you cannot edit or remove templates. Likewise, although you can create as many custom filters as necessary, you cannot create new templates.

Because templates cannot be modified, they provide a constant method of referring back to the original filtering actions applied by Websense-defined filters. For example, the Default category and protocol filter templates apply the same actions as the original Default category and protocol filters. This means that you can always restore the original Websense filtering configuration by creating filters that use the template defaults.

For instructions on using a template to create a new filter, see [Creating a category filter, page 50](#), or [Creating a protocol filter, page 52](#).

Configuring Websense filtering settings

Related Topics:

- ◆ [Filtering categories and protocols, page 40](#)
- ◆ [Block Pages, page 91](#)
- ◆ [Password override, page 73](#)
- ◆ [Account override, page 74](#)
- ◆ [Using Bandwidth Optimizer to manage bandwidth, page 236](#)
- ◆ [Filtering based on keyword, page 222](#)

Use the **Settings > General > Filtering** page to establish basic settings for a variety of filtering features.

Under **Bandwidth Optimizer**, enter the information needed to filter Internet usage based on available bandwidth. For more information about bandwidth-based filtering, see [Using Bandwidth Optimizer to manage bandwidth, page 236](#).



Note

In Websense Web Security Gateway Anywhere environments, hybrid filtering does not use Bandwidth Optimizer settings. No bandwidth-based restrictions are enforced on requests passing through the hybrid service.

1. To specify an **Internet connection speed**, do one of the following:
 - Select a standard speed from the drop-down list.
 - Enter the network speed in kilobits per second in the text field.
2. Enter the default thresholds to use when bandwidth-based filtering is used. Note that when the thresholds are set, but no category or protocol filters enforce bandwidth-base filtering, no bandwidth usage restriction occurs.
 - **Network:** When total network traffic reaches this percentage of total available bandwidth, start enforcing bandwidth-based filtering, as configured in active filters.
 - **Protocol:** When traffic for a specific protocol (like HTTP or MSN Messenger) reaches this percentage of total available bandwidth, start restricting access to that protocol, as configured in active filters.
3. (*Websense Web Security Gateway*) Content Gateway can collect information about bandwidth consumed by HTTP traffic and protocols that tunnel over HTTP for use in reporting. To enable this option, mark **Include bandwidth data collected by Websense Content Gateway**.

Use the **General Filtering** section to determine how users are filtered when multiple group policies apply, specify keyword search options, and set password override, continue, and quota session behavior.

1. To determine how users are filtered when multiple group policies apply, mark or clear **Use most restrictive group policy** (see [Filtering order, page 85](#)).
 - When the option is selected, the policy that applies the most restrictive filtering setting is used. In other words, if one applicable group policy blocks access to a category and another permits access, the user’s request for a site in that category is blocked.
 - When the option is not selected, the most permissive setting is used.
2. Select one of the following **Keyword search options** (see [Filtering based on keyword, page 222](#)).

CGI only	Blocks sites when keywords appear in CGI query strings (after the “?” in a Web address). Example: search.yahoo.com/search?p=test Websense software does not search for keywords before the “?” when this is selected.
URL only	Blocks sites when keywords appear in the URL. If the requested address contains a CGI query string, Websense software searches for keywords up to the “?”.
URL and CGI	Blocks sites when keywords appear anywhere in the address. If a CGI query string is present, Websense software searches for keywords both before and after the “?”.
Disable keyword blocking	Use with caution. Disable keyword blocking turns off all keyword blocking, even if Block keywords is selected in a category filter.

3. In the **Password override timeout** field, enter the maximum number of seconds (up to 3600, default 60) that a user can access sites in all categories after selecting password override (see [Password override, page 73](#)).
4. In the **Continue timeout** field, enter the maximum time in seconds (up to 3600, default 60) that a user who clicks Continue can access sites in categories governed by the Confirm action (see [Filtering actions, page 46](#)).
5. In the **Account override timeout** field, enter the maximum time in minutes (up to 3600, default 5) that a user is filtered by the policy assigned to the override account (see [Account override, page 74](#)).
6. In the **Quota session length** field, enter the interval (up to 60 minutes, default 10) during which users can visit sites in quota-limited categories (see [Using quota time to limit Internet access, page 48](#)).
A session begins when the user clicks the Use Quota Time button.
7. Enter the **Default quota time per day** (up to 240 minutes, default 60) for all users.

To change the quota time for individual users, go to the **Policies > Clients** page. As you make changes to the quota session length and the default quota time per day, the **Default quota sessions per day** is calculated and displayed.

Use the **Block Messages** section to enter the URL or path to the alternative HTML block page you created for the top frame of browser-based block messages (see [Creating alternate block messages, page 100](#)).

- ◆ Separate pages can be used for the different protocols: **FTP**, **HTTP** (including **HTTPS**), and **Gopher**.
- ◆ Leave these fields blank to use the default block message provided with the Websense software, or a customized version of that message (see [Customizing the block message, page 96](#)).
- ◆ In Websense Web Security Gateway Anywhere environments, custom block messages are not applied to hybrid filtering requests.

Under **Search Filtering**, select **Enable search filtering** to have Websense software activate a setting built into certain search engines so thumbnail images and other explicit content associated with blocked sites are not displayed in search results (see [Search filtering, page 48](#)).

The search engines for which this feature is supported are displayed below the check box.

When you have finished configuring Filtering settings, click **OK** to cache the changes. Changes are not implemented until you click **Save All**.

3

Clients

You can customize how Websense software filters requests from specific users or machines by adding them as **clients** in TRITON - Web Security. If you are using on-premises Web filtering, clients can be:

- ◆ **Computers:** Individual machines in your network, defined by IP address.
- ◆ **Networks:** Groups of machines, defined collectively as an IP address range.
- ◆ **Directory clients:** User, group, or domain (OU) accounts in a supported directory service.



Note

In Websense Web Security Gateway Anywhere deployments, hybrid filtering can apply policies to users or groups, and to filtered locations (see [Working with hybrid filtering clients](#), page 76).

Initially, Websense software filters all clients in the same manner, using the **Default** policy (see [The Default policy](#), page 80). Once you add a client to the Clients page in TRITON - Web Security, you can assign that client a specific filtering policy.

When multiple policies could apply, such as when one policy is assigned to the user and another is assigned to the machine, Websense software determines which policy to enforce as follows:

1. Apply the policy assigned to the **user** making the request. If that policy has no filters scheduled at the time of the request, use the next applicable policy.
2. If there is no user-specific policy, or the policy has no active filters at the time of the request, look for a policy assigned to the **computer** (first) or **network** (second) from which the request was made.
3. If there is no computer or network-specific policy, or the policy has no active filters at the time of the request, look for a policy assigned to any **group** to which the user belongs. If the user belongs to multiple groups, Websense software considers all group policies that apply (see [Filtering order](#), page 85).
4. If there is no group policy, look for a policy assigned to the user's **domain** (OU).
5. If no applicable policy is found, or the policy does not enforce a category filter at the time of the request, enforce the **Default** policy for the role to which the client has been assigned.

For more detailed information about how Websense software filters requests, see [Filtering a site](#), page 86.

For information about how the hybrid service applies filtering policies to clients, see [Filtering order](#), page 85.

Working with clients

Related topics:

- ◆ [Clients](#), page 61
- ◆ [Working with computers and networks](#), page 63
- ◆ [Working with users and groups](#), page 64
- ◆ [Adding a client](#), page 71
- ◆ [Changing client settings](#), page 72

Use the **Policy Management > Clients** page to view information about existing clients, add, edit, or delete clients, or move clients to a delegated administration role.

If you are a delegated administrator, add clients to the Clients page from your managed clients list. This allows you to apply policies to the clients. See [Adding a client](#), page 71, for instructions.

Clients are divided into 3 groups:

- ◆ **Directory**, which includes users, groups, and domains (OUs) from your directory service (see [Working with users and groups](#), page 64).
- ◆ **Networks**, IP address ranges within the filtered network that can be governed by a single policy (see [Working with computers and networks](#), page 63).
- ◆ **Computers**, individual machines in the filtered network, identified by IP address (see [Working with computers and networks](#), page 63).

Click the plus sign (+) next to the client type to see a list of existing clients of the selected type. Each client listing includes:

- ◆ The client name, IP address, or IP address range.
- ◆ The **policy** currently assigned to the client. The **Default** policy is used until you assign another policy (see [Internet Filtering Policies](#), page 79).
- ◆ Whether or not the client can use a **password override** (see [Password override](#), page 73) or **account override** (see [Account override](#), page 74) option to view or attempt to view blocked sites.
- ◆ Whether the client has a custom amount of **quota time** allotted (see [Using quota time to limit Internet access](#), page 48).

To find a specific client, browse the appropriate node in the tree.

To edit client policy, password override, quota time, and authentication settings, select one or more clients in the list, and then click **Edit**. See [Changing client settings, page 72](#), for more information.

To add a client, or to apply a policy to a managed client who does not currently appear on the Clients page, click **Add**, and then go to [Adding a client, page 71](#), for more information.

If you have created delegated administration roles (see [Delegated Administration and Reporting, page 285](#)), Super Administrators can move their clients to other roles. First mark the check box next to the client entry, and then click **Move to Role**. When a client is moved to a delegated administration role, the policy and filters applied to the client are copied to the role. See [Moving clients to roles, page 75](#), for more information.

If you have configured Websense software to communicate with an LDAP-based directory service, the **Manage Custom LDAP Groups** button appears in the toolbar at the top of the page. Click this button to add or edit groups based on an LDAP attribute (see [Working with custom LDAP groups, page 69](#)).

To remove a client from TRITON - Web Security, select the client and click **Delete**.

Working with computers and networks

Related topics:

- ◆ [Working with clients, page 62](#)
- ◆ [Working with users and groups, page 64](#)
- ◆ [Adding a client, page 71](#)
- ◆ [Assigning a policy to clients, page 85](#)

In TRITON - Web Security, a **computer** is the IP address (for example, 10.201.3.1) associated with a filtered machine. A **network** is the IP address range (for example, 10.201.3.2 - 10.201.3.44) that corresponds to a group of filtered machines.



Note

In Websense Web Security Gateway Anywhere deployments, hybrid filtering does not apply policies to individual computer and network clients. See [Working with hybrid filtering clients, page 76](#), for information about applying policies to filtered locations.

You can assign policies to computer and network clients just as you would to user, group, or domain clients.

- ◆ Assign a policy to a **computer**, for example, that does not require users to log on, or that can be accessed by users with guest accounts.
- ◆ Assign a policy to a **network** to apply the same filtering policy to several machines at once.

When you assign a policy to a computer or network, that policy is enforced regardless of who is logged on to the filtered machine, **unless** you have assigned a policy to the logged-on user. When on-premises filtering components are used, the computer or network policy takes precedence over any **group** policies that may apply to the user. (In Websense Web Security Gateway Anywhere deployments, hybrid filtering applies the group policy before applying a computer or network policy. See [Working with hybrid filtering clients](#), page 76.)

Working with users and groups

Related topics:

- ◆ [Working with clients](#), page 62
- ◆ [Directory services](#), page 65
- ◆ [Working with custom LDAP groups](#), page 69
- ◆ [Working with computers and networks](#), page 63
- ◆ [Adding a client](#), page 71
- ◆ [Assigning a policy to clients](#), page 85

In order to apply policies to individual users and groups in your network, configure Websense software to access your directory service to obtain directory object (user, group, and domain [OU]) information.

Websense software can communicate with Windows Active Directory in mixed or native mode, and with Novell eDirectory or Oracle (formerly Sun Java) Directory Server Enterprise Edition accessed via Lightweight Directory Access Protocol (LDAP).

- ◆ When you use an LDAP-based directory service, duplicate user names are not supported. Ensure that the same user name does not appear in multiple domains.
- ◆ If you are using Windows Active Directory or Oracle (formerly Sun Java) Directory Server, user names with blank passwords are not supported. Make sure that all users have passwords assigned.

Websense User Service conveys information from the directory service to Policy Server and Filtering Service for use in applying filtering policies. Websense, Inc., recommends installing User Service on a Windows machine (though it can reside on a Linux machine).

To configure Websense software to communicate with your directory service, see [Directory services](#).

Directory services

A directory service is a tool that stores information about a network's users and resources. Before you can add user clients (users, groups, domains, or organizational units) in TRITON - Web Security, you must configure Websense software to retrieve information from your directory service.

Use the **Settings > General > Directory Services** page to identify the directory service used in your network. You can configure settings for only one type of directory service per Policy Server.



Note

In Websense Web Security Gateway Anywhere deployments, information from the Directory Services page is also used to populate the Hybrid Configuration > Shared User Data page. This supports user and group-based filtering by the hybrid service. See [Send user and group data to the hybrid service, page 189](#).

First select a directory service from the Directories list. The selection that you make determines which settings appear on the page.

See the appropriate section for configuration instructions:

- ◆ [Windows Active Directory \(Mixed Mode\), page 66](#)
- ◆ [Windows Active Directory \(Native Mode\), page 66](#)
- ◆ [Novell eDirectory and Oracle \(Sun Java\) Directory Server, page 67](#)



Warning

In Websense Web Security Gateway Anywhere deployments, hybrid filtering supports Windows Active Directory (Native Mode), Oracle Directory Server, and Novell eDirectory.

Once configuration is complete, User Service communicates with the directory service to enable user and group-based filtering. User Service caches the user and group information that it collects for up to 3 hours. If you make changes to user, group, or OU entries in the directory service, use the **Clear Cache** button under User Service Cache to force User Service to refresh its user and group mappings immediately. Note that user-based filtering may slow down for a brief period while the cache is being recreated.

If you plan to allow administrators to use their network accounts to log on to TRITON - Web Security, you must also configure directory service communication on the TRITON Settings > User Directory page. The same directory must be used to authenticate all administrative users. See the TRITON Settings Help for details.

Windows Active Directory (Mixed Mode)

If your directory service is Active Directory in mixed mode, no further configuration is necessary.

In rare circumstances, if you are using another directory service, you may need to supply additional information on this screen. This occurs only when:

- ◆ DC Agent is being used for transparent identification (see [DC Agent, page 255](#))
and
- ◆ User Service runs on a Linux machine

If this matches your configuration, provide the administrative credentials listed under Active Directory (Mixed Mode). You must also configure Websense software to communicate with a Windows Internet Name Server (WINS) to resolve domain names to domain controller IP addresses. See [User Service on Linux, page 415](#).

If your installation does not use this configuration, the administrative credential fields are disabled.

Windows Active Directory (Native Mode)

Windows Active Directory stores user information in one or more global catalogs. The global catalog lets individuals and applications find objects (users, groups, and so on) in an Active Directory domain.

In order for Websense software to communicate with Active Directory in Native Mode, you must provide information about the global catalog servers in your network.

1. Click **Add**, next to the Global catalog servers list. The Add Global Catalog Server page appears.
2. Use the **Server IP or name** field to identify the global catalog server:
 - If you have multiple global catalog servers configured for failover, enter the DNS domain name.
 - If your global catalog servers are not configured for failover, enter the IP address or host name (if name resolution is enabled in your network) of the server to add.
3. Enter the **Port** that Websense software should use to communicate with the global catalog (by default, **3268**).
4. Optionally, enter the **Root context** that Websense software should use to search for user information. If you supply a value, it must be a valid context in your domain.
 - If you have specified a communications port of 3268 or 3269, you do not need to supply a root context.
 - If the specified port is 389 or 636, you must provide a root context.

- If the Root context field is left blank, Websense software begins searching at the top level of the directory service.

**Note**

Avoid having the same user name in multiple domains. If Websense software finds duplicate account names for a user, the user cannot be identified transparently.

5. Specify which administrative account Websense software should use to retrieve user name and path information from the directory service. This account must be able to query and read from the directory service, but does not need to be able to make changes to the directory service, or be a domain administrator.

Select **Distinguished name by components** or **Full distinguished name** to specify how you prefer to enter the account information.

- If you selected Distinguished name by components, enter the **Display name**, account **Password**, **Account folder**, and **DNS domain name** for the administrative account. Use the common name (cn) form of the administrative user name, and not the user ID (uid) form.

**Note**

The **Account folder** field does not support values with the organizational unit (ou) tag (for example, *ou=Finance*). If your administrative account name contains an ou tag, enter the full distinguished name for the administrative account.

- If you selected Full distinguished name, enter the distinguished name as a single string in the **User distinguished name** field (for example, *cn=Admin, cn=Users, ou=InfoSystems, dc=company, dc=net*), and then supply the **Password** for that account.
6. Click **OK**.
 7. Repeat the process above for each global catalog server.
 8. Click **Advanced Directory Settings**, and then go to [Advanced directory settings](#), page 68.

Novell eDirectory and Oracle (Sun Java) Directory Server

To retrieve information from the directory service, Websense software requires the distinguished name, root context, and password for a user account with administrative privileges.

1. Enter the IP address of the directory server machine in the **Server IP** field.
2. Enter the **Port** number that Websense software will use to communicate with the directory. The default is 389.

3. If your directory requires administrator privileges for read-only access, enter the **Administrator distinguished name** and **Password**.
4. Optionally, enter the **Root Context** that Websense software should use when searching for user information. For example, *o=domain.com*.
Narrowing the context increases speed and efficiency in retrieving user information.

**Note**

Avoid having the same user name in multiple domains. If Websense software finds duplicate account names for a user, the user cannot be identified transparently.

5. Click **Advanced Directory Settings**, and then go to [Advanced directory settings](#), page 68.

Advanced directory settings

Related topics:

- ◆ [Windows Active Directory \(Native Mode\)](#), page 66
- ◆ [Novell eDirectory and Oracle \(Sun Java\) Directory Server](#), page 67

These settings can be used to define:

- ◆ How Websense software searches the directory service to find user, group, and domain information
- ◆ Whether Websense software uses an encrypted connection to communicate with the directory service
- ◆ Which character set Websense software uses to encode LDAP information

Configure these settings as needed for any LDAP-based directory service.

1. If you use custom object class types (attribute names) in your directory service, check **Use custom filters**. The default filter strings appear in the Filters fields.
2. Edit the existing filter strings, substituting object class types specific to your directory. For example, if your directory uses an object class type such as **dept** instead of **ou**, insert a new value in the Domain search filter field.

Attributes are always strings used in searching the directory service contents. Custom filters provide the functionality described here.

Attribute	Description
User logon ID attribute	Identifies user logon names
First name attribute	Identifies the user's given name

Attribute	Description
Last name attribute	Identifies the user's surname
Group attribute	Identifies the group's name
MemberOf attribute	Specifies that the user or group is a member of another group. If you are using Novell eDirectory, this corresponds to the groupMembership attribute.
User search filter	Determines how User Service searches for users
Group search filter	Determines how User Service searches for groups
Domain search filter	Determines how User Service searches for domains and organizational units
User's group search filter	Determines how User Service associates users with groups

3. To secure communications between Websense software and your directory service, check **Use SSL**.
4. To determine which character set Websense software uses to encode LDAP information, select **UTF-8** or **MBCS**.
MBCS, or multibyte character set, is commonly used for encoding East Asian languages such as Chinese, Japanese, and Korean.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Working with custom LDAP groups

Related topics:

- ◆ [Working with users and groups, page 64](#)
- ◆ [Directory services, page 65](#)
- ◆ [Adding or editing a custom LDAP group, page 70](#)

Use the **Manage Custom LDAP Groups** page to manage custom groups based on attributes defined in your directory service. This option is available only if you have

configured Websense software to communicate with an LDAP-based directory service.



Important

When you add custom LDAP groups to TRITON - Web Security, the group definitions are stored by the active Policy Server, and do not affect other Policy Server instances. To add custom LDAP groups to multiple Policy Servers, use TRITON - Web Security to connect to each Policy Server and enter the information.

If you add custom LDAP groups, and then either change directory services or change the location of the directory server, the existing groups become invalid. You must add the groups again, and then define each as a client.

- ◆ To add a group, click **Add** (see [Adding or editing a custom LDAP group, page 70](#)).
- ◆ To change an entry in the list, click on its group name (see [Adding or editing a custom LDAP group](#)).
- ◆ To remove an entry, first select it, and then click **Delete**.

When you are finished making changes to custom LDAP groups, click **OK** to cache the changes and return to the previous page. Changes are not implemented until you click **Save All**.

Adding or editing a custom LDAP group

Use the **Add Custom LDAP Group** page to define a group in TRITON - Web Security based on any attribute you have defined in your directory service. Use the **Edit Custom LDAP Group** page to make changes to an existing definition.



Important

If you add custom LDAP groups, and then either change directory services or change the location of the directory server, the existing groups become invalid. You must add the groups again, and then define each as a client.

1. Enter or change the **Group name**. Use a descriptive name that clearly indicates the purpose of the LDAP group.
Group names are case-insensitive, and must be unique.
2. Enter or change the description that defines this group in your directory service. For example:

(WorkStatus=parttime)

In this example, **WorkStatus** is a user attribute that indicates employment status, and **parttime** is a value indicating that the user is a part-time employee.

3. Click **OK** to return to the Manage Custom LDAP Groups page. The new or revised entry appears in the list.
4. Add or edit another entry, or click **OK** to cache changes and return to the previous page. Changes are not implemented until you click **Save All**.

Adding a client

Related topics:

- ◆ [Working with clients, page 62](#)
- ◆ [Working with computers and networks, page 63](#)
- ◆ [Working with users and groups, page 64](#)
- ◆ [Searching the directory service, page 72](#)
- ◆ [Changing client settings, page 72](#)

Use the **Policy Management > Clients > Add Clients** page to add user, group, computer, and network clients to TRITON - Web Security so that you can assign them a policy.

If you are logged on to a delegated administration role, you can only add clients that appear in your managed clients list. In the process of adding managed clients to the Clients page, you must assign them a policy.

1. Identify one or more clients:
 - To add a user, group, or domain (OU) client, browse the **Directory** tree to find entries in your directory service. If you are using an LDAP-based directory service, you can also click **Search** to enable a directory search tool (see [Searching the directory service, page 72](#)).
 - To add a computer or network client, enter an **IP address** or **IP address range**. No two network definitions can overlap, but a network client can include an IP address identified separately as a computer client. In the case of such an overlap, the policy assigned to the computer takes precedence over the policy assigned to the network.
2. Click an arrow button (>) to add each client to the **Selected Clients** list.
To remove an entry from the Selected Clients list, select the client, and then click **Remove**.
3. Select a **Policy** to assign to all clients in the Selected Clients list.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

The clients are added to the appropriate list on the **Policy Management > Clients** page. To change the policy assigned to one or more clients, or to configure additional client settings, select each client entry, and then click **Edit**. See [Changing client settings, page 72](#), for more information.

Searching the directory service

If you have configured Websense software to communicate with an LDAP-based directory service, you can use a search function to identify the directory clients you want to add in TRITON - Web Security. Search is also available for adding managed clients and administrators to delegated administration roles.

To search a directory service to retrieve user, group, and OU information:

1. Click **Search**.
2. Enter all or part of the user, group, or organizational unit **Name**.
3. Use the **Type** list to indicate the type of directory entry (user, group, OU, or all) that you want to find.

In a large directory service, selecting **All** may cause the search to take a very long time.

4. Browse the **Search context** tree to specify which portion of the directory to search. A more precise context helps to speed the search.
5. Click **Go**.
A list of search results is displayed.
6. Select one or more entry in search results, and then click the right arrow (>) to add each selection as a client or administrator.
7. Click **New Search** to enter another set of search criteria.
8. Click **Browse** to return to browsing the directory.
9. When you are finished making changes, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Changing client settings

Use the **Policy Management > Clients > Edit Client** page to change policy and authentication settings for one or more clients. If you select multiple clients before clicking **Edit**, the configuration changes that you make on the **Edit Client** page are applied to all of the selected clients.

1. Select a **Policy** to apply to the selected clients. The Default policy governs clients until another policy is assigned.
2. Under **Block Page Override Options**, indicate whether this client has the option to override (or attempt to override) a block page to view a requested site.

- Mark **Enable password override** to enable the selected clients to enter a password that you specify to access any blocked site for the time period configured on the Settings > General > Filtering page (60 seconds, by default). See [Password override, page 73](#).

Also enter and confirm the password.

You might enable this option for specific users who sometimes need access to sites not generally permitted by your organization's acceptable use policy.

To remove a client's password override privileges, click **Off**.

- Mark **Enable account override** to enable the selected clients to enter a network logon name and password to attempt to access a blocked site by having a different policy applied to the request. If the request is permitted by the new policy, the user can access the site for the time period configured on the Settings > General > Filtering page (5 minutes, by default). See [Account override, page 74](#).

You might enable this option for shared machines (like kiosk machines) typically governed by an IP-address-based policy that allows users to log on via a guest account. Users then have the option to enter their network credentials on the block page to see if their usual policy permits access to a site blocked on the shared machine.

If the user's policy also blocks the site, the user receives a second block page.

3. To allocate a custom amount of **Quota Time** to the selected clients, click **Custom**, and then enter the number of minutes of quota time to assign.

To revert to the default quota settings, click **Default**.

4. Click **OK** to cache your changes and return to the Clients page. Changes are not implemented until you click **Save All**.

The new client settings appear as part of the client listing on the **Policy Management > Clients** page.

Password override

Password override lets users with valid passwords access sites blocked by Websense software. Password override can be granted to individual clients (users, groups, computers, or networks, but not domains [OUs]).

When an administrator enables the password override option, he or she also creates a password. When clients with password override privileges request a blocked site, the

Websense block page includes a password field. The clients can then enter the password to access blocked sites for a limited amount of time.

**Note**

The password override option should not be used when individual clients are managed by multiple Policy Servers.

The timing information associated with this feature is not shared among Policy Servers, and affected clients could be granted more or less Internet access than you intend.

Configure how long clients with password override privileges can access blocked sites per password entry on the **Settings > General > Filtering** page (see [Configuring Websense filtering settings](#), page 57).

Grant password override privileges to specific clients via the **Policy Management > Clients** page (see [Adding a client](#), page 71, or [Changing client settings](#), page 72).

Account override

Account override allows users to change the credentials used to apply a policy to a request.

If, for example, users access the Internet from a kiosk machine, or from a machine where they log on using a local account, rather than a network account, administrators can associate account override permissions with the computer or network (IP-address-based) client.

Account override permissions can also be given to directory clients (users, groups, and domains [OUs]).

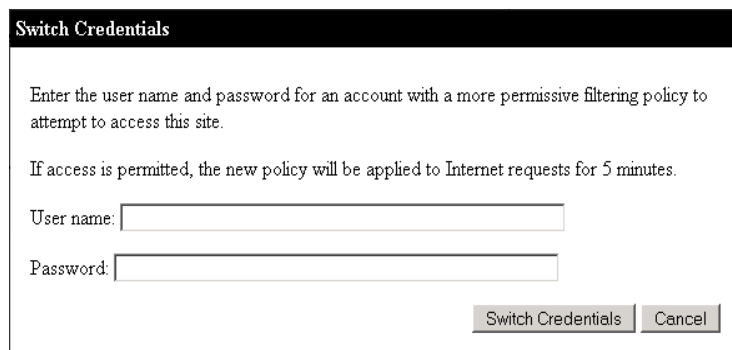
**Note**

The account override option should not be used when individual clients are managed by multiple Policy Servers.

The timing information associated with this feature is not shared among Policy Servers, and affected users could be granted more or less Internet access than you intend.

When user requests are blocked by the current policy, and account override permissions are assigned to the client being filtered (whether that is an IP address or a

directory client), the block page includes an **Enter New Credentials** button. The user can then provide a user name and password.



Once the user clicks **Switch Credentials**, Websense software identifies the policy assigned to the new account, then applies that policy to the request.

- ◆ If the new policy permits the request, the user can access the site.
- ◆ If the new policy blocks the request, the user sees another block page.

In other words, unlike password override, using the account override option does not guarantee access to a blocked site. Instead, it changes the policy used to filter the request.

The new policy is applied to additional requests on that machine for the time period specified on the **Settings > General > Filtering** page (5 minutes, by default). See [Configuring Websense filtering settings, page 57](#).

If, after successfully switching credentials, the user wants to leave the machine before the account override period has ended, the override session can be ended manually by entering the following URL:

```
http://<Filtering_Service_IP_address>:15871/cgi-bin/  
cancel_useraccount_overrider.cgi
```

You may want to configure this URL as a browser bookmark on machines where the account override option is used.

Moving clients to roles

Super Administrators can use the **Move Client To Role** page to move one or more clients to a delegated administration role. Once a client has been moved, that client appears in the Managed Clients list and on the Clients page in the target role.

- ◆ The policy applied to the client in the Super Administrators role and the filters that it enforces are copied to the delegated administration role.
- ◆ Delegated administrators can change the policies applied to their managed clients.
- ◆ Filter Lock restrictions do not affect clients managed by Super Administrators, but do affect managed clients in delegated administration roles.

- ◆ If a group, domain, or organizational unit is added to a role as a managed client, delegated administrators in that role can assign policies to individual users in the group, domain, or organizational unit.
- ◆ If a network (IP address range) is added to a role as a managed client, delegated administrators in that role can assign policies to individual computers in that network.
- ◆ The same client cannot be moved to multiple roles.

To move the selected clients to a delegated administration role:

1. Use the **Select role** drop-down list to select a destination role.
2. Click **OK**

A pop-up dialog box indicates that the selected clients are being moved. The move process may take a while.

3. Changes are not implemented until you click **Save All**.

If delegated administrators in the selected role are logged on with policy access during the move process, they will have to log out of TRITON - Web Security and log on again to see the new clients in their Managed Clients list.

Working with hybrid filtering clients

In Websense Web Security Gateway Anywhere deployments, the hybrid service can provide filtering for Internet requests originating from external IP addresses (locations) that you configure, and for requests from users in unrecognized locations (off-site users, for example) that log on to the hybrid service.

Hybrid filtering can apply policies (created in TRITON - Web Security) to:

- ◆ Users, groups, and domains (OUs) defined in a supported, LDAP-based directory service
This requires that Websense Directory Agent be installed and configured (see [Identification of hybrid filtering users, page 276](#)).
- ◆ Filtered locations, identified on the Hybrid Configuration > Filtered Locations page. A location is identified by the external IP address, IP address range, or subnet of one or more firewall or gateway machines.

Hybrid filtering does **not** apply policies to individual client machines in your network

Directory clients (users, groups, and OUs) filtered by the hybrid service are identified on the Policy Management > Clients page in TRITON - Web Security, just like those filtered by on-premises components.

Applying a policy to a filtered location is similar to applying a policy to a computer or network client:

1. Add the location to the Settings > Hybrid Configuration > Filtered Locations page (see [Define the locations filtered by hybrid filtering, page 175](#)).

2. Add the IP address or range that appears on the Filtered Locations page as a computer or network client on the Policy Management > Clients page (see *Working with computers and networks*, page 63).
3. Apply a policy to the IP address or range.

As with on-premises filtering, any time no user, group, or location policy applies, the Default policy is used.

4

Internet Filtering Policies

Related topics:

- ◆ [Internet Usage Filters, page 39](#)
- ◆ [Clients, page 61](#)
- ◆ [The Default policy, page 80](#)
- ◆ [Working with policies, page 81](#)
- ◆ [Filtering order, page 85](#)

Policies govern user Internet access. A policy is made up of:

- ◆ Category filters, used to apply actions (permit, block) to Web site categories (see [Filtering categories and protocols, page 40](#))
- ◆ Limited access filters, used to permit access to only a restricted list of Web sites (see [Restricting users to a defined list of Internet sites, page 212](#))
- ◆ Protocol filters, used to apply actions to Internet protocols (see [Filtering categories and protocols, page 40](#))
- ◆ A schedule that determines when each category or limited access filter and protocol filter is enforced

A new Websense software or appliance installation includes 3 predefined policies:

- ◆ **Default** filters Internet access for all clients not governed by another policy. Websense software begins enforcing this policy as soon as a subscription key is entered (see [The Default policy, page 80](#)).
- ◆ **Unrestricted** provides unlimited access to the Internet. This policy is not applied to any clients by default.
- ◆ **Example - Standard User** shows how multiple category and protocol filters can be applied in a policy to provide different degrees of filtering restriction at different times. This policy is used in the New User Quick Start tutorial to demonstrate the process of editing a policy and applying it to clients.

Use any of these policies as is, edit them to suit your organization, or create your own policies.

The Default policy

Related topics:

- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Working with policies, page 81](#)
- ◆ [Filtering order, page 85](#)

When you install Websense software, the **Default** policy begins monitoring Internet usage as soon as you enter your subscription key. Initially, the Default policy permits all requests.



Note

When you upgrade from an earlier Websense software version, existing policy settings are preserved. After upgrading, review your policies to ensure that they are still appropriate.

As you create and apply your own filtering policies, the Default policy continues to act as a safety net, filtering Internet access for any clients not governed by another policy.

In a new installation, the Default policy must provide Internet filtering coverage (enforce a combination of category or limited access filters and, if applicable, protocol filters) 24 hours a day, 7 days a week.



Important

Those upgrading from an earlier version of Websense software may have a Default policy that does not cover all time periods. You are not required to change your Default policy. If, however, you do edit the policy, Websense software will not allow you to save the changes until all time periods are covered.

Edit the Default policy as needed to suit the needs of your organization. The Default policy cannot be deleted.

Working with policies

Related topics:

- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Creating a policy, page 82](#)
- ◆ [Editing a policy, page 83](#)
- ◆ [Internet Usage Filters, page 39](#)
- ◆ [Refine Filtering Policies, page 211,](#)

Use the **Policy Management > Policies** page to review existing policy information. This page also serves as a launch point for adding, editing, and deleting policies, copying policies to delegated administration roles (Super Administrators only), and printing detailed information about your policy configuration.

The Policies page includes a list of existing policies. The list includes a name and description for each policy, as well as the number of user, network, and computer clients to whom that policy has been assigned.

- ◆ To add a policy, click **Add**, and then see [Creating a policy, page 82](#), for more information.
- ◆ To edit a policy, click the policy name in the list, and then see [Editing a policy, page 83](#), for more information.
- ◆ To delete a policy, mark the check box next to the policy name, and then click **Delete**.
- ◆ To see which clients are filtered by the policy, click a number in the Users, Networks, or Computers column. The client information appears in a pop-up dialog box.

To print a list of all of your policies and their components, including filters, custom categories and protocols, keywords, custom URLs, and regular expressions, click **Print Policies To File**. This feature creates a detailed spreadsheet of policy information in Microsoft Excel format. It is intended to provide a convenient way for human resources specialists, managers, and others with supervisory authority to review filtering policy information.

If you have created delegated administration roles (see [Delegated Administration and Reporting, page 285](#)), Super Administrators can copy policies that they have created

to other roles for use by delegated administrators. The filters enforced by the policy are also copied.

**Note**

Because delegated administrators are governed by the Filter Lock, when the Permit All filters are copied, the copy is given a new name, and Filter Lock restrictions are applied. Unlike the original filter, the copied filter can be edited.

To copy policies to another role, first mark the check box next to the policy name, and then click **Copy to Role**. This process may take up to several minutes. See [Copying filters and policies to roles, page 216](#), for more information.

Creating a policy

Related topics:

- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Working with policies, page 81](#)
- ◆ [Editing a policy, page 83](#)
- ◆ [Working with filters, page 49](#)
- ◆ [Restricting users to a defined list of Internet sites, page 212](#)

Use the **Policy Management > Policies > Add Policy** page to create a new, custom policy.

1. Enter a unique **Policy name**. The policy name must be between 1 and 50 characters long, and cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
Policy names can include spaces, dashes, and apostrophes.
2. Enter a **Description** for the policy. The description should be clear and detailed to help with policy management in the long term.
The character restrictions that apply to policy names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).
3. To use an existing policy as the foundation for the new policy, mark the **Base on existing policy** check box, and then select a policy from the drop-down list.
To start with an empty policy, leave the check box unmarked.
4. Click **OK** to cache your changes and go to the Edit Policy page.
Use the Edit Policy page to finish defining the new policy. See [Editing a policy, page 83](#).

Editing a policy

Related topics:

- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Working with policies, page 81](#)
- ◆ [Creating a policy, page 82](#)
- ◆ [Working with filters, page 49](#)
- ◆ [Restricting users to a defined list of Internet sites, page 212](#)

Use the **Policy Management > Policies > Edit Policy** page to make changes to an existing policy, or to finish defining a new policy.

Use the top portion of the page to edit the policy name and description:

- ◆ Click **Rename** to change the policy name.
- ◆ Simply type in the **Description** field to change the filter description.

Under the policy description, the **Clients** field lists how many clients of each type (directory, computer, and network) are currently filtered by this policy. To see which clients are governed by the policy, click the link corresponding to the appropriate client type.

To assign this policy to additional clients, click **Apply to Clients** in the toolbar at the top of the page, and then see [Assigning a policy to clients, page 85](#).

Use the **Policy Definition** area to define which filters this policy applies at different times:

1. To add a time block to the schedule, click **Add**.
2. Use the **Start** and **End** columns in the Schedule table to define the time period that this time block covers.

To define filtering for a period that spans midnight (for example, 5 p.m. to 8 a.m.), add two time blocks to the schedule: one that covers the period from the start time until midnight, and one that covers the period from midnight to the end time.

The **Example - Standard User** policy, included with your Websense software, demonstrates how to define a filtering period that spans midnight.

3. Use the **Days** column to define which days of the week are included in this time block. To select days from a list, click the down arrow in the right portion of the column. When you are finished selecting days, click the up arrow.
4. Use the **Category / Limited Access Filter** column to select a filter to enforce during this time block.

To add a new filter to enforce in this policy, select **Create category filter** or **Create limited access filter**. See [Creating a category filter, page 50](#), or [Creating a limited access filter, page 213](#), for instructions.

- Use the **Protocol Filter** column to select a protocol filter to enforce during this time block.

To add a new filter to enforce in this policy, select **Create protocol filter**. See [Creating a protocol filter, page 52](#), for instructions.

- Repeat steps 1 through 5 to add additional time blocks to the schedule.

When any time block in the schedule is selected, the bottom portion of the Edit Policies page shows the filters enforced during that time block. Each filter listing includes:

- ◆ The filter type (category filter, limited access filter, or protocol filter)
- ◆ The filter name and description
- ◆ The filter contents (categories or protocols with actions applied, or a list of sites permitted)
- ◆ The number of policies that enforce the selected filter
- ◆ Buttons that can be used to edit the filter

When you edit a filter on this page, the changes affect every policy that enforces the filter. Before editing a filter that is enforced by multiple policies, click the **Number of policies using this filter** link to see exactly which policies will be affected.

The buttons that appear at the bottom of the filter listing depend on the filter type:

Filter Type	Buttons
category filter	<ul style="list-style-type: none"> ◆ Use the Permit, Block, Confirm, or Quota button to change the action applied to the selected categories (see Filtering actions, page 46). ◆ To change the action applied to a parent category and all of its subcategories, first change the action applied to the parent category, and then click Apply to Subcategories. ◆ To enable keyword blocking, file type blocking, or blocking based on bandwidth, click Advanced.
limited access filter	<ul style="list-style-type: none"> ◆ Use the Add Sites and Add Expressions button to add permitted URLs, IP addresses, or regular expressions to the filter (see Restricting users to a defined list of Internet sites, page 212). ◆ To remove a site from the filter, mark the check box next to the URL, IP address, or expression, and then click Delete.
protocol filter	<ul style="list-style-type: none"> ◆ Use the Permit or Block button to change the action applied to the selected protocols (see Filtering actions, page 46). ◆ To change the action applied to all protocols in a protocol group, change the action applied to any protocol in the group, and then click Apply to Group. ◆ To log data for the selected protocol, or to enable blocking based on bandwidth, click Advanced.

When you finish editing a policy, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Assigning a policy to clients

Related topics:

- ◆ [Internet Filtering Policies, page 79](#)
- ◆ [Creating a policy, page 82](#)
- ◆ [Editing a policy, page 83](#)
- ◆ [Clients, page 61](#)
- ◆ [Adding a client, page 71](#)

Use the **Policies > Edit Policy > Apply Policy to Clients** page to assign the selected policy to clients.

The Clients list shows all of the available directory, computer, and network clients, as well as the policy currently assigned to each client.

Mark the check box next to each client to be filtered by the selected policy, and then click **OK** to return to the Edit Policy page. Click **OK** again to cache your changes.

Click **Save All** to prompt Websense software to begin using the new policy to filter requests from the selected clients.

Filtering order

Multiple criteria, applied in a specific order, are used to determine whether to permit, block, or limit requested Internet data.

For each request it receives, Websense Web Security solutions:

1. Verify subscription compliance, making sure that the subscription is current and the number of subscribed clients has not been exceeded.
2. Determine which policy applies, searching in this order:
 - For users filtered by the on-premises service:
 - a. Policy assigned to the **user**
 - b. Policy assigned to the **IP address** (computer or network) of the machine being used
 - c. Policies assigned to **groups** the user belongs to
 - d. Policies assigned to the user's **domain (OU)**
 - e. The **Default** policy
 - (*Websense Web Security Gateway Anywhere*) For users filtered by the hybrid service:

- a. Policy assigned to the **user**
- b. Policy assigned to **groups** the user belongs to
- c. Policy assigned to the user's **domain (OU)**
- d. Policy assigned to the external **IP address** (filtered location) from which the request originates
- e. The **Default** policy

The first applicable policy found is used.

3. Filter the request according to the policy's restrictions.


In some cases, a user belongs to more than one group or domain, and no higher-priority policy applies. In these cases, the Websense Web security solution checks the policies assigned to each of the user's groups.

- ◆ If all the groups have the same policy, Websense software filters the request according to that policy.
 - ◆ If one of the groups has a different policy, Websense software filters the request according to the **Use more restrictive blocking** selection on the **Settings > General > Filtering** page.
 - If **Use more restrictive blocking** is checked, and any of the applicable policies blocks access to the requested category, Websense software blocks the site.
 - If the option is not checked, and any of the applicable policies permits access to the requested category, Websense software permits the site.
- If one of the applicable policies enforces a limited access filter, the **Use more restrictive blocking** option can have different effects than expected. See [Limited access filters and filtering precedence, page 212](#).
- ◆ If one of the groups has a different policy, and any of the potentially applicable policies enforces file type blocking, the file type blocking settings are ignored.

Filtering a site

Websense Web security software evaluates policy restrictions as follows to determine whether the requested site should be permitted or blocked. (For Websense Web

Security Gateway Anywhere deployments, note that the logic shown here applies to the on-premises software, but not the hybrid service.)

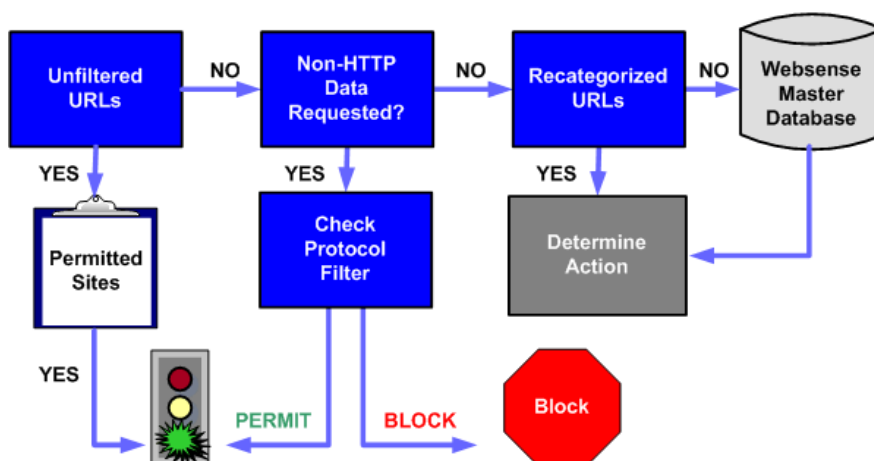


- Determines which **category filter** or **limited access filter** the policy enforces for the current day and time.
 - If the active category filter is **Permit All**, permit the site.
 - If the active category filter is **Block All**, block the site.
 - If the filter is a **limited access filter**, check whether the filter contains the URL or IP address. If so, permit the site. If not, block the site.
 - If any other category filter applies, continue to Step 2.

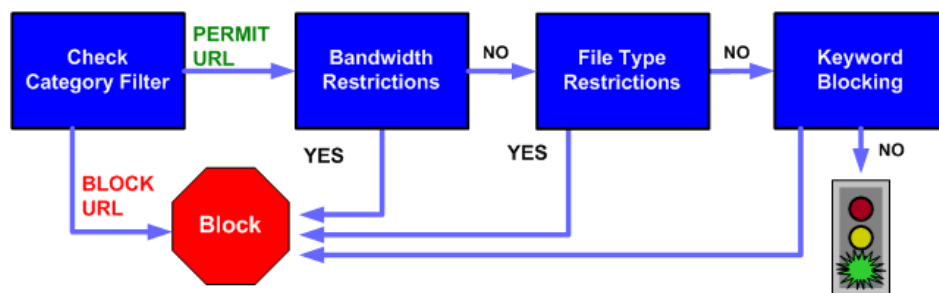


Note

Websense software filters URLs accessed from an Internet search engine's cache like any other URLs. URLs stored this way are filtered according to policies active for their URL categories. Log records for cached URLs show the entire cached URL, including any search engine parameters.

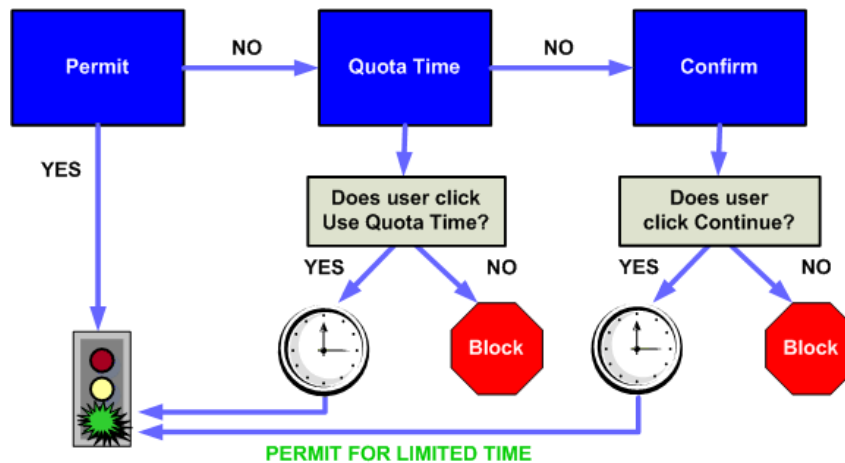


2. Tries to match the site to an entry in the **Unfiltered URLs** list.
 - If the URL appears on the list, permit the site.
 - If the URL does not appear on the list, continue to Step 3.
3. Checks the active **protocol filter** and determines whether any non-HTTP protocols are associated with the request.
 - If so, apply protocol filtering settings to data that may be transmitted.
 - If not, continue to Step 4.
4. Tries to match the site to an entry in the **Recategorized URLs** list.
 - If a match is made, identify the category for the site and go to Step 6.
 - If a match is not made, continue to Step 5.
5. Tries to match the site to an entry in the **Master Database**.
 - If the URL appears in the Master Database, identify the category for the site and continue to Step 6.
 - If a match is not made, categorize the site as Miscellaneous/Uncategorized and continue to Step 6.



6. Checks the active category filter and identifies the action applied to the category containing the requested site.
 - If the action is **Block**, block the site.
 - If any other action is applied, continue to Step 7.
7. Checks for **Bandwidth Optimizer** settings in the active category filter (see [Using Bandwidth Optimizer to manage bandwidth, page 236](#)).
 - If current bandwidth usage exceeds any configured limits, block the site.
 - If current bandwidth usage does not exceed the specified limits, or no bandwidth-based action applies, proceed to Step 8.
8. Checks for **file type** restrictions applied to the active category (see [Managing traffic based on file type, page 238](#)).
 - If the site contains files whose extensions are blocked, block access to those files. If the site itself is comprised of a blocked file type, block access to the site.
 - If the site does not contain files whose extensions are blocked, go to Step 9.

9. Checks for blocked **keywords** in the URL and CGI path, if keyword blocking is enabled (see *Filtering based on keyword*, page 222).
 - If a blocked keyword is found, block the site.
 - If a blocked keyword is not found, continue to Step 10.



10. Handles the site according to the action applied to the category.
 - **Permit**: Permit the site.
 - **Limit by Quota**: Display the block message with an option to view the site using quota time or go back to the previous page.
 - **Confirm**: Display the block message with the option to view the site for work purposes.

Websense software proceeds until the requested site is either blocked or explicitly permitted. At that point, Websense software does not attempt any further filtering. For example, if a requested site belongs to a blocked category and contains a blocked keyword, Websense software blocks the site at the category level without checking the keyword filter. Log Server then logs the request as blocked because of a blocked category, not because of a keyword.



Note

Users with password override privileges can access Internet sites regardless of why the site was blocked.

5

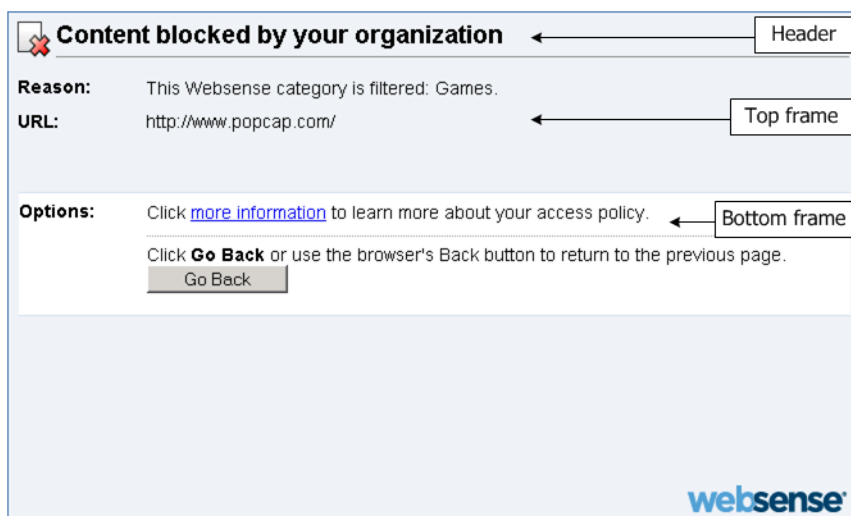
Block Pages

Related topics:

- ◆ *Blocking graphical advertisements*, page 92
- ◆ *Blocking embedded pages*, page 93
- ◆ *Protocol block messages*, page 94
- ◆ *Working with block pages*, page 95
- ◆ *Creating alternate block messages*, page 100
- ◆ *Using an alternate block page on another machine*, page 101
- ◆ *Determining why a request was blocked*, page 102

When Websense Web Security blocks a Web site, it displays a block page in the client's browser.

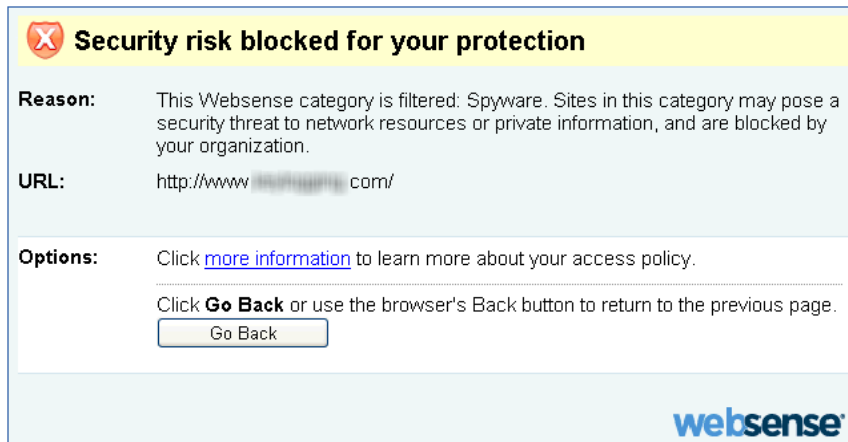
Block pages are constructed from HTML files, by default, made up of 3 main sections.



- ◆ The **header** explains that the site is blocked.
- ◆ The **top frame** contains a block message showing the requested URL and the reason the URL was blocked.

- ◆ The **bottom frame** presents any options available to the user, such as the option to go back to the previous page, or to click a Continue or Use Quota Time button to view the site.

If the site is blocked because it belongs to a category in the Security Risk class (see [Risk classes, page 43](#)), a special version of the block page is displayed.



Default block page files are included with your Websense software. You can use these default files or create your own custom versions.



Note

In Websense Web Security Gateway Anywhere deployments, changes to the on-premises block pages do not affect hybrid filtering block pages. See [Customizing hybrid block pages, page 185](#).

- ◆ Customize the default files to change the block message (see [Working with block pages, page 95](#)).
- ◆ Configure Websense software to use block messages (default or custom) hosted on a remote Web server (see [Using an alternate block page on another machine, page 101](#)).

Blocking graphical advertisements

In some cases, Websense software displays a very small, blank image file (BlockImage.gif) instead of a standard or security block page. This occurs when:

- ◆ The Advertisements category is blocked, and
- ◆ A site tries to display an image (like a GIF or JPG file) hosted at a URL in the Advertisements category.

Advertisements are often displayed in frames or iframes on a page that also displays non-advertisement information. In this case, graphical advertisements typically appear as white (empty) boxes on the page. The rest of the site content displays normally.

In some cases, an entire site may be made up of advertisement images. In this case, the user will see a blank Web page in the browser instead of a standard block message. Users can tell that the site has been blocked because of the URL, which is something like this:

```
http://<Filtering Service IP address>:15871/cgi-bin/  
blockpage.cgi?ws-session=<session number>
```

If you would prefer to show an image other than the default, 1-pixel block image, simply replace the default file:

1. Navigate to the block page directory on the Filtering Service machine (C:\Program Files\WebSense\Web Security\BlockPages\Images or /opt/WebSense/BlockPages/Images, by default).
2. Make a backup copy of the original **blockImage.gif** file.
3. Name your image **blockImage.gif** and copy it to the Images directory (overwriting the original file).

Blocking embedded pages

Most Web pages contain content from multiple sources (ad servers, streaming video sites, social networking applications, image hosting services, and so on). Some sites aggregate content, pulling pieces from multiple sites into a single presentation.

In these instances, users may request sites that contain a mix of permitted and blocked content.

When a frame or iframe within a larger page contains blocked content, Websense software displays a standard or security block page within that frame. When the frame is small, however, the end user might be able to see only a tiny portion of the page (perhaps not even the full block icon), and not understand why the content is blocked.

To address this issue, users can mouse over whatever portion of the block page is visible to see a tooltip-style popup with a brief block message. Clicking the message causes the full block page to appear in a separate window.

To return to browsing the permitted content of the original page, users should close the window showing the block page. Due to browser restrictions, clicking the Back button on a block page opened from within a frame does not have any effect.

If, when the block page is displayed in a new window, it offers a Use Quota Time or Continue option, clicking the button:

1. Closes the new (pop-up) window.
2. Displays the previously blocked content (and only that content) in the original browser window.

To see the original page, including the previously blocked content, do either of the following:

- ◆ Re-enter the site URL.
- ◆ Use the browser Back button to return to the site, then refresh the page.

Protocol block messages

Related topics:

- ◆ [Working with block pages, page 95](#)
- ◆ [Creating alternate block messages, page 100](#)
- ◆ [Using an alternate block page on another machine, page 101](#)

When a user or application requests a blocked, non-HTTP protocol, Websense software typically displays a protocol block message.

When, however, a user requests a blocked FTP, HTTPS, or Gopher site from within a browser, and the request passes through a proxy, an HTML-based block page displays in the browser, instead.

If an application requests the blocked protocol, the user may also receive an error message from the application, indicating that it cannot run. Application error messages are not generated by Websense software.

Some system configuration may be required to display protocol block messages on Windows machines:

- ◆ To display a protocol block message on client machines running Windows NT, XP, or 200x, the Windows Messenger service must be enabled. This service is disabled by default. You can use the Windows Services dialog box to find out if the service is running on a given machine (see [The Windows Services dialog box, page 458](#)).
- ◆ To display protocol block messages on a Windows 98 machine, you must start **winpopup.exe**, located in the Windows directory. Run the application from the command prompt, or configure it to launch automatically by copying it to the Startup folder.

Protocol block messages are not displayed on Linux machines. HTML block pages display regardless of operating system.

If protocol filtering is enabled, Websense software filters protocol requests whether or not the protocol block messages are configured to display on client machines.

Working with block pages

Related topics:

- ◆ [Protocol block messages, page 94](#)
- ◆ [Customizing the block message, page 96](#)
- ◆ [Creating alternate block messages, page 100](#)
- ◆ [Using an alternate block page on another machine, page 101](#)

The files used to create Websense block pages are stored in the **Websense\BlockPages\<language code>\Default** directory.



Note

In Websense Web Security Gateway Anywhere deployments, these block pages are applied only to users filtered by the on-premises software. To customize the pages provided by hybrid filtering, see [Customizing hybrid block pages, page 185](#).

There are 2 primary HTML files used to construct block pages:

- ◆ **master.html** constructs the information frame for the block page, and uses one of the following files to display appropriate options in the bottom frame

File Name	Contents
blockFrame.html	Text and button (Go Back option) for sites in blocked categories.
continueFrame.html	Text and buttons for sites in categories to which the Confirm action is applied.
quotaFrame.html	Text and buttons for sites in categories to which the Quota action is applied.
moreInfo.html	Content for the page that appears when a user clicks the More information link on the block page.

- ◆ **block.html** contains the text for the top frame of the block message, which explains that access is restricted, lists the requested site, and describes why the site is restricted.

In addition, several supporting files are used to supply the text content, styles, and button functionality used in block pages:

File Name	Description
blockStyle.css	Cascading style sheet containing most block page styles
master.css	Cascading style sheet containing styles for block page pop-ups (like the account override pop-up)
popup.html	When an embedded page is blocked (see Blocking embedded pages, page 93), this file is used to display the full-sized block page pop-up.
block.inl	Provides tools used in constructing the block frame of the block page
blockframe.inl	Provides additional information for standard block pages
continueframe.inl	Provides additional information for the block frame when users have a “Continue” option
quotaframe.inl	Provides additional information for the block frame when users have a “Use Quota Time” option
base64.js	JavaScript file used to support credential encryption when users have an “Account Override” option. This file should not be changed or removed.
master.js	JavaScript file used in construction of a standard block page
security.js	JavaScript file used in construction of a security block page
messagefile.txt	Contains text strings used in block pages
WebsenseCopyright.txt	Copyright information for Websense block pages
master.wml	WML file with basic blocking information

In deployments that include Web DLP components, an additional file, **policyViolationDefaultPage.html**, provides block page content when Websense Data Security components block content from being posted to or downloaded from the Web.

Customizing the block message

Related topics:

- ◆ [Changing the size of the message frame, page 98](#)
- ◆ [Changing the logo that displays on the block page, page 98](#)
- ◆ [Using block page content variables, page 99](#)
- ◆ [Reverting to the default block pages, page 100](#)

You can make a copy of the default block page files, and then use the copy to customize the top frame of the block page that users receive.

- ◆ Change the appearance of the block page to use your organization's logo, colors, and style.
- ◆ Add information about your organization's Internet use policies.
- ◆ Provide a method for contacting an administrator about Internet use policies.

To create your own, custom block pages:

1. Navigate to the Websense block page directory. For English:
`<installation path>\BlockPages\en\Default`
2. Copy the block page files to the custom block page directory. For English:
`<installation path>\BlockPages\en\Custom`

**Note**

Do **not** modify the original block message files in the **BlockPages\en\Default** directory. Copy them to the **BlockPages\en\Custom** directory and then modify the copies.

3. Open the file in a text editor, such as Notepad or vi.

**Warning**

Use a plain text editor to edit block message files. Some HTML editors modify HTML code, which could corrupt the files and cause problems displaying the block messages.

4. Modify the text. The files contain comments that guide you in making changes.
Do not modify the tokens (enclosed by \$* and *\$ symbols), or the structure of the HTML code. These enable Websense software to display specific information in the block message.
5. Some block page HTML files use hard-coded paths to reference the support files used to construct the page. If you have modified the stylesheet used to format the block pages (**blockStyle.css**) or the JavaScript file used to construct security block pages (**security.js**), make sure that you also update the path to those files in your custom HTML files. For example:

```
<link rel="stylesheet" href="/en/Custom/blockStyle.css
      type="text">
```
6. Save the file.
7. Restart Filtering Service (see [Stopping and starting Websense services](#), page 329, for instructions).

Changing the size of the message frame

Depending on what information you want to provide in the block message, the default width of the block message and height of the top frame may not be appropriate. To change these size parameters in the **master.html** file:

1. Copy **master.html** from the **Websense\BlockPages\en\Default** directory to **Websense\BlockPages\en\Custom**.
2. Open the file in a text editor, such as Notepad or vi (not an HTML editor).
3. To change the width of the message frame, edit the following line:

```
<div style="border: 1px solid #285EA6;width: 600px...">
```

Change the value of the **width** parameter as required.

4. To cause the top frame of the message to scroll, in order to show additional information, edit the following line:

```
<iframe src="$*WS_BLOCKMESSAGE_PAGE*$*WS_SESSIONID*$" ...  
    scrolling="no" style="width:100%; height: 6em;">
```

Change the value of the **scrolling** parameter to **auto** to display a scroll bar when message text exceeds the height of the frame.

You can also change the value of the **height** parameter to change the frame height.

5. Save and close the file.
6. Restart Filtering Service to implement the change (see [Stopping and starting Websense services](#), page 329).

Changing the logo that displays on the block page

The **master.html** file also includes the HTML code used to display to a Websense logo on the block page. To display your organization's logo instead:

1. Copy the block page files from the **Websense\BlockPages\en\Default** directory to **Websense\BlockPages\en\Custom**, if they have not already been copied.
2. Copy an image file containing your organization's logo to the same location.
3. Open **master.html** in a text editor, such as Notepad or vi (not an HTML editor), and edit the following line to replace the Websense logo with your organization's logo:

```

```

- Replace **wslogo_block_page.png** with the name of the image file containing your organization's logo.
 - Replace the values of the **title** parameter to reflect name of your organization.
4. Save and close the file.
 5. Restart Filtering Service to implement the change (see [Stopping and starting Websense services](#), page 329).

Using block page content variables

Content variables control the information displayed on HTML block pages. The following variables are included with the default block message code.

Variable Name	Content Displayed
WS_DATE	Current date
WS_USERNAME	Current user name (excluding domain name)
WS_USERDOMAIN	Domain name for the current user
WS_IPADDR	IP address of the requesting source machine
WS_WORKSTATION	Machine name of the blocked computer (if no name is available, IP address is displayed)

To use a variable, insert the variable name between the `$* *$` symbols in the appropriate HTML tag:

```
<p id="UserName">*$WS_USERNAME*$</p>
```

Here, `WS_USERNAME` is the variable.

The block message code includes additional variables, described below. You may find some of these variables useful in constructing your own, custom block messages. When you see these variables in Websense-defined block message files, however, please do **not** modify them. Because Filtering Service uses these variables when processing blocked requests, they must remain in place.

Variable Name	Purpose
WS_URL	Displays the requested URL
WS_BLOCKREASON	Displays why the site was blocked (i.e., which filtering action was applied)
WS_ISSECURITY	Indicates whether the requested site belongs to any of the categories in the Security Risk class. When TRUE, the security block page is displayed.
WS_PWOVERRIDECGIDATA	Populates an input field in the block page HTML code with information about use of the Password Override button
WS_QUOTA_CGIDATA	Populates an input field in the block page HTML code with information about use of the Use Quota Time button
WS_PASSWORDOVERRIDE_BEGIN, WS_PASSWORDOVERRIDE_END	Involved in activating password override functionality

Variable Name	Purpose
WS_MOREINFO	Displays detailed information (shown after the More information link is clicked) about why the requested site was blocked
WS_POLICYINFO	Indicates which policy governs the requesting client
WS_MOREINFOCGIDATA	Sends data to Filtering Service about use of the More information link
WS_QUOTATIME	Displays the amount of quota time remaining for the requesting client
WS_QUOTAINTERVALTIME	Displays quota session length configured for the requesting client
WS_QUOTABUTTONSTATE	Indicates whether the Use Quota Time button is enabled or disabled for a particular request
WS_SESSIONID	Acts as an internal identifier associated with a request
WS_TOPFRAMESIZE	Indicates the size (as a percentage) of the top portion of a block page sent by a custom block server, if one is configured
WS_BLOCKMESSAGE_PAGE	Indicates the source to be used for a block page's top frame
WS_CATEGORY	Displays the category of the blocked URL
WS_CATEGORYID	The unique identifier for the requested URL category

Reverting to the default block pages

If users experience errors after you implement customized block messages, you can restore the default block messages as follows:

1. Delete all the files from the **Websense\BlockPages\en\Custom** directory. By default, Websense software will return to using the files in the Default directory.
2. Restart Filtering Service (see [Stopping and starting Websense services](#), page 329).

Creating alternate block messages

Related topics:

- ◆ [Working with block pages](#), page 95
- ◆ [Customizing the block message](#), page 96

You can create your own HTML files to supply the text that appears in the top frame of the block page. Use existing HTML files, create alternate files from scratch, or make copies of **block.html** to use as a template.

- ◆ Create different block messages for each of 3 protocols: HTTP, FTP, and Gopher.
- ◆ Host the files on the Websense machine, or on your internal Web server (see [Using an alternate block page on another machine, page 101](#)).

After creating alternate block message files, you must configure Websense software to display the new messages (see [Configuring Websense filtering settings, page 57](#)). During this process, you can specify which message is used for each of the configurable protocols.

Using an alternate block page on another machine

Related topics:

- ◆ [Working with block pages, page 95](#)
- ◆ [Customizing the block message, page 96](#)
- ◆ [Creating alternate block messages, page 100](#)

Instead of using Websense block pages and customizing just the message in the top frame, you can create your own HTML block pages and host them on an internal Web server.



Note

It is possible to store block pages on an external Web server. If, however, that server hosts a site listed in the Master Database, and that site is in a blocked category, the block page itself is blocked.

Some organizations use alternate, remote block pages to hide the identity of the Websense server machine.

The remote block page can be any HTML file; it does not need to follow the format of the default Websense block pages. Using this method to create block pages, however, does prevent you from using the Continue, Use Quota Time, and Password Override functions available with Websense-defined block pages (default or custom).

When the files are in place, edit the **eimserver.ini** file to point to the new block page.

1. Stop the Websense Filtering Service and Policy Server services, in that order (see [Stopping and starting Websense services, page 329](#)).
2. On the Filtering Service machine, navigate to the Websense **bin** directory (see [Where is the Websense "bin" directory?, page 457](#)).
3. Create a backup copy of the **eimserver.ini** file and store it in another directory.

4. Open **aimserver.ini** file in a text editor, and locate the **[WebsenseServer]** section (at the top of the file).
5. Enter either the host name or the IP address of the server hosting the block page in the following format:

```
UserDefinedBlockPage=http://<host name or IP address>
```

The protocol portion of the URL (http://) is required.
6. Save the file and close the text editor.
7. Restart the Websense Policy Server and Filtering Service, in that order.

When the services have started, users receive the block page hosted on the alternate machine.

Determining why a request was blocked

If you want to investigate why a request was blocked, information is available in the block page source code.

- ◆ If the block page was sent by Filtering Service (for users filtered by the appliance or on-premises software), click **More information**. Next, right-click anywhere in the message text (for example, “Your Websense policy blocks this page at all times”) and select **View Source**. See [Request blocked by Filtering Service, page 102](#).
- ◆ If the block page was sent by hybrid filtering (in Websense Web Security Gateway Anywhere environments), right-click anywhere in the block message and select **View Source**. See [Request blocked by hybrid filtering, page 103](#).

Request blocked by Filtering Service

The HTML source for the more information block page shows information about who requested the site, and what criteria were used to filter the request. Specifically, it shows:

- ◆ The user name and source IP address of the request (if available), and the time (in the format HH:MM) that the request was made.
- ◆ Which policy is being applied to the request, and whether the policy is assigned to the user, group, domain, computer (individual IP address), or network (IP address range).

If more than one group policy could apply, the message also states whether the **Use more restrictive blocking** setting is in use. See [Configuring Websense filtering settings, page 57](#).

- ◆ What aspect of the policy caused the request to be blocked (for example, category or limited access filter, file type, keyword, bandwidth usage).
- ◆ The name of the role in which the policy was assigned.
- ◆ What resource was used to categorize the site (Websense Master Database, real-time database update, a regular expression included in a real-time database

update, custom URL, keyword, Websense Web Security Gateway scanning, and so on).

For example:

```
User Name: WinNT://Test/tester1 Source IP Address:
10.12.132.17 Current Time: 15:30
```

```
This network (10.12.132.0 to 10.12.132.255) is filtered by
policy: role-8**Default. The policy includes a category or
limited access filter for the current time.
```

```
This policy is associated with role: Super Administrator.
```

```
The request was categorized by: Master database.
```

Here, the request is filtered by a policy (Default) applied to the network (IP address range) in which the user's machine is located. The policy assignment was performed in the Super Administrator role, and the requested site was categorized by the Master Database.

Request blocked by hybrid filtering

The HTML source for the block page sent by hybrid filtering shows information about how the requested site was categorized, and how a policy was applied to the request. Specifically, it shows:

- ◆ The name of the role in which the policy was assigned. See [Delegated administration roles, page 286](#).
- ◆ The category assigned to the site.
- ◆ The policy or policies assigned to the request.
- ◆ If file type blocking was used, which file type applies.
- ◆ The protocol (HTTP, HTTPS, or FTP over HTTP) used to make the request.
- ◆ What resource was used to categorize the site (Websense Master Database, real-time database update, a regular expression included in a real-time database update, custom URL, keyword, Websense Web Security Gateway scanning, and so on).
- ◆ If a problem occurred that prevented the hybrid service from reporting why a request was blocked, or if the hybrid service experienced an error when the block page was being displayed, the **Exception reason** field displays an explanation and numeric error code. If the problem recurs, Websense Technical Support can use the error code in troubleshooting the issue.

For example:

```
Role: Super Administrator
Category: Peer-to-Peer File Sharing
Policy: Default
Domain:
Group:
```

FileType:

Network:

Protocol: http

Category Reason String: Master database

Exception reason:

Here, the request is filtered by a policy (Default) in the Super Administrator role that blocks the Peer-to-Peer File Sharing category. The requested HTTP site was categorized by the Master Database.

6

Use Reports to Evaluate Filtering

Related topics:

- ◆ [Presentation reports](#), page 107
- ◆ [Investigative reports](#), page 127
- ◆ [Accessing self-reporting](#), page 151
- ◆ [Real-Time Monitor](#)

TRITON - Web Security provides several reporting tools for use in evaluating the effectiveness of your filtering policies. (Log Server, a Windows-only component, must be installed to enable all reporting features except Real-Time Monitor.)



Note

In organizations that use delegated administration, reporting features may not be available to all administrators. See [Delegated Administration and Reporting](#), page 285.

- ◆ The **Today** page appears when you open TRITON - Web Security. It shows the operating status of Websense software, and can display charts of filtering activities in the network since midnight. (See [Today: Health, Security, and Value Since Midnight](#), page 24.)
- ◆ The **History** page shows charts of filtering activities in the network for up to 30 days, depending on the amount of information in the Log Database. These charts do not include today's activities. (See [History: Last 30 Days](#), page 27.)
- ◆ **Presentation reports** offer a list of predefined report templates and custom reports. Some output tabular reports; some combine a bar chart and a table. Copy any report template and apply your own report filter to create a custom report. See [Presentation reports](#), page 107, for complete details.
- ◆ **Investigative reports** let you browse through log data interactively. The main page shows a summary-level bar chart of activity by risk class. Click the different elements on the page to update the chart or get a different view of the data.

See [Investigative reports, page 127](#), for details on the many ways you can view Internet use data.

- ◆ **Real-Time Monitor** provides insight into current Internet filtering activity in your network, showing the URLs being requested and the action applied to each request. In Websense Web Security Gateway and Web Security Gateway Anywhere deployments, the monitor also shows which sites were scanned by Content Gateway. If a site is dynamically recategorized based on scanning results, both the original category and current category are shown.

See [Real-Time Monitor, page 152](#), for more information.



Important

Internet Explorer 8 Compatibility View is **not** supported with the TRITON console. If you experience odd reporting behavior or page layouts in Internet Explorer 8, make sure that Compatibility View button (between the URL and the Refresh button in the browser address bar) is not selected.

What is Internet browse time?

Related topics:

- ◆ [Database jobs, page 361](#)
- ◆ [Configuring Internet browse time options, page 366](#)

You can generate both presentation and investigative reports showing Internet browse time (IBT), the amount of time an individual spent accessing Web sites. No software program can tell the exact amount of time that someone spends actually viewing a particular site once it is open. Someone might open a site, view it for a few seconds, and then take a business call before requesting another site. Someone else might spend several minutes reading each site in detail before moving on to the next one.

Websense software includes a Log Database job to calculate IBT, using a formula based on certain configurable values. This job runs once a day, so browse time information can lag the actual log data. See [Database jobs, page 361](#), for more information.

For browse time calculations, an Internet session begins when a user opens a browser. It continues as long as that user requests additional Web sites at least every 3 minutes. (This default read time threshold is configurable. See [Configuring Internet browse time options, page 366](#).)

The Internet session ends when more than 3 minutes pass before the user requests another site. Websense software calculates the total time of the session, starting with the time of the first request and ending 3 minutes after the last request.


A new session begins if the user makes additional requests after more than 3 minutes. Commonly, a user's browse time consists of multiple sessions each day.

Presentation reports

Related topics:

- ◆ [Copying a presentation report](#), page 109
- ◆ [Working with Favorites](#), page 116
- ◆ [Running a presentation report](#), page 116
- ◆ [Scheduling presentation reports](#), page 118
- ◆ [Viewing the scheduled jobs list](#), page 123

Use the **Reporting > Presentation Reports** page to generate template-based charts and tabular reports in HTML, PDF, or Microsoft Excel (XLS) format.



Available predefined and custom reports are found in the Report Catalog, which organizes them into related groups. The groups that appear are based on your subscription type. Some groups (like Real Time Security Threats and Scanning Activity) require a Websense Web Security Gateway or Gateway Anywhere subscription.

- ◆ Expand a group to see the reports that it includes.
- ◆ Click a report title to see a brief description of the information included in the report.

To run a presentation report:

1. Select the report in the catalog, and then click **Run**. The Run Report page appears.
2. Specify report details as explained in [Running a presentation report, page 116](#).
 - If you run the report in the foreground (do not schedule the report to run), the report is not automatically saved when you close the application used to view the report (a Web browser, Adobe Reader, or Microsoft Excel, for example). You must save the report manually.
 - If you run the report in the background (schedule the report to run immediately), when the report completes, a copy is saved, and a link to the report appears on the Review Reports page.

To use any predefined or custom report in the Report Catalog as a template for a new report:

1. Select a report in the catalog, and then click **Save As**.
2. Enter a name for the new report, and then click either **Save** or **Save and Edit**.
 - If you click **Save**, you are returned to the Presentation Reports page, where the new report appears in the report catalog. To customize the report, select it, and then click **Edit**.
 - If you click **Save and Edit**, you are taken directly to the Edit Report Filter page. The new report is also added to the Report Catalog.
3. Edit the report filter to modify the report. The report filter controls elements such as which users, categories, protocols, and actions are to be included in your custom report.

For instructions, see [Defining the report filter, page 109](#).

To make changes to the report filter for any custom report, select the report, and then click **Edit**. You cannot modify or delete predefined reports.

To delete a custom report, select the report, and then click **Delete**. If a deleted report appears in any scheduled jobs, it will continue to be generated with that job. See [Viewing the scheduled jobs list, page 123](#), for information on editing and deleting scheduled jobs.

Reports that are used frequently can be marked as Favorites to help you find them more quickly. Just select the report, and then click **Favorite** (see [Working with Favorites, page 116](#)). Mark **Show Favorites only** to display only templates that you have marked as favorites in the Report Catalog.

Use the buttons at the top of the page to schedule reports to run later, view scheduled report jobs, and view and manage reports created by the scheduler.

- ◆ Click **Scheduler** to define a job containing one or more reports to be run at a specific time or on a repeating schedule. See [Scheduling presentation reports, page 118](#).

- ◆ Click **Job Queue** to see and manage a list of existing scheduled jobs, along with the status of each job. See [Viewing the scheduled jobs list, page 123](#).
- ◆ Click **Review Reports** to see and manage a list of reports that were successfully scheduled and run. See [Reviewing scheduled presentation reports, page 125](#).

Copying a presentation report

Related topics:

- ◆ [Presentation reports, page 107](#)
- ◆ [Defining the report filter, page 109](#)
- ◆ [Running a presentation report, page 116](#)

Use the **Save As New Report** page to create an editable copy of any existing report template.

1. Replace the **Report catalog name** with a name that will make it easy to identify the new report. (The default name is the name of the original report template, with a number appended to indicate that this is a copy.)

The name must be from 1 to 85 characters, and cannot duplicate another report name.

2. Click a button to indicate how you want to proceed.

Save	Makes an exact copy of the selected report template, and returns to the Report Catalog.
Save and Edit	Copies the selected report template, and opens the Edit Report Filter page, where you can modify the elements of the report. See Defining the report filter, page 109 .
Cancel	Returns to the Report Catalog without creating a copy of the original report template.

Defining the report filter

Related topics:

- ◆ [Copying a presentation report, page 109](#)
- ◆ [Running a presentation report, page 116](#)

Report filters let you control what information is included in a presentation report. For example, you might choose to limit a report to selected clients, categories, risk classes, or protocols, or even selected filtering actions (permit, block, and so forth). You also can give a new name and description for the entry in the Report Catalog, change the

report title, specify a custom logo to appear, and set other general options through the report filter.



Note

Using a custom logo requires some preparation before you define the report filter. You must create the image in a supported format and place the file in the appropriate location. See [Customizing the report logo, page 114](#).

The specific options available in the filter depend on the report selected. For instance, if you selected a report of group information, such as Top Blocked Groups by Requests, you can control which groups appear in the report but you cannot choose individual users.

The filter for report templates cannot be changed. You can edit the filter for a custom report when you create it by choosing **Save and Edit** on the Save As New Report page, or select the report in the Report Catalog at any time and click **Edit**.

The Edit Report Filter page opens, with separate tabs for managing different elements of the report. Select the items you want on each tab, then click **Next** to move to the next tab. For detailed instructions, see:

- ◆ [Selecting clients for a report, page 110](#)
- ◆ [Selecting categories for a report, page 111](#)
- ◆ [Selecting protocols for a report, page 112](#)
- ◆ [Selecting actions for a report, page 113](#)
- ◆ [Setting report options, page 113](#)

On the **Confirm** tab, choose whether to run or schedule the report, and save the report filter. See [Confirming report filter definition, page 115](#).

Selecting clients for a report

Related topics:

- ◆ [Selecting categories for a report, page 111](#)
- ◆ [Selecting protocols for a report, page 112](#)
- ◆ [Selecting actions for a report, page 113](#)
- ◆ [Setting report options, page 113](#)
- ◆ [Confirming report filter definition, page 115](#)

The **Clients** tab of the Presentation Reports > Edit Report Filter page lets you control which clients are included in the report. You can select only one type of client for each report. For example, you cannot select some users and some groups for the same report.

When the report definition specifies a particular client type, you can choose clients of that type or clients that represent a larger grouping. For example, if you are defining a filter for a report based on Top Blocked Groups by Requests, you can select groups, domains, or organizational units for the report, but you cannot select individual users.

No selections are required on this tab if you want to report on all relevant clients.

1. Select a client type from the drop-down list.
2. Set the maximum number of search results from the **Limit search** list.
Depending on the traffic in your organization, there may be large numbers of users, groups, or domains in the Log Database. This option manages the length of the results list, and the time required to display the search results.
3. Enter one or more characters for searching, and then click **Search**.
Use asterisk (*) as a wildcard to signify missing characters. For example, J*n might return Jackson, Jan, Jason, Jon, John, and so forth.
Define your search string carefully, to assure that all desired results are included within the number selected for limiting the search.
4. Highlight one or more entries in the results list, and click the right arrow button (>) to move them to the **Selected** list.
5. Repeat steps 2-4 as needed to conduct additional searches and add more clients to the Selected list.
6. After you are finished making selections, click **Next** to open the Categories tab.
See [Selecting categories for a report](#), page 111.

Selecting categories for a report

Related topics:

- ◆ [Selecting clients for a report](#), page 110
- ◆ [Selecting protocols for a report](#), page 112
- ◆ [Selecting actions for a report](#), page 113
- ◆ [Setting report options](#), page 113
- ◆ [Confirming report filter definition](#), page 115

The **Categories** tab of the Presentation Reports > Edit Report Filter page lets you control the information included in the report on the basis of categories or risk classes. See [Risk classes](#), page 43.

No selections are required on this tab if you want to report on all relevant categories or risk classes.

1. Select a classification: **Category** or **Risk Class**.
Expand a parent category to display its subcategories. Expand a risk class to see a list of the categories currently assigned to that risk class.

If the associated report is for a specific risk class, only the relevant risk class and the categories it represents are available for selection.



Note

If you select a subset of categories for the risk class named in the report, consider modifying the report title to reflect your selections.

2. Mark the check box for each category or risk class to be reported.
Use the **Select All** and **Clear All** buttons below the list to minimize the number of individual selections required.
3. Click the right arrow button (>) to move your selections to the **Selected** list.
When you mark a risk class, clicking the right arrow places all the associated categories into the Selected list.
4. After all selections are complete, click **Next** to open the Protocols tab. See [Selecting protocols for a report, page 112](#).

Selecting protocols for a report

Related topics:

- ◆ [Selecting clients for a report, page 110](#)
- ◆ [Selecting categories for a report, page 111](#)
- ◆ [Selecting actions for a report, page 113](#)
- ◆ [Setting report options, page 113](#)
- ◆ [Confirming report filter definition, page 115](#)

The **Protocols** tab of the Presentation Reports > Report Filter lets you control which protocols are included in the report.

No selections are required on this tab if you want to report on all relevant protocols.

1. Expand and collapse the protocol groups with the icon beside the group name.
2. Mark the check box for each protocol to be reported.
Use the **Select All** and **Clear All** buttons below the list to minimize the number of individual selections required.
3. Click the right arrow button (>) to move your selections to the **Selected** list.
4. After all selections are complete, click **Next** to open the Actions tab. See [Selecting actions for a report, page 113](#).

Selecting actions for a report

Related topics:

- ◆ [Selecting clients for a report, page 110](#)
- ◆ [Selecting categories for a report, page 111](#)
- ◆ [Selecting protocols for a report, page 112](#)
- ◆ [Setting report options, page 113](#)
- ◆ [Confirming report filter definition, page 115](#)

The **Actions** tab of the Presentation Reports > Edit Report Filter page lets you control which precise filtering actions (for example, permitted by limited access filter or blocked by quota) are included in the report. If the report specifies that it applies only to blocked requests, you can select only block-related actions (blocked by file type, blocked by keyword, and so on).

No selections are required on this tab if you want to report on all relevant actions.

1. Expand and collapse the action groups with the icon beside the group name.
2. Mark the check box for each action to be reported.

Use the **Select All** and **Clear All** buttons below the list to minimize the number of individual selections required.

3. Click the right arrow button (>) to move your selections to the **Selected** list.
4. After all selections are complete, click **Next** to open the Options tab. See [Setting report options, page 113](#).

Setting report options

Related topics:

- ◆ [Customizing the report logo, page 114](#)
- ◆ [Selecting clients for a report, page 110](#)
- ◆ [Selecting categories for a report, page 111](#)
- ◆ [Selecting protocols for a report, page 112](#)
- ◆ [Selecting actions for a report, page 113](#)
- ◆ [Setting report options, page 113](#)
- ◆ [Confirming report filter definition, page 115](#)

Use the **Options** tab of the Presentation Reports > Edit Report Filter page to configure several aspects of the report.

1. Optionally modify the **Report catalog name**. The name must contain from 1 to 85 characters.

This name does not appear on the report itself; it is used only for identifying the unique combination of report format and filter in the Report Catalog.

2. Modify the **Report title** that appears on the report. The title can have up to 85 characters.
3. Modify the **Description** to appear in the Report Catalog. The description can have up to 336 characters.

The description should help you identify this unique combination of report format and filter in the Report Catalog.

4. Select a logo to appear on the report.
All supported image files in the appropriate directory are listed. See [Customizing the report logo, page 114](#).
5. Mark the **Save as Favorite** check box to have the report listed as a Favorite.
The Report Catalog shows a star symbol beside Favorite reports. You can select **Show only Favorites** on the Report Catalog page to reduce the number of reports listed, which enables you to move more quickly to a particular report.
6. Mark the **Show only top** check box and then enter a number from 1 to 20 to limit the number of items reported.
This option appears only if the selected report is formatted as a Top N report, designed to show a limited number of items. The item that is limited depends on the report. For example, for a Top Categories Visited report, this entry determines how many categories are reported.
7. After all entries and selections are complete, click **Next** to open the Confirm tab. See [Confirming report filter definition, page 115](#).

Customizing the report logo

By default, presentation reports display the Websense logo in the upper left corner. When you create a custom report and edit its report filter, you can choose a different logo.

1. Create an image file in one of the following formats:
 - .bmp
 - .gif
 - .jfif
 - .jpe
 - .jpg
 - .jpeg
 - .png
 - .tiff
2. Use a maximum of 25 characters for the image file name, including extension.
3. Copy the image file to the appropriate directory:
 - Windows:
`C:\Program Files\Websense\Manager\ReportingTemplates\images`
 - Linux:
`/opt/Websense/Manager/ReportingTemplates/images/`

All supported image files in this directory automatically appear in the drop-down list on the Options tab of the Edit Report Filter page. The image is automatically scaled to fit within the space allocated for the logo. (See [Setting report options](#), page 113.)



Note

Do not delete or move images that are active in report filters. If the logo file is missing, the report cannot be generated.

Confirming report filter definition

Related topics:

- ◆ [Selecting clients for a report](#), page 110
- ◆ [Selecting categories for a report](#), page 111
- ◆ [Selecting protocols for a report](#), page 112
- ◆ [Selecting actions for a report](#), page 113
- ◆ [Setting report options](#), page 113

The **Confirm** tab of the Presentation Reports > Edit Report Filter page displays the name and description that will appear in the Report Catalog, and lets you choose how to proceed.

1. Review the **Name** and **Description**.

If any changes are needed, click **Back** to return to the Options tab, where you can make those changes. (See [Setting report options](#), page 113.)

2. Indicate how you want to proceed:

Option	Description
Save	Saves the report filter and returns to the Report Catalog. See Presentation reports , page 107.
Save and Run	Saves the report filter and opens the Run Report page. See Running a presentation report , page 116.
Save and Schedule	Saves the report filter and opens the Schedule Report page. See Scheduling presentation reports , page 118.

3. Click **Finish** to implement the selection made in step 2.

Working with Favorites

Related topics:

- ◆ [Presentation reports, page 107](#)
- ◆ [Running a presentation report, page 116](#)
- ◆ [Scheduling presentation reports, page 118](#)

You can mark any presentation report, either template or custom, as a Favorite. Use this option to identify the reports you generate most frequently and want to be able to locate quickly in the Report Catalog.

1. On the **Presentation Reports** page, highlight a report that you generate frequently, or want to be able to locate quickly.
2. Click **Favorite**.
A star symbol appears beside Favorite report names in the list, letting you quickly identify them when all reports are shown.
3. Mark the **Show only Favorites** check box above the Report Catalog to limit the list to those marked as Favorites. Clear this check box to restore the full list of reports.

If your needs change and a Favorite report is no longer being used as frequently, you can remove the Favorite designation.

1. Highlight a report that shows the star symbol of a Favorite.
2. Click **Favorite**.
The star symbol is removed from that report name in the Report Catalog. The report is now omitted from the list if you choose **Show only Favorites**.

Running a presentation report

Related topics:

- ◆ [Presentation reports, page 107](#)
- ◆ [Scheduling presentation reports, page 118](#)

Use the **Presentation Reports > Run Report** page to generate a single report immediately. You can also create jobs with one or more reports and schedule them to run once or on a repeating cycle (see [Scheduling presentation reports, page 118](#)).



Note

Before generating a report in PDF format, make sure that Adobe Reader v7.0 or later is installed on the machine from which you are accessing TRITON - Web Security.

Before generating a report in XLS format, make sure that Microsoft Excel 2003 or later is installed on the machine from which you are accessing TRITON - Web Security.

If the appropriate software is not installed, you have the option to save the file.

To run a report:

1. Select the **Start date** and **End date** to define the time period covered in the report.
2. Select an **Output format** for the report.

Format	Description
PDF	Portable Document Format. PDF files are formatted for viewing, and can be opened in Adobe Reader.
HTML	HyperText Markup Language. HTML files are formatted for viewing, and can be opened in a browser.
XLS	Excel spreadsheet. XLS files are formatted for reuse, and can be opened in Microsoft Excel.

3. If you selected a **Top N** report, choose the number of items to be reported.
4. Specify how you want the report to be generated:
 - Select **Schedule the report to run in the background** (default) to have the report run immediately as a scheduled job. Optionally provide an email address to be notified when the report is complete. You can also provide an email address to be notified if the report cannot be generated. (You can also monitor the job queue to see the status of the report.)

- Deselect **Schedule the report to run in the background** to have the report run in the foreground. In this case, the report is not scheduled, and does not appear on the Review Reports page.



Note

If you plan to run multiple reports in the foreground, make sure that you use the embedded **Close** button to close the pop-up window used to display the “generating report” and “report complete” messages. If you use the browser’s close (X) button, subsequent attempts to run reports in the foreground may fail until you navigate away from the Presentation Reports page, come back, and run the report again.

5. Click **Run**.

- If you scheduled the report to run immediately, the completed report is saved automatically and added to the Review Reports list. To view, save, or delete the report, click **Review Reports** at the top of the Presentation Reports page.
- If you ran the report in the foreground, a new browser window appears, displaying report progress. HTML reports appear in the browser window when complete; with PDF or XLS formats, you have a choice of whether to open the report or save it to disk.

With this option, presentation reports does not automatically store a copy of the report. Use the save functionality built into the application used to open the report if you want to save a copy to view later.

6. To print a report, use the print option offered by the application used to display the report.

For best results, generate PDF output for printing. Then, use the print options in Adobe Reader.

Scheduling presentation reports

Related topics:

- ◆ [Presentation reports, page 107](#)
- ◆ [Running a presentation report, page 116](#)
- ◆ [Viewing the scheduled jobs list, page 123](#)

You can run presentation reports as they are needed, or you can use the **Presentation Reports > Scheduler** page to create jobs that define a schedule for running one or more reports.

Reports generated by scheduled jobs are distributed to one or more recipients via email. As you create scheduled jobs, consider whether your email server will be able to handle the size and quantity of the attached report files.

The completed reports are also added to the Presentation Reports > Review Reports page (see [Reviewing scheduled presentation reports](#), page 125).

To access the Scheduler:

- ◆ Click the **Scheduler** button at the top of the Presentation Reports page (above the Report Catalog).
- ◆ When editing a report filter, choose **Save and schedule** in the Confirm tab, and then click **Finish** (see [Defining the report filter](#), page 109).
- ◆ Click the job name link on the Job Queue page to edit a job.
- ◆ Click **Add** on the Job Queue page to create a new job.

The Scheduler page contains several tabs for selecting the reports to run and the schedule for running them. For detailed instructions, see:

- ◆ [Setting the schedule](#), page 120
- ◆ [Selecting reports to schedule](#), page 121
- ◆ [Setting the date range](#), page 122
- ◆ [Selecting output options](#), page 123

After creating jobs, use the Job Queue to review job status and find other helpful information (see [Viewing the scheduled jobs list](#), page 123).

When a scheduled presentation report has run, the report file is sent to recipients as an email attachment called **presentationreport_0**. The number increments, according to the number of reports attached.

Scheduled reports are also automatically saved to the **ReportingOutput** directory on the TRITON - Web Security machine (C:\Program Files\WebSense\ReportingOutput or /opt/WebSense/ReportingOutput, by default). Note that the name of the attachment sent via email does not match the name of the file stored in the ReportingOutput directory. The best way to find a specific report is to use the Review Reports page, which can be searched by date or job name, as well as report name.

Reports are automatically deleted from the Review Reports page and the ReportingOutput directory after the period specified on the Settings > Reporting > Preferences page (5 days, by default). If you want to retain the reports for a longer time, include them in your backup routine or save them in a location that permits long term storage.

An alert is displayed on the Review Reports page for a period of time before the report is deleted (3 days, by default). Use the Settings > Reporting > Preferences page to change this warning period.

Depending on the number of reports you generate daily, report files can occupy considerable amounts of disk space. Be sure there is adequate disk space available on the TRITON - Web Security machine. If the ReportingOutput directory grows too large before the files are automatically deleted, you can delete the files manually.

Websense software generates the report in the format you choose: PDF (Adobe Reader), XLS (Microsoft Excel), or HTML. If you choose HTML format, the report may display in the TRITON - Web Security content pane. Reports displayed in the

content pane cannot be printed or saved to a file. To print or save a report to file, choose the PDF or XLS output format.



Important

To display presentation reports in PDF format, Adobe Reader v7.0 or later must be installed on the machine from which you are accessing TRITON - Web Security.

To display presentation reports in XLS format, Microsoft Excel 2003 or later must be installed on the machine from which you are accessing TRITON - Web Security.

Setting the schedule

Related topics:

- ◆ [Scheduling presentation reports, page 118](#)
- ◆ [Selecting reports to schedule, page 121](#)
- ◆ [Selecting output options, page 123](#)
- ◆ [Setting the date range, page 122](#)

Define a reporting job to occur once or on a repeating cycle on the **Schedule** tab of the Presentation Reports > Scheduler page.



Note

It is advisable to schedule report jobs on different days or at different times, to avoid overloading the Log Database and slowing performance for logging and interactive reporting.

1. Enter a **Job name** that uniquely identifies this scheduled job.
2. Select a **Recurrence Pattern** and **Recurrence Options** for the job. The specific options available depend on the pattern selected.

Pattern	Options
Once	Enter the exact date on which to run the job, or click the icon to select from a calendar.
Daily	No additional recurrence options are available.
Weekly	Mark the check box for each day of the week the job is to run.
Monthly	Enter the dates during the month for running the job. Dates must be a number between 1 and 31, and must be separated by commas (1,10,20). To run the job on consecutive dates each month, enter a start and end date separated by a hyphen (3-5).

- Under **Schedule Time**, set the start time for running the job.

The job begins according to the time on the machine running TRITON - Web Security.



Note

To start generating the scheduled reports today, select a time late enough that you can complete the job definition before the start time.

- Under **Schedule Period**, select a date for starting the job, and an option for ending the job.

Option	Description
No end date	The job continues to run according to the established schedule, indefinitely. To discontinue the job at some time in the future, either edit or delete the job. See Viewing the scheduled jobs list, page 123 .
End after	Select the number of times to run the job. After that number of occurrences, the job does not run again, but it stays in the Job Queue until you delete it. See Viewing the scheduled jobs list, page 123 .
End by	Set the date when the job stops running. It does not run on or after this date.

- Click **Next** to open the Reports tab. See [Selecting reports to schedule, page 121](#).

Selecting reports to schedule

Related topics:

- ◆ [Scheduling presentation reports, page 118](#)
- ◆ [Setting the schedule, page 120](#)
- ◆ [Selecting output options, page 123](#)
- ◆ [Setting the date range, page 122](#)

Use the **Select Report** tab of the Presentation Reports > Scheduler page to choose reports for the job.

- Highlight a report for this job in the Report Catalog tree.
- Click the right arrow (>) button to move that report to the **Selected** list.
- Repeat steps 1 and 2 until all reports for this job appear in the **Selected** list.
- Click **Next** to open the Date Range tab. See [Setting the date range, page 122](#).

Setting the date range

Related topics:

- ◆ [Scheduling presentation reports](#), page 118
- ◆ [Setting the schedule](#), page 120
- ◆ [Selecting reports to schedule](#), page 121
- ◆ [Selecting output options](#), page 123

Use the **Date Range** tab of the Presentation Reports > Scheduler page to set the date range for the job. The options available depend on your selection for **Date range**.

Date range	Description
All Dates	<p>Reports include all dates available in the Log Database. No additional entries are required.</p> <p>When this option is used for repeating jobs, there may be duplicate information on reports in separate runs.</p>
Specific Dates	<p>Choose the exact start (From) and end (To) dates for the reports in this job.</p> <p>This option is ideal for jobs that run only one time. Choosing this option for a repeating schedule results in duplicate reports.</p>
Relative Dates	<p>Use the drop-down lists to choose the number of periods to report (This, Last, Last 2, and so forth), and the type of period (Days, Weeks, or Months). For example, the job might cover the Last 2 Weeks or This Month.</p> <p>Week represents a calendar week, Sunday through Saturday. Month represents a calendar month. For example, This Week produces a report from Sunday through today; This Month produces a report from the first of the month through today; Last Week produces a report for the preceding Sunday through Saturday; and so forth.</p> <p>This option is ideal for jobs that run on a repeating schedule. It lets you manage how much data appears on each report, and minimize duplication of data on reports in separate runs.</p>

After setting the date range for the job, click **Next** to display the Output tab. See [Selecting output options](#), page 123.

Selecting output options

Related topics:

- ◆ [Scheduling presentation reports](#), page 118
- ◆ [Setting the schedule](#), page 120
- ◆ [Selecting reports to schedule](#), page 121
- ◆ [Setting the date range](#), page 122

After you select the reports for a job, use the **Output** tab to select the output format and distribution options.

1. Select the file format for the finished report.

Format	Description
PDF	Portable Document Format. Recipients must have Adobe Reader v7.0 or later to view the PDF reports.
XLS	Excel Spreadsheet. Recipients must have Microsoft Excel 2003 or later to view the XLS reports.

2. Enter email addresses for distributing the report.
Enter each address on a separate line.
3. Mark the **Customize subject and body of email** check box, if desired. Then, enter the custom **Subject** and **Body** text for this job's distribution email.
4. Click **Save Job** to save and implement the job definition, and display the Job Queue page.
5. Review this job and any other scheduled jobs. See [Viewing the scheduled jobs list](#), page 123.

Viewing the scheduled jobs list

Related topics:

- ◆ [Presentation reports](#), page 107
- ◆ [Scheduling presentation reports](#), page 118
- ◆ [Selecting output options](#), page 123
- ◆ [Scheduling investigative reports](#), page 146

The **Presentation Reports > Job Queue** page lists the scheduled jobs created for presentation reports. The list gives status for each job, as well as basic information about the job, such as how frequently it runs. From this page, you can add and delete scheduled jobs, temporarily suspend a job, and more.

(To review scheduled jobs for investigative reports, see *Managing scheduled investigative reports jobs*, page 148.)

The list provides the following information for each job.

Column	Description
Job Name	The name assigned when the job was created.
Status	Indicates whether the job is <ul style="list-style-type: none"> • running • scheduled (waiting for the next scheduled run time) • completed successfully • failed • misfired (did not run at the last scheduled time due to a problem such as low memory or server shutdown)
State	One of the following: <ul style="list-style-type: none"> • ENABLED indicates a job that runs according to the established recurrence pattern. • DISABLED indicates a job that is inactive, and does not run.
Recurrence	The recurrence pattern (Once, Daily, Weekly, Monthly) set for this job.
History	Click the Details link to open the Job History page for the selected job. See <i>Viewing job history</i> , page 125.
Next Scheduled	Date and time for the next run.
Owner	The user name of the administrator who scheduled the job.

Use the options on the page to manage the jobs. Some of the buttons require that you first mark the check box beside the name of each job to be included.

Option	Description
Job name link	Opens the Scheduler page, where you can edit the job definition. See <i>Scheduling presentation reports</i> , page 118.
Add Job	Opens the Scheduler page where you can define a new job. See <i>Scheduling presentation reports</i> , page 118.
Delete	Deletes from the Job Queue all jobs that have been checked in the list. After a job has been deleted, it cannot be restored. To temporarily stop running a particular job, use the Disable button.
Run Now	Starts running the jobs that have been checked in the list immediately. This is in addition to the regularly scheduled runs.

Option	Description
Enable	Reactivates disabled jobs that have been checked in the list. The job begins running according to the established schedule.
Disable	Discontinues running of enabled jobs that are checked in the list. Use this to temporarily suspend the job that you may want to restore in the future.

Viewing job history

Related topics:

- ◆ [Scheduling presentation reports, page 118](#)
- ◆ [Viewing the scheduled jobs list, page 123](#)

Use the **Presentation Reports > Job Queue > Job History** page to view information about recent attempts to run the selected job. The page lists each report separately, providing the following information.

Column	Description
Report Name	Title printed on the report.
Start Date	Date and time the report started running.
End Date	Date and time the report was complete.
Status	Indicator of whether the report succeeded or failed.
Message	Relevant information about the job, such as whether the report was emailed successfully.

Reviewing scheduled presentation reports

Related topics:

- ◆ [Presentation reports, page 107](#)
- ◆ [Running a presentation report, page 116](#)
- ◆ [Scheduling presentation reports, page 118](#)

Use the **Presentation Reports > Review Reports** page to find, access, and delete scheduled reports. By default, reports are listed from oldest to newest.

To view any report in the list, click the report name.

- ◆ If the report is a single PDF or XLS file, you may be given the option to save or open the report. This depends on your browser security settings and the plug-ins installed on your machine.

- ◆ If the report is very large, it may have been saved as multiple PDF or XLS files and stored in a ZIP file. The file is compressed using ZIP format regardless of whether the report was created on a Windows or Linux machine. Save the ZIP file, then extract the PDF or XLS files it contains to view the report content.
- ◆ Hover the mouse pointer over the report icon next to the report name to see if the report is one or multiple files.

To limit the list to reports that will be deleted soon, mark the **Show only reports due to be purged** check box. The length of time that reports are stored is configured on the Settings > Reporting > Preferences page (see [Configuring reporting preferences](#), page 354).

To search the report list, first select an entry from the **Filter by** drop-down list, and then enter all or part of a name or date. You can search by:

- ◆ The report or job name
- ◆ The name of the administrator that scheduled the report (Requestor)
- ◆ The date the report was created (Creation Date)
- ◆ The date the report is due to be deleted (Purge Date)

Enter your search term, and then click **Go**. The search is case-sensitive.

Click **Clear** to remove the current search term, and then either perform a different search or click **Refresh** to display the complete list of reports.

If a recently completed report does not appear on the Review Reports page, you can also click **Refresh** to update the page with the latest data.

To delete a report, click the X to the right of the report file size.

To see the status of a scheduled report job, click **Job Queue** at the top of the page. See [Viewing the scheduled jobs list](#), page 123, for more information about using the job queue.

To schedule a new report job, click **Scheduler** (see [Scheduling presentation reports](#), page 118).

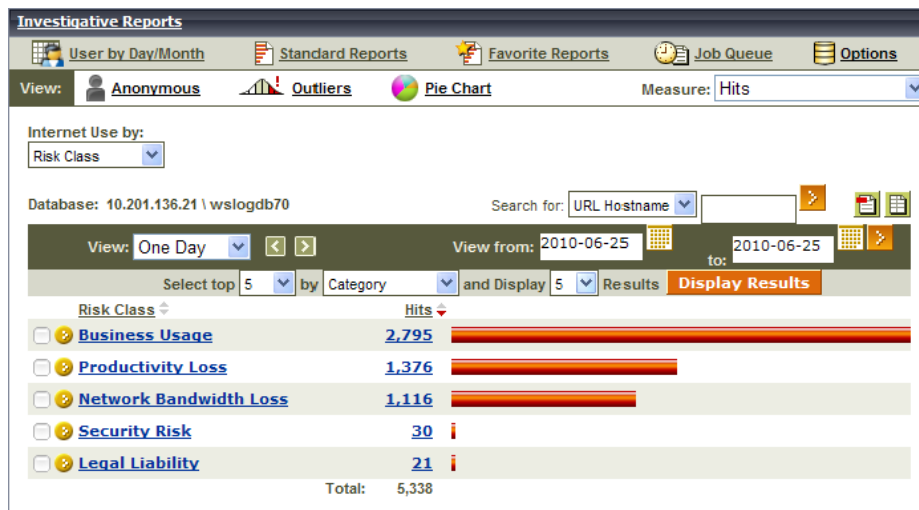
Investigative reports

Related topics:

- ◆ [Summary reports](#), page 129
- ◆ [Multi-level summary reports](#), page 133
- ◆ [Flexible detail reports](#), page 134
- ◆ [User Activity Detail reports](#), page 138
- ◆ [Standard reports](#), page 142
- ◆ [Favorite investigative reports](#), page 143
- ◆ [Scheduling investigative reports](#), page 146
- ◆ [Outliers reports](#), page 149
- ◆ [Output to file](#), page 150
- ◆ [Database connection and report defaults](#), page 368

Use the **Reporting > Investigative Reports** page to analyze Internet filtering activity in an interactive way.



Initially, the main Investigative Reports page shows a summary report of activity by risk class (see [Risk classes](#), page 43).



Work in the summary report view by clicking the available links and elements to explore areas of interest and gain general insight into your organization's Internet usage (see [Summary reports](#), page 129).

Multi-level summary reports (see [Multi-level summary reports](#), page 133) and flexible detail reports (see [Flexible detail reports](#), page 134) let you analyze the information from different perspectives.

Other report views and investigative reports features can be accessed from links at the top of the page. See the table below for a list of links and the features they access. (Not all links are available on all pages.)

Option	Action
User by Day/Month	Displays a dialog box that lets you define a report of a specific user's activity, covering either a day or a month. For more information, see User Activity Detail reports, page 138 .
Standard Reports	Displays a list of predefined reports so you can quickly see a specific combination of data. See Standard reports, page 142 .
Favorite Reports	Lets you save the current report as a Favorite, and displays a list of existing Favorites that you can generate or schedule. See Favorite investigative reports, page 143 .
Job Queue	Displays the list of scheduled investigative reports jobs. See Scheduling investigative reports, page 146 .
Outliers	Displays reports showing Internet usage that is significantly different from average. See Outliers reports, page 149 .
Options	Displays the page for selecting a different Log Database for reporting. The Options page also lets you customize certain reporting features, such as the time period initially shown on summary reports and the default columns for detail reports. See Database connection and report defaults, page 368 .
	Click this button, at the right of the Search fields, to export the current report to a spreadsheet file compatible with Microsoft Excel. You are prompted to either open or save the file. To open the file, Microsoft Excel 2003 or later must be installed. See Output to file, page 150 .
	Click this button, at the right of the Search fields, to export the current report to a PDF file compatible with Adobe Reader. You are prompted to either open or save the file. To open the file, Adobe Reader version 7.0 or later must be installed. See Output to file, page 150 .

Keep in mind that reporting is limited to the information that has been recorded in the Log Database. If you disable logging for user names, IP addresses, or selected categories (see [Configuring how filtered requests are logged, page 355](#)), that information cannot be included. Similarly, if you disable logging for certain protocols (see [Editing a protocol filter, page 53](#)), requests for those protocols are not available. If you want reports to show both the domain name (www.domain.com) and the path to a particular page in the domain (/products/productA) you must log full URLs (see [Configuring how URLs are logged, page 365](#)).

Websense investigative reports are limited by the processor and available memory of the machine running TRITON - Web Security, as well as some network resources. Some large reports may take a very long time to generate. The progress message

includes an option for saving the report as a Favorite so you can schedule it to run at another time. See [Scheduling investigative reports](#), page 146.

Summary reports

Related topics:

- ◆ [Multi-level summary reports](#), page 133
- ◆ [Flexible detail reports](#), page 134
- ◆ [User Activity Detail reports](#), page 138
- ◆ [Standard reports](#), page 142
- ◆ [Favorite investigative reports](#), page 143
- ◆ [Scheduling investigative reports](#), page 146
- ◆ [Outliers reports](#), page 149
- ◆ [Output to file](#), page 150

Initially, the investigative reports page gives a summary report of usage for all users by risk class, showing the current day's activity from the Log Database. The measurement for this initial bar chart is Hits (number of times the site was requested). To configure the time period for this initial summary report, see [Database connection and report defaults](#), page 368.

Use the links and options on the page to quickly change the information reported, or drill down into the report details.

1. Customize the way that results are quantified by selecting one of the following options from the **Measure** list.

Option	Description
Hits	The number of times the URL was requested. Depending on how Log Server is configured, this may be true hits, which logs a separate record for each separate element of a requested site, or it may be visits, which combines the different elements of the site into a single log record.
Bandwidth [KB]	The amount of data, in kilobytes, contained in both the initial request from the user and the response from the Web site. This is the combined total of the Sent and Received values. Keep in mind that some integration products do not send this information to Websense software. Two examples are Check Point FireWall-1 and Cisco PIX Firewall. If your integration does not send this information, and Network Agent is installed, enable Log HTTP requests to enable bandwidth-based reporting. See Configuring NIC settings , page 382.
Sent [KB]	The number of kilobytes sent as the Internet request. This represents the amount of data transmitted, which may be a simple request for a URL, or may be more significant (for example, if the user is registering for a Web site.)

Option	Description
Received [KB]	<p>The number of kilobytes of data received in response to the request, including all text, graphics, and scripts on the page.</p> <p>For sites that are blocked, the number of kilobytes varies according to the software creating the log record. When Network Agent logs the records, the number of bytes received for a blocked site represents the size of the Websense block page.</p> <p>If the log record is created by Websense Security Gateway, as a result of scanning, the kilobytes received represents the size of the page scanned. See Scanning and SSL Bypass Options, page 155, for more information scanning.</p> <p>If another integration product creates the log records, the kilobytes received for a blocked site may be zero (0), may represent the size of the block page, or may be a value obtained from the requested site.</p>
Browse Time	<p>An estimate of the amount of time spent viewing the site. See What is Internet browse time?, page 106.</p>

2. Change the primary grouping of the report by selecting an option from the **Internet Use by** list above the report.

Options vary according to the contents of the Log Database and certain network considerations. For example, if there is only one group or domain in the Log Database, Groups and Domains do not appear in this list. Similarly, if there are too many users (more than 5,000) or groups (more than 3,000), those options do not appear. (Some of these limits can be configured. See [Display and output options](#), page 370.)

3. Click a name in the left column (or the arrow beside the name) to display a list of options, such as by user, by domain, or by action.

The options listed are similar to those listed under Internet Use by, customized to be a meaningful subset of the content currently displayed.



Note

Sometimes an option, such as User or Group, appears in red lettering. In this case, selecting that option may produce a very large report that may be slow to generate. Consider drilling down further into the details before selecting that option.

4. Select one of those options to generate a new summary report showing the selected information for the associated entry.

For example, on a Risk Class summary report, clicking by User under the Legal Liability risk class generates a report of each user's activity in the Legal Liability risk class.

5. Click a new entry in the left column, and then select an option to see more detail about that particular item.
6. Use the arrows beside a column heading to change the report's sort order.

7. Control the summary report with the following options above the chart. Then, delve into related details by clicking the elements of the new report.

Option	Action
Report path (User > Day)	Beside the Internet use by list is a path showing the selections that created the current report. Click any link in the path to return to that view of the data.
View	<p>Select a period for the report: One Day, One Week, One Month, or All. The report updates to show data for the selected period.</p> <p>Use the adjacent arrow buttons to move through the available data, one period (day, week, month) at a time.</p> <p>As you change this selection, the View from fields update to reflect the time period being viewed.</p> <p>The View field displays Custom, instead of a time period, if you choose specific date in the View from fields or through the Favorites dialog box.</p>
View from... to...	<p>The dates in these fields update automatically to reflect the time period being viewed when you make changes in the View field.</p> <p>Alternatively, enter exact start and end dates for the reports, or click the calendar icon to select the desired dates.</p> <p>Click the adjacent right arrow button to update the report after selecting dates.</p>
Pie Chart / Bar Chart	<p>When the bar chart is active, click Pie Chart to display the current summary report as a pie chart. Click the slice label to display the same options that are available when you click an entry in the left column of the bar chart.</p> <p>When the pie chart is active, click Bar Chart to display the current summary report as a bar chart.</p>
Full Screen	Select this option to display the current investigative report in a separate window, without the left and right navigation panes.
Anonymous / Names	<ul style="list-style-type: none"> • Click Anonymous to have reports display an internally-assigned user identification number wherever a user name would have appeared. • When names are hidden, click Names to return to showing user names. <p>Under some circumstances, user names cannot be displayed. For more information, see Configuring how filtered requests are logged, page 355.</p> <p>For more information about hiding user-identifying information, see Anonymizing investigative reports, page 132.</p>
Search for	<p>Select a report element from the list, then enter all or part of a value for the search in the adjacent text box.</p> <p>Click the adjacent arrow button to start the search and display results.</p> <p>Entering a partial IP address, such as 10.5., searches for all subnets, 10.5.0.0 through 10.5.255.255 in this example.</p>

8. Add a subset of information for all or selected entries in the left column by creating a multi-level summary report. See [Multi-level summary reports](#), page 133.
9. Create a tabular report for a specific item in the left column by clicking the adjacent number or measurement bar. This detailed report can be modified to meet your specific needs. See [Flexible detail reports](#), page 134.

Anonymizing investigative reports

If you want to prevent identifying information from appearing in investigative reports, you have several options.

- ◆ The most absolute method is to prevent the logging of user names and source IP addresses. In this case, no user-identifying information is recorded in the Log Database, making it impossible for investigative or presentation reports to include the information. See [Configuring how filtered requests are logged](#), page 355, for instructions.
- ◆ If some administrators need access to reports that include user information, but other administrators should never see user information, use delegated administration roles to control reporting access. You can configure roles to grant access to investigative reports, but hide user names in reports. See [Delegated Administration and Reporting](#), page 285, for details.
- ◆ If you sometimes need to generate reports that contain user information, but sometimes need to generate anonymous reports, use the **Anonymous** option at the top of the Investigative Reports page to hide user names and, optionally, source IP addresses temporarily. See [The Anonymous option](#), page 132, for details.

The Anonymous option

By default, clicking **Anonymous** hides only user names, continuing to show source IP addresses in reports. You can configure Websense software to instead hide both user names and source IP addresses when Anonymous is selected:

1. On the TRITON - Web Security machine, open the **wse.ini** file in a text editor. (By default, this file is located in C:\Program Files\Websense\webroot\Explorer.)
2. Add the following line under the **[explorer]** heading:

```
encryptIP=1
```
3. Save and close the file.

Now, any time you click Anonymous, all user-identifying information is hidden.

When you click **Anonymous**, and then move to a different view of the data, such as detail view or outliers, user names remain hidden in the new report. However, to return to the summary view with the names hidden, you must use the links at the top of the report, not the breadcrumbs in the banner.

Multi-level summary reports

Related topics:

- ◆ [Investigative reports](#), page 127
- ◆ [Summary reports](#), page 129
- ◆ [Flexible detail reports](#), page 134
- ◆ [User Activity Detail reports](#), page 138
- ◆ [Standard reports](#), page 142
- ◆ [Favorite investigative reports](#), page 143
- ◆ [Scheduling investigative reports](#), page 146
- ◆ [Outliers reports](#), page 149
- ◆ [Output to file](#), page 150

Multi-level summary reports show a second level of information to supplement the primary information displayed. For example, if the primary display shows risk classes, you can define a second level to learn which categories have been requested most within each risk class. As another example, if the primary report shows requests for each category, you might show the top 5 categories and the 10 users who made the most requests to each.

Use the settings immediately above the summary report to create a multi-level summary report.



1. In the **Select top** list, choose a number to designate how many primary entries (left column) to report. The resulting report includes the primary entries with the largest values. (This shows the earliest dates if Day is the primary entry.)
Alternatively, mark the check box beside the desired individual entries in the left column to report only those entries. The **Select top** field displays **Custom**.
2. From the **by** list, choose the secondary information to report.
3. In the **Display** field, choose the number of secondary results to report for each primary entry
4. Click **Display Results** to generate the multi-level summary report.
The summary report updates to show only the selected number of primary entries. Below the bar for each primary entry, a list of secondary entries appears.
5. Use the arrows beside a column heading to change the report's sort order.

To return to a single-level summary report, select a different option under **Internet Use by**. Alternatively, click one of the primary or secondary entries, and select an option to generate a new investigative report of that information.

Flexible detail reports

Related topics:

- ◆ [Investigative reports](#), page 127
- ◆ [Summary reports](#), page 129
- ◆ [Multi-level summary reports](#), page 133
- ◆ [Favorite investigative reports](#), page 143
- ◆ [Scheduling investigative reports](#), page 146
- ◆ [Outliers reports](#), page 149
- ◆ [Output to file](#), page 150
- ◆ [Database connection and report defaults](#), page 368
- ◆ [Columns for flexible detail reports](#), page 136

Detail reports give you a tabular view of the information in the Log Database. Access the detail report view from the main page after viewing a summary report for which you want more detail.

You can request a detail view from any row. However, when requesting a detail report based on hits, it is best to start from a row that shows fewer than 100,000 hits. If there are more than 100,000 hits for a particular row, the hits value displays in red to alert you that a detail report may be slow to generate.

Detail report view is considered *flexible* because it lets you design your own report. You can add or delete columns of information, and change the order of the columns displayed. The information is sorted according to order of the columns. You can even reverse the sort order within any column from ascending to descending, or vice versa.

Websense investigative reports are limited by the processor and available memory of the machine running TRITON - Web Security, as well as some network resources. Requests for large reports may time out. When you request a large report, you are given options for generating the report without timeouts.



Important

In any drop-down or values list, some options may appear in red. The red lettering indicates that selecting this option may result in a very large report. It is generally more effective to drill down further into the details before selecting that option.

1. Generate a summary report or multi-level report on the investigative reports main page. (See [Summary reports](#), page 129, or [Multi-level summary reports](#), page 133.)
2. Drill down into the results to focus on the information of immediate interest.

When generating a report on hits, it is best to drill down to an entry that shows fewer than 100,000 hits before opening the detail report view.

- Click the number or the bar on the row that you want to explore in more detail. To include multiple rows in one report, mark the check box for each row before clicking the number or bar on one row.

A pop-up message shows progress while the detail report loads.



Note

If the report takes a long time to generate, consider saving it as a Favorite report by clicking the link in the Loading message, and scheduling it to run later. See [Favorite investigative reports](#), page 143.

- Review the information in the initial report.

The default columns vary, depending on whether you are reporting on hits, bandwidth, or browse time, and on the selections made on the Options page. (See [Database connection and report defaults](#), page 368.)

- Click **Modify Report** at the top of the page.

The **Current Report** list in the Modify Report dialog box shows which columns appear in the current detail report.

- Select a column name in the **Available Columns** or **Current Report** list, and click the right arrow (>) or left arrow (<) buttons to move that column to the other list.

Choose a maximum of 7 columns for the report. The column showing the measure (hits, bandwidth, browse time) from the initial summary report always appears as the right-most column. It does not appear as a choice when modifying the report.

See [Columns for flexible detail reports](#), page 136, for a list of the columns available, and a description of each.

- Select a column name in the **Current Report** list and use the up and down arrow buttons to change the order of the columns.

The column at the top of the Current Report list becomes the left column in the report.

- Click the **Summary** or **Detail** link above the report to toggle between the two displays.

Option	Description
Summary	You must remove the Time column to display a summary report. Summary reports group into a single entry all records that share a common element. The specific element varies, according to the information reported. Typically, the right-most column before the measure shows the summarized element.
Detail	The Detail option displays every record as a separate row. The Time column can be displayed.

9. Click **Submit** to generate the report you defined.
10. Use the following options to modify the displayed report.
 - Use the **View** options above the report to change the time period reported.
 - Click the up or down arrow in a column heading to reverse the sort order for that column, and the associated data.
 - Use the **Next** and **Prev** links above and below the report to display additional pages of the report, if any. By default, each page contains 100 rows, which can be adjusted to fit your needs. See [Display and output options, page 370](#).
 - Click the URL to open the requested Web site in a new window.
11. Click **Favorite Reports** if you want to save the report so that you can generate it again quickly or on a recurring basis (see [Saving a report as a Favorite, page 144](#)).

Columns for flexible detail reports

Related topics:

- ◆ [Flexible detail reports, page 134](#)
- ◆ [Favorite investigative reports, page 143](#)
- ◆ [Scheduling investigative reports, page 146](#)

The table below describes the columns available for detail reports (see [Flexible detail reports, page 134](#)).

Not all columns are available at all times. For example, if the User column is displayed, Group is not available; if Category is displayed, Risk Class is not available.

Column Name	Description
User	Name of the user who made the request. User information must be available in the Log Database to include it on reports. Group information is not available in user-based reports.
Day	Date the Internet request was made.
URL Hostname	Domain name (also called hostname) of the requested site.
Domain	Directory service domain for the directory-based client (user or group, domain, or organizational unit) that made the request.
Group	Name of the group to which the requestor belongs. Individual user names are not given on group-based reports. If the user who requested the site belongs to more than one group in the directory service, the report lists multiple groups in this column.

Column Name	Description
Risk Class	Risk class associated with the category to which the requested site belongs. If the category is in multiple risk classes, all relevant risk classes are listed. See Assigning categories to risk classes , page 353.
Directory Object	Directory path for the user who made the request, excluding the user name. Typically, this results in multiple rows for the same traffic, because each user belongs in multiple paths. If you are using a non-LDAP directory service, this column is not available.
Disposition	Action Websense software took as a result of the request (for example, category permitted or category blocked).
Source Server	IP address of the machine sending requests to Filtering Service. In standalone deployments, this is the Network Agent IP address. In integrated deployments, this is the gateway, firewall, or cache IP address. With Websense Web Security Gateway Anywhere, use this option to identify requests filtered by the hybrid service from both on-site (filtered location) and off-site users.
Protocol	Protocol of the request (for example, HTTP or FTP).
Protocol Group	Master Database group in which the requested protocol falls (for example, Remote Access or Streaming Media).
Source IP	IP address of the machine from which the request was made. With Websense Web Security Gateway Anywhere, you can use this option to review requests coming from a specific hybrid filtered location. See Define the locations filtered by hybrid filtering , page 175.
Destination IP	IP address of the requested site.
Full URL	Domain name and path for the requested site (example: http://www.mydomain.com/products/itemone/). If you are not logging full URLs, this column is blank. See Configuring how URLs are logged , page 365.
Month	Calendar month the request was made.
Port	TCP/IP port over which the user communicated with the site.
Bandwidth	The amount of data, in kilobytes, contained in both the initial request from the user and the response from the Web site. This is the combined total of the Sent and Received values. Keep in mind that some integration products do not send this information to Websense software. Two examples are Check Point FireWall-1 and Cisco PIX Firewall. If your integration does not send this information, and Websense Network Agent is installed, activate the Log HTTP requests option for the appropriate NIC to enable reporting on bandwidth information. See Configuring NIC settings , page 382.
Bytes Sent	Number of bytes sent as the Internet request. This represents the amount of data transmitted, which may be a simple request for a URL, or may be a more significant submission if the user is registering for a Web site, for example.

Column Name	Description
Bytes Received	<p>Number of bytes received from the Internet in response to the request. This includes all text, graphics, and scripts that make up the site.</p> <p>For sites that are blocked, the number of bytes varies according to the software creating the log record. When Websense Network Agent logs the records, the number of bytes received for a blocked site represents the size of the block page.</p> <p>If the log record is created by Websense Security Gateway, as a result of scanning, the bytes received represents the size of the page scanned. See Scanning and SSL Bypass Options, page 155, for more information on scanning.</p> <p>If another integration product creates the log records, the bytes received for a blocked site may be zero (0), may represent the size of the block page, or may be a value obtained from the requested site.</p>
Time	Time of day the site was requested, shown in the HH:MM:SS format, using a 24-hour clock.
Category	Category to which the request was assigned. This may be a category from the Master Database or a custom category.

User Activity Detail reports

Related topics:

- ◆ [Investigative reports, page 127](#)

Click the **User by Day/Month** link to generate a User Activity Detail report for one user. This report gives a graphical interpretation of the user's Internet activity for a single day or a full month.

First, generate a report for a specific user for a selected day. From that report, you can generate a report of the same user's activity for a full month. For detailed instructions, see:

- ◆ [User activity detail by day, page 138](#)
- ◆ [User activity detail by month, page 140](#)

User activity detail by day

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [User Activity Detail reports, page 138](#)
- ◆ [User activity detail by month, page 140](#)

The User Activity Detail by Day report gives a more in-depth view of a specific user's activity on one day.

1. Select **User by Day/Month** at the top of the main page. The User Detail by Day dialog box appears.

2. Enter a user's name, or a portion of the name, in the **Search for user** field, and then click **Search**.

The search displays a scrolling list of up to 100 matching user names from the Log Database.

3. Make a selection from the **Select user** list.
4. In the **Select day** field, either accept the last activity date that appears by default, or choose a different date.

You can type the new date or click the calendar icon to select a date. The calendar selection box indicates the date range covered by the active Log Database.

5. Click **Go to User by Day** to see a detailed report of activity for that user on the requested date.

The initial report shows a timeline of the user's activity in 5-minute increments. Each request appears as an icon, which corresponds to a Websense Master Database category. A single icon represents all custom categories. (The color of the icons corresponds to the risk grouping shown on the User Activity by Month reports. See [User activity detail by month](#), page 140.)

Rest the mouse over an icon to show the exact time, category, and action for the associated request.

Use the controls listed below to modify the report display or to see a legend.

Option	Description
Previous Day / Next Day	Display this user's Internet activity for the previous or next calendar day.
Table View	Displays a list of each requested URL, giving the date and time of the request, the category, and the action taken (blocked, permitted, or other).
Detail View	Displays the initial, graphical view of the report.
Group Similar Hits / View All Hits	Combines into a single row all requests that occurred within 10 seconds of each other and have the same domain, category, and action. This results in a shorter, summarized view of information. The standard time threshold is 10 seconds. If you need to change this value, see Display and output options , page 370. After you click the link, it becomes View All Hits, which restores the original list of each request.
Category View Control	Displays a list of each category in the current report, showing both the category name and the icon representing that category. Control which categories appear in the report by marking the check boxes for the categories to be included. Then, click Accept to update the report according to your selections.

6. Click **User Activity Detail by Month**, above the report, to view the same user's activity for the full month. See [User activity detail by month, page 140](#), for more information.

User activity detail by month

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [User Activity Detail reports, page 138](#)
- ◆ [User activity detail by day, page 138](#)
- ◆ [Category mapping, page 140](#)

While the User Activity Detail by Day report is open, you can switch to see the monthly activity for that user.

1. Open a User Activity Detail by Day report. See [User activity detail by day, page 138](#).
2. Click **User Activity Detail by Month** at the top.
The new report displays a calendar image, with each day's area showing small colored blocks representing the user's Internet activity for that day. Requests to sites in custom categories are shown as gray blocks.
3. Click **Database Category Legend** at the top left to see how the colors represent low to high potential risk for the requested site.
The category assignments are fixed, and cannot be changed. See [Category mapping, page 140](#).
4. Click **Prev** or **Next** to display this user's Internet activity for the previous or the next month.

Category mapping

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [User Activity Detail reports, page 138](#)
- ◆ [User activity detail by month, page 140](#)

The following list identifies which categories are represented by each of the colors on the User Activity by Day and User Activity Detail by Month reports.

Keep in mind that category names in the Master Database are subject to change. Additionally, categories may be added or deleted at any time.

Color	Categories
Gray	Custom Categories Non-HTTP traffic
Dark Blue	Business and Economy and all its subcategories Education , and all its subcategories Health Information Technology , including the Search Engines and Portals, and Web Hosting subcategories Miscellaneous subcategories Content Delivery Networks, Dynamic Content, Images (Media), Image Servers, and Private IP Addresses Productivity/Advertisements
Light Blue	Drugs/Prescribed Medications Government and its Military subcategory Information Technology/URL Translation Sites Miscellaneous , parent category only News and Media , parent category only Special Events
Yellow Green	Abortion and all its subcategories Adult Material/Sex Education Bandwidth , including the subcategories Internet Radio and TV, Personal Network Storage and Backup, and Streaming Media Entertainment , including its subcategory MP3 Games Government/Political Organizations Information Technology/Computer Security Internet Communication/Web-based Email Miscellaneous/File Download Servers Miscellaneous/Network Errors News and Media/Alternative Journals Productivity , including its subcategories Instant Messaging, Message Boards and Clubs, and Online Brokerage and Trading Religion and its subcategories Non-Traditional Religions and Occult and Folklore, and Traditional Religions Security , parent category only Shopping and all its subcategories Social Organizations and all its subcategories Society and Lifestyles , including its subcategories Gay or Lesbian or Bisexual Interest, Hobbies, Personal Web Sites, and Restaurants and Dining Sports and all its subcategories Travel User-Defined Vehicles

Color	Categories
Orange	Adult Material /Nudity Advocacy Groups Bandwidth /Internet Telephony Drugs and its subcategories Abused Drugs, Marijuana, and Supplements and Unregulated Compounds Information Technology /Proxy Avoidance Internet Communication and its subcategory Web Chat Job Search Miscellaneous /Uncategorized Productivity subcategories Freeware and Software Download, and Pay-to-Surf Religion Society and Lifestyles subcategories Alcohol and Tobacco, and Personals and Dating Tasteless Weapons
Red	Adult Material and these subcategories: Adult Content, Lingerie and Swimsuit, and Sex Bandwidth /Peer-to-Peer File Sharing Gambling Illegal or Questionable Information Technology /Hacking Militancy and Extremist Racism and Hate Security subcategories Keyloggers, Malicious Web Sites, Phishing, and Spyware Violence

Standard reports

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Favorite investigative reports, page 143](#)
- ◆ [Scheduling investigative reports, page 146](#)

Standard reports let you display a particular set of information quickly without using the drill-down process.

1. Click the **Standard Reports** link on the main Investigative Reports page.

2. Choose the report containing the desired information. The following reports are available.

Highest Activity Levels

- Which users have the most hits?
- Top 10 users for top 10 visited URLs
- Top 5 users activity in Shopping, Entertainment, and Sports
- Top 5 URLs for the top 5 visited categories

Highest Bandwidth Consumption

- Which groups are consuming the most bandwidth
- Groups consuming most bandwidth in Streaming Media
- Detail URL report on users by Network Bandwidth Loss
- Top 10 groups for Bandwidth categories

Most Time Online

- Which users spent the most time online
- Which users spent the most time on sites in Productivity categories

Most Blocked

- Which users were blocked most?
- Which sites were blocked most?
- Detail URL report on users who were blocked
- Top 10 blocked categories

Highest Security Risk

- Top categories posing a security risk
- Top users of P2P protocol
- Top users of sites in Security categories
- URLs for top 10 machines with spyware activity

Legal Liability

- Legal Liability Risk by Category
 - Top users in Adult categories
-

3. View the report that appears.
4. Save the report as a Favorite if you want to run it on a recurring basis. See [Favorite investigative reports](#), page 143.

Favorite investigative reports

Related topics:

- ◆ [Investigative reports](#), page 127
- ◆ [Scheduling investigative reports](#), page 146

You can save most investigative reports as **Favorites**. This includes reports you generate by drilling down to specific information, standard reports, and detail reports that you have modified to meet your specific needs. Then, run the Favorite report at any time, or schedule it to run on specific days and times.

In organizations that use delegated administration, permission to save and schedule Favorites is set by the Super Administrator. Administrators who are granted this permission can run and schedule only the Favorites they saved; they do not have access to Favorites saved by other administrators.

For detailed instructions on working with Favorite reports, see:

- ◆ [Saving a report as a Favorite, page 144](#)
- ◆ [Generating or deleting a Favorite report, page 144](#)
- ◆ [Modifying a Favorite report, page 145](#)

Saving a report as a Favorite

Related topics:

- ◆ [Favorite investigative reports, page 143](#)
- ◆ [Modifying a Favorite report, page 145](#)

Use the following procedure to save a report as a Favorite.

1. Generate an investigative report with the desired format and information.
2. Click **Favorite Reports**.
3. Accept or modify the name displayed by TRITON - Web Security.
The name may contain letters, numbers and underscore characters (_). No blanks or other special characters can be used.
4. Click **Add**.
The report name is added to the list of Favorites.
5. Select a report on this list, then select an option for managing the report.
Depending on the option you choose, see:
 - [Generating or deleting a Favorite report, page 144](#)
 - [Scheduling investigative reports, page 146](#)

Generating or deleting a Favorite report

Related topics:

- ◆ [Favorite investigative reports, page 143](#)
- ◆ [Modifying a Favorite report, page 145](#)

You can generate a Favorite report at any time, or delete one that has become obsolete.

1. Click **Favorite Reports** to display a list of reports saved as favorites.

**Note**

If your organization uses delegated administration, this list does not include Favorite reports saved by other administrators.

2. Select a report from the list.
3. Do one of the following:
 - Click **Run Now** to generate and display the selected report immediately.
 - Click **Schedule** to schedule a report to run later or on a recurring basis. See [Scheduling investigative reports, page 146](#), for more information.
 - Click **Delete** to remove the report from the Favorites list.

Modifying a Favorite report

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Favorite investigative reports, page 143](#)

You can easily create a new Favorite report that is similar to an existing Favorite report, as follows.

1. Click **Favorite Reports** to display a list of reports saved as favorites.

**Note**

If your organization uses delegated administration, this list does not include Favorite reports saved by other administrators.

2. Select and run the existing Favorite report that most closely resembles the new report you want to create. (See [Generating or deleting a Favorite report, page 144](#).)
3. Modify the displayed report as desired.
4. Click **Favorite Reports** to save the revised display as a Favorite report with a new name. (See [Saving a report as a Favorite, page 144](#).)

Scheduling investigative reports

Related topics:

- ◆ [Favorite investigative reports](#), page 143
- ◆ [Saving a report as a Favorite](#), page 144
- ◆ [Managing scheduled investigative reports jobs](#), page 148

You must save an investigative report as a Favorite before it can be scheduled to run at a later time or on a repeating cycle. When the scheduled report job runs, the resulting reports are sent via email to the recipients you designate. As you create scheduled jobs, consider whether your email server will be able to handle the size and quantity of the attached report files.

Scheduled report files are stored in the following directory:

```
<install_path>\webroot\Explorer\<name>\
```

The default installation path is C:\Program Files\WebSense\Web Security. If the scheduled job has only one recipient, *<name>* is the first portion of the email address (before the @). In the case of multiple recipients, the reports are saved in a directory called Other.



Note

The reports saved from a repeating job use the same file name each time. If you want to save files for longer than a single cycle, be sure to change the file name or copy the file to another location.

Depending on the size and number of reports scheduled, this directory could become very large. Be sure to clear the directory periodically, eliminating unneeded report files.

1. Save one or more reports as Favorites. (See [Saving a report as a Favorite](#), page 144).
2. Click **Favorite Reports** to display a list of reports saved as favorites.



Note

If your organization uses delegated administration roles, this list does not include favorite reports saved by other administrators.

3. Highlight up to 5 reports to run as part of the job.
4. Click **Schedule** to create a scheduled report job, and then provide the information requested on the Schedule Report page.

It is advisable to schedule report jobs on different days or at different times, to avoid overloading the Log Database and slowing performance for logging and interactive reporting.

Field	Description
Recurrence	Select the frequency (Once, Daily, Weekly, Monthly) for running the report job.
Start Date	Choose the day of the week or calendar date for running the job the first (or only) time.
Run Time	Set the time of day for running the job.
Email to	Use the Additional Email Addresses field to add the appropriate addresses to this list. Highlight one or more email addresses to receive the reports in the job. (Be sure to deselect any that should not receive the reports.)
Additional Email Addresses	Enter an email address, and then click Add to put it on the Email to list. The new email address is automatically highlighted with the other selected email addresses.
Customize email subject and body text	Mark this check box to customize your email notification subject line and body text. If this box is not checked, the default subject and body text are used.
Email Subject	Enter the text to appear as the email subject line when scheduled reports are distributed. The default email subject reads: Investigative Reports scheduled job
Email Text	Enter text to be added to the email message for distributing scheduled reports. The email reads as shown below, with your text in place of <CUSTOM TEXT>. Report scheduler generated the attached file or files on <date time>. <CUSTOM TEXT> To view the generated report(s), click on the following link(s). Note: The link will not work if the recipient does not have access to the web server from which the job was sent.
Schedule Job Name	Assign a unique name for the scheduled job. The name identifies this job in the Job Queue. See Managing scheduled investigative reports jobs , page 148.

Field	Description
Output Format	Choose the file format for the scheduled reports: PDF: Portable Document Format files are viewed in Adobe Reader. Excel: Excel spreadsheet files are viewed in Microsoft Excel.
Date Range	Set the date range to be covered by reports in this job. All Dates: all available dates in the Log Database. Relative: Choose a time period (Days, Weeks, or Months) and the specific period to include (This, Last, Last 2, and so on). Specific: set specific dates or a date range for the reports in this job.

5. Click **Next** to display the Schedule Confirmation page.
6. Click **Save** to save your selections and go to the Job Queue page (see [Managing scheduled investigative reports jobs, page 148](#)).

Managing scheduled investigative reports jobs

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Scheduling presentation reports, page 118](#)

When you create a scheduled job for investigative reports, the **Job Queue** page appears, showing the new job and a list of existing scheduled jobs. You can also access the page by clicking the **Job Queue** link on the main investigative reports page.



Note

If your organization uses delegated administration, this page does not show jobs scheduled by other administrators.

The **Schedule Report Detail** section lists each scheduled job in the order it was created showing an overview of the defined schedule and the job status. In addition, the following options are available.

Option	Description
Edit	Displays the schedule defined for this job, and allows you to modify it, as needed.
Delete	Deletes the job and adds an entry to the Status Log section showing the job as Deleted.

The **Status Log** section lists each job that has changed in some way, showing the scheduled start time for the job, the actual end time, and the status.

Click **Clear Status Log** to remove all entries in the Status Log section.

Outliers reports

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Summary reports, page 129](#)

An Outliers report shows which users have the most unusual Internet activity in the database. Websense software calculates the average activity for all users per category, per day, per action (disposition), and per protocol. It then displays the user activity that has the most statistically significant variance from the average. Variance is calculated as the standard deviation from the mean.

1. On the main Investigative Reports page, generate a summary report that displays the information for which you want to see outliers. The report selections underlined and shown in blue beside the Internet Use by field are reflected in the Outliers report.

For example, to view outliers by hits for a particular category, select **Category** in the **Internet Use by** list, and select **Hits** as the **Measure**.



Note

Outliers reports cannot be generated for browse time. If you start from a summary report showing browse time, the Outliers report is based on hits.

2. Click **Outliers**.

The rows are sorted in descending order with the highest variance shown first. Each row shows:



- Total (hits or bandwidth) for the user, category, protocol, day, and action.
 - Average (hits or bandwidth) for all users, for that category, protocol, day, and action.
 - Variance from the average for the user.
3. To see an individual user's activity in this category over time, click the user name. For example, if one user's activity is noticeably high for a certain day, click that user's name to see a report that gives a more in-depth understanding of the user's overall activity.

Output to file

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Printing investigative reports, page 150](#)

After you generating an investigative report, you can use the buttons above the report to save it to a file. The button you click determines the format of the file.

Option	Description
	<p>Saves the report in XLS format.</p> <p>If Microsoft Excel 2003 or later is installed on the machine from which you are accessing TRITON - Web Security, you are prompted to view or save the report. Otherwise, you are prompted to select a directory and file name for the saved report.</p> <p>Use the options in Microsoft Excel to print, save, or email the report.</p>
	<p>Generates a report in PDF format.</p> <p>If Adobe Reader v7.0 or later is installed on the machine from which you are accessing TRITON - Web Security, you are prompted to view or save the report. Otherwise, you are prompted to select a directory and file name for the saved report.</p> <p>Use the options in Adobe Reader to print, save, or email the report.</p>

Printing investigative reports

Related topics:

- ◆ [Investigative reports, page 127](#)
- ◆ [Output to file, page 150](#)

To print investigative reports:

- ◆ Use the Web browser print function while the report is displayed.
- ◆ Create a PDF or XLS file, and then use the print function in Adobe Reader or Microsoft Excel (see [Output to file, page 150](#)).

Although reports have been set up to print successfully from the browser, you may want to test printing to check the result.

User Activity Detail by Month reports are configured to print in landscape mode. All other reports are configured for portrait mode.

When you design your own report (see *Flexible detail reports*, page 134), the column widths differ according to the information included. The page orientation changes to landscape if the report is wider than 8 1/2 inches.

The content of the page is either 7 1/2 inches or 10 inches wide. In the case of A4, the margins are slightly narrower but still within the print range. (The default paper size is Letter, or 8.5 x 11 inches. If you are working with A4 paper, be sure to change this setting in the `wse.ini` file. See *Display and output options*, page 370.)

Accessing self-reporting

Related topics:

- ◆ [Investigative reports](#), page 127
- ◆ [Configuring reporting preferences](#), page 354
- ◆ [Self-reporting](#), page 372

Websense self-reporting allows you to evaluate your own Internet browsing activities and adjust them, as needed, to meet organizational guidelines. It also accommodates government regulations that require organizations to let users see the type of information being collected.

If self-reporting is enabled in your organization, access it from your browser:

1. Enter the URL supplied by your Websense software administrator, or click the Self-Reporting link on the main TRITON - Web Security logon page to access the self-reporting logon page.
2. If **Policy Server** shows a drop-down list, choose the IP address for the Policy Server that logs information on your Internet activity.
Contact your Websense software administrator for assistance.
3. Enter the **User name** and **Password** you use to log on to the network.
4. Click **Log On**.

TRITON - Web Security opens to an investigative report showing your Internet activity by risk class. Click the various links and elements on the page to access other options for alternative views of the information stored on your activity. Use the **Help** system for assistance when working with the reports.

Real-Time Monitor

Related topics:

- ◆ [Use Reports to Evaluate Filtering, page 105](#)
- ◆ [Real-Time Monitor in Multiple Policy Server Deployments, page 154](#)

Use the **Reporting > Real-Time Monitor** page to review current Internet filtering activity in your network.



Important

Internet Explorer 8 Compatibility View is **not** supported with the TRITON console. If Real-Time Monitor does not display correctly in Internet Explorer 8, make sure that Compatibility View button (between the URL and the Refresh button in the browser address bar) is not selected.

Click **Start** to populate the page with data. The page shows recent Internet requests, including:

- ◆ The IP address or name of the **user** who made the request.
 - If user-based filtering is used in your network, and the user name is shown, mouse over an entry to see the IP address.
 - If a user name is longer than 30 characters, a hyphen (“-”) and the last 30 characters of the name are displayed. If you right-click to add a long user name to the search filter, delete the hyphen character from the filter field and click **Show Results** to display matching entries.
- ◆ The **URL** requested.

If the URL is too long to display in the space provided, the field shows the first 30 characters of the URL, a space, a hyphen (“-”), and a space, and then last 20 characters of the URL. Right-click the truncated URL to see the entire string.
- ◆ Whether or not the requested site was recategorized as a result of Content Gateway scanning.
 - The presence of an icon indicates that the site was dynamically recategorized based on the results of scanning. Mouse over the icon to see the original category.
 - No icon indicates that the Master Database category or custom URL category was used. (This includes sites that were scanned by Content Gateway, but not recategorized.)

- ◆ The **Category** assigned to the site.
The actual category used to filter the request is shown, whether that is the Master Database category, the custom URL category, or the category dynamically assigned as a result of scanning.
- ◆ The **Action** (permitted or blocked) applied to the request.
- ◆ The **Time** the request was passed to Real-Time Monitor.
Because Real-Time Monitor receives request information from Usage Monitor in real time, rather than reading the request from the Log Database, the request time shown here may not match the request time that appears in investigative and presentation reports.

To review current data, click **Pause** to prevent the page from continuing to refresh. When you are ready to start monitoring new information, click **Start** again.

Depending on your current settings, Real-Time Monitor holds a set number of records (250, 500, or 1000), and will always display the latest set of available records. When you pause display of new records to review current data, this can mean that the hundreds or thousands of requests that occur while the display is paused will never be displayed in the monitor. (The requests are, however, stored in the Log Database, and will appear in investigative and presentation reports.)

Filtering report data

To filter the data displayed on the screen:

1. Enter all or part of a user name or IP address, URL, category, or action in the **Filter results by** fields. You can also select a time filter to show the past 5, 10, or 15 minutes worth of applicable results.
2. Click **Show Results**.
3. To return to viewing all results, click **Clear Search Filters**.

You can also right-click any entry in the User, URL, Category, or Action fields and select the **Filter by** or **Add...to search filter** option to immediately filter results based on the selected string.

Customizing the report view

To determine how often the monitor refreshes its view, how much data is shown, and how URLs are displayed:

1. Click **Customize** in the toolbar at the top of the Content pane.
2. Select a **Data refresh rate** for the Real-Time Monitor page: 5, 10, 15, 20, 30, or 45 seconds, or 1 minute. The default is 15 seconds.
3. Select the **Number of records shown** in Real-Time Monitor: 250, 500, or 1000.
4. Indicate how URLs should appear in the monitor list. You can choose to show the entire URL, or only the domain portion of the URL (for example, www.google.com).

5. Click **OK** to save your changes and close the dialog box.

Understanding timeout behavior

By default, TRITON Unified Security Center sessions time out after 30 minutes. To run Real-Time Monitor without timing out, click **Full Screen** to open the monitor in a new window. The IP address of the monitored Policy Server appears in the Real-Time Monitor title bar. If you want to monitor multiple Policy Server instances, see [Real-Time Monitor in Multiple Policy Server Deployments](#), page 154, for considerations and instructions.

Real-Time Monitor in Multiple Policy Server Deployments

When you go to the Reporting > Real-Time Monitor page in TRITON - Web Security, Real-Time Monitor shows information for the Policy Server instance to which the management console is currently connected. This means that if you have multiple Policy Servers, when you connect the management console to a new Policy Server instance, Real-Time Monitor starts to display information for a different set of filtering clients.

If you want Real-Time Monitor to continue monitoring traffic for a specific Policy Server, regardless of which Policy Server instance TRITON - Web Security is connected to, click **Full Screen** to open the monitor in a new window. The IP address of the monitored Policy Server is displayed at the top of the screen.

- ◆ Real-Time Monitor receives Internet activity information from Usage Monitor. Each Policy Server must have a Usage Monitor instance associated with it for Real-Time Monitor to show its filtering activity.
- ◆ You can have multiple Real-Time Monitor instances running in full-screen mode, each showing data for a different Policy Server:
 1. Open TRITON - Web Security. It will connect to the central (default) Policy Server.
 2. Go to the Reporting > Real-Time Monitor page and click Full Screen. The IP address of the central Policy Server appears in the title bar.
 3. Return to TRITON - Web Security and use the Policy Server Connection button in the TRITON toolbar to connect to a different Policy Server instance.
 4. Repeat step 2.
 5. Repeat for each additional Policy Server instance in your network.
- ◆ In full screen mode, Real-Time Monitor does not time out.

7

Scanning and SSL Bypass Options

Related topics:

- ◆ [Scanning options, page 157](#)
- ◆ [Content categorization, page 158](#)
- ◆ [Tunneled protocol detection, page 160](#)
- ◆ [Security threats: Content scanning, page 161](#)
- ◆ [Security threats: File scanning, page 162](#)
- ◆ [Advanced options, page 164](#)
- ◆ [Scanning exceptions, page 166](#)
- ◆ [Data files used with scanning, page 168](#)
- ◆ [Reporting on scanning activity, page 168](#)
- ◆ [SSL decryption bypass, page 171](#)

Scanning (advanced analysis) options and SSL decryption bypass features are available with Websense Web Security Gateway and Websense Web Security Gateway Anywhere.

Scanning options support the advanced analysis of Web traffic as it flows through the Content Gateway module (the Websense on-premises proxy). Only sites that are not already blocked, based on the active policy, are analyzed.

- ◆ [Content categorization, page 158](#), categorizes content from URLs that are not in the Websense Master Database and from sites with dynamic Web 2.0 content, as identified by Websense Security Labs. Analysis returns a category for use in filtering.
- ◆ [Tunneled protocol detection, page 160](#), analyzes traffic to discover protocols tunneled over **HTTP** and **HTTPS**. Such traffic is reported to Websense Web filtering for protocol policy enforcement. Analysis includes both inbound and outbound traffic.
- ◆ [Security threats: Content scanning, page 161](#), analyzes:
 - **Inbound content to find security threats** such as malware, viruses, phishing, URL redirection, Web exploits, proxy avoidance, and others.

- **Outbound content** to discover and block malicious content like bot and spyware phone home traffic.
- ◆ *Security threats: File scanning*, page 162, applies 2 methods of inspection to detect security threats. It analyzes files with:
 - **Websense advanced detection techniques** to discover malicious content, such as viruses, Trojan horses, and worms, returning a threat category for policy enforcement.
 - **Traditional anti-virus (AV) definition files** to find virus-infected files.

When either Advanced Detection or Anti-virus Scanning analysis is enabled, you can also optionally analyze:

- **Rich Internet applications**, such as Flash files, to detect and block malicious content.
- **FTP files** to detect and block malicious content.

The **File Type Options** settings determine which types of files are analyzed for malicious content, including executable and unrecognized files. Individual file extensions may also be specified.

- ◆ The **Scanning Timeout, File Size Limit and Content Stripping** Advanced Options apply to all traffic transiting the proxy (*Advanced options*, page 164).

Several presentation reports can provide details about how advanced analysis features protect your network from attempts to access sites containing threats. See *Reporting on scanning activity*, page 168.

SSL decryption bypass options support the specification of Web sites and Web site categories that are **not** subject to decryption and analysis as they flow through the proxy. These options apply only if SSL Manager is enabled on Content Gateway. See *SSL decryption bypass*, page 171.

Scanning exceptions are lists of hostnames or URLs that are always analyzed or never analyzed. The type of analysis to always or never perform is specified per hostname/URL or group of hostnames/URLs. A list of client IP addresses whose content is never analyzed can also be specified. See *Scanning exceptions*, page 166.

Enabling scanning and SSL decryption bypass features

To enable the scanning advanced analysis and SSL decryption bypass features that are available with Websense Web Security Gateway and Gateway Anywhere, an appropriate subscription key must be entered in TRITON - Web Security. You can enter the key:

- ◆ When prompted after logging on
- ◆ On the **Settings > General > Policy Servers** page

Review current key information on the **Settings > General > Account** page.

The key is automatically passed to all Content Gateway instances associated with the current Policy Server. See *Reviewing Policy Server connections*, page 320, and *Managing Content Gateway connections*, page 327, for more information.

For information about configuring advanced analysis options, see [Scanning options, page 157](#). For information about SSL decryption bypass options, see [SSL decryption bypass, page 171](#).

Scanning options

Related topics:

- ◆ [Content categorization, page 158](#)
- ◆ [Tunneled protocol detection, page 160](#)
- ◆ [Security threats: Content scanning, page 161](#)
- ◆ [Security threats: File scanning, page 162](#)
- ◆ [Advanced options, page 164](#)
- ◆ [Scanning exceptions, page 166](#)
- ◆ [Reporting on scanning activity, page 168](#)

Scanning analysis options are available with Websense Web Security Gateway and Websense Web Security Gateway Anywhere. These options control the types of advanced analysis performed on Web traffic as it transits the Content Gateway module (the Websense on-premises proxy).

For an introduction to advanced analysis options and other options related to the Websense proxy, see [Scanning and SSL Bypass Options, page 155](#).

Use the **Settings > Scanning > Scanning Options** page to configure the following:

- ◆ [Content categorization, page 158](#)
- ◆ [Tunneled protocol detection, page 160](#)
- ◆ [Security threats: Content scanning, page 161](#)
- ◆ [Security threats: File scanning, page 162](#)
- ◆ Scanning timeout, file size, and content stripping ([Advanced options, page 164](#))

Basic settings are:

- ◆ **Off.** No analysis.
- ◆ **On.** Advanced security analysis, as modified by selected sub-options.

When the **All content from all sites** or **All files from all sites** options are enabled, all requested Web pages or files are analyzed. The only exceptions are sites listed on the Never Scan list.

In addition to the On/Off settings, analysis is performed or not performed, based on the Always Scan, Never Scan, and client IP exception lists. These lists are maintained

on the **Settings > Scanning > Scanning Exceptions** page. See [Scanning exceptions](#), page 166.



Warning

Sites on the Never Scan list are not analyzed under any circumstances. If a site on the Never Scan list is compromised, scanning options do not analyze and detect the malicious code.

When you have completed your changes on the current page, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Content categorization

Related topics:

- ◆ [Scanning options](#), page 157
- ◆ [Tunneled protocol detection](#), page 160
- ◆ [Security threats: Content scanning](#), page 161
- ◆ [Security threats: File scanning](#), page 162
- ◆ [Advanced options](#), page 164
- ◆ [Scanning exceptions](#), page 166
- ◆ [Reporting on scanning activity](#), page 168

When a URL is requested, content categorization is performed if:

- ◆ The URL has not already been blocked by the active policy
- ◆ The URL is not in the Websense Master Database, or
- ◆ The URL is a Web 2.0 site, as identified by Websense Security Labs

The category that is determined by content categorization is forwarded to Websense filtering software for policy enforcement.

Content categorization provides high value because a significant majority of Web content changes rapidly. In addition, the Internet hosts a large amount of user-generated content, such as that found on social-networking sites. Content categorization analyzes this content at the moment it is needed, when it is requested.

Content categorization can, optionally, include **analysis of URL links embedded in the content**. Such analysis can provide more accurate categorization of certain types of content. For example, a page that otherwise has little or no undesirable content, but that links to sites known to have undesirable content, can itself be more accurately categorized. Link analysis is particularly good at finding malicious links embedded in hidden parts of a page, and in detecting pages returned by image servers that link thumbnails to undesirable sites. For more information about how analysis of link

neighborhoods can improve coverage, read the Websense Security Labs blog post [In Bad Company](#).

The effectiveness of content categorization and link analysis is quantified in several presentation reports. See [Presentation reports, page 107](#), for more information.



Important

If you plan to generate reports of advanced analysis, enable full URL logging (see [Configuring how URLs are logged, page 365](#)). Otherwise, log records include only the domain (www.domain.com) of the site categorized, and individual pages within a site may fit into different categories.

If your site uses WebCatcher to report uncategorized URLs to Websense, Inc. (see [Configuring WebCatcher, page 358](#)), URLs categorized through content categorization are forwarded for inclusion in the Master Database.

To configure content categorization:

1. Go to **Settings > Scanning > Scanning Options**.
2. Select **Off** to disable content categorization.
3. Select **On** (default) to enable content categorization.
4. Select **Analyze links embedded in Web content** to include embedded link analysis in content categorization. Requests that are blocked as a result of link analysis are logged and can be viewed in Scanning Activity presentation reports.
5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Sensitivity level

The algorithms used to perform content categorization are tuned by Websense Security Labs to provide optimal results for most organizations. However, if the Optimized setting does not produce the results you expect, you can adjust the Sensitivity Level to produce more restrictive or more permissive results.

There are 5 sensitivity levels.

- ◆ **Optimized** (the middle setting) is the sensitivity level tuned by Websense Security Labs.
- ◆ **Higher** and **Highest** raise analytic sensitivity.
- ◆ **Lower** and **Lowest** reduce analytic sensitivity.

When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Tunneled protocol detection

Tunneled protocol detection analyzes traffic to discover protocols that are tunneled over HTTP and HTTPS. Traffic that is allowed to tunnel over specific ports is also analyzed. Such traffic is reported to Websense Web filtering for protocol policy enforcement. When tunneled protocol detection is enabled, analysis is performed on both inbound and outbound traffic, regardless of other scanning settings.

HTTP tunneling occurs when applications that use custom protocols for communication are wrapped in HTTP (meaning that standard HTTP request/response formatting is present) in order to use the ports designated for HTTP/HTTPS traffic. These ports are open to allow traffic to and from the Web. HTTP tunneling allows these applications to bypass firewalls and proxies, leaving a system vulnerable.

The tunneled protocol detection feature analyzes HTTP and HTTPS traffic and, when it detects a protocol, forwards it to Websense Web filtering for policy enforcement. At this point, a protocol is blocked or allowed based on policy definitions. This feature can be used to block protocols used for instant messaging, peer-to-peer applications, and proxy avoidance. Note that some applications running over HTTP (for example, Google Video) may not display the protocol block page. See [Filtering categories and protocols](#), page 40, for information about protocol filtering.



Note

Tunneled protocol detection is performed before content categorization. As a result, when a tunneled protocol is identified, protocol policy is enforced and content categorization is not performed.

Use the **Settings > Scanning > Scanning Options** page to enable and configure tunneled protocol detection:

1. Select **Off** to disable tunneled protocol detection.
2. Select **On** (default) to analyze all traffic to detect protocols tunneling over HTTP or HTTPS. Such traffic is reported to Web filtering for policy enforcement.
3. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Use the **Settings > Scanning > Scanning Exceptions** page to specify trusted sites that are never analyzed ([Scanning exceptions](#), page 166).

Security threats: Content scanning

Related topics:

- ◆ [Scanning options](#), page 157
- ◆ [Content categorization](#), page 158
- ◆ [Security threats: File scanning](#), page 162
- ◆ [Scanning exceptions](#), page 166
- ◆ [Reporting on scanning activity](#), page 168

Content Scanning performs advanced analysis on Web page content to discover security threats and malicious code in HTTP and HTTPS (when SSL Manager is enabled).

Use the **Settings > Scanning > Scanning Options** page to enable and configure content analysis.

1. Select **Off** to disable all advanced analysis of content.
2. Select **On** (default) to enable advanced analysis of content for all uncategorized sites and all sites with elevated risk profiles, including Web 2.0 sites, as identified by Websense Security Labs.
3. Select **Scan outbound Web content** to analyze outbound content for security threats like bot and spyware phone home traffic. This option is available only when content scanning is enabled.

When bot or phone home traffic is detected, it is categorized and blocked. This traffic is also logged, so you can run a report to obtain a list of the infected computers in your network.

4. Select **Scan content from all sites** to perform advanced analysis on all inbound content. This option is more resource intensive.
5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Use the **Settings > Scanning > Scanning Exceptions** page to specify untrusted or trusted sites that are always analyzed or never analyzed ([Scanning exceptions](#), page 166).

Security threats: File scanning

Related topics:

- ◆ [Scanning options](#), page 157
- ◆ [Content categorization](#), page 158
- ◆ [Security threats: Content scanning](#), page 161
- ◆ [Advanced options](#), page 164
- ◆ [Scanning exceptions](#), page 166
- ◆ [Reporting on scanning activity](#), page 168

File scanning inspects and analyzes files that users attempt to download or open remotely for viruses and other malicious content. This advanced file analysis returns a category to Websense filtering for policy enforcement.

There are 2 types of file analysis. They can be used together.

- ◆ **Advanced detection** applies techniques developed by Websense to discover known and emerging threats, including viruses, Trojan horses, worms, and others.
- ◆ **Anti-virus scanning** uses anti-virus definition files to identify virus-infected files.

You can configure the specific types of files to analyze by clicking **File Type Options**.



Note

If file analysis is configured to include multimedia files, sometimes when the streaming media is buffered and analyzed, the connection to the server times out. In such cases, one remedy is to create an analysis exception for that site. See [Scanning exceptions](#).

Use the **Settings > Scanning > Scanning Exceptions** page to specify untrusted or trusted sites that are always analyzed or never analyzed ([Scanning exceptions](#), page 166).

Use the **Settings > Scanning > Scanning Options** page to enable and configure file analysis.

Advanced Detection

1. Select **Off** to disable file analysis.
2. Select **On** (default) to enable file analysis on all files from uncategorized sites and all files from sites with elevated risk profiles, including Web 2.0 sites, as identified by Websense Security Labs.

3. Select **Scan all files from all sites** to perform advanced analysis on all inbound files. This option is more resource intensive.

Anti-virus Scanning

1. Select **Off** to disable anti-virus analysis.
2. Select **On** (default) to enable anti-virus analysis on all files from uncategorized sites and all files from sites with elevated risk profiles, including Web 2.0 sites, as identified by Websense Security Labs.
3. Select **Scan all files from all sites** to apply anti-virus analysis to all inbound files. This option is more resource intensive.

Rich Internet application scanning

Select **Scan rich Internet applications** to analyze Flash files for malicious content.

FTP file scanning

Select **Scan FTP files** to perform advanced analysis on files that are downloaded and uploaded with the FTP protocol. (FTP over HTTP file downloads and uploads are subject to the HTTP/HTTPS file analysis settings.) To be meaningful, this option requires that Content Gateway be configured to proxy FTP traffic. See Content Gateway Manager Help.



Note

The **Scan rich Internet applications** and **Scan FTP files** options are available only when advanced detection is enabled. When the advanced detection file-scanning feature is turned off, the rich Internet application scanning analysis feature is disabled and the check box is cleared. When advanced detection is re-enabled, the rich Internet application scanning analysis option returns to the enabled state and the check box is selected.

File Type Options

1. To specify the types of files to analyze, click **File Type Options**. As a best practice, analyze all suspicious files, as identified by Websense Security Labs, and all executable and unrecognized files.
2. To always analyze files having a specific extension, select **Files with the following extensions**, enter the extension in the entry field and click **Add**.
To remove an extension from the list, click on the extension to select it, and click **Delete**.

When you are done configuring file analysis options, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Several presentation reports provide details about attempts to download files containing security risks. These reports are listed in the Report Catalog only after analysis activity has detected sites whose activity has changed since it was assigned a Master Database category. See *Presentation reports*, page 107, for more information.

See *Managing traffic based on file type*, page 238, for information about blocking files based on type and URL category.

Advanced options

Related topics:

- ◆ [Scanning options](#), page 157
- ◆ [Content categorization](#), page 158
- ◆ [Security threats: Content scanning](#), page 161
- ◆ [Security threats: File scanning](#), page 162
- ◆ [Scanning exceptions](#), page 166
- ◆ [Reporting on scanning activity](#), page 168

Use these options to:

- ◆ Set the analysis time limit
- ◆ Set the file size analysis limit
- ◆ Enable stripping of specific types of code from HTML content

These settings apply to all incoming traffic.

Scanning Timeout

Each advanced analysis of content or of a file consumes a variable amount of time that cannot be determined before analysis begins. By default, to ensure a good user experience, analysis is limited to 1.5 seconds (1500 milliseconds). To adjust the timeout, select **Custom** and enter a value within the range 500 - 10000 (milliseconds).

File size limit

To set the maximum file size of files to be analyzed (default 10 MB), select **Custom** and enter the size in megabytes. Files larger than the specified size are not analyzed.

Content stripping

Threats to your system can be hiding in **active content** sent via Web pages. Active content is content that is embedded in the HTML page that performs actions, such as running an animation or a program.

The content stripping options make it possible to specify that content in particular scripting languages (ActiveX, JavaScript, or VB Script) be stripped from incoming

Web pages. If content stripping is enabled, all content in the specified scripting languages is removed from sites flagged as containing dynamic content or appearing on the Always Scan list (see *Scanning options*, page 157).

Content is removed only after the advanced analysis options have categorized the site and Websense filtering software has determined which policy applies.

**Warning**

Web pages that rely on active content that has been stripped do not function as expected. To permit full access to sites that require active content, disable content stripping or add the sites to the Never Scan list.

The user requesting a page with active content does not receive any notification that content has been removed.

Use the **Settings > Scanning > Scanning Options > Advanced Options** area to set content stripping options.

1. In the **Advanced Options > Content Stripping** area, select the types of scripting languages to be removed from incoming Web pages.
To disable content stripping for a selected language, clear the associated check box.
2. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

**Warning**

Content stripping can result in some content being garbled and unreadable. You can reduce the number of such occurrences by making a small change to the Content Gateway configuration.

- 1) Open Content Gateway Manager and go to the **Configure > Protocols > HTTP > Privacy** tab.
 - 2) In the **Remove Headers > Remove Others** field, add: Accept-Encoding
 - 3) Click **Apply** and restart Content Gateway.
-

Scanning exceptions

Related topics:

- ◆ [Scanning options, page 157](#)
- ◆ [Content categorization, page 158](#)
- ◆ [Security threats: Content scanning, page 161](#)
- ◆ [Security threats: File scanning, page 162](#)
- ◆ [Advanced options, page 164](#)

Scanning exceptions are lists of trusted or untrusted sites (hostnames and URLs) that are **never scanned** (never analyzed) or **always scanned** (always analyzed). The type of analysis to never or always perform is specified per hostname or URL, or group of hostnames and URLs.

You can also create a list of trusted client IP addresses whose content is never analyzed.

For an introduction to scanning analysis options, see [Scanning and SSL Bypass Options, page 155](#).

Use the **Always Scan** and **Never Scan** lists to refine the behavior of content categorization, tunneled protocol detection, security threats (content analysis and file analysis), and content stripping.

- ◆ When Content Categorization, Content Scanning, or File Scanning options are **On**, sites on the **Always Scan** list are always analyzed, and sites on the **Never Scan** list are never analyzed (see [Scanning options, page 157](#)).
- ◆ When the Tunneled Protocol Detection option is **On** or **Scan all incoming content from all sites** is selected, sites on the **Never Scan** list are never analyzed.

Use the Never Scan list with caution. If a site on the list is compromised, Websense Web Security Gateway does not analyze the site and cannot detect the security problem.

Hostname/URL Exceptions

To add sites to the Always Scan or Never Scan lists:

1. Click the **Add Hostname/URL** button.

You can specify a site in several ways, and you can specify more than one hostname or URL at a time.

- You can enter a simple hostname, for example, **thissite.com**. Be sure to enter both the hostname and the extension (**thissite.com** and **thissite.net** are distinct hosts).
- Sites with multiple extensions are supported. For example, **www.bbc.co.uk**

- You can use the wild card “*” to match the first subdomain. For example, ***.yahoo.com**.
 - You can enter a complete or partial hostname or URL. The leading HTTP or HTTPS is not required. An exact match is performed on the specified string. For example, **www.example.com/media/**
Or, for example, **www.youtube.com/watch?v=**
2. After entering a single or group of hostnames/URLs, select the scanning analysis options that apply to all of the sites you have entered. You can select one or more options.
To apply different options to different sites, enter the names separately.
A site can appear in only 1 of the 2 lists. You cannot, for example, specify that the same site should never be analyzed for tunneled protocols and always analyzed for content categorization.
 3. To delete a site from a list, select the site and click **Delete**.
 4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

To change the scanning analysis options associated with a site:

1. Select the site in the list and adjust the options.
2. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Client Exceptions

Use the Client Exceptions list to identify trusted users (client IP addresses) whose content is never analyzed.

To add an IP address to the list:

Click in the **Enter clients** box and enter an IP address or IP address range. For example, 10.201.67.245, or 10.201.67.245 - 10.201.67.250.

Click the right arrow (>) to move the address to the list.

To edit an entry:

Select the entry in the list and click **Edit**.

Make the desired changes and click **OK**.

To delete an entry:

Select the entry from the list and click **Delete**.

When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Data files used with scanning

Scanning analysis uses a set of data files to support the investigation. These files are updated regularly by Websense Security Labs and made available on the Websense download server. Websense Content Gateway checks for updated scanning analysis data files at regular intervals. The name and version of each file is displayed in Content Gateway Manager **Monitor > MyProxy > Summary**.

Data file updates occur independently of Websense Master Database updates (including real-time database updates and Real-Time Security Updates).

Every time the **./WCGAdmin start** command is run, a data file check and download is performed. If the download fails, a new download is attempted every 15 minutes until a successful download results.

The default interval for database update checks is 15 minutes. This is the recommended setting. Longer intervals increase the window of vulnerability to emerging, *zero day* exploits.

You can change the polling interval by editing the **PollInterval** value in the **/opt/bin/downloadservice.ini** file on the Content Gateway machine. After editing the **downloadservice.ini** file, you must stop and restart Content Gateway from the command line.

- ◆ To stop, enter: **/opt/WCG/WCGAdmin stop**
- ◆ To restart, enter: **/opt/WCG/WCGAdmin start**

Reporting on scanning activity

Related topics:

- ◆ [Scanning options, page 157](#)
- ◆ [Content categorization, page 158](#)
- ◆ [Security threats: File scanning, page 162](#)
- ◆ [Content stripping, page 164](#)

After you install Websense Content Gateway and enter a key that enables the advanced scanning analysis features, you can see and analyze the effects of these features on the History page and with presentation and investigative reports.

On the **History** page, 2 charts tally requests to Web 2.0 sites over the past 30 days:

- ◆ Top 5 Web 2.0 Categories by Requests
- ◆ Top 5 Web 2.0 Sites by Bandwidth

These charts are not displayed by default. To display these charts, on the History page, under the **History: Last 30 Days** header, click **Customize** and select the charts you want to display from the **Charts for History Page** list. See [Customize the History page, page 30](#).

On the **Presentation Reports** page, the **Scanning Activity** group contains reports that focus on Web 2.0 browsing and analysis activity, including recategorization that results from content categorization. There is also a report that tracks page blocks that result from link analysis.



Important

Enable full URL logging (see [Configuring how URLs are logged, page 365](#)) to ensure that reports of advanced scanning analysis are meaningful. Otherwise, reports can display only the domain (www.domain.com) of the site categorized, even though individual pages within the site may fall into different categories, or be recategorized for different reasons.

You can copy a security or scanning analysis report template to create a custom report. You can then edit the report filter to refine the information included when you generate that custom report.

Some security threat reports include a **Threat ID** column. You can click the individual threat ID to open a Websense Security Labs Web page that describes the type of threat identified.

Other presentation reports can contain information on advanced analysis activities, as well as standard filtering activities. For example, the Detail of Full URLs by Category report, found in the Internet Activity group of the Report Catalog, provides a detailed listing of each URL accessed within each category. To make a report that is specific to advanced scanning analysis, copy the Detail of Full URLs by Category report, and edit the report filter for the new custom report. On the Actions tab, select only permitted and blocked actions that relate to scanning analysis. On the Options tab, change the report catalog name and report title to identify this as an advanced analysis report. For example, you might change the name and title to Advanced Analysis: Detail of Full URLs by Category.

Investigative reports can also be used to gain insight into advanced analysis activities.

1. In the **Internet use by** drop-down list, select **Action**.
2. In the resulting report, click an action, such as **Category blocked real time**, to show a list of drill-down options.
3. Click the desired drill-down option, such as **Category** or **User**.
4. Click the **Hits value** or the bar on any row to see related detail.
5. Click **Modify Report**, at the top of the page, to add the **Full URL** column to the report.

See *Investigative reports*, page 127, for details on using all the investigative reports features.

How scanning is logged

There are important differences in the way that standard Web filtering activity and advanced analysis activity are logged.

For standard Web filtering, you have several options to reduce the size of the Log Database.

- ◆ Enable **visits** to log only one record for each Web site requested. See *Log Server Configuration utility*, page 357.
- ◆ Enable **consolidation** to combine into a single log record multiple requests with certain common elements. See *Log Server Configuration utility*, page 357.
- ◆ Disable **full URL logging** to log only the domain name (www.domain.com) for each request, and not the path to the specific page in the domain (/products/productA). See *Configuring how URLs are logged*, page 365.



Note

If your organization needs reports that include the full URL of each site visited, you should leave full URL logging enabled. Otherwise, reports will include only the domain (www.domain.com) of the site categorized, even though individual pages within the site may fall into different categories, or be recategorized for different reasons.

- ◆ Configure **selective category logging** to limit logging to only those categories that are required for your organization. See *Configuring how filtered requests are logged*, page 355.



Note

Enabling **visits**, **consolidation**, or **selective category logging**, will impact the accuracy of Internet Browse Time.

Scanning (advanced analysis) features, however, are bound only partially by these settings. When scanning analyzes a site, it creates 2 separate log records.

- ◆ **Web filter records** take advantage of any size reduction settings that have been implemented, and are available for all Web filter reports.
- ◆ **Scanning (advanced analysis) records** ignore most size reduction settings. Every separate hit is logged, requests to all categories are logged, and no records are consolidated. A scanning analysis record is generated regardless of whether the site is blocked or permitted as a result of analysis. Only the setting for full URL logging is honored for advanced analysis records.

If you have enabled any Log Database size reduction options, the numbers that appear in scanning analysis reports may **not** match those that appear in standard filtering reports, even when the reports are configured for the same users, time periods, and categories. For example, if you have chosen to log visits, and a user requests a site analyzed by scanning features, that user request appears as one visit in standard filtering reports, but may show as multiple hits in advanced analysis reports.

To see comparable data for standard and advanced analysis filtering, **disable** the Log Database size reduction settings. Because this may result in a very large and fast-growing database, make sure that the Log Database machine has adequate hard disk, processing, and memory capacity.

See [Reporting Administration, page 349](#), for more information on configuring size reduction settings. See [Presentation reports, page 107](#), and [Investigative reports, page 127](#), for information on generating reports.

SSL decryption bypass

When Content Gateway is configured to use SSL Manager to handle encrypted traffic, Category Bypass settings can be used to specify categories of Web sites for which decryption and inspection are bypassed. You can also maintain a list of destination hostnames or IP addresses for which SSL decryption is not performed.

A predefined **Privacy Category** group includes categories that may be subject to regulatory requirements. Default privacy categories include the following:

- ◆ Education
- ◆ Financial Data Services
- ◆ Government
- ◆ Health
- ◆ Online Brokerage and Trading
- ◆ Prescription Drugs

Traffic that involves Web sites in these categories may include personal identification information that should not be decrypted. In order to avoid liability for inspecting this type of information, you may want to specify some or all of these categories for decryption bypass. End users can determine that the Web site they are viewing is not decrypted by verifying that the certificate is the original for that site.

Use **Settings > SSL Decryption Bypass** to select the default privacy categories for SSL decryption bypass:

1. Click the **Select Privacy Categories** button. Check boxes for the Web site categories that constitute the default group are selected in the Category Bypass box.
2. Click the arrow to the right of the category tree to add the privacy categories to the **Categories selected for SSL decryption bypass** box.

You can create your own set of categories for SSL decryption bypass. On the **SSL Decryption Bypass** page, specify individual Web site categories for which decryption is not performed:

1. Select a category or subcategory for bypass by clicking its check box.
2. Click the arrow to right of the category tree to enter the selected category into the **Categories selected for SSL decryption bypass** box.

To clear your selections from the category tree, click the **Clear All** button.

To remove a category or subcategory from the list of categories designated for bypass, select the category and click the **Remove** button.

To identify an individual destination hostname or IP address for SSL decryption bypass:

1. Enter the destination hostname or IP address in the **Enter hostnames or IP addresses** box, one entry per line.
2. Click the arrow to right of the **Enter hostnames or IP addresses** box to add the hostname or IP address to the list of entries for which SSL decryption is not performed.

To modify a hostname or IP address, click the **Edit** button and modify the entry in the **Edit URL or Hostname** dialog box. Click **Update** to save your changes or **Cancel** to close the dialog box without saving your changes.

To remove an entry from this bypass list, select the hostname or IP address and click the **Delete** button.

When you are finished configuring SSL decryption bypass, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

8

Configure Hybrid Filtering

Websense Web Security Gateway Anywhere offers a flexible, comprehensive Web security solution that lets you combine on-premises and hybrid (in-the-cloud) filtering as needed to manage Internet activity for your organization. You decide which method to use for which users.

An organization might use the robust on-premises software to provide Web security for the main office or campus, while smaller regional offices or satellite locations send their Internet requests through the hybrid service. Hybrid filtering is also useful for users who are off-network, such as telecommuters, those who travel for business, and so on (see [Applying hybrid filtering to off-site users](#), page 206).

With Web Security Gateway Anywhere you define clients and create policies for on-premises and hybrid filtering in the same user interface—TRITON - Web Security—and configuration, reporting, and management are centralized.

To use hybrid filtering:

1. [Activate your hybrid filtering account](#), page 174
2. [Define the locations filtered by hybrid filtering](#), page 175
3. [Specify sites not filtered by hybrid filtering](#), page 180 (if any)
4. [Configure user access to hybrid filtering](#), page 183
5. [Identification of hybrid filtering users](#), page 276
6. [Send user and group data to the hybrid service](#), page 189

In order to ensure that the hybrid service has current policy, user, and group information, and that the on-premises reporting software has data from users filtered by the hybrid service, see [Schedule communication with hybrid filtering](#), page 196.

Activate your hybrid filtering account

Related topics:

- ◆ [Define the locations filtered by hybrid filtering, page 175](#)
- ◆ [Specify sites not filtered by hybrid filtering, page 180](#)
- ◆ [Configure user access to hybrid filtering, page 183](#)
- ◆ [Send user and group data to the hybrid service, page 189](#)
- ◆ [Schedule communication with hybrid filtering, page 196](#)

Before you can configure the hybrid service to start filtering Internet requests for your organization, you must activate your hybrid account by submitting a contact email address. This creates a connection between the on-premises and hybrid portions of Websense Web Security Gateway Anywhere.

Use the Hybrid Filtering section of the **Settings > General > Account** page to provide the contact email address and country for your Websense filtering administrators (see [Configuring your account information, page 32](#)).

The email address is typically an alias monitored by the group responsible for managing your Websense software. It is very important email sent to this account be received and acted upon promptly.

- ◆ Websense Technical Support uses this address to send out notifications about urgent issues affecting hybrid filtering.
- ◆ If there is a configuration problem with your account, failure to respond to an email message from Technical Support in a timely fashion could lead to service interruptions.
- ◆ Should certain rare problems occur, the email address is used to send the information needed to allow Sync Service to resume contact with the hybrid service.
- ◆ This email address is **not** used to send marketing, sales, or other, general information.

The country you enter provides the system with time zone information.

Once you have activated hybrid filtering for your account, you can specify which locations (identified by IP address, IP address range, or subnet) are filtered by the hybrid service, how information is exchanged between the on-premises and hybrid portions of your Web security software, how users filtered by the hybrid service are authenticated, and more.

Define the locations filtered by hybrid filtering

Related topics:

- ◆ [Specify sites not filtered by hybrid filtering, page 180](#)
- ◆ [Configure user access to hybrid filtering, page 183](#)
- ◆ [Schedule communication with hybrid filtering, page 196](#)

Select **Settings > Hybrid Configuration > Filtered Locations** to review, add, or edit information about the locations that contain users who can be filtered by the hybrid service.

A **filtered location** is the IP address, IP address range, or subnet from which browsers connecting to Web filtering appear to be originating. In Web Security Gateway Anywhere deployments, hybrid filtering can be applied to off-site users, regardless of how those users are filtered when they are in-network.

- ◆ For users filtered by hybrid filtering both in and outside the network, enter their in-network location details and specify that the location is filtered using the hybrid service. When off-site users make an Internet request, they are prompted to log on to hybrid filtering so that the appropriate user or group-based policy can be applied.

Because the hybrid service is hosted outside your network, any locations filtered by the hybrid service must be external addresses, visible from the Internet.

Locations filtered by the hybrid service:

- Are public-facing IP addresses for offices filtered by Web Security Gateway Anywhere
- Are often the external address of your Network Address Translation (NAT) firewall
- Could include branch offices, remote sites, or satellite campuses

These locations are NOT:

- IP addresses of individual client machines
- The IP address of any Content Gateway machine used by the on-premises components of Websense Web Security Gateway Anywhere
- ◆ For users filtered by on-premises components (Filtering Service) when they are inside the network, you can configure the browser PAC file to determine whether the user is in-network or off-site before sending an Internet request for filtering.

If you are using the PAC file generated by the hybrid service, this configuration occurs automatically based on the settings that you provide on the Filtered Locations page. Enter the in-network location details for these users and specify whether their on-premise filtering is through a firewall-integrated or transparent proxy (for example, Content Gateway in transparent mode), or an explicit proxy. If Internet requests from in-network machines at a specified location pass through

an explicit proxy, you provide the proxy location (host name or IP address) and port to ensure requests are routed properly for users at that location.

Each location that you define appears in a table that combines a name and description with technical configuration details, including the selected proxy mode, the type of location (single IP address, IP address range, or subnet), and the actual external IP address or addresses from which requests originate.

To edit an existing entry, click the location **Name**, and then see [Editing filtered locations](#), page 177.

To define a new location, click **Add**, and then see [Adding filtered locations](#), page 176.

To remove a location, mark the check box next to the location name, and then click **Delete**.

To add and edit on-premises explicit proxies for use with filtered locations, click **Manage Explicit Proxies**, then see [Managing explicit proxies](#), page 179.

If you have added or edited a location entry, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Adding filtered locations

Use the **Filtered Locations > Add Filtered Location** page to define a location either filtered by the hybrid service (like a branch office, remote site, or satellite campus), or that contains users filtered by the hybrid service when off site.

1. Enter a unique location **Name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Names can include spaces, dashes, and apostrophes.

2. Enter a short **Description** of the location (up to 255 characters). This appears next to the location name on the Filtered Locations page, and should clearly identify the location to any administrator.

The character restrictions that apply to names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. If you are adding a location filtered by the hybrid service, select the **Time zone** of this location. Time zone information is used in applying policies, to ensure that the correct filters are applied at the appropriate time.

Each location filtered by the hybrid service can have a different time zone setting. Locations filtered by transparent or explicit proxies use the time zone of the machine on which Filtering Service is running as the time zone for policy enforcement.

4. In the **Type** field, indicate how you want to define this location: as an **IP address**, an IP address **Range**, or a **Subnet**.

If you are providing a subnet, specify whether you are identifying it by **By bit range (CIDR)** or **By subnet mask**, and then select a bit range or mask.

5. Enter the external IP address, range, or subnet of the firewall or firewalls through which filtered clients at this location access the Internet.
 - For locations filtered by the hybrid service, these are external IP addresses, visible from outside your network, and not internal (LAN) addresses.



Important

Do not enter private IP addresses (in the ranges 10.0.0.0 - 10.255.255.255, 172.16.0.0 - 172.31.255.255, and 192.168.0.0 - 192.168.255.255) to identify locations filtered by the hybrid service. Because these addresses are not visible from outside your network, and are used within multiple local area networks, the hybrid service does not accept private IP addresses as valid entries.

If the proxy mode for this location is Transparent or Explicit, you can enter private IP addresses.

- Do not include the IP address of any Content Gateway machine used by the on-premises components of Websense Web Security Gateway Anywhere.
 - External IP addresses must be unique to your organization, not shared with any other entity, so that the hybrid service is able to associate requests originating from these locations with the policies belonging to your organization.
6. Select the **Proxy mode** for this location: using the **Hybrid** service, a **Transparent** proxy, or an **Explicit** on-premises proxy.

If you select Explicit, there must be at least one proxy defined in the Explicit proxy configuration table. To add a new explicit proxy to the table, click **Add**, select a proxy location and preference order from the popup window, then click **OK**. See [Managing explicit proxies](#), page 179 for more information on the available explicit proxies.

The filtered location uses the first proxy on the list. If that proxy is not available, Web filtering requests from the filtered location are redirected to the next proxy on the list. To change the order, select any proxy on the list and then click **Move Up** or **Move Down** to change its position in the list.

To remove a proxy from the table, mark the check box next to the proxy name, and then click **Delete**. The deleted proxy is no longer available for this filtered location, but can still be selected for other filtered locations.

7. Click **OK** to return to the Filtered Locations page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Editing filtered locations

Use the **Filtered Locations > Edit Filtered Locations** page to change the way a location filtered by the hybrid service is defined.

1. If you make changes to the location **Name**, make sure that the name remains unique. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Names can include spaces, dashes, and apostrophes.

2. Enter or update the **Description** of the location.
The character restrictions that apply to names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).
3. Verify the **Time zone** of this location to ensure accurate policy enforcement.
4. Verify the method used to define this location: by **IP address**, **IP address Range**, or **Subnet**.

If you are entering or editing a subnet, specify whether it is identified by **By bit range (CIDR)** or **By subnet mask**, and then select a bit range or mask.

5. Update the external IP address, range, or subnet of the firewall or firewalls through which filtered clients at this location access the Internet.
 - For locations filtered by the hybrid service, these are external IP addresses, visible from outside your network, and not internal (LAN) addresses.



Important

Do not enter private IP addresses (in the ranges 10.0.0.0 - 10.255.255.255, 172.16.0.0 - 172.31.255.255, and 192.168.0.0 - 192.168.255.255) to identify locations filtered by the hybrid service. Because these addresses are not visible from outside your network, and are used within multiple local area networks, the hybrid service does not accept private IP addresses as valid entries.

If the proxy mode for this location is Transparent or Explicit, you can enter private IP addresses.

- These never include the IP address of the Content Gateway machine. The Content Gateway is used only by the on-premises portion of your Websense software.
 - External IP addresses must be unique to your organization, not shared with any other entity, so that the hybrid service is able to associate requests originating from these locations with the policies belonging to your organization.
If a filtered location is already associated with one organization, and a second organization adds the same IP address as a filtered location, hybrid filtering may be temporarily suspended for the second organization, until the conflict is resolved.
6. Verify the **Proxy mode** for this location: using the **Hybrid** service, a **Transparent** proxy, or an **Explicit** on-premises proxy.

If you select Explicit, the filtered location uses the first proxy on the list defined in the Explicit Proxy Configuration table. If that proxy is not available, Web filtering

requests from the filtered location are redirected to the next proxy on the list. To change the preference order, select any proxy on the list and then click **Move Up** or **Move Down** to change its position in the list.

To view the details of an existing table entry, click the proxy **Name**.

To add a new explicit proxy to the table, click **Add**, select a proxy location and preference order from the popup window, then click **OK**. See [Managing explicit proxies, page 179](#) for more information on the available explicit proxies.

To remove a proxy from the table, mark the check box next to the proxy name, and then click **Delete**. The deleted proxy is no longer available for this filtered location, but can still be selected for other filtered locations.

7. Click **OK** to return to the Filtered Locations page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Managing explicit proxies

Related topics:

- ◆ [Adding an explicit proxy, page 179](#)
- ◆ [Editing an explicit proxy, page 180](#)

Use the **Filtered Locations > Manage Explicit Proxies** page to review, add, and edit the on-premises explicit proxies available for use with filtered locations.

Each explicit proxy that you define appears in a table that displays a proxy name, its IP address or host name, the port number or numbers used for HTTP, SSL, or FTP access, and the filtered locations (if any) that currently reference the proxy.

To edit an existing entry, click the proxy **Name**, and then see [Editing an explicit proxy, page 180](#).

To define a new explicit proxy, click **Add**, and then see [Adding an explicit proxy, page 179](#).

To remove a proxy, mark the check box next to the proxy name, and then click **Delete**.



Note

You cannot delete a proxy that is being used by one or more filtered locations. If you wish to delete a proxy, first edit each filtered location to remove the proxy from the Explicit Proxy Configuration table.

Adding an explicit proxy

When managing explicit proxies, use the **Add Explicit Proxy** page to define an on-premises explicit proxy to be used for your filtered locations.

1. Enter a unique proxy **Name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
Names can include spaces, dashes, and apostrophes.
2. Enter the **IP address or name** of the explicit proxy. This must be in one of the following forms:
 - An IP address (for example 123.45.67.89)
 - A host name (for example my.example.com)The IP address or name can include a port number, for example 123.45.67.89:443.
3. Enter at least one port number for the proxy. This can be an **HTTP port**, an **SSL port**, or an **FTP port**.
4. Click **OK** to return to the Manage Explicit Proxies page.

Editing an explicit proxy

When managing explicit proxies, use the **Edit Explicit Proxy** page to change the details of a configured on-premises explicit proxy.

1. If you make changes to the proxy **Name**, make sure the name remains unique. The name must be between 1 and 50 characters long, and cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
Names can include spaces, dashes, and apostrophes.
2. Verify or update the **IP address or name** of the explicit proxy. This must be in one of the following forms:
 - An IP address (for example 123.45.67.89)
 - A host name (for example my.example.com)The IP address or name can include a port number, for example 123.45.67.89:443.
3. If you change the proxy port, enter at least one port number. This can be an **HTTP port**, an **SSL port**, or an **FTP port**.
4. Click **OK** to return to the Manage Explicit Proxies page.

Specify sites not filtered by hybrid filtering

Related topics:

- ◆ [Define the locations filtered by hybrid filtering, page 175](#)
- ◆ [Send user and group data to the hybrid service, page 189](#)
- ◆ [Schedule communication with hybrid filtering, page 196](#)

Select **Settings > Hybrid Configuration > Unfiltered Destinations** to review, add, or edit information about target sites to which you want to grant clients unfiltered access. Clients can access these sites directly, without sending the request to either the hybrid service or an on-premises explicit proxy in a filtered location, if used. Typical unfiltered destinations include organizational webmail sites, internal IP addresses, and Microsoft update sites.



Tip

As a best practice, add your organization's webmail address as an unfiltered destination. This ensures that:

- ◆ You can access messages from Technical Support in situations that cause your proxy or the hybrid service to block all requests.
- ◆ Off-site users who have forgotten (or not created) their hybrid filtering password can retrieve it via email.

Destinations listed here are added to the Proxy Auto-Configuration (PAC) file that defines how filtered users' browsers connect to the hybrid service (see [Configure user access to hybrid filtering, page 183](#)). By default, the PAC file excludes all non-routable and multicast IP address ranges from filtering. Therefore, if you are using private IP address ranges defined in RFC 1918 or RFC 3330, you need not enter them here.

Each unfiltered destination that you define appears in a table that combines a name and description with technical configuration details, including how the destination is defined (as an IP address, domain, or subnet), and the actual IP address, domain, or subnet that users can access directly.

To edit an existing entry, click the location **Name**, and then see [Editing unfiltered destinations, page 182](#).

To define a new location, click **Add**, and then see [Adding unfiltered destinations, page 181](#).

To remove an unfiltered destination, mark the check box next to the destination name, and then click **Delete**.

If you have added or edited an unfiltered destination entry, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Adding unfiltered destinations

Use the **Unfiltered Destinations > Add Unfiltered Destination** page to define a site or set of sites that users can access directly, without sending a request to the hybrid service or an on-premises explicit proxy.

1. Enter a unique destination **Name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Names can include spaces, dashes, and apostrophes.

2. Enter a short **Description** of the destination. This appears next to the unfiltered destination name on the Unfiltered Destinations page, and should clearly identify the unfiltered target site or sites to any administrator.

The character restrictions that apply to names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. In the **Type** field, indicate how you want to define this destination: as an **IP address**, **Domain**, or **Subnet**.
If you are providing a subnet, specify whether you are identifying it by **By bit range (CIDR)** or **By subnet mask**, and then select a bit range or mask.
4. Enter the IP address, domain, or subnet that you want users to be able to access without sending the request to the hybrid service or an on-premises explicit proxy.
5. Select the **Proxy** type that this unfiltered destination applies to.

- Select **Hybrid** to enable all hybrid filtering users to access the destination directly without sending a request to the hybrid service.
- Select **Explicit** to enable all users in filtered locations using an on-premises explicit proxy to access the destination directly.
- Select **Hybrid and Explicit** to enable all users filtered by the hybrid service and an on-premises explicit proxy from a filtered location to access the destination directly.

6. Click **OK** to return to the Unfiltered Destinations page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Editing unfiltered destinations

Use the **Unfiltered Destinations > Edit Unfiltered Destination** page to change the definition of an existing site or set of sites that users can access directly, without sending a request to the hybrid service.

1. If you make changes to the destination **Name**, make sure that the name remains unique. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Names can include spaces, dashes, and apostrophes.

2. Enter or update the destination **Description**.

The character restrictions that apply to names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. Verify the method used to define this destination: **IP address**, **Domain**, or **Subnet**.

If you are entering or editing a subnet, specify whether you are identifying it by **By bit range (CIDR)** or **By subnet mask**, and then select a bit range or mask.

4. Update the IP address, domain, or subnet that you want users to be able to access without sending the request to the hybrid service.
5. Update the **Proxy** type that this unfiltered destination applies to.

- Select **Hybrid** to enable all hybrid filtering users to access the destination directly without sending a request to the hybrid service.
 - Select **Explicit** to enable all users in filtered locations using an on-premises explicit proxy to access the destination directly.
 - Select **Hybrid and Explicit** to enable all users filtered by the hybrid service and an on-premises explicit proxy from a filtered location to access the destination directly.
6. Click **OK** to return to the Unfiltered Destinations page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Configure user access to hybrid filtering

To use hybrid filtering, you must configure how users connect to and are filtered by the hybrid service. To do so, select **Settings > Hybrid Configuration > User Access**.

The **Proxy Auto-Configuration (PAC) File** section shows the URL from which browsers on filtered users' machines retrieve the PAC file (see [What is a PAC file?](#), page 187). The PAC file defines which requests the browsers send to the hybrid service, and which are sent directly to the target site (see [Specify sites not filtered by hybrid filtering](#), page 180). The PAC file also contains information about filtered locations, and the proxy configuration for any locations that filter their users through an explicit or transparent proxy when on-premises, so that traffic can be routed properly at all locations.



Note

The exact mechanism for configuring a user's browser to use the PAC file depends on the browser and your network environment. For example, if you are using Microsoft Active Directory and Internet Explorer or Mozilla Firefox, you might want to automate the process by using group policies.

Use the **Availability** section to specify whether all Internet requests should be permitted or blocked when the hybrid service is unable to access policy information for your organization.

Under **Time Zone**, use the drop-down list to select a default time zone to use when applying policies in the following situations:

- ◆ For users connecting to the hybrid service from an IP address that is not part of an existing filtered location (see [Define the locations filtered by hybrid filtering](#), page 175)
The default time zone is used, for example, by off-site users, or for other users that self-register with the hybrid service.
- ◆ Whenever time zone information is not available for a filtered location

Use the **Custom End User Block Page** section to define a customized logo and text for block pages displayed by the hybrid service (see [Customizing hybrid block pages](#), page 185).

Use the **HTTPS Notification Pages** section to enable users making HTTPS requests to view the appropriate Websense notification pages (see [Enabling HTTPS notification pages](#), page 186).

If hybrid filtering uses directory data collected by Websense Directory Agent to identify users, you can configure hybrid passwords for user accounts on the **Hybrid Configuration > Shared User Data** page in TRITON - Web Security (see [Send user and group data to the hybrid service](#), page 189). If your organization does not use directory data collected by Directory Agent to identify users connecting to the hybrid service from outside filtered locations, you can let users **self-register** for the service. This allows users with email accounts associated with domains that you specify under **Registered Domains** to identify themselves to the hybrid service.

Users requesting Internet access from an unrecognized IP address are prompted to self-register. The domain portion of the user's email address is used to associate the user with your organization so that the proper Default policy is applied.

Users who cannot be associated with an organization are filtered by the hybrid service Default policy.

- ◆ Click **Add** to add a domain (see [Adding domains](#), page 184).
- ◆ Click a domain entry to edit the domain or its attributes (see [Editing domains](#), page 185).

You can also apply hybrid filtering to off-site users connecting from unknown IP addresses, regardless of how those users are filtered when they are in-network or connecting from a filtered location. Under Off-site Users, mark **Enable hybrid filtering of off-site users**.

If you clear this check box, any user connecting from an unknown IP address will not be filtered.

See [Applying hybrid filtering to off-site users](#), page 206 for more information.

Adding domains

Use the **User Access > Add Domain** page to identify the domains and subdomains (if any) belonging to your organization. This makes it possible for users with email addresses in the specified domains to self-register (authenticate themselves) for hybrid filtering. This is typically enabled only in organizations that do not use Directory Agent to send user information to hybrid filtering.

The hybrid service is unable to provide user name information about self-registered users to the on-premises components for use in reporting. Only the IP address from which the request originated is logged.

1. Enter a **Domain** name (in the format **sampledomain.org**) belonging to your organization.

2. Enter a clear **Description** of the domain as a point of reference to simplify hybrid filtering administration.
3. If you want users with email addresses in both the domain and its subdomains (like **university.edu** and **humanities.university.edu**) to be able to self-register, mark **Include subdomains**.
4. Click **OK** to return to the User Access page.
5. Click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Editing domains

Use the **User Access > Edit Domain** page to make changes to the domain entries that allow users to self-register for hybrid filtering.

1. Verify the domain **Name** and make changes, if necessary.
2. Update the **Description** as needed.
3. To change whether or not email addresses in subdomains are considered valid, mark or clear **Include subdomains**.
4. Click **OK** to return to the User Access page.
5. Click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Customizing hybrid block pages

When the hybrid service denies access to a resource, it serves a default block page. You can either use the default page, or modify the page text to suit your needs. For example, you could:

- ◆ Add information about your organization's Internet use policies.
- ◆ Provide a method for contacting Human Resources or a Websense administrator about Internet use policies.
- ◆ Add your organization's logo.

Customizing the logo

If you want to customize the logo that appears on a hybrid filtering block page, create a directory named **logo** in the Websense **ssdata** directory (by default, `\Program Files\Websense\Web Security\bin\ssdata\` on Windows, or `/opt/websense/bin/ssdata/` on Linux). Then place your logo file in that directory.

The logo must be a JPEG, GIF, or PNG file. If a file with one of these extensions exists in the **logo** directory, Sync Service detects it and sends the data to the hybrid service. The file must be greater than 0KB and smaller than 50KB for Sync Service to send it. Sync Service also detects when there is a newer version of the file and updates the version on the hybrid service. If there are multiple valid files in this directory, Sync Service uses the most recent file.

The **Hybrid Service** page displays the date and time that Sync Service sent a customized block page logo to the hybrid service (see *Monitor communication with the hybrid service*, page 197).

To stop using a customized logo file, delete the file from the **logo** directory.



Note

Clearing **Use a custom block page title and message** on the **Hybrid Configuration > User Access** page does not automatically remove the customized logo from your block pages. The logo file must be deleted from the logo directory for Sync Service to stop pushing the file to the hybrid service.

Customizing the text

1. On the **Hybrid Configuration > User Access** page, mark **Use a custom block page title and message**.
2. Enter the page **Title** and **Message**. This must be in plain text, with no HTML tags.
3. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Enabling HTTPS notification pages

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

If you install the Websense SSL certificate for hybrid filtering, the hybrid proxy can establish SSL channels with newer browsers (Internet Explorer 8 or later, and Firefox 3.5 or later) in order to serve notification pages to the user – for example, a block page if the SSL site is in a category that requires a notification, or the appropriate page if authentication is required.

To preserve performance, only HTTPS traffic is diverted in this manner; HTTP traffic goes through the proxy to the requested site.

To ensure hybrid filtering users can see the notification pages when browsing with HTTPS, you need a root certificate on each client machine that can act as a Certificate Authority for SSL requests to the hybrid proxy.

**Note**

End users using Websense Authentication Service require this root certificate to ensure seamless authentication to HTTPS sites. If the certificate is not installed for Authentication Service users, they must authenticate using NTLM identification or manual authentication, depending on the settings on the Hybrid User Identification page. See [Deploying Websense Authentication Service, page 280](#).

To install the hybrid root certificate on all clients using hybrid filtering:

1. On the **Hybrid Configuration > User Access** page, click **View Hybrid SSL Certificate**.
2. Save the certificate file to a location of your choice.
3. Deploy the SSL certificate to your hybrid filtering users with your preferred administration or deployment method, for example Microsoft Group Policy Object (GPO) or a third-party deployment tool.

Once you have distributed the certificate, mark **Use the hybrid SSL certificate to display a notification page for HTTPS requests when required**, then click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

What is a PAC file?

A Proxy Auto-Configuration file is a JavaScript function definition that a browser calls to determine how to handle requests. The PAC file used to enable hybrid filtering contains a number of global settings and allows you to configure sites (for example, intranet sites or organizational Web mail) that users can access directly, without sending the request to the hybrid service (see [Specify sites not filtered by hybrid filtering, page 180](#)).

If you want to use hybrid filtering on client machines, you must configure browser settings on each of the clients to point to the URL hosting the PAC file. This URL is displayed on the **Hybrid Configuration > User Access** page in TRITON - Web Security (see [Configure user access to hybrid filtering, page 183](#)).

The exact mechanism for configuring a browser to use the PAC file depends on the browser and network environment. For example, if you are using Microsoft Active Directory and Internet Explorer or Mozilla Firefox, you have the option to automate the process via group policies. Users can also be instructed to set up their browsers manually.

- ◆ For Microsoft Internet Explorer 8 or 9, go to **Tools > Internet Options** and click the **Connections** tab. Click **LAN Settings**, and then mark **Use automatic configuration script**. Enter the PAC file URL in the **Address** field.

- ◆ For Mozilla Firefox 3.x, go to **Tools > Options**, click the **Advanced** icon, and then select the **Network** tab. Under Connection, click **Settings**, and then select **Automatic proxy configuration URL**. Enter the PAC file URL in the blank field.

The default PAC file is supplied by Websense, and comprises default settings from the hybrid service and any changes you make on the Hybrid Configuration pages. If you want to customize the PAC file, create a directory named **pac** in the Websense **ssdata** directory (by default, \Program Files\Websense\Web Security\bin\ssdata\pac on Windows, or /opt/websense/bin/ssdata/pac on Linux). Then you have the following options:

- ◆ To use your own PAC file, create a file named *websense.pac* and place it in the Websense **pac** directory.
- ◆ To add a customized fragment to the default PAC file, place the JavaScript fragment in a file named *customfinal.pac*, and put it in the Websense **pac** directory. This fragment is appended to the default PAC file, replacing the token `_CUSTOMFINALPAC_`.

**Note**

The customized *websense.pac* file must contain the following function:

```
function FindProxyForURL(url, host) {}
```

If this function is not in the file, it will be rejected by the hybrid service.

If either of these files exists in the **pac** directory, Sync Service detects it and sends the data to the hybrid service. The file must be greater than 0KB and smaller than 50KB for Sync Service to send it. Sync Service also detects when there is a newer version of the PAC file or fragment and updates the version on the hybrid service.

The recommended state for custom PAC files is to set up a custom file or a custom fragment, not both. If both files exist in the **pac** directory, we recommend you decide whether a full customized PAC file or a customized fragment suits your needs better, and delete the other file from the directory.

To stop using a customized PAC file or fragment, delete the file or fragment from the **pac** directory.

The **Hybrid Service** page displays the type of PAC file you are using, and lists the date and time that Sync Service last sent a customized file or fragment to the hybrid service (see *Monitor communication with the hybrid service*, page 197).

If you are unfamiliar with PAC files, it is useful to search the Internet for basic information. Wikipedia has a good introductory article, and a good Web site for more information and several example PAC files is <http://www.findproxyforurl.com/>.

Send user and group data to the hybrid service

If your organization uses a supported, LDAP-based directory service—Windows Active Directory (Native Mode), Oracle (Sun Java) Directory Server, or Novell eDirectory—you can collect user and group data and send it to the hybrid service. This is accomplished using 2 Websense components:

- ◆ **Websense Directory Agent** - collects user and group information from Directory Server and collates it for hybrid filtering.
- ◆ **Websense Sync Service** - Transports policy, reporting, custom PAC file information, and user/group data between the on-premises and hybrid systems.

When hybrid filtering is configured properly, the information from Directory Agent can be used to apply user- and group-based filtering.

If your organization uses Windows Active Directory in mixed mode, user and group data cannot be collected and sent to the hybrid service.

If hybrid filtering uses directory data collected by Directory Agent to identify users, you have 2 options:

- ◆ Configure the hybrid service to automatically create a hybrid logon password for all user accounts sent by Directory Agent. Passwords are sent to each user's email address in staggered intervals to avoid a sudden influx of email messages.
- ◆ Have users request their own password the first time they connect to the hybrid service from outside a filtered location. In order for the process to succeed, users must provide an email address that matches an account sent by Directory Agent. The password is then sent to that email address.

**Note**

For this reason, be sure that your organization's webmail address has been added as an unfiltered destination. See [Specify sites not filtered by hybrid filtering](#), page 180.

Configure Directory Agent settings for hybrid filtering

Select **Settings > Hybrid Configuration > Shared User Data** to review and edit your current Directory Agent configuration, and to configure Directory Agent to communicate with Sync Service.

The table near the top of the page lists the Active Directory global catalogs identified on the **Settings > General > Directory Services** page. Add or remove global catalog servers, or change the directory service used by Websense software, on that page.



Note

If you remove an Active Directory server from the Directory Services page, you should also do a manual reset to ensure the server is fully removed from Directory Agent settings.

To perform a manual reset on appliance-based deployments, contact Websense Technical Support.

To perform a manual reset on software-based deployments, delete all files in the Websense **bin/snapshots** directory (see *Where is the Websense "bin" directory?*, page 457). Then go to **Settings > Hybrid Configuration > Scheduling**, and click **Send** under Send Update Now.

To refine the way that Directory Agent searches the directory and packages results for the hybrid service, click an IP address or host name in the table. See *Configure how data is gathered for hybrid filtering*, page 191.

To view root contexts that have been added for hybrid users in the global catalog servers, click **View** under Context in the table. See *Adding and editing root contexts*, page 193.

To have the hybrid service generate passwords for all user accounts that it sees, scroll down to the **Generate User Passwords** section and mark **Automatically generate and email passwords**.

In order for Directory Agent data to be sent to the hybrid service:

1. Scroll to the **Synchronize User Data** section.
2. Verify the **Name or IP address** of the Sync Service machine and the **Port** used for Sync Service communication (by default, 55832).

In most configurations, these fields are populated automatically, but can be updated manually, if needed.

3. Click **Test Connection** to verify that Directory Agent can send data to Sync Service. The test may take a minute or more.
 - If the connection is made, a success message is displayed.
 - If the connection cannot be made, verify the IP address or host name of the Sync Service machine and the communication port. Also verify that the Sync Service machine is on, that Sync Service is running, and that your network firewall permits connections on the Sync Service port.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configure how data is gathered for hybrid filtering

Use the **Shared User Data > Active Directory (Native Mode)** page to refine the way that Directory Agent searches the selected directory server and packages user and group information for the hybrid service.

1. Under Root Context for Hybrid Filtering Users, click **Add** to provide a **Root Context** to use when gathering user and group data from the directory. Narrow the context to increase speed and efficiency. See [Adding and editing root contexts](#), page 193.

It is best to provide contexts that include only users filtered by the hybrid service.

If you are using Active Directory and have multiple Directory Agent instances, make sure that each has a unique, non-overlapping root context. Especially watch out for this if:

- Multiple Directory Agent instances are configured to connect to domain controllers that all manage the same Active Directory server.
- One Directory Agent instance is configured to communicate with an Active Directory parent domain and another instance is configured to communicate with an Active Directory child domain (a separate global catalog server).

You can further refine the data that is sent to the hybrid service by defining patterns, or search filters, used to remove duplicate or otherwise unwanted entries from the directory search results. See [Optimizing search results](#), page 195, for more information.

Oracle (Sun Java) Directory Server and hybrid filtering

If your organization uses Oracle (Sun Java) Directory Server, select **Settings > Hybrid Configuration > Shared User Data** to refine the way that Directory Agent searches the directory and packages user and group information for the hybrid service.



Important

To use any version of Sun Java System Directory or Oracle Directory Server to send user and group information to the hybrid service, a Directory Agent configuration change is required.

Open the **das.ini** file (located in the Websense **bin** directory on the Directory Agent machine) and locate the following section:

```
# Enable next two parameters if your DS
is Sun Java
# GroupMembershipAttribute=uniqueMember
# MemberOfAttribute=memberOf
```

Enable the `GroupMembershipAttribute` and `MemberOfAttribute` parameters by removing the `#` symbol from the beginning of those lines, then save the file and restart Directory Agent.

1. Under Root Context for Hybrid Filtering Users, click **Add** to provide a **Root Context** to use when gathering user and group data from the directory. Narrow the context to increase speed and efficiency. See [Adding and editing root contexts](#), page 193.

It is best to provide a context that includes only users filtered by the hybrid service.

2. Under Synchronize User Data, verify the **Name or IP address** of the Sync Service machine and the **Port** used for Sync Service communication (by default, 55832).
These fields are populated automatically, but can be updated manually, if needed.
3. Click **Test Connection** to verify that Directory Agent can send data to Sync Service. The test may take a minute or more.
 - If the connection is made, a success message is displayed.
 - If the connection cannot be made, verify the IP address or host name of the Sync Service machine and the communication port. Also verify that the Sync Service machine is on, that Sync Service is running, and that your network firewall permits connections on the Sync Service port.

You can further refine the data that is sent to the hybrid service by defining patterns, or search filters, used to remove duplicate or otherwise unwanted entries from the directory search results. See [Optimizing search results](#), page 195, for more information.

Novell eDirectory and hybrid filtering

If your organization uses Novell eDirectory, select **Settings > Hybrid Configuration > Shared User Data** to refine the way that Directory Agent searches the directory and packages user and group information for the hybrid service.

1. Under Root Context for Hybrid Filtering Users, click **Add** to provide a **Root Context** to use when gathering user and group data from the directory. Narrow the context to increase speed and efficiency. See [Adding and editing root contexts](#), page 193.

It is best to provide a context that includes only users filtered by the hybrid service.

2. Under Synchronize User Data, verify the **Name or IP address** of the Sync Service machine and the **Port** used for Sync Service communication (by default, 55832).

These fields are populated automatically, but can be updated manually, if needed.

3. Click **Test Connection** to verify that Directory Agent can send data to Sync Service. The test may take a minute or more.
 - If the connection is made, a success message is displayed.
 - If the connection cannot be made, verify the IP address or host name of the Sync Service machine and the communication port. Also verify that the Sync Service machine is on, that Sync Service is running, and that your network firewall permits connections on the Sync Service port.

You can further refine the data that is sent to the hybrid service by defining patterns, or search filters, used to remove duplicate or otherwise unwanted entries from the directory search results. See [Optimizing search results](#), page 195, for more information.

Adding and editing root contexts

Use the **Settings > Hybrid Configuration > Shared User Data > Add Context** page to refine the way that Directory Agent searches your user directory and packages user and group information for the hybrid service.

You can add multiple root contexts for your user directory. It is best to provide contexts that include only users filtered by the hybrid service: for example, you might have hybrid users in multiple OUs. Alternatively, if you want to synchronize all users in a number of specific groups, then you can enter a context for each group where each context is the fully qualified group name.

By default, Directory Agent uses the user and group filters defined under [Advanced directory settings](#) on the **Settings > General > Directory Services** page. If required, you can customize these filters for each hybrid filtering context, for example to include only users that are members of a “hybrid” group.

You can also choose to exclude certain contexts from the Directory Agent search. You might want to do this if you have a particular context that is not required or could

cause problems with the hybrid service, for example an administrator group that has multiple email addresses in a record.

1. Enter a root **Context** to use when gathering user and group data from the directory. Narrow the context to increase speed and efficiency.
2. To verify that the context that you have entered uses valid syntax and exists in the directory, click **Test Context**.
3. Under Directory Search, indicate how far below the root context Directory Agent looks for users and groups.
 - Select **None** to limit searches to the root context only.
 - Select **One Level** to limit searches to the root context and one level below.
 - Select **All Levels** to expand searches to the root context and all levels below.
4. If you selected One Level or All Levels for group searches, mark **Include all users in selected groups, regardless of context** if you want to ensure that all users are included from the groups found in the directory search, even if some of those users are in a different context. This option is selected by default for Active Directory.
5. To fine-tune the search filters that Directory Agent uses for this context, click **Customize Search Filters**.
6. Mark **Customize search filters**, and edit the user and group search filters as required.
7. If required, click **Add** under Exclude Contexts to enter a context within the directory search that should not be sent to the hybrid service.
8. Under Directory Search, indicate how far below the excluded context Directory Agent looks for users and groups.
 - Select **None** to limit searches to the root context only.
 - Select **One Level** to limit searches to the root context and one level below.
 - Select **All Levels** to expand searches to the root context and all levels below.

Note that the user and group levels for an excluded context cannot be greater than the defined levels for the root context. For example, if the root context's Directory Search level for either users or groups is set to None, the corresponding users or groups search level for the excluded context will also be set to None and cannot be changed.
9. To verify that the context that you have entered uses valid syntax, exists in the directory, and does not have any user or group dependencies with the root context, click **Test Context**.
10. Click **OK** to save the excluded context.

When you are finished, click **OK** to close the Add Context page and update the Root Context for Hybrid Filtering Users table. You must also click **OK** on the Shared User Data page to cache the change.

Optimizing search results

Optimizing search results further refines the data that is sent to the hybrid service by defining patterns, or search filters, used to remove duplicate or otherwise unwanted entries from the directory search results. It also provides a way to modify the **mail** attribute for directory entries collected by Directory Agent before they are sent to the hybrid service.

If, for example, the **mail** attribute in your directory service has a partial or internal email address reference, you could use a search filter to replace that partial or internal information with external information, usable by the hybrid service. This would be useful for those who configure the hybrid service to automatically create passwords for users so that they can connect to hybrid filtering when they are off site (see [Configuring hybrid filtering for off-site users, page 207](#)).

Any search filters that you create in TRITON - Web Security are applied to the directory data collected by Directory Agent before that data is sent to the hybrid service.

Click **Optimize Search Results** to see the current search filters, or to create new search filters using wildcards or regular expressions. There are 2 types of search filters: one to filter user entries and one to filter group entries.

- ◆ To create a new search filter, click **Add** under the appropriate table.
- ◆ To edit an existing search filter, click the associated **Find String**.

A pop-up dialog box prompts you to edit or enter:

- ◆ **Find string:** The text to search for in the original directory data collected by Directory Agent.
- ◆ **Replace string:** The new text that you want to substitute for the original text in data sent to the hybrid service.

When you are finished, click **OK** to close the dialog box and update the Filter User Results or Filter Group Results table. You must also click **OK** on the Shared User Data page to cache the change.

At this time, Directory Agent applies the search filters that you create only to the **mail** attribute.

Schedule communication with hybrid filtering

Select **Settings > Hybrid Configuration > Scheduling** to specify how frequently directory data collected by Directory Agent is sent to the hybrid service, and how often reporting data is retrieved.



Note

Policy data is collected whenever you click **Save All** in TRITON - Web Security, and sent to the hybrid service at 15 minute intervals by default. If you have made an important update to your policy data, and want to send user and group information right away, click **Send** under Send Policy Data Now.

To configure how often directory information is sent to the hybrid service:

1. Under **Send User Data**, select one or more days of the week to send user and group information to the hybrid service. If you are using directory information to identify users, you must send Directory Agent data at least once a week.
2. Enter start and end times to define the time period during which Sync Service attempts to send directory data to the hybrid service. Typically, directory data is sent at a period of low traffic in your network.
3. If you have made an important update to your directory service data, and want to send user and group information right away, click **Send** under Send Update Now. If TRITON - Web Security receives confirmation from Sync Service, a success message is displayed. This means that Sync Service will send the data, not that the data has been received by the hybrid service.

To configure whether the hybrid service collects reporting data, and how often Sync Service retrieves the data:



Important

In order for Sync Service to pass hybrid reporting data to Log Server, a hybrid communication port must be configured on the Settings > General > Logging page. See [Configuring how filtered requests are logged, page 355](#), for details.

If you are using distributed logging, Sync Service must be configured to communicate with the central Log Server. Hybrid logging data cannot be passed from remote Log Server instances to the central Log Server.

1. Under Collect and Retrieve Reporting Data, mark **Have the hybrid service collect reporting data for the clients it filters**.

If you clear this check box, log data is not saved for hybrid filtering users. No information about these users' Internet activity will appear in reports.

2. Select one or more days of the week for Sync Service to request reporting data from the hybrid service. You must retrieve data at least once a week.
3. Enter start and end times to define the time period during which Sync Service retrieves data from the hybrid service. You may want to retrieve data at a period of low traffic in your network.
4. Select how often you want Sync Service to request reporting data from the hybrid service within the specified start and end times.

Sync Service cannot download reporting data any more frequently than every 15 minutes. This means that there is a time delay between when hybrid filtering makes Internet requests and when those requests appear in TRITON - Web Security reports.

If you need to route Sync Service traffic to and from the hybrid service through a proxy server or firewall:

1. Under Route Sync Service Traffic, mark **Route Sync Service traffic through a proxy server or firewall**.
2. Enter the IP address or host name of the proxy server or firewall, and specify the port that is to be used.
3. If the specified server requires authentication, enter the user name and password for Sync Service to access it.

When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Monitor communication with the hybrid service

You can view the status of the hybrid service on the **Main > Status > Hybrid Service** page in TRITON - Web Security. This page displays when data was most recently sent to or received from the hybrid service. If an attempt to send or receive data failed, find out when the failure occurred, and which components were involved.

The page lists the date and time that Sync Service last:

- ◆ Connected or attempted to connect to the hybrid service for any reason
- ◆ Sent or attempted to send directory information to the hybrid service
- ◆ Retrieved or attempted to retrieve log (reporting) data from the hybrid service
- ◆ Sent or attempted to send log data to Log Server
- ◆ Sent or attempted to send account information to the hybrid service
- ◆ Sent or attempted to send policy information to the hybrid service

If you have not yet set up the connection between the on-premises and hybrid portions of Websense Web Security Gateway Anywhere, a message explains that "No communication has occurred."

Under Last Directory Agent Sync Results, the page lists:

- ◆ The date and time that Directory Agent last sent data to the hybrid service
- ◆ The total number of users and groups processed by Directory Agent
- ◆ The number of users and groups that were updated in the hybrid service
- ◆ The number of groups filtered out because they contained invalid values
- ◆ The number of users filtered out because they included invalid email addresses
- ◆ The number of new users and groups synchronized with the hybrid service
- ◆ The number of obsolete users and groups removed from the hybrid service

This page also allows you to access authentication method reports from the hybrid service (see [View hybrid service authentication reports](#), page 198), and displays the type of PAC file you are using:

- ◆ The default PAC file from the hybrid service
- ◆ A customized PAC file uploaded from the Websense **pac** directory (see [What is a PAC file?](#), page 187)
- ◆ The default PAC file with an uploaded customized fragment

If you are using a custom file or fragment, the page shows how long the file or fragment has been in use.

If a **Secondary date stamp** is shown for the PAC file, this means that Sync Service has uploaded both a custom PAC file and a custom fragment from the **pac** directory. The recommended state for custom PAC files is to set up a custom file or a custom fragment, not both. To rectify this, go to the **pac** directory (by default, \Program Files\Websense\Web Security\bin\data\pac on Windows, or /opt/websense/bin/data/pac on Linux) and delete either websense.pac or customfinal.pac.

If you are using a customized block page logo, this page displays the date and time the logo file was uploaded to the hybrid service.

View hybrid service authentication reports

Select **View Report** on the **Main > Status > Hybrid Service** page to download reporting data from the hybrid service and see a breakdown of how hybrid users are identified or authenticated with the service.

The report output consists of a pie chart and a table, showing the number of clients using each available authentication method over the last 7 days. **Web Endpoint**, **Authentication Service**, **NTLM identification**, and **Manual authentication** are all set up for clients on the Settings > Hybrid Configuration > Hybrid User Identification page (see [Send user and group data to the hybrid service](#), page 189).

X-Authenticated-User authentication is available if you have deployed one of the following as a downstream chained proxy server:

- ◆ Microsoft® Internet Security and Acceleration (ISA) Server or Forefront™ Threat Management Gateway (TMG) server
- ◆ BlueCoat Proxy SG

The downstream proxy server performs user authentication and forwards requests to the hybrid proxy using the X-Authenticated-User header.

Click an authentication method in the table to see a list of users who have most recently authenticated with that method. You cannot click an authentication method that you have not deployed or that is currently not in use.

Each authentication method report can contain up to 1000 users. The users are listed by user name, email address, and last logon time. Click the arrow buttons at the bottom of the report to view previous or subsequent pages.

Reports displayed in the content pane cannot be printed or saved to a file. To print or save a report to file, click **Export to PDF** or **Export to XLS** to view the report in the appropriate output format.

**Important**

To display authentication reports in PDF format, Adobe Reader v7.0 or later must be installed on the machine from which you are accessing TRITON - Web Security.

To display authentication reports in XLS format, Microsoft Excel 2003 or later must be installed on the machine from which you are accessing TRITON - Web Security.

Each report includes the data and time it was last updated. Updates are not automatic: to download the latest report data from the hybrid service, click **Update**.

9

Filter Users Off Site

Related topics:

- ◆ [Using remote filtering software, page 202](#)
- ◆ [Applying hybrid filtering to off-site users, page 206](#)

In addition to filtering users inside your organization's network, Websense Web security solutions provide options for filtering users when they are outside the network.

- ◆ Install remote filtering software to monitor Internet activity for users outside the network. See [Using remote filtering software, page 202](#).

Remote filtering software is included with Websense Web Security Gateway Anywhere subscriptions, and is available as an option for Websense Web Filter, Websense Web Security, and Websense Web Security Gateway customers.

- ◆ Use hybrid filtering to monitor Internet activity for users outside the network, regardless of how they are filtered when they are in the network. See [Applying hybrid filtering to off-site users, page 206](#).

Hybrid filtering is available only with Websense Web Security Gateway Anywhere.

These methods can be used, for example, to filter users who work from home, users who travel using company laptops, or students who use institutional laptops on and off campus.



Important

With Websense Web Security Gateway Anywhere, you can use remote filtering software for some off-site users and hybrid filtering for others. The hybrid service cannot, however, be used to monitor Internet activity for machines that also have Remote Filtering Client installed.

Using remote filtering software

Related topics:

- ◆ [When server communication fails, page 203](#)
- ◆ [Configuring Remote Filtering settings, page 204](#)

By default, remote filtering software components monitor HTTP, SSL, and FTP traffic and apply a user-based policy or the Default policy. Remote filtering software does not apply policies to IP addresses (computers or network ranges).

- ◆ Bandwidth-based filtering is not applied to remote filtering clients, and bandwidth generated by remote filtering traffic is not included in bandwidth measurements and reports.
- ◆ Remote filtering software can only block or permit FTP and SSL (HTTPS) requests. FTP and HTTPS sites in categories assigned the quota or confirm action are blocked when the user is outside the network.
- ◆ While remote filtering software always monitors HTTP traffic, you can configure it to ignore FTP traffic, HTTPS traffic, or both. See [Configure remote filtering to ignore FTP or HTTPS traffic, page 205](#).

Remote filtering software includes the following components:

- ◆ **Remote Filtering Server** is installed inside your network's outermost firewall, and configured so that filtered machines outside the network can communicate with it.
- ◆ **Remote Filtering Client** is installed on Microsoft Windows machines that are used outside the network.



Notes

Follow the recommendations in the *Deployment and Installation Center* carefully to deploy these components. See the *Remote Filtering Software* technical paper for instructions on installing them.

Configure Network Agent or the integration product that passes traffic to your Websense software **not** to monitor the Remote Filtering Server machine (see [Configuring global settings, page 379](#)).

All communication between Remote Filtering Client and Remote Filtering Server is authenticated and encrypted.

By default, when an HTTP, SSL, or FTP request is made from a machine with Remote Filtering Client installed:

1. The client first determines whether or not it is inside the network by sending a **heartbeat** to the Remote Filtering Server in the DMZ.
2. If the machine is **inside** the network, Remote Filtering Client takes no action. The request is passed to Network Agent or an integration product, and filtered like other in-network Internet activity.
3. If the machine is **outside** the network, Remote Filtering Client communicates with Remote Filtering Server over the configured port (80, by default).
4. Remote Filtering Server then contacts Filtering Service (installed inside the network) to ask what action to apply to the request.
5. Filtering Service evaluates the request and sends a response to Remote Filtering Server.
6. Finally, Remote Filtering Server responds to Remote Filtering Client, either permitting the site or sending the appropriate block message.

When it starts for the first time, Remote Filtering Client creates a log file that tracks its filtering activities, such as entering and leaving the network, failing open or closed, and restarting the client. You control the presence and size of this log file. See [Configuring Remote Filtering settings, page 204](#).

Complete information about planning for, deploying, and configuring remote filtering software is available in the [Remote Filtering Software](#) technical paper, available from the [Websense Knowledge Base](#).

When server communication fails

Related topics:

- ◆ [Using remote filtering software, page 202](#)
- ◆ [Configuring Remote Filtering settings, page 204](#)

Filtering occurs when Remote Filtering Client, outside the network, successfully communicates with Remote Filtering Server.

You can configure what action Remote Filtering Client takes if it cannot contact Remote Filtering Server.

- ◆ By default, Remote Filtering Client permits all HTTP, SSL, and FTP requests while it continues attempting to contact Remote Filtering Server (fails open). When the communication is successful, the appropriate filtering policy is enforced.
- ◆ When Remote Filtering Client is configured to block all requests (fail closed), a timeout value is applied (default 15 minutes). The clock begins running when the remote computer is started. Remote Filtering Client attempts to connect to

Remote Filtering Server immediately and continues cycling through available Remote Filtering Servers until it is successful.

If the user has Web access at startup, no filtering occurs (all requests are permitted) until Remote Filtering Client connects to the Remote Filtering Server.

If Remote Filtering Client cannot connect within the configured timeout period, all Internet access is blocked (fail closed) until connection to Remote Filtering Server can be established.

**Note**

If Remote Filtering Server cannot connect to Filtering Service for any reason, an error is returned to the Remote Filtering Client, and filtering always fails open.

This timeout period allows users who pay for Internet access when travelling to start the computer and arrange for connection without being locked out. If the user does not establish Web access before the 15 minute timeout period expires, Web access cannot be established during that session. When this occurs, the user must restart the computer to begin the timeout interval again.

To change whether requests are permitted or blocked when Remote Filtering Client can't communicate with Remote Filtering Server, and change the timeout value, see [Configuring Remote Filtering settings, page 204](#).

Configuring Remote Filtering settings

Related topics:

- ◆ [When server communication fails, page 203](#)
- ◆ [Configure remote filtering to ignore FTP or HTTPS traffic, page 205](#)
- ◆ [Configure the Remote Filtering Client heartbeat interval, page 206](#)

Use the **Settings > General > Remote Filtering** page to configure options that affect all Remote Filtering Clients associated with this installation.

For detailed information about how remote filtering works, which components are involved, and how to deploy components, see the [Remote Filtering Software](#) technical paper.

1. Select the **Block all requests...** check box to prevent users from accessing the Internet when Remote Filtering Client cannot communicate with Remote Filtering Server (fail closed).

By default, users have unfiltered access to the Internet when Remote Filtering Client cannot communicate with the Remote Filtering Server (fail open).

For more information, see [When server communication fails, page 203](#).

- If you selected the block all requests option, specify a **Timeout interval** (by default, 15 minutes). During the timeout period, all HTTP, SSL, and FTP requests are permitted.

If the Remote Filtering Client cannot communicate with Remote Filtering Server during the timeout interval, all Internet access will be blocked until communication is reestablished.

Selecting **No timeout** may lock out a remote computer before the user can establish Internet connection from a hotel or other pay-for-use-provider.



Warning

Websense, Inc., does not recommend choosing **No timeout** or setting the timeout period to a very low number.

- Select a **Maximum size** for the Remote Filtering Client log file size (in megabytes), up to 10. Select **No log** to disable logging.

This controls the size and existence of the log file the remote computer creates when it is initially disconnected from the Remote Filtering Server. This log file tracks the following events:

- The computer leaves the network
- The computer rejoins the network
- The Remote Filtering Client is restarted
- Fail open condition occurs
- Fail closed condition occurs
- Remote Filtering Client receives a policy update

The computer retains the 2 most recent logs. These logs can be used to troubleshoot connection issues or other problems with remote filtering software.

Configure remote filtering to ignore FTP or HTTPS traffic

You can configure remote filtering software to ignore FTP traffic, HTTPS traffic, or both. HTTP traffic is always monitored.

If you have multiple Remote Filtering Servers, repeat these steps for each instance.

- Navigate to the Websense **bin** directory (see [Where is the Websense "bin" directory?](#), page 457) on the Remote Filtering Server machine.
- Open the **securewispproxy.ini** file in a text editor.
- To disable FTP filtering for this Remote Filtering Server instance, add the following line to the file:

```
FilterFTP=0
```

If you want to later turn FTP filtering back on, change the parameter value from "0" to "1".

- To disable HTTPS filtering for this Remote Filtering Server instance, add the following line to the file:

```
FilterHTTPS=0
```

If you want to later turn HTTPS filtering back on, change the parameter value from “0” to “1”.

5. Save and close the file.
6. Restart the Remote Filtering Server service or daemon.

Configure the Remote Filtering Client heartbeat interval

In order to determine whether it is inside or outside of the network, Remote Filtering Client sends a heartbeat to Remote Filtering Server. If the heartbeat connection succeeds, Remote Filtering Client knows that it is inside the network. By default, Remote Filtering Client continues to send the heartbeat every 15 minutes to ensure that its status has not changed.

If you would prefer that Remote Filtering Client send the heartbeat less frequently once it has determined that it is inside the network, you can increase the heartbeat interval. In this case, Remote Filtering Client will only send a more frequent heartbeat if it registers a change in network.

To change the heartbeat interval:

1. Navigate to the Websense **bin** directory (see [Where is the Websense “bin” directory?](#), page 457) on the Remote Filtering Server machine.
2. Open the **securewisproxy.ini** file in a text editor.
3. Find the **HeartbeatRetryInterval** parameter and change its value. For example:

```
HeartbeatRetryInterval=360
```

In this example, the heartbeat will be sent every 360 minutes, or 6 hours.

- The value can be any number of minutes between 0 and 1440 (24 hours).
 - The default is 15 minutes.
4. Save and close the file.
 5. Restart the Remote Filtering Server service or daemon.

Applying hybrid filtering to off-site users

Related topics:

- ◆ [Configuring hybrid filtering for off-site users](#), page 207
- ◆ [Off-site user self-registration](#), page 207

In Websense Web Security Gateway Anywhere deployments, hybrid filtering can be applied to off-site users, regardless of how those users are filtered when they are in-network.

- ◆ For users filtered by on-premises components (Filtering Service) when they are inside the network, you can configure the browser PAC file to determine whether the user is in-network or off-site before sending an Internet request for filtering. If you are using the PAC file generated by the hybrid service, this configuration occurs automatically based on the settings that you provide in TRITON - Web Security.
- ◆ For users filtered by hybrid filtering both in and outside the network, no PAC file changes are required. When off-site users make an Internet request, they are prompted to log on to hybrid filtering so that the appropriate user or group-based policy can be applied.



Important

While you can use remote filtering software for some off-site users and hybrid filtering for others, the hybrid service cannot be used to monitor Internet activity for machines that also have Remote Filtering Client installed.

Configuring hybrid filtering for off-site users

To configure hybrid filtering for users outside a filtered location:

- ◆ If hybrid filtering uses directory data collected by Websense Directory Agent to identify users, you can either configure the hybrid service to automatically create a hybrid logon password for all user accounts sent by Directory Agent (see [Send user and group data to the hybrid service, page 189](#)), or you can have users request their own password the first time they connect to the hybrid service from outside a filtered location (see [Off-site user self-registration, page 207](#)).
- ◆ If your organization does not use directory data collected by Directory Agent to identify users connecting to the hybrid service, you can let users **self-register** for the service. See [Configure user access to hybrid filtering, page 183](#).
- ◆ Once you have established an identification policy for off-site users, mark **Enable off-site users** on the **Settings > Hybrid Configuration > User Access** page in TRITON - Web Security. See [Configure user access to hybrid filtering, page 183](#).

Off-site user self-registration

If you are not sending directory service data to the hybrid service (in other words, if you have not enabled Directory Agent), users must self-register in order to be filtered properly when they are off site (outside a filtered location).

In order for users to be allowed to self-register, you must first identify the domains associated with your organization on the **Settings > Hybrid Configuration > User Access** page in TRITON - Web Security (see [Configure user access to hybrid filtering, page 183](#)).

Users connecting to hybrid filtering from outside a filtered location are prompted to enter a user name and password, or to register. To register with the hybrid service:

1. The user provides a name and email address.
2. Hybrid filtering then sends a password to the user via email, along with a link that can be used to change the password.
3. The user clicks the link, and is prompted to enter the password.
4. Registration is complete.

When registered users connect to the hybrid service from outside a filtered location, they enter their email address and password. Hybrid filtering then applies your organization's Default policy to their Internet requests.

10

Protect Vital Information

Websense Web Security secures your enterprise from Web-based threats, liability issues, and productivity loss. But what if you want—or are required—to protect sensitive data, such as social security numbers or credit card numbers, from leakage over the Web? Or what if you want to monitor removable media devices, printers, instant messages, copy/paste operations, or email for the such data?

To protect against data loss over the Web, you can deploy Websense Web Security Gateway Anywhere. To protect against data loss over other channels, including the Web, you can purchase Websense Data Protect, Data Monitor, Data Discover, Data Endpoint, or the full Data Security Suite as add-ons to your Web security software.

Websense Web and data security solutions interoperate in fundamental ways, giving the data security software access to user information (collected by User Service) and URL categorization information (from the Master Database).

By combining Web and data security, you can create data loss prevention (DLP) policies that base rules on URL categories. For example, you can define a rule that credit card numbers cannot be posted to known fraud sites. You can also define rules based on users and computers rather than IP addresses. For example, Jane Doe cannot post financial information to FTP sites.

For an end-to-end description of setting up data loss protection over the Web, see the *Deployment and Installation Center*. This covers installation, deployment, and configuration of the various components, including Websense Content Gateway.

For instructions on creating data security policies, see the TRITON - Data Security Help.

11

Refine Filtering Policies

At its simplest, Internet usage filtering requires a single policy that applies one category filter and one protocol filter 24 hours a day, 7 days a week. Websense software offers tools, however, for going far beyond this basic filtering, to achieve precisely the level of granularity you need to manage Internet usage. You can:

- ◆ Create **limited access filters** to block access to all but a specified list of sites for certain users (see [Restricting users to a defined list of Internet sites, page 212](#)).
- ◆ Create **custom categories** to redefine how selected sites are filtered (see [Working with categories, page 218](#)).
- ◆ **Recategorize URLs** to move specific sites from their default, Master Database category to another Websense-defined or custom category (see [Recategorizing URLs, page 226](#)).
- ◆ Define **unfiltered URLs** to allow users to access specific sites, even when the sites are assigned to a blocked category in the active category filter (see [Defining unfiltered URLs, page 225](#)).
- ◆ Implement **bandwidth** restrictions, blocking users from accessing otherwise permitted categories and protocols when bandwidth usage reaches a specified threshold (see [Using Bandwidth Optimizer to manage bandwidth, page 236](#)).

In Websense Web Security Gateway Anywhere environments, bandwidth-based restrictions are not enforced for hybrid filtering users.

- ◆ Define **keywords** used to block sites in otherwise permitted categories when keyword blocking is enabled and activated (see [Filtering based on keyword, page 222](#)).
- ◆ Define **file types** used to block the download of selected types of files from otherwise permitted categories when file type blocking is activated (see [Managing traffic based on file type, page 238](#)).

Restricting users to a defined list of Internet sites

Related topics:

- ◆ [Limited access filters and filtering precedence, page 212](#)
- ◆ [Creating a limited access filter, page 213](#)
- ◆ [Editing a limited access filter, page 214](#)

Limited access filters provide a very precise method of filtering Internet access. Each limited access filter is a list of individual Web sites. Like category filters, limited access filters are added to policies and enforced during a specified time period. When a limited access filter is active in a policy, users assigned that policy can visit only sites in the list. All other sites are blocked.

For example, if the First Grade policy enforces a limited access filter that includes only certain educational and reference sites, students governed by the First Grade policy can visit only those sites, and no others.

When a limited access filter is active, a block page is returned for any requested URL not included in that filter.

Websense software can support up to 2,500 limited access filters containing 25,000 URLs in total.

Limited access filters and filtering precedence

In some cases, more than one filtering policy could apply to a single user. This happens when a user belongs to more than one group, and the groups are governed by different policies. In addition, a URL might both appear in a limited access filter and be defined as an unfiltered URL.

When multiple group policies apply to a user, the **Use more restrictive blocking** setting (see [Filtering order, page 85](#)) determines how the user is filtered. By default, this setting is off.

Websense software determines which filtering setting is less restrictive at the filter level. In cases where a user might be filtered by multiple policies, one of which is enforcing a limited access filter, “less restrictive” may sometimes seem counterintuitive.

When **Use more restrictive blocking** is **OFF**:

- ◆ If the **Block All** category filter and a limited access filter could apply, the limited access filter is always considered less restrictive.
- ◆ If any other category filter and a limited access filter could apply, the category filter is considered less restrictive.

This means that even when the limited access filter permits the site and the category filter blocks the site, the site is blocked.

When **Use more restrictive blocking** is **ON**, a limited access filter is considered more restrictive than any category filter except **Block All**.

The table below summarizes how the **Use more restrictive blocking** setting affects filtering when multiple policies could apply:

	<i>Use more restrictive blocking OFF</i>	<i>Use more restrictive blocking ON</i>
limited access filter + Block All category filter	limited access filter (request permitted)	Block All (request blocked)
limited access filter + permitted category	category filter (request permitted)	limited access filter (request permitted)
limited access filter + blocked category	category filter (request blocked)	limited access filter (request permitted)
limited access filter + Quota/Confirm category	category filter (request limited by quota/confirm)	limited access filter (request permitted)
limited access filter + unfiltered URL	unfiltered URL (request permitted)	limited access filter (request permitted)

Creating a limited access filter

Related topics:

- ◆ [Working with filters, page 49](#)
- ◆ [Restricting users to a defined list of Internet sites, page 212](#)
- ◆ [Editing a limited access filter, page 214](#)

Use the **Add Limited Access Filter** page (accessed via the **Filters** or **Edit Policy** page) to give your new filter a unique name and a description. After creating the filter, enter a list of permitted URLs, assign the filter to a policy, and apply the policy to clients.

1. Enter a unique **Filter name**. The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Filter names can include spaces, dashes, and apostrophes.

2. Enter a short **Description** of the filter. This description appears next to the filter name in the Limited Access Filters section of the Filters page, and should explain the filter's purpose to help administrators manage policies over time.

The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. To see and edit the new filter, click **OK**. To abandon your changes and return to the Filters page, click **Cancel**.

When you create a new limited access filter, it is added to the **Policy Management > Filters > Limited Access Filters** list. Click a filter name to edit the filter.

To finish customizing your new filter, continue with [Editing a limited access filter](#).

Editing a limited access filter

Related topics:

- ◆ [Restricting users to a defined list of Internet sites, page 212](#)
- ◆ [Limited access filters and filtering precedence, page 212](#)
- ◆ [Creating a limited access filter, page 213](#)
- ◆ [Editing a policy, page 83](#)

A limited access filter is a list of Web sites (URLs or IP addresses) and regular expressions, used to identify specific sites that users can access. When the filter is applied to clients, those clients cannot visit any site that is not in the list.



Important

If a site permitted by a limited access filter becomes infected with malicious code, as long as Security categories are blocked, user requests for that site are blocked.

For instructions to change this behavior, see [Prioritizing Security Risk categorization, page 227](#).

Use the **Policy Management > Filters > Edit Limited Access Filter** page to make changes to an existing limited access filter. You can change the filter name and description, see a list of policies that enforce the filter, and manage which sites are included in the filter.

When you edit a limited access filter, the changes affect every policy that enforces the filter.

1. Verify the filter name and description. To change the filter name, click **Rename**, and then enter the new name. The name is updated in all policies that enforce the selected limited access filter.
2. Use the **Policies using this filter** field to see how many policies currently enforce this filter. If 1 or more policies enforce the filter, click **View policies** to list them.
3. Under Add or Remove Sites, enter the URLs and IP addresses that you want to add to the limited access filter. Enter one URL or IP address per line.
 - For HTTP sites, it is not necessary to include the HTTP:// prefix.

- When an HTTP site is filtered according to its Master Database category, Websense software matches the URL with its equivalent IP address. This is not the case for limited access filters. To permit a site's URL and IP address, add both to the filter.
 - For FTP and HTTPS sites, include the prefix and provide the site's IP address, rather than host (domain) name.
4. Click the right arrow (>) to move the URLs and IP addresses to the Permitted sites list.
 5. In addition to adding individual sites to the limited access filter, you can add regular expressions that match multiple sites. To create regular expressions, click **Advanced**.
 - Enter one regular expression per line, and then click the right arrow to move the expressions to the Permitted sites list.
 - To verify that a regular expression matches the intended sites, click **Test**.
 - See [Using regular expressions, page 241](#), for detailed information about using regular expressions for filtering.
 6. Review the URLs, IP addresses, and regular expressions in the **Permitted sites** list.
 - To make changes to a site or expression, select it and click **Edit**.
 - To remove a site or expression from the list, select it and click **Delete**.
 7. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

Adding sites from the Edit Policy page

Related topics:

- ◆ [Restricting users to a defined list of Internet sites, page 212](#)
- ◆ [Limited access filters and filtering precedence, page 212](#)
- ◆ [Creating a limited access filter, page 213](#)
- ◆ [Editing a policy, page 83](#)

Use the **Policies > Edit Policy > Add Sites** page to add sites to a limited access filter.

Enter one URL or IP address per line. If you do not specify a protocol, Websense software automatically adds the **HTTP://** prefix.

When you are finished making changes, click **OK** to return to the Edit Policy page. You must also click **OK** on the Edit Policy page to cache the changes. Changes are not implemented until you click **Save All**.

Changes made to a limited access filter affect all policies that enforce the filter.

Copying filters and policies to roles

Related topics:

- ◆ [Creating a category filter, page 50](#)
- ◆ [Creating a protocol filter, page 52](#)
- ◆ [Creating a limited access filter, page 213](#)
- ◆ [Creating a policy, page 82](#)

Super Administrators can use the **Filters > Copy Filters To Role** and **Policies > Copy Policies To Role** pages to copy one or more filters or policies to a delegated administration role. Once the filter or policy has been copied, delegated administrators can use the filters or policies to filter their managed clients.

- ◆ In the target role, the tag “(Copied)” is added to the end of the filter or policy name. A number is added if the same filter or policy is copied multiple times. For example, “(Copied 2).”
- ◆ Delegated administrators can rename or edit filters or policies that have been copied to their role.
- ◆ Category filters copied to a delegated administration role set the filtering action to Permit for custom categories created in the role. Delegated administrators should update the copied category filters to set the desired action for their role-specific custom categories.
- ◆ Changes made by a delegated administrator to a filter or policy copied to their role by a Super Administrator do not affect the Super Administrator’s original filter or policy, or any other role that received a copy of the filter or policy.
- ◆ Filter Lock restrictions do not affect the Super Administrator’s original filter or policy, but they do affect the delegated administrator’s copy of the filter or policy.
- ◆ Because delegated administrators are affected by Filter Lock restrictions, the Permit All category and protocol filters cannot be copied to a delegated administration role.

To copy a filter or policy:

1. On the Copy Filters to Role or Copy Policies to Role page, verify that the correct policies or filters appear in the list at the top of the page.
2. Use the **Select a role** drop-down list to select a destination role.
3. Click **OK**.

A pop-up dialog box indicates that the selected filters or policies are being copied. The copy process may take a while.

The changes are not implemented until you click **Save All**.

After the copy process is complete, the copied filters or policies will be available to delegated administrators in the selected role the next time they log on to TRITON -

Web Security. If a delegated administrator is logged on to the role with policy access when the filters or policies are copied, they will not see the new filters or policies until they log off and log on again.

Building filter components

Use the **Policy Management > Filter Components** page to access tools used to refine and customize the way that Websense software enforces your organization's Internet access policies. The 4 buttons on the screen are associated with the following tasks:

Edit Categories	<ul style="list-style-type: none"> • Recategorize a URL (see Redefining filtering for specific sites, page 224). For example, if the Shopping category is blocked by your Internet filtering policies, but you want to permit access to specific supplier or partner sites, you could move those sites to a permitted category, like Business and Economy. • Define or edit custom categories (see Creating a custom category, page 221). Create additional subcategories within Websense-defined parent categories, or within the User-Defined parent category, and then assign URLs to the new categories. • Assign keywords to a category (see Filtering based on keyword, page 222). To recategorize and block access to sites whose URLs contain a specific string, first define keywords, and then enable keyword blocking in a category filter. • Create regular expressions (see Using regular expressions, page 241), patterns or templates that can be used to match multiple URLs and assign them to a category.
Edit Protocols	Define or edit custom protocol definitions (see Creating a custom protocol, page 234 , and Editing custom protocols, page 231). For example, if members of your organization use a custom messaging tool, you could create a custom protocol definition to permit use of that tool while blocking other Instant Messaging / Chat protocols.
File Types	Create or edit file type definitions, used to block specific types of files within otherwise permitted categories (see Managing traffic based on file type, page 238).
Unfiltered URLs	Define specific sites to permit for all clients, even when they belong to a blocked category (see Defining unfiltered URLs, page 225). Note that adding a URL to this list does not override the Block All category filter or limited access filters.

Working with categories

Related topics:

- ◆ [Editing categories and their attributes, page 218](#)
- ◆ [Creating a custom category, page 221](#)
- ◆ [Filtering based on keyword, page 222](#)
- ◆ [Redefining filtering for specific sites, page 224](#)

Websense software provides multiple methods for filtering sites that are not in the Master Database, and for changing the way that individual sites in the Master Database are filtered.

- ◆ Create **custom categories** for more precise filtering and reporting.
- ◆ Use **recategorized URLs** to define categories for uncategorized sites, or to change the category for sites that appear in the Master Database.
- ◆ Define **keywords** to recategorize all sites whose URL contains a certain string.

If you want to configure whether or not attempts to access a category are recorded in the Log Database, see [Configuring how filtered requests are logged, page 355](#). If a category is not logged, client requests for that category do not appear in reports.

Editing categories and their attributes

Related topics:

- ◆ [Creating a custom category, page 221](#)
- ◆ [Reviewing all customized category attributes, page 219](#)
- ◆ [Making global category filtering changes, page 220](#)
- ◆ [Filtering based on keyword, page 222](#)
- ◆ [Redefining filtering for specific sites, page 224](#)

Use the **Policy Management > Filter Components > Edit Categories** page to create and modify custom categories, recategorized URLs, and keywords.

The existing categories, both Websense-defined and custom, are listed in the left portion of the content pane. To see current custom settings associated with a category, or to create new custom definitions, first select a category from the list.

To see a list of all custom URLs, keywords, and regular expressions associated with all categories, click **View All Custom URLs / Keywords** in the toolbar at the top of the page. See [Reviewing all customized category attributes, page 219](#), for more information.

- ◆ To create a new category, click **Add**, and then go to [Creating a custom category, page 221](#), for further instructions.
To remove an existing custom category, select the category, and then click **Delete**. You cannot delete Websense-defined categories.
- ◆ To change the name or description of a custom category, select the category and click **Rename** (see [Renaming a custom category, page 220](#)).
- ◆ To change the filtering action associated with a category in all category filters, click **Override Action** (see [Making global category filtering changes, page 220](#)).
- ◆ The **Recategorized URLs** list shows which recategorized sites (URLs and IP addresses) have been assigned to this category.
 - To add a site to the list, click **Add URLs**. See [Recategorizing URLs, page 226](#), for further instructions.
 - To change an existing recategorized site, select the URL or IP address, and then click **Edit**.
- ◆ The **Keywords** list shows which keywords have been associated with this category.
 - To define a keyword associated with the selected category, click **Add Keywords**. See [Filtering based on keyword, page 222](#), for further instructions.
 - To change an existing keyword definition, select the keyword, and then click **Edit**.
- ◆ In addition to URLs and keywords, you can define **Regular Expressions** for the category. Each regular expression is a pattern or template used to associate multiple sites with the category.
To see or create regular expressions for the category, click **Advanced**.
 - To define a regular expression, click **Add Expressions** (see [Using regular expressions, page 241](#)).
 - To change an existing regular expression, select the expression, and then click **Edit**.
- ◆ To delete a recategorized URL, keyword, or regular expression, select the item to remove, and then click **Delete**.

When you are finished making changes on the Edit Categories page, click **OK** to cache the changes and return to the Filter Components page. Changes are not implemented until you click **Save All**.

Reviewing all customized category attributes

Use the **Filter Components > Edit Categories > View All Custom URLs and Keywords** page to review custom URL, keyword, and regular expression definitions. You can also delete definitions that are no longer needed.

The page contains 3 similar tables, one for each category attribute: custom URLs, keywords, or regular expressions. In each table, the attribute is listed next to the name of the category with which it is associated.

To delete a category attribute, mark the appropriate check box, and then click **Delete**.

To return to the Edit Categories page, click **Close**. If you deleted any items on the View All Custom URLs and Keywords page, click **OK** on the Edit Categories page to cache the changes. Changes are not implemented until you click **Save All**.

Making global category filtering changes

Use the **Filter Components > Edit Categories > Override Action** page to change the action applied to a category in all existing category filters. This also determines the default action applied to the category in new filters.

Although this change overrides the action applied to the category in all existing filters, administrators can later edit those filters to apply a different action.

Before changing the filtering settings applied to a category, first verify that the correct category name appears next to **Selected Category**. Next, you can:

1. Chose a new **Action** (Permit, Block, Confirm, or Quota). See [Filtering actions, page 46](#), for more information.
By default, **Do not change current settings** is selected for all options on the page.
2. Specify whether or not to **Block Keywords**. See [Filtering based on keyword, page 222](#), for more information.
3. Specify whether or not to **Block File Types**, and customize blocking settings. See [Managing traffic based on file type, page 238](#), for more information.
4. Under **Advanced Filtering**, specify whether or not to use Bandwidth Optimizer to manage access to HTTP sites, and customize blocking settings. See [Using Bandwidth Optimizer to manage bandwidth, page 236](#), for more information.



Important

Changes made here affect every existing category filter, except **Block All** and **Permit All**.

5. Click **OK** to return to the Edit Categories page (see [Editing categories and their attributes, page 218](#)). The changes are not cached until you click **OK** on the Edit Categories page.

Renaming a custom category

Use the **Filter Components > Edit Categories > Rename Category** page to change the name or description associated with a custom category.

- ◆ Use the **Filter name** field to edit the category name. The new name must be unique, and cannot exceed 50 characters.

The name cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

- ◆ Use the **Description** field to edit the category description. The description cannot exceed 255 characters.

The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

When you are finished making changes, click **OK** to return to the Edit Categories page. The changes are not cached until you click **OK** on the Edit Categories page.

Creating a custom category

Related topics:

- ◆ [Editing categories and their attributes, page 218](#)
- ◆ [Filtering based on keyword, page 222](#)
- ◆ [Redefining filtering for specific sites, page 224](#)

In addition to using the more than 90 Websense-defined categories in the Master Database, you can define your own **custom categories** to provide more precise filtering and reporting. For example, create custom categories like:

- ◆ **Business Travel**, to group sites from approved vendors that employees can use to buy airplane tickets and make rental car and hotel reservations
- ◆ **Reference Materials**, to group online dictionary and encyclopedia sites deemed appropriate for elementary school students
- ◆ **Professional Development**, to group training sites and other resources that employees are encouraged to use to build their skills

Use the **Policy Management > Filter Components > Edit Categories > Add Category** page to add custom categories to any parent category. You can create up to 100 custom categories.

1. Enter a unique, descriptive **Category name**. The name cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

2. Enter a **Description** for the new category.

The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. Select a parent category from the **Add to** list. By default, **All Categories** is selected.
4. Enter the sites (URLs or IP addresses) that you want to add to this category. See [Recategorizing URLs, page 226](#), for more information.
You can also edit this list after creating the category.
5. Enter the keywords that you want to associate with this category. See [Filtering based on keyword, page 222](#), for more information.
You can also edit this list after creating the category.

6. Define a default filtering **Action** to apply to this category in all existing category filters. You can edit this action in individual filters later.

**Note**

Category filters copied to a delegated administration role set the filtering action to Permit for custom categories created in the role. Delegated administrators should update the copied category filters to set the desired action for their role-specific custom categories.

7. Enable any **Advanced Filtering** actions (keyword blocking, file type blocking, or bandwidth blocking) that should be applied to this category in all existing category filters.
8. When you are finished defining the new category, click **OK** to cache changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

The new category is added to the Categories list and custom URL and keyword information for the category is displayed.

Filtering based on keyword

Related topics:

- ◆ [Recategorizing URLs, page 226](#)
- ◆ [Configuring Websense filtering settings, page 57](#)
- ◆ [Creating a category filter, page 50](#)
- ◆ [Editing a category filter, page 51](#)
- ◆ [Working with categories, page 218](#)

Keywords are associated with categories, and then used to offer protection against sites that have not explicitly been added to the Master Database or defined as a custom URL. Three steps are necessary to enable keyword blocking:

1. Enable keyword blocking at a global level (see [Configuring Websense filtering settings, page 57](#)).
2. Define keywords associated with a category (see [Defining keywords, page 223](#)).
3. Enable keyword blocking for the category in an active category filter (see [Editing a category filter, page 51](#)).

When keywords have been defined and keyword blocking is enabled for a specific category, Websense software tries to match the keyword against each requested URL as follows:

- ◆ If the keyword contains only ASCII characters, the keyword is matched against the domain, path, and query portions of a URL.
For example, if you associated the keyword “nba” with the permitted Sports category, the following URLs are blocked:
 - sports.espn.go.com/**nba**/
 - modern**nb**akery.com
 - fashion**nb**ar.com
- ◆ If the keyword contains characters outside the ASCII character set (like “ç”, “й”, or “á”), the keyword is matched against only the path and query portions of the string.
For example, if you associated the keyword “fútbol” with the permitted Sports category:
 - “www.**fútbol**.com” is **permitted** (the domain portion of the URL is not matched)
 - “es.wikipedia.org/wiki/**Fútbol**” is **blocked** (the path portion of the URL is matched)

When a site is blocked by keyword, the site is recategorized according to the keyword match. Reports show the keyword category, rather than the Master Database category, for the site.

Be cautious when defining keywords to avoid unintended overblocking.



Important

Avoid associating keywords with any of the Extended Protection subcategories. Keyword blocking is not enforced for these categories.

When a request is blocked based on a keyword, this is indicated on the Websense block page that the user receives.

Defining keywords

Related topics:

- ◆ [Editing a category filter, page 51](#)
- ◆ [Working with categories, page 218](#)
- ◆ [Filtering based on keyword, page 222](#)
- ◆ [Using regular expressions, page 241](#)

A keyword is a string of characters (like a word, phrase, or acronym) that might be found in a URL. Assign keywords to a category, and then enable keyword blocking in a category filter.

Use the **Policy Management > Filter Components > Edit Categories > Add Keywords** page to associate keywords with categories. If you need to make changes to a keyword definition, use the **Edit Keywords** page.

When you define keywords, be cautious to avoid unintended overblocking. You might, for example, intend to use the keyword “sex” to block access adult sites, but end up blocking search engine requests for words like sextuplets or City of Essex, and sites like msexchange.org (Information Technology), vegasexperience.com (Travel), and sci.esa.int/marsexpress (Educational Institutions).

Enter one keyword per line.

- ◆ Do not include spaces in keywords. URL and CGI strings do not include spaces between words.
- ◆ Include a backslash (\) before special characters such as:
 . , # ? * +
If you do not include the backslash, Websense software ignores the special character.
- ◆ Avoid associating keywords with any of the Extended Protection subcategories. Keyword blocking is not enforced for these categories.

When you are finished adding or editing keywords, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

In order for keyword blocking to be enforced, you must also:

1. Enable keyword blocking via the **Settings > Filtering** page (see [Configuring Websense filtering settings, page 57](#)).
2. Enable keyword blocking in one or more active category filters (see [Editing a category filter, page 51](#)).

Redefining filtering for specific sites

Related topics:

- ◆ [Creating a custom category, page 221](#)
- ◆ [Filtering based on keyword, page 222](#)
- ◆ [Defining unfiltered URLs, page 225](#)
- ◆ [Recategorizing URLs, page 226](#)

With custom URLs, you can:

- ◆ Apply more precise filtering to sites that are not in the Websense Master Database. By default, the action applied to the **Miscellaneous/Uncategorized** category is used to filter these sites.
- ◆ Filter sites differently than their Master Database category.

Websense software looks for custom URL definitions for a site before consulting the Master Database, and therefore filters the site according to the category assigned to the custom URL.

There are 2 types of custom URLs: unfiltered and recategorized.

- ◆ Unfiltered URLs are permitted for all users not governed by the Block All category filter or a limited access filter (see [Defining unfiltered URLs, page 225](#)).
- ◆ Recategorized URLs have been moved from their Master Database category to another Websense-defined or custom category (see [Recategorizing URLs, page 226](#)).

A recategorized URL is not blocked by default. It is filtered according to the action applied to its new category in each active category filter.

When a site is filtered according to its Master Database category, Websense software matches the URL with its equivalent IP address. This is not the case for custom URLs. To change the way a site is filtered, define both its URL and its IP address as a custom URL.

If a site can be accessed via multiple URLs, define each URL that can be used to access the site as a custom URL to ensure that the site is permitted or blocked as intended.

If a site is moved to a new domain, and an HTTP redirect is used to send users to the new URL, the new URL is not automatically filtered the same way as the redirecting site. To make sure that the site is filtered appropriately at its new address, create a new custom URL.

Defining unfiltered URLs

Related topics:

- ◆ [Working with categories, page 218](#)
- ◆ [Redefining filtering for specific sites, page 224](#)
- ◆ [Recategorizing URLs, page 226](#)

Use the **Policy Management > Filter Components > Unfiltered URLs** page to define a list of sites that any user can access, except when governed by the Block All category filter or a limited access filter.



Important

If an unfiltered URL becomes infected with malicious code, as long as Security categories are blocked, user requests for that site are blocked.

For instructions to change this behavior, see [Prioritizing Security Risk categorization, page 227](#).

The **Permitted sites** list in the right portion of the content pane lists the unfiltered sites (URLs and IP addresses) and regular expressions that you have defined (see [Using regular expressions, page 241](#)). Each site is associated with a category.

- ◆ The URL may be associated with its Master Database category, or recategorized.
- ◆ When a user requests access to the unfiltered URL, the request is logged as a permitted custom URL in the category to which it has been assigned.

To add an unfiltered URL:

1. Under **Define Unfiltered URLs**, enter one URL or IP address per line, and then click the right arrow (>).

Websense software does not match a custom URL with its equivalent IP address. To permit both the URL and the IP address for a site, add both to the Unfiltered URLs list.

2. To add regular expressions that match multiple sites, click **Advanced**. Enter one regular expression per line, and then click the right arrow to move the expressions to the Unfiltered URLs list. To verify that a pattern matches the intended sites, click **Test**.

See [Using regular expressions, page 241](#), for detailed information.

3. When you are finished, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

To remove a site from the Unfiltered URLs list, select the URL, IP address, or regular expression, and then click **Delete**.

Recategorizing URLs

Related topics:

- ◆ [Working with categories, page 218](#)
- ◆ [Redefining filtering for specific sites, page 224](#)
- ◆ [Defining unfiltered URLs, page 225](#)

Use the **Policy Management > Filter Components > Edit Categories > Recategorize URLs** page to add individual sites to any category. Make changes to existing recategorized sites on the **Edit URLs** page.



Important

If a site is recategorized into a permitted category, and later becomes infected with malicious code, as long as Security categories are blocked, user requests for that site are blocked.

For instructions to change this behavior, see [Prioritizing Security Risk categorization, page 227](#).

Recategorize URLs to change the way that individual sites are filtered and logged. When you add recategorized sites:

- ◆ Enter each URL or IP address on a separate line.
- ◆ Include the protocol for any non-HTTP site. If the protocol is omitted, Websense software filters the site as an HTTP site.

For HTTPS sites, also include the port number (https://63.212.171.196:443/, https://www.onlinebanking.com:443/).

- ◆ Websense software recognizes custom URLs exactly as they are entered. If the Search Engines and Portals category is blocked, but you recategorize **www.yahoo.com** in a permitted category, the site is permitted only if users type the full address. If a user types images.search.yahoo.com, or just yahoo.com, the site is still blocked. If you recategorize **yahoo.com**, however, all sites with yahoo.com in the address are permitted.

When you are finished adding or editing recategorized sites, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

After saving recategorized URLs, use the **URL Category** tool in the right shortcut pane to verify that the site is assigned to the correct category. See [Using the Toolbox to verify filtering behavior](#), page 242.

Prioritizing Security Risk categorization

By default, when a site is categorized in a Security Risk category, the site is filtered based on its Security Risk classification, even when the site:

- ◆ Is added as a recategorized URL in a permitted category
- ◆ Is listed as an unfiltered URL
- ◆ Appears in a limited access filter



Note

Although the Extended Protection categories are default members of the Security Risk class, because they group sites that are still being analyzed, they receive lower prioritization than other categories. As a result, custom categorization always takes precedence over Extended Protection categorization.

When Filtering Service or the hybrid service assigns a site to a Security Risk class category (based on Master Database category or Content Gateway scanning):

- ◆ If a category filter is in effect, and the security-related category is blocked, the site is blocked.
- ◆ If a limited access filter is in effect, the site is blocked.

Configure which categories are part of the Security Risk class on the **Settings > General > Risk Classes** page in TRITON - Web Security.

If you want to always filter based on custom categorization, regardless of whether a site appears in a Security Risk category (like Malicious Web Sites or Spyware):

1. Navigate to the Websense **bin** directory on the Filtering Service machine (see [Where is the Websense "bin" directory?](#), page 457) and open the **eimserver.ini** file in a text editor.
2. Navigate to the **[FilteringManager]** section and add the following line:

```
SecurityCategoryOverride=OFF
```
3. Save and close the file.
4. Restart Filtering Service.
 - *Windows*: Use the Services dialog box (Start > Administrative Tools > Services) to restart Websense Filtering Service.
 - *Linux*: Use the `/opt/Websense/WebsenseDaemonControl` command to stop and then start Filtering Service.

In Websense Web Security Gateway Anywhere environments, you can also disable this feature for hybrid filtering:

1. Navigate to the Websense **bin** directory on the Sync Service machine (see [Where is the Websense "bin" directory?](#), page 457) and open the **syncservice.ini** file in a text editor.
2. If it does not already exist, add a section called **[hybrid]**, and then add the **SecurityCategoryOverride** parameter, as shown here:

```
[hybrid]
SecurityCategoryOverride=false
```
3. Save and close the file.
4. Restart Sync Service.
 - *Windows*: Use the Services dialog box (Start > Administrative Tools > Services) to restart Websense Sync Service.
 - *Linux*: Use the `/opt/Websense/WebsenseDaemonControl` command to stop and then start Sync Service.

Blocking posts to sites in some categories

By default, if users are permitted access to a category, like Message Boards and Forums, they can both view and post to sites in the category.

You can configure Websense software to block posting to sites in specific categories using the **BlockMessageBoardPosts** configuration parameter.

- ◆ If the parameter is set to **ON**, users are blocked from posting only to sites in the Message Boards and Forums category.

- ◆ The parameter can also take a comma-separated list of category identifiers (in the form **112,122,151**). In this case, users are blocked from posting to sites in any of the listed categories.

To enable this feature for on-premises components:

1. Navigate to the Websense **bin** directory on the Filtering Service machine (see [Where is the Websense “bin” directory?](#), page 457) and open the **eimserver.ini** file in a text editor.

2. Navigate to the **[WebsenseServer]** section and add the following line:

```
BlockMessageBoardPosts=<value>
```

Here, <value> can be either **ON** or a comma-separated list of category identifiers

3. Save and close the file.
4. Restart Filtering Service.
 - *Windows:* Use the Services dialog box (Start > Administrative Tools > Services) to restart Websense Filtering Service.
 - *Linux:* Use the `/opt/Websense/WebsenseDaemonControl` command to stop and then start Filtering Service.

In Websense Web Security Gateway Anywhere environments, to enable this feature for hybrid filtering:

1. Navigate to the Websense **bin** directory on the Sync Service machine (see [Where is the Websense “bin” directory?](#), page 457) and open the **syncservice.ini** file in a text editor.

2. If it does not already exist, add a section called **[hybrid]**, and then add the **BlockMessageBoardPosts** parameter, as shown here:

```
[hybrid]
BlockMessageBoardPosts=<value>
```

Here, <value> is a comma-separated list of category identifiers.

3. Save and close the file.
4. Restart Sync Service.
 - *Windows:* Use the Services dialog box (Start > Administrative Tools > Services) to restart Websense Sync Service.
 - *Linux:* Use the `/opt/Websense/WebsenseDaemonControl` command to stop and then start Sync Service.

Working with protocols

The Websense Master database includes protocol definitions used to filter Internet protocols other than HTTP, HTTPS and FTP. These definitions include Internet applications and data transfer methods such as those used for instant messaging, streaming media, file sharing, file transfer, Internet mail, and other network and database operations.

These protocol definitions can even be used to filter protocols or applications that bypass a firewall by tunneling through ports normally used by HTTP traffic. Instant messaging data, for example, can enter a network whose firewall blocks instant messaging protocols by tunneling through HTTP ports. Websense software accurately identifies these protocols, and filters them according to policies you configure.



Note

In Websense Web Filter and Websense Web Security deployments, Network Agent must be installed to enable protocol-based filtering.

With Websense Web Security Gateway, it is possible to filter non-HTTP protocols that tunnel over HTTP ports without using Network Agent. See [Tunneled protocol detection](#), page 160, for more information.

In addition to using Websense-defined protocol definitions, you can define custom protocols for filtering. Custom protocol definitions can be based on IP addresses or port numbers, and can be edited.

To block traffic over a specific port, associate that port number with a custom protocol, and then assign that protocol a default action of **Block**.

To work with custom protocol definitions, go to **Policy Management > Filter Components**, and then click **Protocols**. See [Editing custom protocols](#), page 231, and [Creating a custom protocol](#), page 234, for details.

Filtering protocols

Related topics:

- ◆ [Working with protocols](#), page 229
- ◆ [Editing custom protocols](#), page 231
- ◆ [Creating a custom protocol](#), page 234
- ◆ [Adding or editing protocol identifiers](#), page 232
- ◆ [Adding to a Websense-defined protocol](#), page 235

When Network Agent is installed or with a Websense Web Security Gateway deployment, Websense software can block Internet content transmitted over particular ports, or using specific IP addresses, or marked by certain signatures, regardless of the

nature of the data. By default, blocking a port intercepts all Internet content entering your network over that port, regardless of source.

**Note**

Occasionally, internal network traffic sent over a particular port may not be blocked, even though the protocol using that port is blocked. The protocol may send data via an internal server more quickly than Network Agent can capture and process the data. This does not occur with data originating outside the network.

When a protocol request is made, Websense software uses the following steps to determine whether to block or permit the request:

1. Determine the protocol (or Internet application) name.
2. Identify the protocol based on the request destination address.
3. Search for related port numbers or IP addresses in custom protocol definitions.
4. Search for related port numbers, IP addresses, or signatures in Websense-defined protocol definitions.

If Websense software is unable to determine any of this information, all content associated with the protocol is permitted.

Editing custom protocols

Related topics:

- ◆ [Working with protocols, page 229](#)
- ◆ [Creating a custom protocol, page 234](#)
- ◆ [Creating a protocol filter](#)
- ◆ [Editing a protocol filter](#)
- ◆ [Working with categories](#)

Use the **Policy Management > Filter Components > Edit Protocols** page to create and edit custom protocol definitions, and to review Websense-defined protocol definitions. Websense-defined protocols cannot be edited.

The Protocols list includes all custom and Websense-defined protocols. Click on a protocol or protocol group to get information about the selected item in the right-hand portion of the content pane.

To add a new, custom protocol, click **Add Protocol**, and then continue with [Creating a custom protocol, page 234](#).

To edit a protocol definition:

1. Select the protocol in the Protocols list. The protocol definition appears to the right of the list.
2. Click **Override Action** to change the filtering action applied to this protocol in all protocol filters (see [Making global protocol filtering changes](#), page 233).
3. Click **Add Identifier** to define additional protocol identifiers for this protocol (see [Adding or editing protocol identifiers](#), page 232).
4. Select an identifier in the list, and then click **Edit** to make changes to the **Port**, **IP Address Range**, or **Transport Method** defined by that identifier.
5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

To delete a protocol definition, select an item in the Protocols list, and then click **Delete**.

Adding or editing protocol identifiers

Use the **Filter Components > Edit Protocols > Add Protocol Identifier** page to define additional protocol identifiers for an existing custom protocol. Use the **Edit Protocol Identifier** page to make changes to a previously-defined identifier.

Before creating or changing an identifier, verify that the correct protocol name appears next to **Selected Protocol**.

When working with protocol identifiers, remember that at least one criterion (port, IP address or transport type) must be unique for each protocol.

1. Specify which **Ports** are included in this identifier.
 - If you select **All Ports**, that criterion overlaps with other ports or IP addresses entered in other protocol definitions.
 - Port ranges are not considered unique if they overlap. For example, the port range 80-6000 overlaps with the range 4000-9000.
 - Use caution when defining a protocol on port 80 or 8080. Network Agent listens for Internet requests over these ports.

You can configure Network Agent to ignore these ports in conjunction with a Websense Web Security Gateway deployment.

Since custom protocols take precedence over Websense protocols, if you define a custom protocol using port 80, all other protocols that use port 80 are filtered and logged like the custom protocol.

2. Specify which **IP Addresses** are included in this identifier.
 - If you select **All external IP addresses**, that criterion overlaps with any other IP addresses entered in other protocol definitions.
 - IP address ranges are not considered unique if they overlap.
3. Specify which **Protocol Transport** method is included in this identifier.
4. Click **OK** to cache your changes and return to the Edit Protocols page. Changes are not implemented until you click **Save All**.

Renaming a custom protocol

Use the **Filter Components > Edit Protocols > Rename Protocol** page to change the name of a custom protocol, or move it to a different protocol group.

- ◆ Use the **Name** field to edit the protocol name. The new name cannot exceed 50 characters.

The name cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

- ◆ To move the protocol to a different protocol group, select the new group from the **In group** field.

When you are finished making changes, click **OK** to return to the Edit Protocols page. You must also click **OK** on the Edit Protocols to cache the changes.

Making global protocol filtering changes

Use the **Filter Components > Edit Protocols > Override Action** page to change the way a protocol is filtered in all existing protocol filters. This also determines the default action applied to the protocol in new filters.

Although this change overrides the filtering action applied in all existing protocol filters, administrators can later edit those filters to apply a different action.

1. Verify that the correct protocol name appears next to **Selected Protocol**.
2. Select a new **Action** (Permit or Block) to apply to this protocol. By default, **No change** is selected. See [Filtering actions, page 46](#), for more information.
3. Specify new **Logging** options. Protocol traffic must be logged to appear in reports and enable protocol usage alerts.
4. Specify whether or not **Bandwidth Optimizer** is used to manage access to this protocol. See [Using Bandwidth Optimizer to manage bandwidth, page 236](#), for more information.



Important

Changes made here affect every existing protocol filter, except **Block All** and **Permit All**.

5. When you are finished, click **OK** to return to the Edit Protocols page (see [Editing custom protocols, page 231](#)). You must also click **OK** on the Edit Protocols page to cache the changes.

Creating a custom protocol

Related topics:

- ◆ [Working with protocols](#), page 229
- ◆ [Filtering protocols](#), page 230
- ◆ [Editing custom protocols](#), page 231
- ◆ [Adding to a Websense-defined protocol](#), page 235

Use the **Filter Components > Protocols > Add Protocol** page to define a new, custom protocol.

1. Enter a **Name** for the protocol.

The name cannot include any of the following characters:

* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

A custom protocol can be assigned the same name as a Websense-defined protocol, in order to extend the number of IP addresses or ports associated with the original protocol. See [Adding to a Websense-defined protocol](#), page 235, for more information.

2. Expand the **Add protocol to this group** drop-down list, and then select a protocol group. The new protocol appears in this group in all protocol lists and filters.
3. Define a unique **Protocol Identifier** (set of **ports**, **IP addresses**, and **transport methods**) for this group. You can add additional identifiers later, from the Edit Protocols page.

Follow these guidelines for creating protocol identifiers:

- At least one criterion (port, IP address or transport type) must be unique for each protocol definition.
- If you select **All Ports** or **All external IP addresses**, that criterion overlaps with any other ports or IP addresses entered in other protocol definitions.
- Port ranges or IP address ranges are not considered unique if they overlap. For example, the port range 80-6000 overlaps with the range 4000-9000.



Note

Use caution when defining a protocol on port 80 or 8080. Network Agent listens for Internet requests over these ports. (In Websense Web Security Gateway deployments, you can configure Network Agent to ignore these ports.)

Since custom protocols take precedence over Websense protocols, if you define a custom protocol using port 80, all other protocols that use port 80 are filtered and logged like the custom protocol.

The following tables provide examples of valid and invalid protocol definitions:

Port	IP Address	Transport Method	Accepted combination?
70	ANY	TCP	Yes - the port number makes each protocol identifier unique.
90	ANY	TCP	

Port	IP Address	Transport Method	Accepted combination?
70	ANY	TCP	No - the IP addresses are not unique. 10.2.1.201 is included in the "ANY" set.
70	10.2.1.201	TCP	

Port	IP Address	Transport Method	Accepted combination?
70	10.2.3.212	TCP	Yes - the IP addresses are unique.
70	10.2.1.201	TCP	

4. Under Default Filtering Action, specify the default action (**Permit** or **Block**) that should be applied to this protocol in all active protocol filters:
 - Indicate whether traffic using this protocol should be **Logged**. Protocol traffic must be logged to appear in reports and enable protocol usage alerts.
 - Indicate whether access to this protocol should be regulated by **Bandwidth Optimizer** (see [Using Bandwidth Optimizer to manage bandwidth, page 236](#)).
5. When you are finished, click **OK** to return to the Edit Protocols page. The new protocol definition appears in the Protocols list.
6. Click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Adding to a Websense-defined protocol

You cannot add a port number or IP address directly to a Websense-defined protocol. You can, however, create a custom protocol with the same name as the Websense-defined protocol, and then add ports or IP addresses to its definition.

When a custom protocol and a Websense-defined protocol have the same name, Websense software looks for protocol traffic at the ports and IP addresses specified in both definitions.

In reports, custom protocol names have a "C_" prefix. For example, if you created a custom protocol for SQL_NET and specified additional port numbers, reports display C_SQL_NET when the protocol uses the port numbers in the custom protocol.

Using Bandwidth Optimizer to manage bandwidth

Related topics:

- ◆ [Working with categories, page 218](#)
- ◆ [Working with protocols, page 229](#)
- ◆ [Configuring the default Bandwidth Optimizer limits, page 237](#)

When you create a category or protocol filter, you can elect to limit access to a category or protocol based on bandwidth usage.

- ◆ Block access to categories or protocols based on total network bandwidth usage.
- ◆ Block access to categories based on total bandwidth usage by HTTP traffic.
- ◆ Block access to a specific protocol based on bandwidth usage by that protocol.



Note

If you have Websense Web Security Gateway Anywhere, be aware that hybrid filtering does not enforce bandwidth-based restrictions.

For example:

- ◆ Block the AOL Instant Messaging protocol if total network bandwidth usage exceeds 50% of available bandwidth, or if current bandwidth usage for AIM exceeds 10% of the total network bandwidth.
- ◆ Block the Sports category when total network bandwidth usage reaches 75%, or when bandwidth usage by all HTTP traffic reaches 60% of available network bandwidth.

Protocol bandwidth usage includes traffic over all ports, IP addresses, or signatures defined for the protocol. This means that if a protocol or Internet application uses multiple ports for data transfer, traffic across all of the ports included in the protocol definition are counted toward that protocol's bandwidth usage total. If an Internet application uses a port not included in the protocol definition, however, traffic over that port is not included in bandwidth usage measurements.

Websense software records bandwidth used by filtered TCP- and UDP-based protocols.

Websense, Inc., updates Websense protocol definitions regularly to ensure bandwidth measurement accuracy.

Network Agent sends network bandwidth data to Filtering Service at a predetermined interval. This ensures that Websense software accurately monitors bandwidth usage, and receives measurements that are closest to an average.

In a Websense Web Security Gateway deployment, the Content Gateway collects data about bandwidth used by HTTP traffic and protocols that tunnel over HTTP. You can specify that this data be used to determine bandwidth-based protocol filtering in the Bandwidth Optimizer settings.

1. In TRITON - Web Security, go to **Settings > General > Filtering**.
2. Select the **Bandwidth Monitoring** check box.
3. When you are finished, click **OK** to cache your change. Changes are not implemented until you click **Save All**.

When bandwidth-based filtering options are active, Websense software begins bandwidth-based filtering 10 minutes after initial configuration, and 10 minutes after each Websense Policy Server restart. This delay ensures accurate measurement of bandwidth data and use of this data in filtering.

When a request is blocked based on bandwidth limitations, the Websense block page displays this information in the **Reason** field. For more information, see [Block Pages, page 91](#).

Configuring the default Bandwidth Optimizer limits

Related topics:

- ◆ [Editing a category filter, page 51](#)
- ◆ [Editing a protocol filter, page 53](#)
- ◆ [Using Bandwidth Optimizer to manage bandwidth, page 236](#)

Before specifying bandwidth settings in policies, verify the default bandwidth thresholds that trigger bandwidth-based filtering settings. The Websense-defined values are:

- ◆ Default bandwidth for network: **50%**
- ◆ Default bandwidth per protocol: **20%**

Default bandwidth values are stored by Policy Server, and enforced by all associated instances of Network Agent.

To change the default bandwidth values:

1. In TRITON - Web Security, go to **Settings > General > Filtering**.
2. Enter the bandwidth usage thresholds that will trigger bandwidth-based filtering, when bandwidth filtering is enabled.
 - When a category or protocol is blocked based on traffic for the entire network, **Default bandwidth for network** defines the default filtering threshold.
 - When a category or protocol is blocked based on traffic for the protocol, the **Default bandwidth per protocol** defines the default filtering threshold.

You can override the default threshold values for each category or protocol in any category or protocol filter.

3. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Any changes to the defaults have the potential to affect any category and protocol filters that enforce Bandwidth Optimizer restrictions.

- ◆ To manage bandwidth usage associated with a particular protocol, edit the active protocol filter or filters.
- ◆ To manage bandwidth usage associated with a particular URL category, edit the appropriate category filter or filters.

When you filter categories based on HTTP bandwidth usage, Websense software measures total HTTP bandwidth usage over all ports specified as HTTP ports for Websense software.

Managing traffic based on file type

When you create a category filter, you can define filtering based on file extensions, restricting access to particular file types from sites in certain categories. For example, permit the category Sports, but block video files from sites in the Sports category.



Note

To implement full filtering for video and audio Internet media, combine protocol-based filtering with file type filtering. In this case, protocol filtering handles streaming media, while file type filtering handles files that can be downloaded and then played.

Websense software provides several predefined file types, or groupings of file extensions used for similar purposes. These file type definitions are maintained in the Master Database, and may be changed as part of the Master Database update process.

You can implement filtering using predefined file types, modify the existing file type definitions, or create new file types. Note, however, that you cannot delete Websense-defined file types, or delete the file extensions associated with them.

When a user requests a site, Websense software first determines the site category, and then checks for filtered file extensions.



Note

When multiple group policies could apply to a user request, file type blocking is not performed.

When a user tries to access a file whose extension is blocked, the **Reason** field on the Websense block page indicates that the file type was blocked. For more information, see [Block Pages, page 91](#).

The standard block page is not displayed if a blocked GIF or JPEG image comprises just a portion of a permitted page. Instead, the image region appears blank. This avoids the possibility of displaying a small portion of a block page in multiple locations on an otherwise permitted page.

File type definitions may contain as many or as few file extensions as are useful for filtering purposes. Websense-defined file types, for example, include the following file extensions:

Audio	Compressed Files		Executables	Video	
.aif	.ace	.mim	.bat	.asf	.mpg
.aifc	.arc	.rar	.exe	.asx	.mpv2
.aiff	.arj	.tar		.avi	.qt
.m3u	.b64	.taz		.ivf	.ra
.mid	.bhx	.tgz		.mlv	.ram
.midi	.cab	.tz		.mov	.wm
.mp3	.gz	.uu		.mp2	.wmp
.ogg	.gzip	.uue		.mp2v	.wmv
.rmi	.hqx	.xxe		.mpa	.wmx
.snd	.iso	.z		.mpe	.wxv
.wav	.jar	.zip			
.wax	.lzh				
.wma					

Any of the file extensions associated with a Websense-defined file type can be added to a custom file type. The file extension is then filtered and logged according to the settings associated with the custom file type.

To view existing file type definitions, edit file types, or create custom file types, go to **Policy Management > Filter Components**, and then click **File Types**. See [Working with file types, page 239](#), for more information.

Working with file types

Related topics:

- ◆ [Managing traffic based on file type, page 238](#)
- ◆ [Editing a category filter, page 51](#)
- ◆ [Filtering a site, page 86](#)

Use the **Policy Management > Filter Components > Edit File Types** page to create and manage up to 32 **file types**. File types are groups of file extensions that can be explicitly blocked in category filters (see *Managing traffic based on file type*, page 238).

- ◆ Click on a file type to see the file extensions associated with that type.
- ◆ To add extensions to the selected file type, click **Add Extension**, and then see *Adding file extensions to a file type*, page 240, for further instructions.
- ◆ To create a new file type, click **Add File Type**, and then see *Adding custom file types*, page 240, for further instructions.
- ◆ To delete a custom file type or extension, select an item, and then click **Delete**.
You cannot delete Websense-defined file types, or delete the file extensions associated with them.
You can, however, add file extensions associated with a Websense-defined file type to a custom file type. The file extension is then filtered and logged according to the settings associated with the custom file type. You cannot add the same extension to multiple custom file types.

When you are finished making changes to file type definitions, click **OK**. Changes are not implemented until you click **Save All**.

Adding custom file types

Use the **Filter Components > Edit File Types > Add File Type** page to define custom file types.

1. Enter a unique **File type name**.
You can create a custom file type with the same name as a Websense-defined file type in order to add additional file extensions to the existing file type.
2. Enter file extensions, one per line, in the **File extensions** list. You do not need to include the dot (“.”) before each extension.
3. Click **OK** to return to the Edit File Types screen. The new file type appears in the File Types list.
4. When you are finished working with file type definitions, click **OK** on the Edit File Types page. Changes are not implemented until you click **Save All**.

Adding file extensions to a file type

Use the **Filter Components > Edit File Types > Add File Extensions** page to add file extensions to the selected file type.

1. Verify that the expected file type name appears next to **Selected file type**.
2. Enter file extensions, one per line, in the **File extensions** list. You do not need to include the dot (“.”) before each extension.
3. Click **OK** to return to the Edit File Types screen. The new file extensions appear in the Custom file extensions list.

4. When you are finished working with file type definitions, click **OK** on the Edit File Types page. Changes are not implemented until you click **Save All**.

Using regular expressions

A **regular expression** is a template or pattern used to match multiple strings, or groups of characters. You can use regular expressions in limited access filters, or to define custom URLs or keywords. Filtering Service then tries to match the general pattern, rather than a specific, single URL or keyword.

Consider this simple regular expression:

```
domain.(com|org|net)
```

This expression pattern matches the URLs:

- ◆ domain.com
- ◆ domain.org
- ◆ domain.net

Use regular expressions with caution. They provide a powerful filtering tool, but can easily result in either unexpected over blocking or under blocking. Also, poorly constructed regular expressions can result in excessive filtering overhead.



Important

Using regular expressions as filtering criteria may increase CPU usage. Tests have shown that with 100 regular expressions, the average CPU usage on the Filtering Service machine increased by 20%.

As with keywords, when non-ASCII characters appear in a regular expression, the expression is matched against only the path and query strings in a URL, and not the domain (“www.domain.com/**path?query**”).

Websense software supports most Perl regular expression syntax, with 2 exceptions. The unsupported syntax is unlikely to be useful for matching strings that could be found in a URL.

Unsupported regular expression syntax includes:

```
(?{code})
```

```
??{code})
```

For further help with regular expressions, see:

en.wikipedia.org/wiki/Regular_expression
www.regular-expressions.info/

Using the Toolbox to verify filtering behavior

The right shortcut pane in TRITON - Web Security includes a **Toolbox** that allows you to perform quick checks of your filtering setup.

Click a tool name to access the tool. Click the name again to see the list of tools. For more information about using a tool, see:

- ◆ [URL Category, page 242](#)
- ◆ [Check Policy, page 242](#)
- ◆ [Test Filtering, page 243](#)
- ◆ [URL Access, page 243](#)
- ◆ [Investigate User, page 243](#)

You can also click **Support Portal** to access the Websense Technical Support Web site in a new browser tab or window. From the Support Portal, you can use the Knowledge Base to access tutorials, tips, articles, and documentation.

URL Category

To find out how a site is currently categorized:

1. Click **URL Category** in the Toolbox.
2. Enter a URL or IP address.
3. Click **Go**.

The site's current category is displayed in a popup window. If your organization has recategorized the URL, the new category is shown.

The site's categorization may depend on which version of the Master Database (including real-time updates) you are using.

Check Policy

Use this tool to determine which policies apply to a specific client. The results are specific to the current day and time.

1. Click **Check Policy** in the Toolbox.
2. To identify a directory or computer client, enter either:
 - A fully qualified user name
To browse or search the directory to identify the user, click **Find User** (see [Identifying a user to check policy or test filtering, page 244](#)).
 - An IP address
3. Click **Go**.

The name of one or more policies is displayed in a popup window. Multiple policies are displayed only when no policy has been assigned to the user, but policies have

been assigned to multiple groups, domains, or organizational units to which the user belongs.

Even if multiple policies are shown, only one policy is enforced for a user at any given time (see [Filtering order](#), page 85).

Test Filtering

To find out what happens when a specific client requests a particular site:

1. Click **Test Filtering** in the Toolbox.
2. To identify a directory or computer client, enter either:
 - A fully qualified user name
To browse or search the directory to identify the user, click **Find User** (see [Identifying a user to check policy or test filtering](#), page 244).
 - An IP address
3. Enter the URL or IP address of the site you want to check.
4. Click **Go**.

The site category, the action applied to the category, and the reason for the action are displayed in a popup window.

URL Access

To see whether users have attempted to access a site in the past 2 weeks, including today:

1. Click **URL Access** in the Toolbox.
2. Enter all or part of the URL or IP address of the site you want to check.
3. Click **Go**.

An investigative report shows whether the site has been accessed, and if so, when.

You might use this tool after receiving a security alert to find out if your organization has been exposed to phishing or virus-infected sites.

Investigate User

To review a client's Internet usage history for the last 2 weeks, excluding today:

1. Click **Investigate User** in the Toolbox.
2. Enter all or part of a user name or computer IP address.
3. Click **Go**.

An investigative report shows the client's usage history.

Identifying a user to check policy or test filtering

Use the **Find User** page to identify a user (directory) client for the Check Policy or Test Filtering tool.

The page opens with the **User** option selected. Expand the **Directory Entries** folder to browse the directory, or click **Search**. The search feature is available only if you are using an LDAP-based directory service.

To search the directory to find a user:

1. Enter all or part of the user **Name**.
2. Expand the **Directory Entries** tree and browse to identify a search context.
You must click a folder (DC, OU, or CN) in the tree to specify the context. This populates the field below the tree.
3. Click **Search**. Entries matching your search term are listed under **Search Results**.
4. Click a user name to select a user, or click **Search Again** to enter a new search term or context.
To return to browsing the directory, click **Cancel Search**.
5. When the correct fully qualified user name appears in the **User** field, click **Go**.

If you are using the Test Filtering tool, make sure that a URL or IP address appears in the **URL** field before you click **Go**.

To identify a computer client instead of a user, click **IP address**.

12

User Identification

To apply policies to users and groups, Websense software must be able to identify the user making a request, given the originating IP address. Various identification methods are available:

- ◆ An integration device or application identifies and authenticates users, and then passes user information to Websense software. For more information, see the *Deployment and Installation Center*.
- ◆ A Websense transparent identification agent works in the background to communicate with a directory service and identify users (see [Transparent identification](#)).
- ◆ Websense software prompts users for their network credentials, requiring them to log on when they open a Web browser (see [Manual authentication, page 247](#)).

In Websense Web Security Gateway Anywhere environments, the hybrid service must likewise be able to identify users to apply user and group based policies. It does not use information provided by User Service or transparent identification agents. Instead, the following methods are available:

- ◆ A component called Websense Directory Agent collects the information used to identify users (see [Identification of hybrid filtering users, page 276](#)).
- ◆ Websense Web Endpoint is installed on client machines to provide transparent authentication, enforce use of hybrid filtering, and pass authentication details to the hybrid service.
- ◆ Websense Authentication Service provides transparent authentication via a virtual machine hosted on your network that communicates with your directory service.

Transparent identification

Related topics:

- ◆ [Manual authentication, page 247](#)
- ◆ [Configuring user identification methods, page 248](#)

In general, **transparent identification** describes any method that Websense software uses to identify users in your directory service without prompting them for logon information. This includes integrating Websense software with a device or application that provides user information for use in filtering, or using optional Websense transparent identification agents.

- ◆ Websense *DC Agent*, [page 255](#), is used with a Windows-based directory service. The agent periodically queries domain controllers for user logon sessions and polls client machines to verify logon status. It runs on a Windows server and can be installed in any domain in the network.
- ◆ Websense *Logon Agent*, [page 258](#), identifies users as they log on to Windows domains. The agent runs on a Linux or Windows server, but its associated Logon Application runs only on Windows machines.
- ◆ Websense *RADIUS Agent*, [page 261](#), can be used in conjunction with either Windows- or LDAP-based directory services. The agent works with a RADIUS server and client to identify users logging on from remote locations.
- ◆ Websense *eDirectory Agent*, [page 265](#), is used with Novell eDirectory. The agent uses Novell eDirectory authentication to map users to IP addresses.

For instructions on installing each agent, see the *Deployment and Installation Center*. Agent can be used alone, or in certain combinations (see [Configuring multiple agents](#), [page 271](#)).

Both general user identification settings and specific transparent identification agents are configured in TRITON - Web Security. Go to the **Settings > General > User Identification** page.

See [Configuring user identification methods](#), [page 248](#), for detailed configuration instructions.

In some instances, Websense software may not be able to obtain user information from a transparent identification agent. This can occur if more than one user is assigned to the same machine, or if a user is an anonymous user or guest, or for other reasons. In these cases, you can prompt the user to log on via the browser (see [Manual authentication](#), [page 247](#)).

Transparent identification of remote users

In certain configurations, Websense software can transparently identify users logging on to your network from remote locations:

- ◆ If you have deployed the Websense Remote Filtering Server and Remote Filtering Client, Websense software can identify any off-site user logging on to a cached domain using a domain account. For more information, see [Filter Users Off Site](#), [page 201](#).
- ◆ If you have deployed DC Agent, and remote users directly log on to named Windows domains in your network, DC Agent can identify these users (see [DC Agent](#), [page 255](#)).

- ◆ If you are using a RADIUS server to authenticate users logging on from remote locations, RADIUS Agent can transparently identify these users so you can apply filtering policies based on users or groups (see [RADIUS Agent, page 261](#)).

Manual authentication

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Setting authentication rules for specific machines, page 249](#)
- ◆ [Secure manual authentication, page 252](#)
- ◆ [Configuring user identification methods, page 248](#)

Transparent identification is not always available or desirable in all environments. For organizations that do not use transparent identification, or in situations when transparent identification is not available, you can still filter based on user and group-based policies using **manual authentication**.

Manual authentication prompts users for a user name and password the first time they access the Internet through a browser. Websense software confirms the password with a supported directory service, and then retrieves policy information for that user.

You can configure Websense software to enable manual authentication any time transparent identification is not available (see [Configuring user identification methods, page 248](#) and [Configure user access to hybrid filtering, page 183](#)).

You can also create a list of specific machines with custom authentication settings on which users are prompted to log on when they open a browser (see [Setting authentication rules for specific machines, page 249](#)).

When manual authentication is enabled, users may receive HTTP errors and be unable to access the Internet if:

- ◆ They make 3 failed attempts to enter a password. This occurs when the user name or password is invalid.
- ◆ They click **Cancel** to bypass the authentication prompt.

When manual authentication is enabled, users who cannot be identified are prevented from browsing the Internet.

Configuring user identification methods

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Working with users and groups, page 64](#)

Use the **Settings > General > User Identification** page to manage when and how Websense software attempts to identify users in the network in order to apply user- and group-based policies.

- ◆ Configure Policy Server to communicate with transparent identification agents.
- ◆ Review and update transparent identification agent settings.
- ◆ Set a global rule to determine how Websense software responds when users cannot be identified by a transparent identification agent or integration device.
- ◆ Identify machines in your network to which global user identification rules do not apply, and specify whether and how users of those machines should be authenticated.

If you are using Websense transparent identification agents, the agents are listed under **Transparent Identification Agents**:

- ◆ **Server** shows the IP address or name of the machine hosting the transparent identification agent.
- ◆ **Port** lists the port that Websense software uses to communicate with the agent.
- ◆ **Type** indicates whether the specified instance is a DC Agent, Logon Agent, RADIUS Agent, or eDirectory Agent. (See [Transparent identification, page 245](#), for an introduction to each type of agent.)

To add an agent to the list, select the agent type from **Add Agent** drop-down list. Click one of the following links for configuration instructions:

- ◆ [Configuring DC Agent, page 256](#)
- ◆ [Configuring Logon Agent, page 259](#)
- ◆ [Configuring RADIUS Agent, page 263](#)
- ◆ [Configuring eDirectory Agent, page 267](#)

To remove an agent instance from the list, mark the checkbox next to the agent information in the list, and then click **Delete**.

Under **Additional Authentication Options**, specify the default response of Websense software when users are not identified transparently (by an agent or integration):

- ◆ Click **Apply computer or network policy** to ignore user and group-based policies in favor of computer and network-based policies, or the Default policy.

- ◆ Click **Prompt user for logon information** to require users to provide logon credentials when they open a browser. User and group-based policies can then be applied (see [Manual authentication, page 247](#)).
- ◆ Specify the default domain **Context** that Websense software should use any time a user is prompted for log on credentials. This is the domain in which users' credentials are valid.

If you use the Exceptions list to specify any machines on which users are prompted for logon information, you must provide a default domain context, even if the global rule is to apply a computer or network-based policy.

After establishing the general rule that determines when and how users are identified by Websense software, you can create exceptions to the rule.

For example, if you use a transparent identification agent or integration product to identify users, and have enabled manual authentication to prompt users for their credentials when they cannot be identified transparently, you can identify specific machines on which:

- ◆ Users who cannot be identified are never be prompted for their credentials. In other words, when transparent identification fails, manual authentication is not attempted, and the computer or network policy, or the Default policy, is applied.
- ◆ User information is always ignored, even when it is available, and users are always prompted for their credentials.
- ◆ User information is always ignored, even when it is available, and users are never prompted for their credentials (the computer or network policy, or the Default policy, is always applied).

To create an exception, click **Exceptions**, and then see [Setting authentication rules for specific machines, page 249](#).

When you are finished making changes on this page, click **OK** to save your changes. To avoid saving changes, click **Cancel**.

Setting authentication rules for specific machines

Related topics:

- ◆ [Configuring user identification methods, page 248](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Secure manual authentication, page 252](#)

Selective authentication lets you determine whether users requesting Internet access from a specific client machine (identified by IP address) are prompted to provide their logon credentials via the browser. This can be used to:

- ◆ Establish different authentication rules for a machine in a public kiosk than for employees of the organization supplying the kiosk.

- ◆ Ensure that users of an exam-room computer in a medical office are always identified before getting Internet access.

Machines with special user identification settings applied are listed on the **Settings > User Identification** page. Click **Exceptions** to establish specific user identification settings for some machines in your network, or see if special settings have been defined for a specific machine.

To add a machine to the list, click **Add**, and then see [Defining exceptions to user identification settings, page 250](#), for further instructions.

When you are finished adding machines or network ranges to the list, click **OK**. Changes are not implemented until you click **Save All**.

Defining exceptions to user identification settings

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Configuring user identification methods, page 248](#)

Use the **User Identification > Add IP Addresses** page to identify machines to which specific user identification rules should be applied.

1. Enter an **IP address** or **IP address range** to identify machines to which to apply a specific authentication method, and then click the right-arrow button to add them to the **Selected** list.

If the same rules should be applied to multiple machines, add them all to the list.

2. Select an entry in the **User identification** drop-down list to indicate whether Websense software should attempt to identify users of these machines transparently.
 - Select **Try to identify user transparently** to request user information from a transparent identification agent or integration device.
 - Select **Ignore user information** to avoid using any transparent method to identify users.
3. Indicate whether users should be prompted to provide logon credentials via the browser. This setting applies when user information is not available, either because other identification failed, or because user information was ignored.
 - Select **Prompt user for logon information** to require users to provide logon credentials.

If **Try to identify user transparently** is also selected, users receive a browser prompt only if they are not identified transparently.
 - Select **Apply computer or network policy** to ensure that users are never required to provide logon credentials.

If **Try to identify user transparently** is also selected, users whose credentials can be verified transparently are filtered by the appropriate user-based policy.

4. Click **OK** to return to the User Identification page.
5. When you are finished updating the Exceptions list, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Revising exceptions to user identification settings

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Configuring user identification methods, page 248](#)

Use the **Settings > User Identification > Edit IP Addresses** page to make changes to entries in the Exceptions list. Changes made on this page affect all machines (identified by IP address or range) that appear in the Selected list.

1. Select an entry in the **User identification** drop-down list to indicate whether Websense software should attempt to identify users of these machines transparently.
 - Select **Try to identify user transparently** to request user information from a transparent identification agent or integration device.
 - Select **Ignore user information** to avoid using any transparent method to identify users.
2. Indicate whether users should be prompted to provide logon credentials via the browser. This setting applies when user information is not available, either because transparent identification failed, or because transparent identification was ignored.
 - Select **Prompt user for logon information** to require users to provide logon credentials.

If **Try to identify user transparently** is also selected, users receive a browser prompt only if they are not identified transparently.
 - Select **Apply computer or network policy** to ensure that users are never prompted to provide logon credentials.

If **Try to identify user transparently** is also selected, users whose credentials can be verified transparently are filtered by the appropriate user-based policy.
3. Click **OK** to return to the User Identification page.
4. When you are finished updating the Exceptions list, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Secure manual authentication

Related topics:

- ◆ [Configuring user identification methods](#), page 248
- ◆ [Manual authentication](#), page 247
- ◆ [Setting authentication rules for specific machines](#), page 249
- ◆ [Activating secure manual authentication](#), page 253

Websense secure manual authentication uses Secure Sockets Layer (SSL) encryption to protect authentication data being transmitted between client machines and Websense software. An SSL server built into Filtering Service provides encryption of user names and passwords transmitted between client machines and Filtering Service. By default, secure manual authentication is disabled.



Note

Secure manual authentication cannot be used with remote filtering software. The Remote Filtering Server can not serve block pages to clients if it is associated with a Filtering Service instance that has secure manual authentication enabled.

To enable this functionality, you must perform the following steps:

1. Generate SSL certificates and keys, and place them in a location accessible by Websense software and readable by Filtering Service (see [Generating keys and certificates](#), page 252).
2. Enable secure manual authentication (see [Activating secure manual authentication](#), page 253) and secure communication with the directory service.
3. Import certificates into the browser (see [Accepting the certificate within the client browser](#), page 254).

Generating keys and certificates

Related topics:

- ◆ [Manual authentication](#), page 247
- ◆ [Setting authentication rules for specific machines](#), page 249
- ◆ [Secure manual authentication](#), page 252
- ◆ [Activating secure manual authentication](#), page 253
- ◆ [Accepting the certificate within the client browser](#), page 254

A certificate consists of a public key, used to encrypt data, and a private key, used to decipher data. Certificates are issued by a Certificate Authority (CA). You can generate a certificate from an internal certificate server, or obtain a client certificate from any third-party CA, such as VeriSign.

The CA issuing the client certificate must be trusted by Websense software. Typically, this is determined by a browser setting.

- ◆ For answers to common questions about private keys, CSRs, and certificates, see http://httpd.apache.org/docs/2.2/ssl/ssl_faq.html#aboutcerts.
- ◆ To learn more about generating your own private key, CSR, and certificate, see www.akadia.com/services/ssh_test_certificate.html.

There are many tools that you can use to generate a self-signed certificate, including the OpenSSL toolkit (available from openssl.org).

Regardless of the method you choose for generating the certificate, use the following general steps.

1. Generate a private key (**server.key**).
2. Generate a Certificate Signing Request (CSR) with the private key.



Important

When prompted for the CommonName, enter the IP address of the Filtering Server machine. If you skip this step, client browsers will display a security certificate error.

3. Use the CSR to create a self-signed certificate (**server.crt**).
4. Save the **server.crt** and **server.key** files in a location that Websense software can access, and where they can be read by Filtering Service.

Activating secure manual authentication

Related topics:

- ◆ [Manual authentication, page 247](#)
- ◆ [Setting authentication rules for specific machines, page 249](#)
- ◆ [Secure manual authentication, page 252](#)
- ◆ [Generating keys and certificates, page 252](#)
- ◆ [Accepting the certificate within the client browser, page 254](#)

1. Stop Websense Filtering Service (see [Stopping and starting Websense services, page 329](#)).
2. Navigate to the Websense installation directory on the Filtering Service machine (by default, **C:\Program Files\Websense\bin** or **/opt/Websense/bin/**).

3. Locate **eimserver.ini** and make a backup copy of the file in another directory.
4. Open the original INI file in a text editor.
5. Find the **[WebsenseServer]** section, and then add the line:

```
SSLManualAuth=on
```

6. Below the previous line, add the following:

```
SSLCertFileLoc=[path]
```

Replace **[path]** with the full path to the SSL certificate, including the certificate file name (for example, C:\secmanauth\server.crt).

7. Also add:

```
SSLKeyFileLoc=[path]
```

Replace **[path]** with the full path to the SSL key, including the key file name (for example, C:\secmanauth\server.key).

8. Save and close **eimserver.ini**.
9. Start Websense Filtering Service.

After starting, Filtering Service listens for requests on the default secure HTTP port (**15872**).

The preceding steps ensure secure communication between the client machine and Websense software. To also secure communication between Websense software and the directory service, make sure that **Use SSL** is selected on the **Settings > Directory Services** page. See [Advanced directory settings, page 68](#), for details.

Accepting the certificate within the client browser

Related topics:

- ◆ [Manual authentication, page 247](#)
- ◆ [Setting authentication rules for specific machines, page 249](#)
- ◆ [Secure manual authentication, page 252](#)
- ◆ [Generating keys and certificates, page 252](#)
- ◆ [Activating secure manual authentication, page 253](#)

The first time you try to browse to a Web site, the browser will display a warning about the security certificate. To avoid seeing this message in the future, install the certificate in the certificate store.

Microsoft Internet Explorer

1. Open the browser and go to a Web site.
A warning appears, stating that there is a problem with the site's security certificate.
2. Click **Continue to this website (not recommended)**.
If you receive an authentication prompt, click **Cancel**.

3. Click the **Certificate Error** box to the right of the address bar (at the top of the browser window), and then click **View certificates**.
4. On the General tab of the Certificate dialog box, click **Install Certificate**.
5. Select **Automatically select the certificate store based on the type of certificate**, and then click **Next**.
6. Click **Finish**.
7. When asked whether to install the certificate, click **Yes**.

Users will no longer receive certificate security warnings related to Filtering Service on this machine.

Mozilla Firefox (Version 3.x and later)

1. Open the browser and go to a Web site.
A warning message appears.
2. Click **Or you can add an exception**.
3. Click **Add Exception**.
4. Make sure that **Permanently store this exception is selected**, and then click **Confirm Security Exception**.

Users will no longer receive certificate security warnings related to Filtering Service on this machine.

DC Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Configuring DC Agent, page 256](#)
- ◆ [Configuring different settings for an agent instance, page 272](#)

Websense DC Agent runs on Windows and detects users in a Windows network running NetBIOS, WINS, or DNS networking services.

DC Agent and User Service gather network user data and send it to Websense Filtering Service. Several variables determine the speed of data transmission, including the size of your network and the amount of existing network traffic.

To enable transparent identification with DC Agent:

1. Install DC Agent. For more information, see the *Deployment and Installation Center*.



Note

Run DC Agent using domain administrator privileges. The domain admin account must also be a member of the Administrators group on the DC Agent machine.

This is required for DC Agent to retrieve user logon information from the domain controller. If you cannot install DC Agent with such privileges, configure administrator privileges for these services after installation. For more information, see [Websense software is not applying user or group policies, page 399](#).

2. Configure DC Agent to communicate with other Websense components and with domain controllers in your network (see [Configuring DC Agent](#)).
3. Use TRITON - Web Security to add users and groups to filter (see [Adding a client, page 71](#)).

Websense software can prompt users for identification if DC Agent is unable to identify users transparently. For more information, see [Manual authentication, page 247](#).

Configuring DC Agent

Related topics:

- ◆ [Transparent identification](#)
- ◆ [Manual authentication](#)
- ◆ [Configuring user identification methods](#)
- ◆ [DC Agent](#)
- ◆ [Configuring multiple agents](#)

Use the **User Identification > DC Agent** page to configure a new instance of DC Agent, as well as to configure the global settings that apply to all instances of DC Agent.

To add a new instance of DC Agent, first provide basic information about where the agent is installed, and how Filtering Service should communicate with it. These settings may be unique to each agent instance.

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

**Note**

Machine names must start with an alphabetical character (a-z), not a numeric or special character.

Machine names containing certain extended ASCII characters may not resolve properly. If you are using a non-English version of Websense software, enter an IP address instead of a machine name.

2. Enter the **Port** that DC Agent should use to communicate with other Websense components. The default is 30600.
3. To establish an authenticated connection between Filtering Service and DC Agent, check **Enable Authentication**, and then enter a **Password** for the connection.

Next, customize global DC Agent communication and troubleshooting, domain controller polling, and computer polling settings. By default, changes that you make here affect all DC Agent instances. Settings marked with an asterisk (*), however, can be overridden in an agent's configuration file to customize the behavior of that agent instance (see [Configuring different settings for an agent instance](#), page 272).

1. Under DC Agent Communication, enter the **Communications port** to be used for communication between DC Agent and other Websense components. The default is 30600.

Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30601.

2. Under Domain Controller Polling, mark **Enable domain controller polling** to enable DC Agent to query domain controllers for user logon sessions.

You can specify which domain controllers each instance of DC Agent polls in the agent's configuration file. See [Configuring multiple agents](#), page 271, for details.

3. Use the **Query interval** field to specify how often (in seconds) DC Agent queries domain controllers.

Decreasing the query interval may provide greater accuracy in capturing logon sessions, but also increases overall network traffic. Increasing the query interval decreases network traffic, but may also delay or prevent the capture of some logon sessions. The default is 10 seconds.

4. Use the **User entry timeout** field to specify how frequently (in hours) DC Agent refreshes the user entries in its map. The default is 24 hours.
5. Under Computer Polling, check **Enable computer polling** to enable DC Agent to query computers for user logon sessions. This may include computers that are outside the domains that the agent already queries.

DC Agent uses WMI (Windows Management Instruction) for computer polling. If you enable computer polling, configure the Windows Firewall on client machines to allow communication on port **135**.

6. Enter a **User map verification interval** to specify how often DC Agent contacts client machines to verify which users are logged on. The default is 15 minutes.

DC Agent compares the query results with the user name/IP address pairs in the user map it sends to Filtering Service. Decreasing this interval may provide greater user map accuracy, but increases network traffic. Increasing the interval decreases network traffic, but also may decrease accuracy.

7. Enter a **User entry timeout** period to specify how often DC Agent refreshes entries obtained through computer polling in its user map. The default is 1 hour.

DC Agent removes any user name/IP address entries that are older than this timeout period, and that DC Agent cannot verify as currently logged on. Increasing this interval may lessen user map accuracy, because the map potentially retains old user names for a longer time.



Note

Do not make the user entry timeout interval shorter than the user map verification interval. This could cause user names to be removed from the user map before they can be verified.

8. Click **OK** to immediately save and implement your changes.

Logon Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Configuring Logon Agent, page 259](#)
- ◆ [Configuring different settings for an agent instance, page 272](#)

Websense Logon Agent identifies users in real time, as they log on to domains. This eliminates the possibility of missing a user logon due to a query timing issue.

Logon Agent (also called Authentication Server) can reside on a Windows or Linux machine. The agent works with the Websense logon application (LogonApp.exe) on Windows client machines to identify users as they log on to Windows domains.

In most cases, using either DC Agent or Logon Agent is sufficient, but you can use both agents together. In this case, Logon Agent takes precedence over DC Agent. DC Agent only communicates a logon session to Filtering Service in the unlikely event that Logon Agent has missed one.

Install Logon Agent, and then deploy the logon application to client machines from a central location. For more information, see the *Deployment and Installation Center*.

After installation, configure the agent to communicate with client machines and with the Websense Filtering Service (see *Configuring Logon Agent*).

**Note**

If you are using Windows Active Directory (Native Mode) and User Service is installed on a Linux machine, see *User Service on Linux*, page 415, for additional configuration steps.

Configuring Logon Agent

Related topics:

- ◆ [Transparent identification](#), page 245
- ◆ [Manual authentication](#), page 247
- ◆ [Configuring user identification methods](#), page 248
- ◆ [Logon Agent](#), page 258
- ◆ [Configuring multiple agents](#), page 271

Use the **User Identification > Logon Agent** page to configure a new instance of Logon Agent, as well as to configure the global settings that apply to all instances of Logon Agent.

To add a new instance of Logon Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

**Note**

Machine names must start with an alphabetical character (a-z), not a numeric or special character.

Machine names containing certain extended ASCII characters may not resolve properly. If you are using a non-English version of Websense software, enter an IP address instead of a machine name.

2. Enter the **Port** that Logon Agent should use to communicate with other Websense components. The default is 30602.

3. To establish an authenticated connection between Filtering Service and Logon Agent, check **Enable Authentication**, and then enter a **Password** for the connection.
4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global Logon Agent communications settings. By default, changes that you make here affect all Logon Agent instances.

1. Under Logon Agent Communication, enter the **Communications port** that should be used for communication between Logon Agent and other Websense components. The default is 30602.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30603.
3. Under Logon Application Communication, specify the **Connection port** that the logon application uses to communicate with Logon Agent. The default is 15880.
4. Enter the **Maximum number of connections** that each Logon Agent instance allows. The default is 200.

If your network is large, you may need to increase this number. Increasing the number does increase network traffic.

5. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

To configure the default settings that determine how user entry validity is determined, you must first determine whether Logon Agent and the client logon application will operate in **persistent mode** or **nonpersistent mode** (default).

Nonpersistent mode is activated by including the /NOPERSIST parameter when launching **LogonApp.exe**. (More information is available in the **LogonApp_ReadMe.txt** file, which is included with your Logon Agent installation.)

- ◆ In persistent mode, the logon application contacts Logon Agent periodically to communicate user logon information.

If you are using persistent mode, specify a **Query interval** to determine how frequently the logon application communicates logon information.



Note

If you change this value, the change does not take effect until the previous interval period has elapsed. For example, if you change the interval from 15 minutes to 5 minutes, the current 15-minute interval must end before the query starts occurring every 5 minutes.

- ◆ In nonpersistent mode, the logon application sends user logon information to Logon Agent only once for each logon.

If you are using nonpersistent mode, specify a **User entry expiration** time period. When this timeout period is reached, the user entry is removed from the user map.

When you are finished making configuration changes, click **OK** to save your settings.


RADIUS Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Processing RADIUS traffic, page 262](#)
- ◆ [Configuring the RADIUS environment, page 262](#)
- ◆ [Configuring RADIUS Agent, page 263](#)
- ◆ [Configuring the RADIUS client, page 264](#)
- ◆ [Configuring the RADIUS server, page 265](#)
- ◆ [Configuring different settings for an agent instance, page 272](#)

Websense RADIUS Agent lets you apply user and group-based policies using authentication provided by a RADIUS server. RADIUS Agent enables transparent identification of users who access your network using a dial-up, Virtual Private Network (VPN), Digital Subscriber Line (DSL), or other remote connection (depending on your configuration).

RADIUS Agent works together with the RADIUS server and RADIUS client in your network to process and track Remote Access Dial-In User Service (RADIUS) protocol traffic. This enables you to assign particular filtering policies to users or groups that access your network remotely, as well as to local users.



When you install RADIUS Agent, the agent integrates with existing Websense components. However, RADIUS Agent, your RADIUS server, and your RADIUS client must be configured appropriately (see [Configuring RADIUS Agent, page 263](#)).

Processing RADIUS traffic

The Websense RADIUS Agent acts as a proxy that forwards RADIUS messages between a RADIUS client and a RADIUS server (or multiple clients and servers).

RADIUS Agent does not authenticate users directly. Instead, the agent identifies remote users and associates them with IP addresses so a RADIUS server can authenticate those users. Ideally, the RADIUS server passes authentication requests to an LDAP-based directory service.

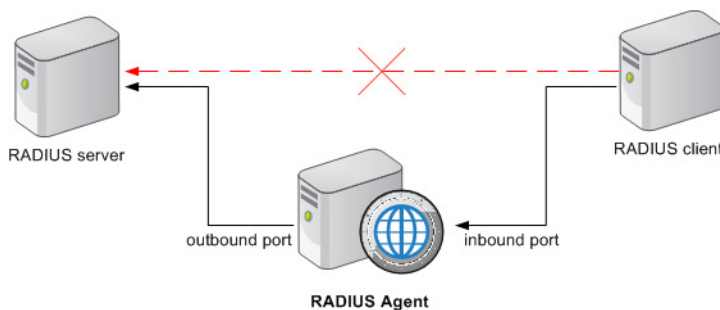
RADIUS Agent stores user name-to-IP-address pairings in a user map. If your RADIUS client supports accounting (or user logon tracking), and accounting is enabled, RADIUS Agent gleans more detail about user logon sessions from the RADIUS messages it receives.

When properly configured, Websense RADIUS Agent captures and processes all RADIUS protocol packets of these types:

- ◆ **Access-Request:** Sent by a RADIUS client to request authorization for a network access connection attempt.
- ◆ **Access-Accept:** Sent by a RADIUS server in response to an Access-Request message; tells the RADIUS client that the attempted connection is authorized and authenticated.
- ◆ **Access-Reject:** Sent by a RADIUS server in response to an Access-Request message; tells the RADIUS client that the attempted connection is rejected.
- ◆ **Accounting-Stop-Request:** Sent by a RADIUS client to tell the RADIUS server to stop tracking user activity.

Configuring the RADIUS environment

Websense RADIUS Agent serves as a proxy between a RADIUS client and a RADIUS server. This diagram shows a simplified view of how using RADIUS Agent differs from a standard RADIUS setup.



RADIUS Agent and the RADIUS server should be installed on separate machines. The agent and server cannot have the same IP address, and must use different ports.

After installing RADIUS Agent, configure the RADIUS Agent in TRITON - Web Security (see [Configuring RADIUS Agent](#), page 263). You must also:

- ◆ Configure the RADIUS client (typically a Network Access Server [NAS]) to communicate with RADIUS Agent instead of directly with your RADIUS server.
- ◆ Configure the RADIUS server to use RADIUS Agent as a proxy (see the RADIUS server documentation). If you have multiple RADIUS servers, configure each one separately.

**Note**

If you use Lucent RADIUS Server and RRAS, you must configure the RADIUS server to use Password Authentication Protocol (PAP), and the RRAS server to accept only PAP requests. For more information, see the related product documentation.

Configuring RADIUS Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Configuring user identification methods, page 248](#)
- ◆ [RADIUS Agent, page 261](#)
- ◆ [Configuring multiple agents, page 271](#)

Use the **User Identification > RADIUS Agent** page to configure a new instance of RADIUS Agent, as well as to configure the global settings that apply to all instances of RADIUS Agent.

To add a new instance of RADIUS Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

**Note**

Machine names must start with an alphabetical character (a-z), not a numeric or special character.

Machine names containing certain extended ASCII characters may not resolve properly. In non-English environments, enter an IP address instead of a name.

2. Enter the **Port** that RADIUS Agent should use to communicate with other Websense components. The default is 30800.

3. To establish an authenticated connection between Filtering Service and RADIUS Agent, check **Enable Authentication**, and then enter a **Password** for the connection.
4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global RADIUS Agent settings. By default, changes that you make here affect all RADIUS Agent instances. Settings marked with an asterisk (*), however, can be overridden in an agent's configuration file to customize the behavior of that agent instance (see [Configuring different settings for an agent instance, page 272](#)).

1. Enter the **Communications port** used for communication between RADIUS Agent and other Websense components. The default is 30800.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30801.
3. Under RADIUS Server, enter the **RADIUS server IP or name**. RADIUS Agent forwards authentication requests to the RADIUS server, and must know the identity of this machine.
4. If Microsoft RRAS is in use, enter the IP address of the **RRAS machine**. Websense software queries this machine for user logon sessions.
5. Enter the **User entry timeout** interval, used to determine how often RADIUS Agent refreshes its user map. Typically, the default query value (24 hours) is best.
6. Use the **Authentication Ports** and **Accounting Ports** settings to specify which ports RADIUS Agent uses to send and receive authentication and accounting requests. For each type of communication, you can specify which port is used for communication between:
 - RADIUS Agent and the RADIUS server
 - RADIUS Agent and the RADIUS client
7. When you are finished, click **OK** to immediately save your settings.

Configuring the RADIUS client

Your RADIUS client must be configured to transmit authentication and accounting requests to the RADIUS server via RADIUS Agent.

Modify your RADIUS client configuration so that:

- ◆ The RADIUS client sends authentication requests to machine and port on which RADIUS Agent listens for authentication requests. This is the **Authentication Port** specified during RADIUS Agent configuration.
- ◆ The RADIUS client sends accounting requests to the machine and port on which RADIUS Agent listens for accounting requests. This is the **Accounting Port** specified during RADIUS Agent configuration.

The exact procedure for configuring a RADIUS client differs by client type. For details, see your RADIUS client documentation.

**Note**

The RADIUS client should include the attributes **User-Name** and **Framed-IP-Address** in authentication and accounting messages it sends. RADIUS Agent uses the values of these attributes to interpret and store user name/IP address pairs. If your RADIUS client does not generate this information by default, configure it to do so (see the RADIUS client documentation).

Configuring the RADIUS server

To enable proper communication between Websense RADIUS Agent and your RADIUS server:

- ◆ Add the IP address of the RADIUS Agent machine to your RADIUS server's client list. For instructions, see your RADIUS server documentation.
- ◆ Define shared secrets between the RADIUS server and all RADIUS clients that use the agent to communicate with the RADIUS server. Shared secrets are usually specified as authentication security options.

Configuring a shared secret for RADIUS clients and the RADIUS server provides secure transmission of RADIUS messages. Typically, the shared secret is a common text string. For instructions, see your RADIUS server documentation.

**Note**

The RADIUS server should include the attributes **User-Name** and **Framed-IP-Address** in authentication and accounting messages. RADIUS Agent uses the values of these attributes to interpret and store user name/IP address pairs. If your RADIUS server does not generate this information by default, configure it to do so (see the RADIUS server documentation).

eDirectory Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Configuring eDirectory Agent, page 267](#)
- ◆ [Configuring different settings for an agent instance, page 272](#)

Websense eDirectory Agent works together with Novell eDirectory to transparently identify users so Websense software can filter them according to policies assigned to users, groups, domains, or organizational units.

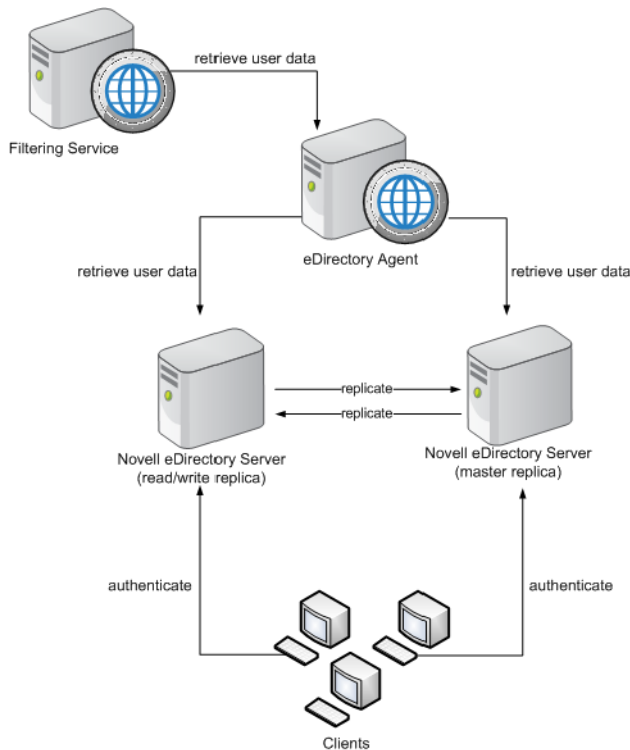
eDirectory Agent gathers user logon session information from Novell eDirectory, which authenticates users logging on to the network. The agent then associates each authenticated user with an IP address, and records user name-to-IP-address pairings to a local user map. eDirectory Agent then communicates this information to Filtering Service.



Note

From a Novell client running Windows, multiple users can log on to a single Novell eDirectory server. This associates one IP address with multiple users. In this scenario, eDirectory Agent's user map only retains the user name/IP address pairing for the last user logged on from a given IP address.

One instance of Websense eDirectory Agent can support one Novell eDirectory master, plus any number of Novell eDirectory replicas.



Special configuration considerations

- ◆ If you have integrated Cisco Content Engine v5.3.1.5 or higher with Websense software:

- Run the following Websense services on the same machine as Cisco Content Engine:

Websense eDirectory Agent
Websense User Service
Websense Filtering Service
Websense Policy Server

- Ensure that all Novell eDirectory replicas are added to the **wsedir.ini** file on the same machine.
- Delete the **eDirAgent.bak** file.

Run Websense Reporting Tools services on a machine **separate** from Cisco Content Engine and Websense software.

- ◆ Websense software supports using NMAS with eDirectory Agent. To use eDirectory Agent with NMAS enabled, eDirectory Agent must be installed on a machine that is also running the Novell client.

Configuring eDirectory Agent

Related topics:

- ◆ [Transparent identification, page 245](#)
- ◆ [Manual authentication, page 247](#)
- ◆ [Configuring user identification methods, page 248](#)
- ◆ [eDirectory Agent, page 265](#)
- ◆ [Configuring eDirectory Agent to use LDAP, page 269](#)
- ◆ [Configuring multiple agents, page 271](#)

Use the **User Identification > eDirectory Agent** page to configure a new instance of eDirectory Agent, as well as to configure the global settings that apply to all instances of eDirectory Agent.

To add a new instance of eDirectory Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.



Note

Machine names must start with an alphabetical character (a-z), not a numeric or special character.

Machine names containing certain extended ASCII characters may not resolve properly. In non-English environments, enter an IP address instead of a name.

2. Enter the **Port** that eDirectory Agent should use to communicate with other Websense components. The default is 30700.
3. To establish an authenticated connection between Filtering Service and eDirectory Agent, check **Enable Authentication**, and then enter a **Password** for the connection.
4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global eDirectory Agent communication settings. By default, changes that you make here affect all eDirectory Agent instances. Settings marked with an asterisk (*), however, can be overridden in an agent's configuration file to customize the behavior of that agent instance (see [Configuring different settings for an agent instance](#), page 272).

1. Enter the default **Communications port** used for communication between eDirectory Agent and other Websense components. The default is 30700.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30701.
3. Under eDirectory Server, specify a **Search base** (root context) for eDirectory Agent to use as a starting point when searching for user information in the directory.
4. Provide the administrative user account information that eDirectory Agent should use to communicate with the directory:
 - a. Enter the **Administrator distinguished name** for a Novell eDirectory administrative user account.
 - b. Enter the **Password** used by that account.
 - c. Specify a **User entry timeout** interval to indicate how long entries remain in the agent's user map.

This interval should be approximately 30% longer than a typical user logon session. This helps prevent user entries from being removed from the map before the users are done browsing.

Typically, the default value (24 hours) is recommended.

**Note**

In some environments, instead of using the User entry timeout interval to determine how frequently eDirectory Agent updates its user map, it may be appropriate to query the eDirectory Server at regular intervals for user logon updates. See [Enabling full eDirectory Server queries](#), page 270.

5. Add the eDirectory Server master, as well as any replicas, to the **eDirectory Replicas** list. To add an eDirectory Server master or replica to the list, click **Add**, and then follow the instructions in [Adding an eDirectory server replica](#), page 269.

When you are finished making configuration changes, click **OK** to save your settings.

Adding an eDirectory server replica

One instance of the Websense eDirectory Agent can support one Novell eDirectory master, plus any number of Novell eDirectory replicas running on separate machines.

eDirectory Agent must be able to communicate with each machine running a replica of the directory service. This ensures that the agent gets the latest logon information as quickly as possible, and does not wait for eDirectory replication to occur.

Novell eDirectory replicates the attribute that uniquely identifies logged-on users only every 5 minutes. Despite this replication time lag, eDirectory Agent picks up new logon sessions as soon as a user logs on to any eDirectory replica.

To configure eDirectory Agent installation to communicate with eDirectory:

1. In the Add eDirectory replica screen, enter the IP address or name for eDirectory **Server** (master or replica).
2. Enter the **Port** that eDirectory Agent uses to communicate with the eDirectory machine. The valid values are **389** (standard port) and **636** (SSL port).
3. Click **OK** to return to the eDirectory Agent page. The new entry appears in the eDirectory Replicas list.
4. Repeat the process for any additional eDirectory server machines.
5. Click **OK** to cache changes, and then click **Save All**.
6. Stop and start eDirectory Agent so that the agent can begin communicating with the new replica. See [Stopping and starting Websense services, page 329](#), for instructions.

Configuring eDirectory Agent to use LDAP

Websense eDirectory Agent can use Netware Core Protocol (NCP) or Lightweight Directory Access Protocol (LDAP) to get user logon information from Novell eDirectory. By default, eDirectory Agent on Windows uses NCP. On Linux, eDirectory Agent must use LDAP.

If you are running eDirectory Agent on Windows, but want the agent to use LDAP to query Novell eDirectory, set the agent to use LDAP instead of NCP. Generally, NCP provides a more efficient query mechanism.

To set eDirectory Agent on Windows to use LDAP:

1. Ensure that you have at least one Novell eDirectory replica containing all directory objects to monitor and filter in your network.
2. Stop the Websense eDirectory Agent service (see [Stopping and starting Websense services, page 329](#)).
3. Navigate to the eDirectory Agent installation directory (by default, **\Program Files\Websense\bin**), and then open the **wsedir.ini** file in a text editor.
4. Modify the **QueryMethod** entry as follows:

```
QueryMethod=0
```

This sets the Agent to use LDAP to query Novell eDirectory. (The default value is 1, for NCP.)

5. Save and close the file.
6. Restart the Websense eDirectory Agent service.

Enabling full eDirectory Server queries

In small networks, you can configure Websense eDirectory Agent to query the eDirectory Server for all logged-on users at regular intervals. This allows the agent to detect both newly logged-on users and users who have logged off since the last query, and to update its local user map accordingly.



Important

Configuring eDirectory Agent to use full queries is not recommended for larger networks, because the length of time required to return query results depends on the number of logged on users. The more logged-on users there are, the higher the performance impact.

When you enable full queries for eDirectory Agent, the **User entry timeout** interval is not used, because users who have logged off are identified by the query. By default, the query is performed every 30 seconds.

Enabling this feature increases eDirectory Agent processing time in 2 ways:

- ◆ Time needed to retrieve the names of logged-on users each time a query is performed
- ◆ Time required to process user name information, remove obsolete entries from the local user map, and add new entries based on the most recent query

eDirectory Agent examines the entire local user map after each query, rather than identifying only new logons. The time required for this process depends on the number of users returned by each query. The query process can therefore affect both eDirectory Agent and Novell eDirectory Server response times.

To enable full queries:

1. On the eDirectory Agent machine, navigate to the Websense **bin** directory (see *Where is the Websense "bin" directory?*, page 457).
2. Locate the file **wsedir.ini** and make a backup copy in another directory.
3. Open **wsedir.ini** in a text editor (like Notepad or vi).
4. Go to the **[eDirAgent]** section of the file and find the following entry:

```
QueryMethod=<N>
```

Make a note of the QueryMethod value, in case you want to revert to the default setting later.

5. Update the **QueryMethod** value as follows:

- If the current value is 0 (communicate with the directory via LDAP), change the value to **2**.
- If the current value is 1 (communicate with the directory via NCP), change the value to **3**.

**Note**

If changing this query value slows system performance, return the QueryMethod entry to its previous value.

6. If the default query interval (30 seconds) is not appropriate for your environment, edit the **PollInterval** value appropriately.
Note that the interval time is set in **milliseconds**.
7. Save and close the file.
8. Restart the Websense eDirectory Agent service (see [Stopping and starting Websense services](#), page 329).

Configuring multiple agents

Related topics:

- ◆ [DC Agent](#), page 255
- ◆ [Logon Agent](#), page 258
- ◆ [RADIUS Agent](#), page 261
- ◆ [eDirectory Agent](#), page 265

It is possible to combine multiple transparent identification agents within the same network. If your network configuration requires multiple agents, it is best to install each agent on a separate machine. In some cases, however, you can configure Websense software to work with multiple agents on a single machine.

The following transparent identification agent combinations are supported:

Combination	Same host?	Same network?	Configuration required
Multiple DC Agents	No	Yes	Ensure that all instances of DC Agent can communicate with Filtering Service.
Multiple RADIUS Agents	No	Yes	Configure each instance to communicate with Filtering Service.

Combination	Same host?	Same network?	Configuration required
Multiple eDirectory Agents	No	Yes	Configure each instance to communicate with Filtering Service.
Multiple Logon Agents	No	Yes	Configure each instance to communicate with Filtering Service.
DC Agent + RADIUS Agent	Yes	Yes	Install these agents in separate directories. Configure each agent to communicate with Filtering Service using a different communication port.
DC Agent + eDirectory Agent	No	No	Websense software does not support communication with both Windows and Novell directory services in the same deployment. However, you can have both agents installed, with only 1 active agent.
DC Agent + Logon Agent	Yes	Yes	Configure both agents to communicate with Filtering Service. By default, each agent uses a unique port, so port conflicts are not an issue unless these ports are changed.
eDirectory Agent + Logon Agent	No	No	Websense software does not support communication with both Windows and Novell directory services in the same deployment. However, you can have both agents installed, with only 1 active agent.
RADIUS Agent + eDirectory Agent	Yes	Yes	Configure each agent to communicate with Filtering Service using a different communication port.
DC Agent + Logon Agent + RADIUS Agent	Yes	Yes	Though this combination is rarely required, it is supported. Install each agent in a separate directory. Configure all agents to communicate with Filtering Service using different communication ports.

Configuring different settings for an agent instance

The TRITON - Web Security transparent identification agent configuration settings are global, and apply to all instances of the agent you have installed. If you have

multiple instances of any agent, however, you can configure one instance independently of the others.

Unique settings you specify for a particular agent instance override the global settings in the Settings dialog box. Settings that can be overridden are marked with an asterisk (*).

1. Stop the transparent identification agent service (see [Stopping and starting Websense services](#), page 329).
2. On the machine running the agent instance, navigate to the agent installation directory and open the appropriate file in a text editor:
 - for DC Agent: **transid.ini**
 - for Logon Agent: **authserver.ini**
 - for eDirectory Agent: **wsedir.ini**
 - for RADIUS Agent: **wsradius.ini**
3. Locate the parameter to change for this agent instance (see [INI file parameters](#), page 274).

For example, you can enable an authenticated connection between this agent instance and other Websense services. To do this, enter a value for the **password** parameter in the INI file:

```
password= [xxxxxxx]
```

4. Modify any other values as desired.
5. Save and close the INI file.
6. If you made a change to **DC Agent** settings, you must remove 2 files from the Websense **bin** directory (see [Where is the Websense "bin" directory?](#), page 457):
 - a. Stop all Websense services on the DC Agent machine (see [Stopping and starting Websense services](#), page 329).
 - b. Delete the following files:


```
Journal.dat
XidDcAgent.bak
```

 These files are recreated when you start the Websense DC Agent service.
 - c. Restart the Websense services (including DC Agent), and then skip to **step 8**.
7. Restart the transparent identification agent service.
8. Update the agent settings in TRITON - Web Security:
 - a. Go to **Settings > General > User Identification**.
 - b. Under **Transparent Identification Agents**, select the agent and then click **Edit**.



Note

If you modified the **port** value for this agent instance, remove and then re-add the agent. First select the existing agent entry and click **Delete**, and then click **Add Agent**.

- c. Verify the **Server IP or name** and **Port** this agent instance uses. If you specified a unique port number in the INI file, ensure that your entry matches that value.
- d. If you specified a unique authentication password in the INI file, ensure that the **Password** entry shown here is correct.
- e. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

INI file parameters

TRITON - Web Security field label	.ini parameter name	Description
Communications port (<i>all agents</i>)	port	The port over which the agent communicates with other Websense services.
Diagnostic Port (<i>all agents</i>)	DiagServerPort	The port over which the agent troubleshooting tool listens for data from the agent.
Password (<i>all agents</i>)	password	The password the agent uses to authenticate connections to other Websense services. Specify a password to enable authentication.
Query interval (<i>DC Agent</i>)	QueryInterval	The interval at which DC Agent queries domain controllers.
Server IP or name Port (<i>eDirectory Agent</i>)	Server=IP:port	The IP address and port number of the machine running eDirectory Agent.
Search base (<i>eDirectory Agent</i>)	SearchBase	The root context of the Novell eDirectory server.
Administrator distinguished name (<i>eDirectory Agent</i>)	DN	The name of the administrative user for Novell eDirectory server.
Password (<i>eDirectory Agent</i>)	PW	The password for the Novell eDirectory server administrative user.
RADIUS server IP or name	RADIUSHost	The IP address or name of your RADIUS server machine.
RRAS machine IP (Windows Only) (<i>RADIUS Agent</i>)	RRASHost	The IP address of the machine running RRAS. Websense queries this machine for user logon sessions.
Authentication Ports: Between RADIUS Agent and RADIUS server	AuthOutPort	The port on which the RADIUS server listens for authentication requests.

Authentication Ports: Between RADIUS clients and RADIUS Agent	AuthInPort	The port over which RADIUS Agent accepts authentication requests.
Accounting Ports: Between RADIUS Agent and RADIUS server	AccOutPort	The port over which the RADIUS server listens for RADIUS accounting messages.
Accounting Ports: Between RADIUS clients and RADIUS Agent	AccInPort	The port over which RADIUS Agent accepts accounting requests.

Configuring an agent to ignore certain user names

You can configure a transparent identification agent to ignore logon names that are not associated with actual users. This feature is often used to deal with the way that some Windows 200x and XP services contact domain controllers in the network.

For example, **user1** logs on to the network, and is identified by the domain controller as **computerA/user1**. That user is filtered by a Websense policy assigned to **user1**. If a service starts up on the user's machine that assumes the identity **computerA/ServiceName** to contact the domain controller, this can cause filtering problems. Websense software treats **computerA/ServiceName** as a new user with no policy assigned, and filters this user by the computer policy, or by the **Default** policy.

To address this issue:

1. Stop the agent service (see [Stopping and starting Websense services](#), page 329).
2. Navigate to the `\Websense\bin\` directory, and open the **ignore.txt** file in a text editor.
3. Enter each user name on a separate line. Do not include wildcard characters, such as `“*“`:

```
maran01
WindowsServiceName
```

Websense software ignores these user names, regardless of which machine they are associated with.

To prompt Websense software to ignore a user name within a specific domain, use the format **username, domain**.

```
aperez, engineering1
```

4. When you are finished, save and close the file.
5. Restart the agent service.

The agent ignores the specified user names, and Websense software does not consider these names in filtering.

Identification of hybrid filtering users

Related topics:

- ◆ [Websense Directory Agent, page 281](#)
- ◆ [When users are not identified, page 283](#)
- ◆ [Working with hybrid filtering clients, page 76](#)

Select **Settings > Hybrid Configuration > Hybrid User Identification** to configure how users are identified by the hybrid service, and to test and configure users' connections to the service. You can configure multiple authentication or identification options for your hybrid users if required.

To ensure that the appropriate per-user or per-group policy is applied to hybrid filtering users, whether from a filtered location or when off-site, Websense Web Security Gateway Anywhere provides options for identifying hybrid filtering users transparently:

- ◆ **Websense Web Endpoint** is installed on client machines to provide transparent authentication, enforce use of hybrid filtering, and pass authentication details to the hybrid service. See [Deploying Web Endpoint, page 277](#).
- ◆ **Websense Authentication Service** provides clientless transparent authentication via a gateway hosted on your network. See [Deploying Websense Authentication Service, page 280](#).

If you do not deploy either Web Endpoint or Authentication Service, the hybrid service can identify users transparently or manually when they connect to hybrid filtering.

- ◆ Users can only be identified transparently if they are logging on from a known IP address, defined as a **filtered location** (see [Define the locations filtered by hybrid filtering, page 175](#)).
- ◆ The hybrid service can be configured to automatically generate passwords for all users whose information is collected by Directory Agent (see [Configure user access to hybrid filtering, page 183](#)).
- ◆ If you do not enable transparent identification, off site users are prompted for an email address and password when they open a browser and connect to the Internet. Other hybrid filtering users are filtered based on their IP address.

Indicate how the hybrid service should identify users requesting Internet access. These options are also used as a fallback if either the endpoint or Authentication Service fails.

- ◆ Mark **Always authenticate users on first access** to enable transparent or manual authentication when users first connect to hybrid filtering.
Internet Explorer and Firefox can be used for transparent user identification. Other browsers will prompt users for logon information.

If Directory Agent is sending data to the hybrid service, using NTLM to identify users is recommended.

- ◆ Mark **Use NTLM to identify users when possible** to use directory information gathered by Directory Agent to identify users transparently, if possible.

When this option is selected, the hybrid service uses NTLM to identify the user if the client supports it, and otherwise provides a logon prompt.

If you do not select this option, users who could not be identified via another means see a logon prompt when accessing the Internet. Basic authentication is used to identify users who receive a logon prompt. Advise end users **not** to use the same password for hybrid filtering that they use to log on to the network.



Note

When NTLM is used to identify users, **do not** use self-registration (configured on the User Access page under Registered Domains).

- ◆ Specify whether or not a Welcome page is displayed when users who have not been identified via NTLM open a browser to connect to the Internet. The Welcome page:
 - Provides a simple selection of common search engines to get the user started
 - Is used mainly by those who connect to the hybrid service from outside a filtered location (while working from home or traveling, for example)

When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Once you have set up hybrid filtering and configured user browsers to access the PAC file, you can use the links provided under **Verify End User Configuration** to make sure that end user machines have Internet access and are correctly configured to connect to the hybrid service.

If your hybrid filtering account has not been verified (which may mean that no email address has been entered on the Settings > General > Account page), the URLs are not displayed.

Deploying Web Endpoint

Websense Web Endpoint is a piece of software that gets installed on a client machine. It enforces the use of the hybrid service for Web filtering, and passes authentication information to the hybrid proxies, enabling secure transparent authentication.

You can deploy Web Endpoint in the following ways:

- ◆ Download the installation files and install them manually on the selected client machines using your preferred distribution method. For example, you might push the endpoint out using Microsoft Group Policy Object (GPO) or similar deployment software.

- ◆ Deploy the endpoint to some or all of your hybrid filtering users directly from the hybrid service. Each user will be asked to install the endpoint software on their machine.

If a user does not install the endpoint, he or she is authenticated according to the options you have selected on the User Identification page. Websense Authentication Service is used if configured; otherwise the hybrid service falls back to NTLM identification or finally basic authentication. The user is again asked to install the endpoint next time they start a browsing session.



Note

Web Endpoint can only be deployed on Windows operating systems.

To deploy Web Endpoint to end users:

1. On the **Settings > Hybrid Configuration > Hybrid User Identification** page, mark **Enable installation and update of Web Endpoint on client machines**.

This setting defines whether the endpoint, and automatic updates (if you select that option) will be deployed to the client machines that you define. If you uncheck this option at a later date, there will be no further new deployments of the endpoint or automatic updates of existing installations. However, the installed endpoints will continue to work unless they are uninstalled from the client machines.

2. Choose your deployment method:

- To deploy the endpoint manually, using your preferred distribution method, click **Deploy Web Endpoint Manually**.

If you are using GPO to distribute the endpoint, note the GPO command displayed on screen and use it to configure your deployment script. See [Distributing Web Endpoint via GPO, page 279](#).

Click **View Web Endpoint Files** to view the endpoint versions suitable for your client machines. Select a client operating system, then click on a version of the endpoint to download. You can also view a PDF of the release notes for each version by clicking a release notes link. Click **Close** when done.

- To deploy the endpoint directly from the hybrid service, mark **Deploy Web Endpoint from hybrid service proxies**.

Choose whether the endpoint is deployed to all users that are filtered through the hybrid service, or only to off-site users.

You can provide a customized message that appears to end users at the beginning of the endpoint download and installation process. The message can be used to reassure the user that the download is company-approved, and to provide any further information they may need. To customize the message, click **Advanced Settings**, then enter your organization name and the message you want to display. Click **View Sample Page** to see what will appear to the end user.

The sample page also contains the default text that is always displayed to the end user at the beginning of the download.

3. Mark **Automatically update endpoint installations when a new version is released** if you want to ensure that all endpoints on your client machines always have the latest version when it is available from the hybrid service.
4. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Distributing Web Endpoint via GPO

Follow the steps below to deploy endpoint clients through an Active Directory group policy object (GPO):

1. Create a shared folder (create a folder and turn on sharing in the Properties menu).
2. Create a batch file (.bat) in the shared folder, for example “installmsi.bat”. This can be done in any text editor.

Type the following msixec command into the batch file and save it.

```
msiexec /package "\\path\WebSense Endpoint.msi" /quiet /norestart WSCONTEXT=xxxx
```

Where `path` is the path to the installer that you downloaded from the portal, and `WSCONTEXT=xxxx` is the GPO script command noted from the Hybrid User Identification page.

3. Test your batch file manually to make sure it runs on other workstations. You can do this by opening the server path to the file on a workstation and attempting to run the file. If the file does not run, check your permissions.
4. Open the Group Policy Management Console (GPMC).
5. Create a new (or open an existing) GPO on the organization unit (OU) in which your computer accounts reside. To create a new GPO:
 - a. In the console tree, right-click **Group Policy Objects** in the forest and domain in which you want to create a Group Policy object (GPO).
 - b. Click **New**.
 - c. In the **New GPO** dialog box, specify a name for the new GPO, and the click **OK**.
6. Open **Computer Configuration > Windows Settings > Scripts**, and double-click **Startup** in the right pane of the screen.
7. Click **Add**.
8. In the **Script Name** field type the full network path and filename of the script batch file you created in step 2.
9. Click **OK**.
10. Close the GPMC.
11. Run the **gpupdate /force** command at the command prompt to refresh the group policy.

The application should be installed on startup. The client may not be fully functional until a reboot occurs.

Deploying Websense Authentication Service

Websense Authentication Service is an on-network virtual machine that authenticates user identity, attributes, and roles with enterprise directories. All communications between components are secured.

When Authentication Service is installed on your network, clients connecting to the hybrid proxy from a filtered location are redirected to Authentication Service. Once Authentication Service has authenticated a user against your directory service, they are directed back to the proxy and the appropriate policy is applied. Clients who have authenticated once do not then have to authenticate again for subsequent Web browsing sessions.

Off-site users must enter their email address and network password whenever they connect to the hybrid proxy. These credentials are verified with Authentication Service and the appropriate policy is applied for that session.

For more information, see the *Authentication Service Installation and Configuration Guide*.

To download and deploy Websense Authentication Service:

1. On the **Settings > Hybrid Configuration > Hybrid User Identification** page, click **Authentication Service Files** to view the available Authentication Service downloads.

Click on a file name to download that version. You can also view a PDF of the release notes for each version by clicking a release notes link. Click **Close** when done.
2. Install Authentication Service, following the instructions in the *Authentication Service Installation and Configuration Guide*.
3. On the **Settings > Hybrid Configuration > User Access** page, download and install the hybrid SSL certificate to ensure seamless authentication to HTTPS sites. If the certificate is not installed for Authentication Service users, they must authenticate using NTLM identification or manual authentication, depending on the settings on the Hybrid User Identification page. See [Enabling HTTPS notification pages](#), page 186.
4. Once Authentication Service is installed and configured and the SSL certificate is installed on clients, copy the metadata URL from the Configuration > Federation page in the Authentication Service Console and enter it in the Metadata URL field on the Hybrid User Identification page.
5. Test that Authentication Service is set up correctly by deploying the temporary PAC file URL displayed under Test Authentication Service to a client machine. An end user on the client machine should be able to browse Web sites according to their policy with no requests for credentials.

If you receive an error page while browsing with the test PAC file, check you have followed all of the steps in the *Authentication Service Installation and Configuration Guide*.

6. Mark **Enable Authentication Service** to activate Authentication Service for all client machines. Ensure that the client used to test the implementation in step 6 is reverted to the standard PAC file.
7. Under Session Timeout, define how often users' credentials are revalidated for security reasons. The options are 1, 7, 14, or 30 days.
8. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Websense Directory Agent

In Websense Web Security Gateway Anywhere environments, an interoperability component called **Websense Directory Agent** is required if you want to enable user, group, and domain (OU) based filtering through the hybrid service.

Directory Agent must be installed on a machine from which it can communicate with:

- ◆ Your supported LDAP-based directory service (Windows Active Directory [Native Mode], Oracle Directory Server, or Novell eDirectory)

If your organization uses Windows Active Directory in mixed mode, user and group data cannot be collected and sent to the hybrid service.

- ◆ Websense Sync Service

Directory Agent can be installed on the same machine as other Websense components, including Sync Service and User Service.

After deployment, use TRITON - Web Security to configure Directory Agent to collect data from your directory service (see [Send user and group data to the hybrid service, page 189](#)). Once configured, Directory Agent collects user and group data from your directory service and sends it to Sync Service in LDIF format.

At scheduled intervals (see [Schedule communication with hybrid filtering, page 196](#)), Sync Service sends the user and group information collected by Directory Agent to the hybrid service. Sync Service compresses large files before sending them.

Directory Agent and User Service

Related topics:

- ◆ [Identification of hybrid filtering users, page 276](#)
- ◆ [Working with users and groups, page 64](#)
- ◆ [Directory services, page 65](#)
- ◆ [Send user and group data to the hybrid service, page 189](#)

Although Directory Agent collects directory information independently, it has one important dependency on User Service. At installation, Directory Agent must connect to a Policy Server instance that has a User Service associated with it. Directory Agent

can be configured to communicate only with the directory that this User Service instance is configured to use.

In other words, in a distributed deployment, if you have multiple Policy Servers, each with an associated User Service, and the User Service instances connect to different directory servers, you must associate Directory Agent with the Policy Server whose User Service connects to the directory that you want to use for hybrid filtering user identification.

- ◆ You can have multiple Directory Agent instances.
- ◆ Each Directory Agent instance must be associated with a different Policy Server.
- ◆ All Directory Agent instances must connect to a single Sync Service. (A deployment can have only one Sync Service instance.)

You must configure the Sync Service connection manually for all supplemental Directory Agent instances. (Communication is configured automatically for the Directory Agent instance that connects to the same Policy Server as Sync Service.) To do this:

1. When you log on to TRITON - Web Security, select the appropriate Policy Server instance for the Directory Agent that you want to configure.
2. Go to the Settings > Hybrid Configuration > Shared User Data page.
3. Under Synchronize User Data, verify the **Name or IP address** of the Sync Service machine and the **Port** used for Sync Service communication (by default, 55832).
4. Click **Test Connection** to verify that Directory Agent can send data to Sync Service. The test may take a minute or more.
 - If the connection is made, a success message is displayed.
 - If the connection cannot be made, verify the IP address or host name of the Sync Service machine and the communication port. Also verify that the Sync Service machine is on, that Sync Service is running, and that your network firewall permits connections on the Sync Service port.
5. Click **OK** to cache your changes, and then click **Save All** to implement them.

Directory Agent configuration can not be performed until there is a supported User Service configuration. Changes to User Service configuration may also require you to update your Directory Agent configuration.

- ◆ User Service configuration is performed on the Settings > General > Directory Services page (see [Working with users and groups](#), page 64).
- ◆ Directory Agent configuration is performed on the Settings > Hybrid Configuration > Shared User Data page (see [Send user and group data to the hybrid service](#), page 189).

You can configure Directory Agent to use a different root context than User Service, and to process its directory data differently than User Service. Also, with Windows Active Directory, if User Service is configured to communicate with multiple global catalog servers, Directory Agent can communicate with all of them.

Note that if you have multiple Directory Agent instances, each instance must use a unique, non-overlapping root context.

When users are not identified

Related topics:

- ◆ [Identification of hybrid filtering users](#), page 276
- ◆ [Working with hybrid filtering clients](#), page 76

If you do not choose to deploy Directory Agent, Web Endpoint, or Authentication Service, or disable user identification, only 3 types of policies can be applied to users:

- ◆ The policy applied to the external IP address from which the user connects.
This IP address must be defined as a filtered location.
- ◆ Your organization's Default policy, if the request originates from outside a filtered location, or if no computer or network policy has been applied to the filtered location.
- ◆ The hybrid filtering Default policy, if the user's connection cannot be associated with your organization.
This is a rare case, that should occur only if there is a configuration problem with your hybrid filtering account.

User and group policies cannot be applied to self-registered users. Self-registered users are always filtered by the Default policy (see [Off-site user self-registration](#), page 207).

13

Delegated Administration and Reporting

Delegated administration provides an effective way to distribute responsibility for Web security configuration, Internet access management, Internet activity reporting, and compliance auditing to multiple individuals. For example:

- ◆ Allow individual managers to set policies and run reports on users in their teams.
- ◆ Give local administrators for regional offices or campuses policy management permissions, as well as some access to local configuration options, but limit reporting access to protect end-user privacy.
- ◆ Ensure that Human Resources can run Internet activity reports on some or all clients, identified by user name or IP address.
- ◆ Grant auditors access to view all configuration and policy management screens in TRITON - Web Security without the ability to save changes.

The sections that follow detail the main concepts of delegated administration, and then provide specific configuration and implementation instructions.

- ◆ *The fundamentals of delegated administration*, page 285
- ◆ *Preparing for delegated administration*, page 291
- ◆ *Managing delegated administration roles*, page 296
- ◆ *Updating delegated administration roles*, page 304
- ◆ *Performing delegated administrator tasks*, page 307
- ◆ *Enabling network accounts*, page 312

The fundamentals of delegated administration

Related topics:

- ◆ *Delegated administration roles*, page 286
- ◆ *Delegated administrators*, page 287
- ◆ *Delegated administration and reporting permissions*, page 288
- ◆ *Administrators in multiple roles*, page 290
- ◆ *Multiple administrators accessing TRITON - Web Security*, page 290

Before setting up delegated administration for your organization, there are 3 main concepts to understand:

- ◆ **Roles** are containers used to group **administrators** and clients. There are 3 types of roles. See [Delegated administration roles](#), page 286.
- ◆ **Administrators** are individuals or groups given responsibility for configuring TRITON - Web Security settings, managing policies for clients, running Internet activity reports, or auditing the system. An administrator's set of responsibilities is determined by the **role** and **permissions** that the administrator is assigned. See [Delegated administrators](#), page 287.
- ◆ **Permissions** determine what responsibilities (like creating policies or running reports) an **administrator** has within a **role**. The available permissions change based on which type of role an administrator is assigned to. See [Delegated administration and reporting permissions](#), page 288.

Delegated administration roles

A **role** groups clients—users, groups, domains (OUs), computers, and networks—with one or more administrators.

- ◆ Clients in a delegated administration role are referred to as **managed clients**.
- ◆ Administrators can perform different tasks (like managing policies or running reports) for managed clients in their role, based on their **permissions**.

TRITON - Web Security includes one predefined role: Super Administrator. Although it is not shown, **admin**, the Global Security Administrator account, is a member of this role. The admin account cannot be deleted, nor can its permissions be changed.



Important

You cannot delete the Super Administrator role or the admin account.

Administrators assigned to the Super Administrator role have the ability to create roles, assign administrators and managed clients to roles, and determine the permissions for administrators in the role. Global Security Administrators can add administrators to the Super Administrator role.

Super Administrators can create 2 types of delegated administration and reporting roles:

- ◆ **Policy management and reporting:** User policies are managed by administrators in the role. Administrators in the role can optionally also run reports.
- ◆ **Investigative reporting:** Administrators can run investigative reports showing Internet activity for only managed clients in the role. Client policies are managed in one or more other roles.

Define as many additional roles as are appropriate for the organization. For example:

- ◆ Create a role for each department, with the department manager as administrator and the department members as managed clients.
- ◆ In a geographically distributed organization, create a role for each location and assign all the users at the location as managed clients of that role. Then, assign one or more individuals at the location as administrators.

Delegated administrators

Administrators are the individuals who can access TRITON - Web Security. Depending on their permissions, they may be able to:

- ◆ Log on and view the Status > Today page, but take no other actions.
- ◆ Access all configuration and management features of TRITON - Web Security, but save no changes.
- ◆ Run reports on specific groups of clients, or on all clients.
- ◆ Manage policies for specific groups of clients.
- ◆ Have full configuration access to all features of TRITON - Web Security.

The specific permissions available depend on the administrator's role type (Super Administrator, policy management and reporting, or investigative reporting). See [Delegated administration roles, page 286](#).

Global Security Administrators (like **admin**) define administrator accounts in TRITON Settings. These accounts may either be network logon accounts (defined in a supported directory service) or local accounts, used only to access TRITON. Once an account has been defined, the Global Security Administrator assigns each one a level of logon access to one or more TRITON modules.

The levels of Web Security access that can be granted to administrators are:

- ◆ **Access and account management**, which grants unconditional Super Administrator permissions (see [Delegated administration and reporting permissions, page 288](#)).
- ◆ **Access**, which allows the administrator to log on and view limited portions of the Status > Today, History, and Alerts pages only. Super Administrators can add those administrators to roles to allow them some level of additional policy management access, reporting access, or both.

Any administrator account that has been granted access to the Web Security module appears on the Delegated Administration > View Administrator Accounts page. These accounts are also listed on the Delegated Administration > Edit Role > Add Administrators page.

Only administrators that have already been granted Web Security access via TRITON Settings can be added to roles.

Delegated administration and reporting permissions

The permissions available to an administrator depend on the type of role to which that administrator is assigned (Super Administrator, policy management and reporting, or investigative reporting).

Super Administrator permissions

Administrators in the Super Administrator role can be granted the following permissions:

- ◆ **Unconditional Super Administrators** can:
 - Access all system configuration settings for the Websense installation (managed via the Settings tab).
 - Add or remove administrators in the Super Administrator role.
 - Create or edit the Filter Lock that blocks certain categories and protocols for all users managed by delegated administration roles. See [Creating a Filter Lock](#), page 293.
 - Manage policies for clients in the Super Administrator role, including the Default policy that applies to all clients not assigned another policy in any role.
 - Create and run reports on all clients, regardless of which role they are assigned to.
 - Access Real-Time Monitor.
 - Review the audit log, which records administrator access to and actions within TRITON - Web Security.
- ◆ **Conditional Super Administrators** can be given policy, reporting, or Real-Time Monitor, or auditor permissions, or a combination.
 - **Policy** permissions allow conditional Super Administrators to:
 - Create and edit delegated administration roles, filtering components, filters, and policies, and to apply policies to clients that are not managed by any other role.
 - Access database download, directory service, user identification, and Network Agent configuration settings. Conditional Super Administrators with reporting permissions can also access configuration settings for the reporting tools.
 - Create and edit delegated administration roles, but not to delete roles or remove the administrators or managed clients assigned to them.
 - **Reporting** permissions allow conditional Super Administrators to:
 - Access graphs on the Today and History pages.
 - Run investigative and presentation reports on all users.

If an administrator is granted reporting permissions only, the Create Policy, Recategorize URL, and Unblock URL options in the Common Tasks list are not shown. Additionally, the Check Policy tool does not appear in the Toolbox.

- **Real-Time Monitor** permissions allow Super Administrators to monitor all Internet filtering activity for each Policy Server associated with TRITON - Web Security.

Only one administrator at a time can log on to a role with **policy** permissions. Therefore, if an administrator is logged on to the Super Administrator role to perform policy or configuration tasks, other Super Administrators can log on with only reporting, auditor, or status monitor permissions in the role. Super Administrators also have the option to select a different role to manage.

To switch to another role after logon, go to the **Role** drop-down list in the Web Security toolbar and select a role.

Policy Management and Reporting permissions

Delegated administrators in policy management and reporting roles can be given any combination of the following permissions:

- ◆ Policy permissions allow delegated administrators to create and manage filter components (including custom categories and recategorized URLs), filters (category, protocol, and limited access), and policies for their managed clients.

Filters created by delegated administrators are restricted by the Filter Lock, which may designate some categories and protocols as **blocked and locked**. These categories and protocols cannot be permitted by delegated administrators.

Only one administrator at a time can log on to a role with policy permissions. Therefore, if an administrator is logged on to a role to perform policy tasks, other administrators in the role can log on with auditing (read-only), reporting, or Real-Time Monitor permissions only. Administrators who have been assigned to multiple roles also have the option to select a different role to manage.

To switch to another role after logon, go to the **Role** drop-down list in the banner and select a role.

- ◆ Reporting permissions can be granted in either of 2 general categories: report on **all clients**, or report on **only managed clients** in the role.
 - Any delegated administrator with reporting permissions can be given access to the Today and History pages, investigative reports, and the Settings page used to manage the Log Database.
 - Delegated administrators with the option to report on all clients can also be given access to presentation reports.
- ◆ **Real-Time Monitor** permissions allow administrators to monitor all Internet filtering activity for each Policy Server associated with TRITON - Web Security.

Investigative Reporting permissions

Administrators in investigative reporting roles can create investigative reports for managed clients in their role. (Clients' policies are managed in other roles.) They can also use the URL Category, URL Access, and Investigate User tools.

These administrators do not have access to presentation reports or Real-Time Monitor, but can optionally be allowed to view charts on the Status > Today and History pages.

Auditors

Any Super Administrator or delegated administrator account (except admin) can be granted **Auditor** permissions. An auditor can see all of the TRITON - Web Security features and functions that other administrators in the role can access, but cannot save any changes.

Instead of the OK and Cancel buttons that allow other administrators to cache or discard changes, Auditors are given a single Back button. The Save All button is disabled.

Administrators in multiple roles

Related topics:

- ◆ [Delegated administration roles, page 286](#)
- ◆ [Delegated administrators, page 287](#)
- ◆ [Delegated administration and reporting permissions, page 288](#)

Depending on the needs of your organization, the same administrator may be assigned to multiple roles. Administrators assigned to multiple roles must choose a single role to manage at logon.

After logon, your permissions are as follows:

- ◆ **Policy:** you can add and edit filters and policies for the role selected during logon, and apply policies to that role's managed clients.
- ◆ **Reporting:** you have the combined reporting permissions of all your roles. For example, suppose you are assigned to 3 roles, with reporting permissions as follows:
 - Role 1: no reporting
 - Role 2: investigative reporting only
 - Role 3: report on all clients, full access to all reporting features

In this situation, regardless of which role you choose during logon, you are permitted to view reports on the Today and History pages, and report on all clients, using all reporting features.

If you are logged on for reporting only, the Role field in the banner bar indicates whether you have Full Reporting (report on all clients) or Limited Reporting (report on managed clients only) permissions.

Multiple administrators accessing TRITON - Web Security

Administrators in **different** roles can access TRITON - Web Security simultaneously to perform whatever activities their role permissions allow. Since they manage different clients, they can create and apply policies without conflict.

The situation is different if administrators with policy permissions in the **same** role try to connect at the same time. Only **one administrator at a time** can log on with policy permissions in the shared role. If a second administrator tries to log on with policy permissions while another administrator logged on, the second administrator is given a choice:

- ◆ Log on with read-only access (similar to temporary auditor permissions).
When this option is selected, the Role drop-down box shows “Role Name - [Read-Only]” as the current role, and offers the option of switching to “Role Name” (without any modifiers). This makes it possible to access the role with policy permissions when the role is no longer locked.
- ◆ Log on for reporting only, if the administrator has reporting permissions.
- ◆ Log on to a different role, if the administrator is assigned to any other roles.
- ◆ Log on to view only the Status pages until the role becomes available (Limited Status access).
- ◆ Try again later, after the first administrator logs off.

Administrators who are not using their policy permissions can do one of the following to unlock the role and allow another administrator to log on to manage policies:

- ◆ If generating reports, select **Release Policy Permissions** from the **Role** drop-down list.
When this option is selected, policy management features are hidden from the logged-on administrator, but reporting features remain active.
- ◆ If monitoring system performance, select **Status Monitor** from the **Role** drop-down list.

Administrators in Status Monitor mode can access the Status > Today, History, and Alerts pages, as well as Real-Time Monitor (if applicable). Their session does not time out.

If an administrator in Status Monitor mode tries to go to a page other than Today, History, Alerts, or Real-Time Monitor, he or she is prompted to log on again.

Preparing for delegated administration

Related topics:

- ◆ [The fundamentals of delegated administration, page 285](#)
- ◆ [Creating a Filter Lock, page 293](#)
- ◆ [Preparing delegated administrators, page 295](#)
- ◆ [Managing delegated administration roles, page 296](#)

Before creating delegated administration roles, there are 2 key planning and setup tasks for the Super Administrator to perform:

- ◆ Review and edit the Filter Lock, which blocks specified categories and protocols for managed clients in all delegated administration roles. By default, the Filter Lock blocks and locks several categories, so it is important to check the default settings against the requirements of your organization. (See [Creating a Filter Lock, page 293.](#))
 - Filter Lock restrictions are automatically enforced for all filters created in or copied to a delegated administration role, and cannot be modified by the delegated administrator.
 - Delegated administrators can apply any filtering action to categories and protocols **not** blocked and locked in the Filter Lock.
 - Changes to the Filter Lock are implemented for all managed clients as soon as the changes are saved. Delegated administrators who are working in TRITON - Web Security when the changes take effect will not see the changes in their filters until the next time they log on.
 - Filter Lock restrictions do not apply to clients managed by the Super Administrator role.
- ◆ Determine which Super Administrator policies and filters will be copied to each new role that you plan to create, and make adjustments to existing policies as needed.
 - By default, each role is created with a single Default policy, created from the Default category and protocol filter (**not** the Default policy) currently configured for the Super Administrator role.
 - Optionally, you can instead copy all policy objects (policies, filters, custom categories, and custom URLs) from the Super Administrator role to the new role. The delegated administrator then starts with a complete set of policies and policy components.
 - Copies of policies and filters in a delegated administration role are subject to the Filter Lock, and are therefore not identical to the same policies and filters in the Super Administrator role.
 - When the Unrestricted policy is copied, the policy **and** filter names are changed to reflect the fact that they are subject to the Filter Lock, and no longer permit all requests.

Copying Super Administrator policy objects to a new role can take a very long time, depending on how much information must be copied.

Once these planning steps are completed, each of the following delegated administration components must be put into place:

1. A Global Security Administrator creates administrator accounts on the TRITON Settings > Administrators page, and grant the accounts the appropriate level of Web Security access.
2. A Super Administrator creates delegated administration roles on the Policy Management > Delegated Administration page, then adds administrators and managed clients to the roles. See [Managing delegated administration roles, page 296.](#)

- The Super Administrator notifies the delegated administrators that they have been granted administrative access to TRITON - Web Security, and explains their level of permissions. See *Preparing delegated administrators*, page 295.

Creating a Filter Lock

Related topics:

- ◆ [Locking categories, page 293](#)
- ◆ [Locking protocols, page 294](#)

The **Policy Management > Filter Lock** page lets you specify categories and protocols that are blocked for all managed clients in delegated administration roles. Any category or protocol that is blocked in the Filter Lock is considered **blocked and locked**.

- ◆ Click the **Categories** button to block and lock specific categories or category elements (keywords and file types). See *Locking categories*, page 293.
- ◆ Click the **Protocols** button to block and lock protocols, or to specify protocols that are always logged. See *Locking protocols*, page 294.

Locking categories

Related topics:

- ◆ [Creating a Filter Lock, page 293](#)
- ◆ [Locking protocols, page 294](#)

Use the **Policy Management > Filter Lock > Categories** page to select the categories to be blocked and locked for all members of delegated administration roles. You also can block and lock keywords and file types for a category.

- Select a category in the tree.
Delegated administration roles do not have access to custom categories created by the Super Administrators. Therefore, custom categories do not appear in this tree.
- Set the restrictions for this category in the box that appears beside the category tree.

Option	Description
Lock category	Blocks and locks access to sites in this category.
Lock keywords	Blocks and locks access based on keywords defined for this category in each role.

Option	Description
Lock file types	Blocks and locks the selected file types for sites in this category. Be sure to mark the check box for each file type to be blocked and locked. Custom file types created by the Super Administrator are included on this list because they are available to delegated administration roles.
Apply to Subcategories	Applies the same settings to all subcategories of this category.

You can block and lock selected elements for all categories at once, if appropriate. Select **All Categories** in the tree, and then select the elements to be blocked for all categories. Then, click **Apply to Subcategories**.

3. When you are finished making changes, click **OK** to cache the changes and return to the Filter Lock page. Changes are not implemented until you click **Save All**.

Locking protocols

Related topics:

- ◆ [Creating a Filter Lock, page 293](#)
- ◆ [Locking categories, page 293](#)

Use the **Policy Management > Filter Lock > Protocols** page to block and lock access to or lock logging of selected protocols for all clients managed by delegated administration roles.



Note

Protocol logging is associated with protocol usage alerts. You cannot generate usage alerts for a protocol unless it is set for logging in at least one protocol filter. Enabling the **Lock protocol logging** option through the Filter Lock assures that usage alerts can be generated for the protocol. See [Configuring protocol usage alerts, page 337](#).

1. Select a protocol in the tree.
Delegated administration roles do not have access to custom protocols created by the Super Administrator. Therefore, custom protocols do not appear in this tree.

- Set the restrictions for this protocol in the box that appears beside the protocol tree.

Option	Description
Lock protocol	Blocks and locks access to applications and Web sites using this protocol.
Lock protocol logging	Logs information about access to this protocol, and prevents delegated administrators from disabling logging.
Apply to Group	Applies the same settings to all protocols in the group.

When you are finished making changes, click **OK** to cache the changes and return to the Filter Lock page. Changes are not implemented until you click **Save All**.

Preparing delegated administrators

Related topics:

- ◆ [The fundamentals of delegated administration, page 285](#)
- ◆ [Preparing for delegated administration, page 291](#)
- ◆ [Performing delegated administrator tasks, page 307](#)

After assigning individuals as administrators in any administrative role, make sure to give them the following information:

- ◆ The URL for logging on to TRITON - Web Security. By default:

```
https://<TRITON_location>:9443/triton/
```

 Substitute the IP address or host name of the machine running TRITON - Web Security.
- ◆ What Policy Server to select after logon, if applicable. In an environment with multiple Policy Server instances, administrators can select the Policy Server to use from the Web Security toolbar. They must select the Policy Server that is configured to communicate with the directory service that authenticates their managed clients.
- ◆ Whether to use their network logon account or a local Websense account when logging on to TRITON - Web Security. If administrators log on with local accounts, provide the user name and password.
- ◆ Their permissions: to create and apply policies to clients in the role, generate reports, create policies and generate reports, or audit administrator tasks without implementing changes.

Advise administrators who have both policy and reporting permissions to consider what activities they plan to perform during the session. If they only plan to generate reports, recommend that they go to the **Role** field in the banner, and choose **Release Policy Permissions**. This frees the policy permissions for the

role, enabling another administrator to access TRITON - Web Security and manage policy for that role.

- ◆ How to find the list of clients managed by their role. Administrators can go to Policy Management > Delegated Administration, and then click their role name to display the Edit Role page, which includes a list of managed clients.
- ◆ Limitations imposed by the Filter Lock, if any categories or protocols have been blocked and locked.
- ◆ The tasks that are generally performed by administrators. See [Performing delegated administrator tasks, page 307](#).

Be sure to notify delegated administrators when you add or change custom file types and protocols. These components automatically appear in filters and policies for all roles, so it is important for those administrators to know when changes have been made.

Managing delegated administration roles

Related topics:

- ◆ [The fundamentals of delegated administration, page 285](#)
- ◆ [Preparing for delegated administration, page 291](#)
- ◆ [Managing role conflicts, page 304](#)

The **Policy Management > Delegated Administration** page offers different options, depending on whether it is viewed by a Super Administrator or a delegated administrator.

Super Administrators see a list of all the roles currently defined, and have the following options available.

Option	Description
Add	Click to add a new role. See Adding roles, page 297 .
Role	Click a role name to view or configure the role. See Editing roles, page 298 .
Delete	Mark the check box next to a role name, then click the button to delete the selected roles. Available to unconditional Super Administrators only. See Delete roles, page 305 , for information about how a role's clients are managed after the role is deleted.
Advanced	Click to access the Manage Role Priority function.

Option	Description
Manage Role Priority	Click to specify which role's policy settings are used when the same client exists in multiple groups that are managed by different roles. See Managing role conflicts , page 304.
View Administrator Accounts	Click to see the local and network administrator accounts with TRITON - Web Security access, and review their permission level and role assignments. See Reviewing administrator accounts , page 312.

Delegated administrators see only the roles in which they are administrators, and have access to more limited options.

Option	Description
Role	Click to view the clients assigned to the role, and the specific reporting permissions granted. See Editing roles , page 298.

Adding roles

Related topics:

- ◆ [Preparing for delegated administration](#), page 291
- ◆ [Managing delegated administration roles](#), page 296
- ◆ [Editing roles](#), page 298

Use the **Delegated Administration > Add Role** page to provide a name and description for the new role.

1. Enter a **Name** for the new role.

The name must be between 1 and 50 characters long, and cannot include any of the following characters:

* < > ' { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,

Role names can include spaces and dashes.

2. Enter a **Description** for the new role.

The description may be up to 255 characters. The character restrictions that apply to role names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).

3. Specify the **Role Type**:

- A **Policy management and reporting** role allows administrators the ability to create filters and policies and apply them to manage clients. Administrators in these roles may also be given permission to report on managed clients or all clients.

If you select this role type, also indicate whether or not to **Copy all Super Administrator policies, filters, and filter components to the new role**. If you select this option, the process of creating the role may take several minutes.

If you do not copy all Super Administrator policies to the role, a Default policy is created for the role that enforces the Super Administrator Default category and protocol filters.

- An **Investigative reporting** role allows administrators to report on their managed clients only, using the investigative reports tool. Managed clients in an investigative reporting role may also be added to a policy management and reporting role.
4. Click **OK** to display the **Edit Role** page and define the characteristics of this role. See [Editing roles](#), page 298.
 - If you created a policy management and reporting role, the new role is added to the Role drop-down list in the Web Security toolbar the next time you log on.
 - If you created an investigative reporting role, the name does not appear in the role drop-down. This reflects the fact that reporting permissions are cumulative (see [Administrators in multiple roles](#), page 290).

Editing roles

Related topics:

- ◆ [Managing delegated administration roles](#), page 296
- ◆ [Adding roles](#), page 297
- ◆ [Managing role conflicts](#), page 304

Delegated administrators can use the **Delegated Administration > Edit Role** page to view the list of clients managed by their role, and the specific reporting permissions granted.

Super Administrators can use this page to select the administrators and clients for a role, and to set administrator permissions, as described below. Only unconditional Super Administrators can delete administrators and clients from a role.

1. Change the role **Name** and **Description**, as needed.
The name of the Super Administrator role cannot be changed.
2. Add or remove administrators for this role (Super Administrators only).

Item	Description
User Name	Administrator's user name.
Account Type	Indicates whether the user is defined in the network directory service (Directory) or unique to the TRITON console (Local).

Item	Description
Reporting	Give the administrator permission to use reporting tools.
Real-Time Monitor	Give the administrator permission to monitor all Internet filtering activity for any Policy Server.
Policy	Give the administrator permission to create filters and policies, and apply policies to the role's managed clients. In the Super Administrator role, administrators with policy permission can also manage certain Websense configuration settings. See Super Administrator permissions , page 288.
Auditor	Give the administrator permissions to see all of the features and functions available to other administrators in the role, but without the ability to save changes. The check boxes for other permissions are disabled when Auditor permissions are selected.
Add	Open the Add Administrators page. See Adding Administrators , page 301.
Delete	Remove the selected administrators from the role. <ul style="list-style-type: none"> • Available to unconditional Super Administrators only. • Unconditional Super Administrator accounts can only be removed from the TRITON Settings > Administrators page.

3. Add and delete **Managed Clients** for the role.

Changes can be made by Super Administrators only. Delegated administrators can view the clients assigned to their role.

Item	Description
<Name>	Displays the name of each client explicitly assigned to the role. Administrators in the role must add the clients to the Clients page before policies can be applied. See Performing delegated administrator tasks , page 307.
Add	Opens the Add Managed Clients page. See Adding managed clients , page 302.
Delete	Available to unconditional Super Administrators only, this button removes from the role any clients marked in the managed clients list. Some clients cannot be deleted directly from the managed clients list. See Delete managed clients , page 306, for more information.

4. Use the **Reporting Permissions** area to select the features available to administrators in this role who have reporting access.

- a. Choose the general level of reporting permissions:

Option	Description
Report on all clients	Select this option to give administrators permission to generate reports on all network users. Use the remaining options in the Reporting Permissions area to set the specific permissions for administrators in this role.
Report on managed clients only	Select this option to limit administrators to reporting on the managed clients assigned to this role. Then, select the investigative reports features these administrators can access. Administrators limited to reporting on managed clients only cannot access presentation reports or user-based reports on the Today and History pages.

- b. Mark the check box for each reporting feature that appropriate administrators in the role are permitted to use.

Option	Description
Access presentation reports	Enables access to presentation reports features. This option is available only when administrators can report on all clients. See Presentation reports , page 107.
View reports on Today and History pages	Enables display of charts showing Internet activity on these pages. See Today: Health, Security, and Value Since Midnight , page 24 and History: Last 30 Days , page 27. If this option is deselected, administrators can view only the Health Alert and Value areas of the Today page, and the Value Estimates on the History page.
Access investigative reports	Enables access to basic investigative reports features. When this option is selected, additional investigative reports features can be selected, also. See Investigative reports , page 127.
View user names in investigative reports	Enables administrators in this role to view user names, if they are logged. See Configuring how filtered requests are logged , page 355. Deselect this option to show only system-generated identification codes, instead of names. This option is available only when administrators are granted access to investigative reports.
Save investigative reports as favorites	Enables administrators in this role to create favorite investigative reports. See Favorite investigative reports , page 143. This option is available only when administrators are granted access to investigative reports.

Option	Description
Schedule investigative reports	Enables administrators in this role to schedule investigative reports to run at a future time or on a repeating cycle. See Scheduling investigative reports, page 146 . This option is available only when administrators are granted permissions to save investigative reports as favorites.
Manage the Log Database	Enables administrators to access the Settings > Log Database page. See Log Database administration settings, page 361 .

- When you are finished making changes, click **OK** to cache the changes and return to the Delegated Administration page. Changes are not implemented until you click **Save All**.

Adding Administrators

Related topics:

- ◆ [Delegated administrators, page 287](#)
- ◆ [Editing roles, page 298](#)

Super Administrators can use the **Delegated Administration > Edit Role > Add Administrators** page to specify which individuals are administrators for a role.



Note

Administrators can be added to multiple roles. These administrators must choose a role during logon. In this situation, the administrator receives the combined reporting permissions for all roles.

Delegated administrators have significant control over the Internet activities of their managed clients. To ensure that this control is handled responsibly and in accordance with your organization's acceptable use policies, Super Administrators should use the Audit Log page to monitor changes made by administrators. See [Viewing and exporting the audit log, page 328](#).

- If you plan to assign network accounts as delegated administrators, make sure you are logged on to the Policy Server whose Settings > General > Directory Service configuration (see [Directory services, page 65](#)) matches the TRITON Settings > User Directory configuration.

If you are adding only local accounts as administrators, you can be logged on to any Policy Server.

2. Under **Local Accounts**, mark the check box for one or more users, and then click the right arrow button to move the highlighted users to the **Selected** list.
3. Under **Network Accounts**, mark the check box for one or more users, and then click the right arrow (>) button to move them to the **Selected** list.

**Note**

Custom LDAP groups cannot be added as administrators.

4. Set the **Permissions** for administrators in this role.

Option	Description
Administrator: Policy Management	Let administrators in this role apply policies to their managed clients. This also grants access to certain Websense configuration settings.
Administrator: Reporting	Grant administrators access to reporting tools. Use the Edit Role page to set the specific reporting features permitted.
Administrator: Real-Time Monitor	Allow administrators to monitor Internet filtering traffic in real time. See Real-Time Monitor , page 152.
Auditor	Give the administrator access to view all features available to other administrators in the role, without the ability to save changes.

5. When you are finished making changes, click **OK** to return to the Edit Role page.
6. Click **OK** on the Edit Role page to cache your changes. Changes are not implemented until you click **Save All**.

Adding managed clients

Related topics:

- ◆ [Managing delegated administration roles](#), page 296
- ◆ [Editing roles](#), page 298

Managed clients are the users and computers assigned to a role, whose policies are set by the role's administrators. Directory clients (users, groups, and domains [OUs]), computers (individual IP addresses), and networks (IP address ranges) can all be defined as managed clients.

Super Administrators can use the **Delegated Administration > Edit Role > Add Managed Clients** page to add as many clients to a role as needed. Each client can be assigned to only one policy management and reporting role.

If you assign a network range as managed client in one role, you cannot assign individual IP addresses within that range to any other role. Additionally, you cannot

specifically assign a user, group, or domain (OU) to 2 different roles. However, you can assign a user to one role, and then assign to a different role a group or domain (OU) of which the user is a member.

**Note**

If a group is a managed client in one role, and that role's administrator applies a policy to each member of the group, individual users in that group cannot later be assigned to another role.

When adding managed clients, consider which client types to include.

- ◆ If you add IP addresses to a role, administrators for that role can report on **all** activity for the specified machines, regardless of who is logged on.
- ◆ If you add users to a role, administrators can report on all activity for those users, regardless of the machine where the activity occurred.

Administrators are not automatically included as managed clients in the roles they administer, since that would enable them to set their own policy. To allow administrators to view their own Internet usage, enable self-reporting (see [Self-reporting, page 372](#)).

If your organization has deployed multiple Policy Servers, and the Policy Servers communicate with different directories, be sure to select the Policy Server connected to the directory containing the clients you want to add.

**Note**

Best practices indicate that all directory clients in the same role be defined in the same directory.

1. Select clients for the role:
 - Under **Directory**, mark the check box for one or more users.
If your environment uses Active Directory (Native Mode) or another LDAP-based directory service, you can search the directory to find specific user, group, or domain (OU) names. See [Searching the directory service, page 72](#).
 - Under **Computer**, enter the IP address of a computer to be added to this role.
 - Under **Network**, enter the first and last IP addresses in a range of computers to be added as a unit.
2. Click the right arrow (>) button adjacent to the client type to move the clients to the **Selected** list.
3. When you are finished making changes, click **OK** to return to the Edit Role page.
4. Click **OK** on the Edit Role page to cache your changes. Changes are not implemented until you click **Save All**.

Managing role conflicts

Related topics:

- ◆ [Managing delegated administration roles, page 296](#)
- ◆ [Adding managed clients, page 302](#)

Directory services allow the same user to belong to multiple groups. As a result, a single user may exist in groups that are managed by different delegated administration roles. The same situation exists with domains (OUs).

Additionally, it is possible for a user to be managed by one role, and belong to a group or domain (OU) that is managed by a different role. If the administrators for both of these roles are logged on simultaneously, the administrator responsible for the user could apply policy to that user at the same time as the administrator responsible for the group applies policy to the individual members of the group.

Use the **Delegated Administration > Manage Role Priority** page to tell Websense software what to do if different policies apply to the same user because of an overlap. When a conflict occurs, Websense software applies the filtering policy from the role that appears highest on this list.

1. Select any role on the list, except Super Administrator.



Note

The Super Administrator role is always first on this list. It cannot be moved.

2. Click **Move Up** or **Move Down** to change its position in the list.
3. Repeat steps 1 and 2 until all roles have the desired priority.
4. When you are finished making changes, click **OK** to cache the changes and return to the Delegated Administration page. Changes are not implemented until you click **Save All**.

Updating delegated administration roles

Related topics:

- ◆ [Delete roles, page 305](#)
- ◆ [Delete managed clients, page 306](#)

Policies and managed clients are typically added to a role when the role is created.

- ◆ Delegated administrators with policy permissions can edit existing policies and create new policies within the role that they manage.
- ◆ As new members join the organization, a Super Administrator can add them to existing roles (see [Editing roles](#), page 298).

Super Administrators can also move clients (see [Moving clients to roles](#), page 75) and policies ([Copying filters and policies to roles](#), page 216) from the Super Administrator role to an existing delegated administration role at any time.

- ◆ When a client is moved to a delegated administration role, the policy applied in the Super Administrator role is also copied. During this copy process, the filters are updated to enforce the restrictions of the Filter Lock, if any.

In the target role, the tag “(Copied)” is added to the end of the filter or policy name. Administrators for that role can readily identify the new item and update it appropriately.

Encourage administrators in the role to rename the filters and policies, and to edit them as needed, to clarify their settings and to minimize duplicates. These changes can simplify future maintenance efforts.

After the client is moved to the new role, only an administrator in that role can modify the client’s policy or the filters it enforces. Changes in the original policy or filters in the Super Administrator role do not affect copies of the policy or filters in delegated administration roles.

- ◆ When policies and filters are copied to a delegated administration role directly, the same constraints are enforced that apply when filters and policies are copied as part of moving a client.
 - Filter Lock restrictions are implemented during the copy.
 - Permit All category and protocol filters are renamed, and become editable filters subject to the Filter Lock.
 - Copied filters and policies are identified in the role by the (Copied) tag in the name.

Consider editing policy descriptions before starting the copy, to assure that they are meaningful to the administrators in the target roles.

Delete roles

On the **Delegated Administration** page, unconditional Super Administrators can delete any roles that have become obsolete.

Deleting a role also removes all clients that the role’s administrators have added to the Clients page. After the role is deleted, if those clients belong to any networks, groups, or domains managed by other roles, they are governed by the appropriate policy applied in those roles (see [Filtering order](#), page 85). Otherwise, they are governed by the Super Administrator’s Default policy.

1. On the **Delegated Administration** page, mark the check box beside each role to be deleted.



Note

You cannot delete the Super Administrator role.

2. Click **Delete**.
3. Confirm the delete request to remove the selected roles from the Delegated Administration page. Changes are not permanent until you click **Save All**.
The deleted role is cleared from Role drop-down list in the banner the next time you log on to TRITON - Web Security.

Delete managed clients

Clients cannot be deleted directly from the managed clients list (Delegated Administration > Edit Role) if:

- ◆ the administrator has applied a policy to the client
- ◆ the administrator has applied a policy to one or more members of a network, group, or domain (OU)

There may also be problems if the Super Administrator is connected to a different Policy Server than the one that communicates with the directory service containing the clients to be deleted. In this situation, the current Policy Server and directory service do not recognize the clients.

An unconditional Super Administrator can assure that the appropriate clients can be deleted, as follows.

1. Open the **Policy Server** list in the Web Security toolbar and make sure that you are connected to the Policy Server that communicates with the appropriate directory. You must be logged on with unconditional Super Administrator permissions.
2. Open the **Role** list in the Web Security toolbar, and select the role from which managed clients are to be deleted.
3. Go to **Policy Management > Clients** to see a list of all the clients to which the delegated administrator has explicitly assigned a policy.
This may include both clients that are specifically identified on the role's managed clients list, and clients who are members of networks, groups, domains, or organizational units on the managed clients list.
4. Delete the appropriate clients.
5. Click **OK** to cache the changes.
6. Open the **Role** list in the banner, and select the **Super Administrator** role.
7. Go to **Policy Management > Delegated Administration > Edit Role**.
8. Delete the appropriate clients from the managed clients list, and then click **OK** to confirm the delete request.

9. Click **OK** on the Edit Role page to cache the changes. Changes are not implemented until you click **Save All**.

Managing Super Administrator clients

Clients who are not specifically assigned to a delegated administration role are managed by Super Administrators. There is no Managed Clients list for the Super Administrator role.

To apply policies to these clients, add them to the Policy Management > Clients page. See [Adding a client, page 71](#). Clients who have not been assigned a specific policy are governed by the Super Administrator Default policy.

There may be times when you cannot add clients to the Clients page. This can occur when the client is a member of a network, group, or domain (OU) that is assigned to another role. If the administrator of the other role has applied a policy to individual members of the network or group, those clients cannot be added to the Super Administrator role.

Performing delegated administrator tasks

Delegated administrators who have **policy** permissions can perform the following tasks.

- ◆ [View your user account, page 307](#)
- ◆ [View your role definition, page 308](#)
- ◆ [Add clients to the Clients page, page 309](#)
- ◆ [Create policies and filters, page 310](#)
- ◆ [Apply policies to clients, page 311](#)

Reporting permissions can be granted at a granular level. The specific reporting permissions granted to your role determine which of the following tasks are available to administrators with reporting permissions. See [Generate reports, page 311](#).

View your user account

Related topics:

- ◆ [Performing delegated administrator tasks, page 307](#)
- ◆ [View your role definition, page 308](#)
- ◆ [Add clients to the Clients page, page 309](#)
- ◆ [Create policies and filters, page 310](#)
- ◆ [Apply policies to clients, page 311](#)

If you log on to TRITON - Web Security with network credentials, password changes are handled through your network directory service. Contact your system administrator for assistance.

If you have been assigned a local user name and password, view information about your account and change your password within the TRITON console.

1. Click **TRITON Settings** in the TRITON toolbar, just under the banner.
The My Account page opens.
2. To change your password, first enter your current password, then enter and confirm a new password.
 - The password must be between 4 and 255 characters.
 - Strong passwords are recommended: 8 characters or longer, including at least one uppercase letter, lowercase letter, number, and special character (such as hyphen, underscore, or blank).Click **OK** to save and implement the change.
3. To see a list of roles that you can administrator, go to the TRITON - Web Security Policy Management > Delegated Administration > View Administrator Accounts page.
 - If you are assigned to manage only one role, its name appears in the list.
 - If you are assigned to manage multiple roles, click **View** next to your user name to see them listed.
4. When you are finished, click **Close** to return to the Delegated Administration page.

View your role definition

Related topics:

- ◆ [Performing delegated administrator tasks, page 307](#)
- ◆ [View your user account, page 307](#)
- ◆ [Add clients to the Clients page, page 309](#)
- ◆ [Create policies and filters, page 310](#)
- ◆ [Apply policies to clients, page 311](#)

Open the Delegated Administration page and click your role name to display the Edit Role page, which lists the role's managed clients. This page also shows the reporting features available to administrators who have reporting permissions in this role.

Administrators who have only reporting permissions are unable to view this page. Only the specified reporting features are available to these administrators.

Add clients to the Clients page

Related topics:

- ◆ [Performing delegated administrator tasks, page 307](#)
- ◆ [View your user account, page 307](#)
- ◆ [View your role definition, page 308](#)
- ◆ [Create policies and filters, page 310](#)
- ◆ [Apply policies to clients, page 311](#)

Super Administrators assign managed clients to a role, but delegated administrators must add them to the Clients page before applying policies. See [Adding a client, page 71](#), for instructions.

When clients are added to a managed clients list, they are immediately filtered by a policy in the role.

- ◆ Clients previously assigned a policy within the Super Administrator role are governed by a copy of that policy in the new role. The Move to Role process automatically copies the applicable policy.
- ◆ Clients not previously assigned a policy are filtered by the new role's Default policy. Initially, this Default policy enforces a Default category and protocol filter copied from the Super Administrator role.

Any client listed on the Delegated Administration > Edit Role page for your role can be added to the Clients page and assigned a policy. For groups, domains (OUs), and networks assigned to the role, you can also add:

- ◆ Individual users who members of the group or domain/organizational unit
- ◆ Individual computers that are members of the network

Because a user may be part of multiple groups, domains, or organizational units, adding individuals from a larger client grouping has the potential to create conflicts when different roles manage groups, domains, or organizational units with common members. If administrators in different roles access TRITON - Web Security at the same time, they might add the same client (individual member of a group, for instance) to their Clients page. In that situation, Internet filtering for that client is governed by the priority established for each role. See [Managing role conflicts, page 304](#).

Create policies and filters

Related topics:

- ◆ [Performing delegated administrator tasks](#), page 307
- ◆ [View your user account](#), page 307
- ◆ [View your role definition](#), page 308
- ◆ [Add clients to the Clients page](#), page 309
- ◆ [Apply policies to clients](#), page 311

When your role was created, it automatically inherited the current Default category filter and protocol filter from the Super Administrator role. A role-specific Default policy was created that enforces the inherited Default category and protocol filters. (This role-specific Default policy is automatically applied to any client added to the role until another policy is assigned.)

The Super Administrator may have copied other policies and filters to your role, as well.

In addition to policies and filters, you also inherit any custom file types and protocols created by the Super Administrator.

You can edit inherited policies and filters. Changes you make affect your role only. Any changes the Super Administrator later makes to the original policies and filters do not affect your role.



Note

Changes the Super Administrator makes to file types and protocols automatically affect the filters and policies in your role.

When a Super Administrator informs you of changes to these components, review your filters and policies to be sure they are handled appropriately.

You can also create as many new filters and policies as you need. Filters and policies created by a delegated administrator are available only to administrators logged on to your role. For instructions on creating policies, see [Working with policies](#), page 81. For instructions on creating filters, see [Working with filters](#), page 49.

You can edit filter components for your role, with some limitations.

- ◆ **Categories:** add custom categories, and edit both Master Database and custom categories, defining recategorized URLs and keywords for use within their role; change the action and advanced filtering option applied by default in category filters they create. (Changes to a category's default action are implemented only if the category is not locked by the Filter Lock.)

- ◆ **Protocols:** change the action and advance filtering options applied by default in protocol filters they create. (Changes to a protocol's default action are implemented only if the protocol is not locked by the Filter Lock.) Delegated administrators cannot add or delete protocol definitions.
- ◆ **File types:** view the file extensions assigned to each file type. Delegated administrators cannot add file types or change the extensions assigned to a file type.
- ◆ **Unfiltered URLs:** add URLs and add regular expressions that represent sites to be permitted for all managed clients in their role only.

For more information, see [Building filter components](#), page 217.

If a Super Administrator has implemented Filter Lock restrictions, there may be categories or protocols that are automatically blocked, and cannot be changed in the filters you create and edit.

Apply policies to clients

Related topics:

- ◆ [Performing delegated administrator tasks](#), page 307
- ◆ [View your user account](#), page 307
- ◆ [View your role definition](#), page 308
- ◆ [Add clients to the Clients page](#), page 309
- ◆ [Create policies and filters](#), page 310

After creating a policy, you can apply that policy directly to clients who have already been added to the Clients page by clicking the **Apply to Clients** button. See [Assigning a policy to clients](#), page 85.

Alternatively, you can go to the Clients page and add the clients who should be governed by this policy. See [Working with clients](#), page 62.

Generate reports

If you have reporting permissions, the specific reporting options available are set by the Super Administrator. To learn which features you can use, go to the Delegated Administration page and click the role name. The Edit Role page shows the reporting features for which you have permissions. See:

- ◆ [Presentation reports](#), page 107
- ◆ [Investigative reports](#), page 127
- ◆ [Real-Time Monitor](#), page 152

Reviewing administrator accounts

Use the **Delegated Administration > View Administrator Accounts** page to:

- ◆ See a list of local and network accounts that have been given Web Security access by a Global Security administrator.
- ◆ Check the level of permissions assigned to each account.
- ◆ See a list of roles associated with each account.

If an account has been added to a single role as an administrator, that role is listed to the right of the account name. If the account can be used to manage multiple roles, click **View** to see the roles listed

Delegated administrators see account information for only their own account, and not for all accounts.

When you are finished reviewing administrator accounts, click **Close** to return to the Delegated Administration page.

Enabling network accounts

Global Security Administrators can use the **TRITON Settings > User Directory** page to enter the directory service information needed to allow administrators to log on to TRITON - Web Security with their network credentials.

This task is done **in addition to** the configuration done by Web Security Super Administrators to define the directory service used to identify user and group clients.



Note

Client directory service information is configured on the Settings > Directory Services page (see [Directory services](#), page 65).

TRITON administrators' network credentials must be authenticated against a single directory service. If your network includes multiple directories, a trusted relationship must exist between the directory specified in TRITON Settings and the others.

If it is not possible to define a single directory service for use with the TRITON Unified Security Center, consider creating local accounts for administrators.

Specific instructions for defining the directory used to authenticate administrator logons can be found in the TRITON Settings Help.

14

Web Security Server Administration

Related topics:

- ◆ [Websense Web Security product components, page 314](#)
- ◆ [Working with Policy Server, page 320](#)
- ◆ [Viewing and exporting the audit log, page 328](#)
- ◆ [Stopping and starting Websense services, page 329](#)
- ◆ [Alerting, page 332](#)
- ◆ [Backing up and restoring your Websense data, page 340](#)

Internet usage filtering requires interaction between several Websense Web Security software components:

- ◆ User requests for Internet access are received by Network Agent, Content Gateway, or an integrated third-party product or device (integration).
- ◆ The requests are sent to Filtering Service for processing.
- ◆ Filtering Service communicates with Policy Server and Policy Broker to apply the appropriate policy in response to the request.

A single Policy Database holds client, filter, policy, and general configuration information, whether there is a single Policy Server or multiple Policy Servers.

Each instance of TRITON - Web Security is associated with that Policy Database, and can be used to configure any Policy Server associated with that database.

Because the policy configuration performed in TRITON - Web Security is stored in the central database, policy information is automatically available to all Policy Servers associated with that Policy Database.

Websense Web Security product components

Related topics:

- ◆ [Filtering components, page 315](#)
- ◆ [Reporting components, page 317](#)
- ◆ [User identification components, page 318](#)
- ◆ [Working with Policy Server, page 320](#)
- ◆ [Stopping and starting Websense services, page 329](#)
- ◆ [Reviewing current system status, page 339](#)

Websense Web Security software is made up of several components that work together to provide user identification, Internet filtering, and reporting capabilities. This section provides an overview of each component to help you understand and manage your filtering environment.

The primary Websense components include:

- ◆ Policy Database
- ◆ Policy Broker
- ◆ Policy Server
- ◆ Filtering Service
- ◆ Network Agent
- ◆ Master Database
- ◆ TRITON - Web Security
- ◆ Usage Monitor
- ◆ User Service
- ◆ Log Server
- ◆ Log Database

Websense software also includes optional transparent identification agents:

- ◆ DC Agent
- ◆ RADIUS Agent
- ◆ eDirectory Agent
- ◆ Logon Agent

Additional optional components include:

- ◆ Content Gateway
- ◆ Remote Filtering Server
- ◆ Remote Filtering Client

The following components support interoperability between Websense solutions:

- ◆ Linking Service
- ◆ Directory Agent
- ◆ Sync Service

When Websense software is integrated with a firewall, proxy, caching device, or similar product, a filtering plug-in may also be installed.

Filtering components

Component	Description
Policy Database	Stores Websense software settings and policy information. Installed automatically with Policy Broker.
Policy Broker	Manages requests from Websense components for policy and general configuration information.
Policy Server	<ul style="list-style-type: none"> ◆ Identifies and tracks the location and status of other Websense components. ◆ Stores configuration information specific to a single Policy Server instance. ◆ Communicates configuration data to Filtering Service, for use in filtering Internet requests. <p>Configure Policy Server settings in TRITON - Web Security (see Working with Policy Server, page 320).</p> <p>Policy and most configuration settings are shared between Policy Servers that share a Policy Database (see Working in a multiple Policy Server environment, page 322).</p>
Filtering Service	<p>Provides Internet filtering in conjunction Network Agent or a third-party integration product. When a user requests a site, Filtering Service receives the request and determines which policy applies.</p> <ul style="list-style-type: none"> ◆ Filtering Service must be running for Internet requests to be filtered and logged. ◆ Each Filtering Service instance downloads its own copy of the Websense Master Database. <p>Configure filtering and Filtering Service behavior in TRITON - Web Security (see Internet Usage Filters, page 39, and Configuring Websense filtering settings, page 57).</p>
Network Agent	<ul style="list-style-type: none"> ◆ Enhances filtering and logging functions ◆ Enables protocol management ◆ Enables filtering in a stand-alone environment <p>For more information, see Network Configuration, page 375.</p>

Component	Description
Master Database	<ul style="list-style-type: none"> • Includes more than 36 million Web sites, sorted into more than 90 categories and subcategories • Contains more than 100 protocol definitions for use in filtering protocols <p>Download the Websense Master Database to activate Internet filtering, and make sure that the database is kept up to date. If the Master Database is more than 2 weeks old, no filtering can occur. See <i>The Websense Master Database</i>, page 34, for more information.</p>
TRITON - Web Security	<p>Serves as the configuration, management, and reporting interface for Websense software.</p> <p>Use TRITON - Web Security to define and customize Internet access policies, configure Websense software components, report on Internet filtering activity, and more.</p> <p>TRITON - Web Security is made up of the following services:</p> <ul style="list-style-type: none"> • Websense - TRITON Web Security • Websense Web Reporting Tools • Websense Explorer Report Scheduler • Websense Information Service for Explorer • Websense Reporter Scheduler <p>See <i>Working in TRITON - Web Security</i>, page 18, for more information.</p>
Usage Monitor	<ul style="list-style-type: none"> • Enables alerting based on Internet usage. • Provides Internet usage information to Real-Time Monitor. <p>Usage Monitor tracks URL category access (shown in Real-Time Monitor) and protocol access, and generates alert messages according to the alerting behavior you have configured. See <i>Alerting</i>, page 332, and <i>Real-Time Monitor</i>, page 152, for more information.</p>
Content Gateway	<ul style="list-style-type: none"> • Provides a robust proxy and cache platform. • Can analyze the content of Web sites and files in real time to categorize previously uncategorized sites. • Enables protocol management <p>See <i>Scanning and SSL Bypass Options</i>, page 155.</p> <p>As part of a Websense Web Security Gateway deployment, also:</p> <ul style="list-style-type: none"> • Analyzes HTML code to find security threats (for example, phishing, URL redirection, Web exploits, and proxy avoidance). • Inspects file content to assign a threat category (for example, viruses, Trojan horses, or worms). • Strips active content from certain Web pages.

Component	Description
Remote Filtering Client	<ul style="list-style-type: none"> • Resides on client machines outside the network firewall. • Identifies the machines as clients to be filtered, and communicates with Remote Filtering Server. <p>See Filter Users Off Site, page 201, for more information.</p>
Remote Filtering Server	<ul style="list-style-type: none"> • Allows filtering of clients outside a network firewall. • Communicates with Filtering Service to provide Internet access management of remote machines. <p>See Filter Users Off Site, page 201, for more information.</p>

Reporting components

Component	Description
Log Server	<p>Logs Internet request data, including:</p> <ul style="list-style-type: none"> • The request source • The category or protocol associated with the request • Whether the request was permitted or blocked • Whether keyword blocking, file type blocking, quota allocations, bandwidth levels, or password protection were applied <p>With Network Agent and some integration products, Log Server also stores information about the amount of bandwidth used.</p> <p>Log Server is a Windows-only component that must be installed to enable all reporting features of TRITON - Web Security.</p> <p>After installing Log Server, configure Filtering Service to pass logging data to the correct location (see Configuring how filtered requests are logged, page 355).</p>
Log Database	Stores Internet request data collected by Log Server for use by Websense reporting tools.
Real-Time Monitor	<p>Displays current URL filtering activity, including:</p> <ul style="list-style-type: none"> • Request source (user name or IP address) • URL (full or domain only) • Category (Master Database, custom URL, or dynamic, based on Content Gateway scanning) • Whether the request was permitted or blocked • Time of the request <p>Real-Time Monitor is made up of 3 services:</p> <ul style="list-style-type: none"> • Websense RTM Client • Websense RTM Server • Websense RTM Database <p>See Real-Time Monitor, page 152.</p>

User identification components

Component	Description
User Service	<ul style="list-style-type: none"> • Communicates with your directory service. • Conveys user-related information, including user-to-group and user-to-domain relationships, to Policy Server and Filtering Service, for use in applying filtering policies. <p>If you have installed and configured a Websense transparent identification agent (see Transparent identification, page 245), User Service helps to interpret user logon session information, and uses this information to provide user name-to-IP-address associations to Filtering Service.</p> <p>When you add users and groups as Websense clients (see Adding a client, page 71), User Service provides name and path information from the directory service to TRITON - Web Security.</p> <p>For information about configuring directory service access, see Directory services, page 65.</p>
DC Agent	<ul style="list-style-type: none"> • Offers transparent user identification for users in a Windows-based directory service. • Communicates with User Service to provide up-to-date user logon session information to Websense software for use in filtering. <p>For more information, see DC Agent, page 255.</p>
Logon Agent	<ul style="list-style-type: none"> • Provides unsurpassed accuracy in transparent user identification in Linux and Windows networks. • Does not rely on a directory service or other intermediary when capturing user logon sessions. • Detects user logon sessions as they occur. <p>Logon Agent communicates with the logon application on client machines to ensure that individual user logon sessions are captured and processed directly by Websense software.</p> <p>For more information, see Logon Agent, page 258.</p>
eDirectory Agent	<ul style="list-style-type: none"> • Works with Novell eDirectory to transparently identify users. • Gathers user logon session information from Novell eDirectory, which authenticates users logging on to the network. • Associates each authenticated user with an IP address, and then works with User Service to supply the information to Filtering Service. <p>For more information, see eDirectory Agent, page 265.</p>
RADIUS Agent	<p>Enables transparent identification of users who use a dial-up, Virtual Private Network (VPN), Digital Subscriber Line (DSL), or other remote connection to access the network.</p> <p>For more information, see RADIUS Agent, page 261.</p>

Interoperability components

Component	Description
Directory Agent	In Websense Web Security Gateway Anywhere deployments, collects user and group information from a supported directory service for use in filtering by the hybrid service.
Filtering Plug-In	When Websense software is integrated with certain firewall, proxy, cache, or similar products, a filtering plug-in may be installed to enable communication between Filtering Service and the integration.
Linking Service	In Websense Web Security Gateway Anywhere deployments, or in environments that combine Websense data and Web security solutions, gives data security software access to Master Database categorization information and user and group information collected by User Service.
Sync Service	In Websense Web Security Gateway Anywhere deployments: <ul style="list-style-type: none"> • Sends policy updates and user and group information to the hybrid service. • Receives reporting data from the hybrid service.

Understanding the Policy Database

Websense Policy Database stores both policy data (including clients, filters, filter components, and delegated administration settings) and global settings configured in TRITON - Web Security. Settings specific to a single Policy Server instance (like its Filtering Service and Network Agent connections) are stored separately.

In multiple Policy Server environments, a single Policy Database holds policy and general configuration data for all Policy Server instances.

1. At startup, each Websense component requests applicable configuration information from the Policy Database via the Policy Broker.
2. Running components frequently check for changes to the Policy Database.
3. The Policy Database is updated each time administrators make changes in TRITON - Web Security and click Save All.
4. After a change to the Policy Database, each component requests and receives the changes that affect its functioning.

Back up the Policy Database on a regular basis to safeguard important configuration and policy information. See [Backing up and restoring your Websense data](#), page 340, for more information.

Working with Policy Server

Policy Server communicates with Filtering Service to aid in policy enforcement and is responsible for identifying other Websense software components and tracking their location and status.

When you log on to TRITON - Web Security, you are logging onto a graphical interface to Policy Server.

- ◆ You cannot log on to TRITON - Web Security until it is configured to communicate with Policy Server.
- ◆ If your Websense software installation includes multiple Policy Servers, you can choose between Policy Server instances at logon time.
- ◆ You can add and remove Policy Server instances within TRITON - Web Security.

Communication between TRITON - Web Security and one Policy Server instance is established during TRITON - Web Security installation.

Most environments require only one Policy Server. A single Policy Server can communicate with multiple Filtering Service and Network Agent instances for load balancing. In very large organizations (10,000+ users), however, it may help to install multiple instances of Policy Server. If you install additional Policy Servers, add each instance to TRITON - Web Security (see [Reviewing Policy Server connections](#), page 320).

Reviewing Policy Server connections

Use the **Settings > General > Policy Servers** page to review Policy Server information for all Policy Server instances associated with TRITON - Web Security.

If you have multiple Policy Server instances that share a subscription key, you can create one instance as the primary Policy Server, and the others as secondary Policy Servers. This may save you some time in configuring the Policy Server connections.

To see the secondary Policy Server instances associates with a primary Policy Server in the list, click the “+” symbol next to the Policy Server name or IP address.

For each listed Policy Server, the table shows both a short description and a communication port. Primary Policy Server entries also include the subscription key associated with the instance and its secondaries, as well as the subscription level (for example, Web Security or Web Security Gateway) of the key.

- ◆ Click **Add** to associate an additional Policy Server with this TRITON - Web Security instance. See [Adding or editing Policy Server instances](#), page 321.
- ◆ Click a Policy Server IP address or name to edit configuration information for the selected instance. See [Adding or editing Policy Server instances](#), page 321.

TRITON - Web Security is associated with one Policy Server instance at installation time. This becomes the base Policy Server for the TRITON - Web Security installation, and its IP address and description cannot be changed.

You can, however, change the subscription key associated with the base Policy Server.

- ◆ Mark one or more Policy Server entries and click **Delete** to remove the connection between TRITON - Web Security and the selected Policy Server.

Clicking Delete removes the Policy Server instance from TRITON - Web Security, but does not uninstall or stop the Websense Policy Server service. You cannot delete the base Policy Server instance.

After adding or editing a Policy Server connection, click **OK** on the Policy Servers page to cache your changes. Changes are not implemented until you click **Save All**.

Adding or editing Policy Server instances

Use the Settings > General > **Add Policy Server** or **Edit Policy Server** page to associate a new Policy Server instance with TRITON - Web Security, or to update configuration information for an existing Policy Server.

1. Enter or edit the **IP address or name** and communication **Port** for the Policy Server instance. The default port is **55806**.
2. Enter or update the **Description** of the selected Policy Server instance.
You cannot change the description for the base Policy Server.
3. Indicate whether this is a **Primary** or **Secondary** Policy Server.
 - A primary Policy Server has a different subscription key than other Policy Server instances associated with TRITON - Web Security.
 - A secondary Policy Server uses the same subscription key as another Policy Server that has already been associated with TRITON - Web Security.
4. If this is a **secondary** Policy Server, select the IP address of the primary Policy Server from which the secondary should get its key, and then click **OK** to return to the Policy Servers page.

You must also click **OK** on the Policy Servers page to cache your changes. Changes are not implemented until you click **Save All**.

5. If this is a **primary** Policy Server, indicate whether to **Use the current subscription key** registered to the new instance or **Enter a subscription key**.
 - If you are editing an existing entry, the current subscription key and subscription type are displayed below the radio buttons.
 - Click **Verify Policy Server** to make sure that TRITON - Web Security can communicate with the new Policy Server. If you have selected “Use the current subscription key,” and the connection is successful, the subscription key is displayed.
 - If you are not sure whether the new Policy Server instance already has a key registered, you can either select the option to enter the key manually, or click **Verify Policy Server** to see if TRITON - Web Security finds an existing key for the instance.
6. Click **OK** to return to the Policy Servers page. You must click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Working in a multiple Policy Server environment

In some distributed environments with a large number of users, it may be appropriate to install multiple Policy Servers. This entails some special considerations.

- ◆ If you implement a configuration that allows the same client to be managed by different Policy Servers, depending on current load, do **not** implement time-based policy actions:

- Password Override
- Account Override
- Confirm
- Quota

The timing information associated with these features is not shared among Policy Servers, and clients could be granted more or less Internet access than you intend.

Remember that the Default policy is enforced whenever no other policy applies to a client. If clients can be governed by more than one Policy Server, you may want to make sure that the Default policy does not enforce category filters that apply time-based actions.

- ◆ Because policy information is stored in the Policy Database, policy changes are automatically shared between all Policy Servers when you click **Save All**.
- ◆ Many global configuration settings (like risk class definitions and alerting options) are also shared between Policy Servers.
- ◆ Configuration settings that are specific to a single Policy Server (like its Filtering Service and Network Agent connections) are stored locally by each Policy Server and not distributed.

To switch between Policy Servers in TRITON - Web Security to review or configure settings that apply to a single Policy Server instance:

1. In the Websense banner, expand the **Policy Server** list and select an IP address.
2. If there are unsaved changes to the current Policy Server instance, a warning prompt appears. Do one of the following:
 - Click **Cancel** to remain logged on to the current Policy Server so that you can save your changes.
 - Click **OK** to abandon the changes and log on to the new Policy Server.If there are no unsaved changes, you are taken directly to the logon screen.
3. At the logon screen, enter a user name and password to log on to the selected Policy Server, and then click **Log On**.

Changing the Policy Server IP address

Before changing the IP address of the Policy Server machine, **stop all Websense services** on the machine. If TRITON - Web Security is also installed on the machine, this includes the both the Websense TRITON - Web Security and the Websense Web Reporting Tools services.

After changing the IP address, you must manually update Websense configuration files used by TRITON - Web Security, Policy Server, and other Websense services before filtering resumes.

Step 1: Update TRITON - Web Security configuration

Update TRITON - Web Security to use the new IP address to connect to Policy Server.

1. On the TRITON - Web Security machine, stop the **Websense Web Reporting Tools** and **Websense TRITON - Web Security** services (if necessary).

If TRITON - Web Security and Policy Server are installed on this same machine, these services should already be stopped.

2. Navigate to the following directory:

```
websense\Web Security\tomcat\conf\Catalina\localhost\
```

3. Locate the **mng.xml** file, and then make a backup copy of the file in another directory.
4. Open **mng.xml** in a text editor (like Notepad or vi) and replace each instance of the old Policy Server IP address with the new one.
The Policy Server IP address appears twice: as the **ps/default/host** value and the **psHosts** value.
5. When you are finished, save and close the file.

Do not restart the TRITON - Web Security services until you have completed the remaining configuration updates in this section.

Step 2: Update Policy Server configuration

Update the Policy Server configuration file, and the initialization file used to configure communication between Websense components.

1. If you have not already done so, stop all Websense services on the Policy Server machine (see *Stopping and starting Websense services*, page 329).
2. Navigate to the Websense **bin** directory (see *Where is the Websense "bin" directory?*, page 457).
3. Locate the **config.xml** file, and then make a backup copy of the file in another directory.
4. Open **config.xml** in a text editor and replace each instance of the old Policy Server IP address with the new one.
5. When you are finished, save and close the file.
6. In the **bin** directory, locate the **websense.ini** file, and then make a backup copy in another directory.
7. Open **websense.ini** in a text editor and replace each instance of the old Policy Server IP address with the new one.
8. When you are finished, save and close the file.

Step 3: Verify the Log Database connection

Use the Windows ODBC Data Source Administrator on the Policy Server machine to verify the ODBC connection to the Log Database.

1. Go to **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.
2. On the **System DSN** tab, select the appropriate data source name (by default, **wslogdb70**), and then click **Configure**.
3. Verify that the correct database server machine is selected, and then click **Next**.
4. Enter the credentials used to connect to the database, and then click **Next**.
5. Accept the defaults on the next 2 screens, and then click **Test Data Source**.



Note

If the test fails, check the database server machine name and try again.

If the machine name is correct, but the test continues to fail, verify that the correct connection port is being used, and that the firewall allows communication on the selected port.

Step 4: Restart Websense services

1. Reboot the Policy Server machine. Make sure that all Websense services on the machine restart normally.
2. If the TRITON - Web Security used to configure this Policy Server is installed on another machine, restart the **Websense Web Reporting Tools** and **Websense TRITON - Web Security** services on that machine.



Note

If TRITON - Web Security is installed on the same machine as Policy Server, administrators must use the new IP address to log on.

Working with Filtering Service

Filtering Service is the Websense software component that works with Network Agent, Content Gateway, or a third-party integration product to filter Internet activity. When a user requests a site, Filtering Service receives the request, determines which policy applies, and uses the applicable policy to determine how the site is filtered.

Each Filtering Service instance downloads its own copy of the Websense Master Database to use in determining how to filter Internet requests.

Filtering Service also sends information about Internet activity to Log Server, so that it can be recorded and used for reporting.

When you log on to TRITON - Web Security, a **Filtering Service Summary** on the Status > Today page lists the IP address and current status of each Filtering Service instance associated with the current Policy Server. Click a Filtering Service IP address for more detailed information about the selected Filtering Service.

Review Filtering Service details

Use the **Status > Today > Filtering Service Details** page to review the status of an individual Filtering Service instance. The page lists:

- ◆ The Filtering Service IP address
- ◆ Whether or not the selected instance is running
- ◆ The Filtering Service version
This should match your Websense software version, including any hotfixes that have been applied.
- ◆ The operating system of the Filtering Service machine
- ◆ The Websense software platform
This indicates whether Websense software is running in standalone mode or integrated with Content Gateway or a third-party product.
- ◆ The IP address and status of any Network Agent instances with which the selected Filtering Service communicates.
- ◆ The IP address and status of any Content Gateway instances with which the selected Filtering Service communicates

Click **Close** to return to the Today page.

Review Master Database download status

Each Filtering Service instance in your network downloads its own copy of the Master Database. When you are working in TRITON - Web Security, the Health Alert Summary on the Status > Today page displays a status message when a Master Database download is in progress, or if a download attempt fails.

For detailed information about recent or ongoing database downloads, click **Database Download** on the Today page toolbar. The Database Download page includes an entry for each Filtering Service instance associated with the current Policy Server.

Initially, the Database Download page displays a quick download summary, showing where the database was downloaded, which database version was downloaded, and whether the download was successful. From this summary view, you can:

- ◆ Initiate a database download for a single Filtering Service (click **Update**).
- ◆ Initiate database downloads for all listed Filtering Service instances (click **Update All**).
- ◆ Cancel one or all ongoing updates.

Click an IP address in the list on the right to review more detailed database download status for the selected Filtering Service.

- ◆ If the selected Filtering Service has encountered download problems, a recommendation for addressing the problem may be displayed.
- ◆ To manually initiate a database download for the selected Filtering Service, click **Update**.

During database download, the status screen shows detailed progress information for each stage of the download process. Click **Close** to hide progress information and continue working in TRITON - Web Security.

Resuming Master Database downloads

If a Master Database download is interrupted, Websense software attempts to resume the download automatically. If Filtering Service is able to reconnect to the download server, the download resumes from where it was interrupted.

You can manually restart a failed or interrupted download. This does not resume the download from the point of interruption, but instead restarts the process from the beginning.

1. In TRITON - Web Security, go to **Status > Today** and click **Database Download**.
2. Click **Stop All Updates** to stop the interrupted process.
3. Select a Filtering Service instance and click **Update**, or click **Update All**, to restart the download process from the beginning.

Working with Content Gateway

Related topics:

- ◆ [Working with Filtering Service, page 324](#)
- ◆ [Working with Policy Server, page 320](#)
- ◆ [Managing Content Gateway connections, page 327](#)

Content Gateway is a Linux-only Websense software component that provides high-performance Web proxy services in Websense Web Security Gateway and Gateway Anywhere deployments. Content Gateway is also used as a proxy by Websense Data Security and Email Security Gateway solutions.

In Websense Web Security Gateway and Gateway Anywhere deployments, Content Gateway provides:

- ◆ Real-time content scanning and Web site classification to protect the network from malicious Web content. This is especially valuable for Web 2.0 sites, whose multiple sources and dynamic nature limit the usefulness of static categorization.

- ◆ Advanced file scanning to discover and block infected and malicious files from being uploaded or downloaded
- ◆ Detection of inbound and outbound protocols tunneled over HTTP and HTTPS and apply protocol filtering

Content Gateway works with Filtering Service to filter Internet requests based on both:

- ◆ Static categorization by the Master Database or custom URL definitions
- ◆ Dynamic recategorization resulting from content scanning and analysis

At installation, Content Gateway establishes communication with a Policy Server instance. This connection:

- ◆ Allows Policy Server to pass subscription key information to Content Gateway, mitigating the need to maintain keys in 2 management consoles
- ◆ Provides TRITON - Web Security with information about Filtering Service connections to Content Gateway
- ◆ Is used to populate the Settings > General > Content Gateway Access page in TRITON - Web Security, and makes it possible to launch Content Gateway Manager from within TRITON

Managing Content Gateway connections

Use the **Settings > General > Content Gateway Access** page to review configuration and status information for Content Gateway instances associated with the current Policy Server, or to launch Content Gateway Manager for a selected instance.

When a Content Gateway instance is registered with a Policy Server, the Content Gateway Access page is automatically updated with IP address, host name, and status information for that Content Gateway. This information appears in one of 3 tables:

- ◆ If the Content Gateway is part of a cluster, a table is displayed with the cluster name as its title. All Content Gateway instances in the cluster are listed. If there are multiple clusters, multiple tables will appear.
- ◆ If the Content Gateway is not clustered, it is shown in the Unclustered Content Gateway instances table.
- ◆ If Policy Server cannot communicate with a Content Gateway instance, it appears in the Not Responding table. This table is only displayed when Policy Server cannot communicate with a registered Content Gateway instance.

To launch Content Gateway Manager for any listed instance, click the corresponding link in the **IP Address** column of the table.

To update the description of an instance, to make it easier to manage Content Gateway connections, mark the radio button next to an instance IP address and click **Edit Description**.

If a Content Gateway instance appears in the **Not Responding** table because the instance has been uninstalled or relocated, mark the radio button next to the instance name and click **Delete**.

After editing Content Gateway descriptions or deleting obsolete entries, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Viewing and exporting the audit log

Websense software provides an audit trail showing which administrators have accessed TRITON - Web Security, as well as any changes made to policies and settings. This information is available only to Super Administrators who are granted policy permissions (see [Super Administrator permissions](#), page 288).

Delegated administrators have significant control over the Internet activities of their managed clients. Monitoring their changes through the audit log enables you to ensure that this control is handled responsibly and in accordance with your organization's acceptable use policies.

Use the **Status > Audit Log** page to view the audit log, and to export selected portions of it to an Excel spreadsheet (XLS) file, if desired.

Audit records are saved for 60 days. To preserve audit records longer than 60 days, use the export option to export the log on a regular basis. Exporting does not remove records from the audit log.

When the Audit Log page opens, the most recent records are shown. Use the scroll bar and the paging buttons above the log to view older records.

The log displays the following information. If an item is truncated, click the partial entry to display the full record in pop-up dialog box.

Column	Description
Date	Date and time of the change, adjusted for time zones. To assure consistent data in the audit log, be sure all machines running Websense components have their date and time settings synchronized.
User	User name of the administrator who made the change.
Server	IP address or name of machine running the Policy Server affected by the change. This appears only for changes that affect the Policy Server, such as changes made on the Settings tab.
Role	Delegated administration role affected by the change. When a change affects a client explicitly assigned as a managed client in the delegated administrator's role, that change shows as affecting the Super Administrator role. If the change affects a client that is a member of a network range, group, domain or organizational unit assigned to the role, the change shows as affecting the delegated administrator's role.
Type	Configuration element that was changed, such as policy, category filter, or logon/logoff.

Column	Description
Element	Identifier for the specific object changed, such as the category filter name or role name.
Action	Type of change made, such as add, delete, change, log on, and so on.
Previous	Value before the change.
Current	New value after the change.

Not all items are shown for all records. For example, the role is not displayed for logon and logoff records.

To export audit log records:

1. Select a time period from the **Export range** list.
Choose **Last 60 days** to export the entire audit log file.
2. Click **Go**.
If Microsoft Excel is installed on the machine running TRITON - Web Security, the exported file opens. Use options in Excel to save or print the file.
If Microsoft Excel is not installed on the machine running TRITON - Web Security, follow the on-screen instructions to either locate the software or save the file.

Stopping and starting Websense services

Websense services are configured to start each time the machine restarts. However, in some cases you need to stop or start one or more product components separately from a machine restart.



Note

If Filtering Service is in the process of downloading the Master Database, it does not stop running until the download is complete.

When you stop **all Websense services**, always end with the policy services, in the order shown:

1. Websense Policy Server
2. Websense Policy Broker
3. Websense Policy Database

Note that unless a problem specifically pertains to Policy Broker or the Policy Database, it is rarely necessary to restart these services. Avoid restarting these services when possible.

When you start all Websense services, always start with the policy services, in the reverse of the shutdown order (starting with Policy Database and ending with Policy Server).

When you stop the services associated with **Real-Time Monitor**:

- ◆ Also stop the TRITON - Web Security services (Websense TRITON - Web Security and Websense Web Reporting Tools).
- ◆ Stop the Real-Time Monitor services in the order shown:
 1. Websense RTM Client
 2. Websense RTM Server
 3. Websense RTM Database

Start the Real-Time Monitor services in the reverse of shutdown order (starting with RTM Database and ending with RTM Client).

Windows

1. Open the Windows Services dialog box (**Start > Settings > Control Panel > Administrative Tools > Services**).
2. Right-click the Websense service name, and then select **Stop** or **Start**.

Linux

On Linux machines, there are 2 tools that can be used to stop and start daemons:

- ◆ The **WebsenseAdmin** script starts, stops, and restarts **all** daemons on the machine.
- ◆ The **WebsenseDaemonControl** script starts and stops **individual** daemons.



Warning

Do not use the **kill** command to stop a Websense service, as it may corrupt the service.

To use the WebsenseAdmin script to start or stop all daemons:

1. Go to the **/opt/Websense** directory.
2. Check the status of the Websense services with the following command:

```
./WebsenseAdmin status
```
3. Stop, start, or restart all Websense services with the commands:
 - `./WebsenseAdmin stop`
 - `./WebsenseAdmin start`
 - `./WebsenseAdmin restart`

To use the WebsenseDaemonControl script to start or stop a daemon:

1. Go to the `/opt/Websense` directory.
2. Enter the following command:

```
./WebsenseDaemonControl
```

A list of installed components is displayed, showing whether each process is running or stopped.
3. Enter the letter associated with a component to start or stop the associated process. To refresh the list, enter **R**.
4. When you are finished, enter **Q** or **X** to exit the tool.

V-Series appliance

On Websense V-Series appliances, use Appliance Manager to stop, start, and restart Websense services.

To restart services:

1. Go to the Status > General page. This page is displayed by default when you log on to Appliance Manager.
2. Scroll to the Network Agent section and click **Restart Module**.
3. When the Network Agent module has restarted, go to the Web Security section and click **Restart Module**.

To stop the services (perhaps while performing a maintenance task):

1. Scroll to the Network Agent section of the Status > General page and click **Stop Services**.
2. In the Websense Web Security section, also click **Stop Services**.
3. When you are ready start the services again:
 - a. Go to the Websense Web Security section and click **Start Services**.
 - b. Go to the Network Agent section and click **Start Services**.

Websense Web Security installation directories

The Websense Web Security installation directory depends on the machine operating system, the operating system version, and whether you have a new installation or have upgraded existing components.

On **Windows** machines, the default installation directories are:

- ◆ For most upgraded components:
`C:\Program Files\Websense\`
- ◆ For new components on 32-bit platforms:
`C:\Program Files\Websense\Web Security\`

- ◆ For new components on Windows Server 2008 R2:
C:\Program Files (x86)\ Websense\Web Security\

On **Linux** machines, the default installation directory is:

/opt/Websense/

Alerting

Related topics:

- ◆ [Flood control, page 333](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Configuring system alerts, page 335](#)
- ◆ [Configuring category usage alerts, page 336](#)
- ◆ [Configuring protocol usage alerts, page 337](#)

To facilitate tracking and management of both Websense software and client Internet activity, Super Administrators can configure alerts to be sent when selected events occur.

- ◆ **System alerts:** Notification regarding Web Security events relating to subscription status and Master Database activity, as well as Content Gateway events, including loss of contact to a domain controller, log space issues, and more.
- ◆ **Usage alerts:** Notification when Internet activity for particular categories or protocols reaches configured thresholds.

Alerts can be sent to selected recipients via email, on-screen pop-up messages (Windows **net send** messaging), or SNMP messages.



Note

On-screen pop-up alerts cannot be sent to Linux or Windows 2008 machines.

Pop-up alerts **can** be sent from a Linux machine running Policy Server to Windows machines, provided that the Samba client is installed on the Linux machine. See the *Deployment and Installation Center*.

Usage alerts can be generated for both Websense-defined and custom categories or protocols.

Flood control

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Configuring general alert options, page 333](#)

There are built-in controls for usage alerts to avoid generating excessive numbers of alert messages. Use the **Maximum daily alerts per usage type** setting to specify a limit for how many alerts are sent in response to user requests for particular categories and protocols. See [Configuring general alert options, page 333](#), for more information.

You can also set threshold limits for each category and protocol usage alert. For example, if you set a threshold limit of 10 for a certain category, an alert is generated after 10 requests for that category (by any combination of clients). See [Configuring category usage alerts, page 336](#), and [Configuring protocol usage alerts, page 337](#), for more information.

Suppose that the maximum daily alerts setting is 20, and the category alert threshold is 10. Administrators are only alerted the first 20 times category requests exceed the threshold. That means that only the first 200 occurrences result in alert messages (threshold of 10 multiplied by alert limit of 20).

Configuring general alert options

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Configuring system alerts, page 335](#)
- ◆ [Configuring category usage alerts, page 336](#)
- ◆ [Configuring protocol usage alerts, page 337](#)

Websense software can notify administrators of various kinds of system events, such as updates to Master Database categories and subscription issues, as well as Internet usage that exceeds defined thresholds.

Use the **Settings > Alerts > Enable Alerts** page to select and configure the desired notification methods, as described below. Then, use the other pages in the Settings > Alerts section to enable the alerts you want to receive.

1. Enter a number in the **Maximum daily alerts per usage type** field to limit the total number of alerts generated daily for each category and protocol usage alert.
For example, you might configure usage alerts to be sent every 5 times (threshold) someone requests a site in the Sports category. Depending on the number of users and their Internet use patterns, that could generate hundreds of alerts each day.

If you enter 10 as the maximum daily alerts per usage type, only 10 alert messages are generated each day for the Sports category. In this example, these messages alert you to the first 50 requests for Sports sites (5 requests per alert multiplied by 10 alerts).

2. Mark the **Enable email alerts** check box to deliver alerts and notifications by email. Then, configure these email settings.

SMTP server IP or name	IP address or host name for the SMTP server through which email alerts should be routed.
From email address	Email address to use as the sender for email alerts.
Administrator email address (To)	Email address of the primary recipient of email alerts.
Recipient email addresses (Cc)	Email address for up to 50 additional recipients. Each address must be on a separate line.

3. Mark the **Enable pop-up alerts** check box to display pop-up messages on specific computers. Then, enter the IP address or machine name for up to 50 **Recipients**, each on a separate line.



Note

On-screen pop-up alerts cannot be sent to Linux or Windows 2008 machines.

Pop-up alerts **can** be sent from a Linux machine running Policy Server to Windows machines, provided that the Samba client is installed on the Linux machine. See the *Deployment and Installation Center*.

4. Mark the **Enable SNMP alerts** check box to deliver alert messages through an SNMP Trap system installed in your network. Then, provide information about your SNMP Trap system.

Community name	Name of the trap community on your SNMP Trap server.
Server IP or name	IP address or name of the SNMP Trap server.
Port	Port number SNMP messages use.

5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring system alerts

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Reviewing current system status, page 339](#)

TRITON - Web Security displays detailed system health and status information via the **Status > Alerts** page, described in [Reviewing current system status, page 339](#).

To assure that administrators are notified of significant system events, configure Websense system alerts to be distributed by email, pop-up message, or through your SNMP trap system.

Websense Web Security Gateway and Gateway Anywhere administrators have the option to enable system alerts for both Web Security events (related to subscription and database download issues) and Content Gateway events for a variety of issues.

Use the **Settings > Alerts > System** page to specify which alerts to send, and select the methods used to send each notification.

To enable an alert, mark one or more check box to the right of the message summary to indicate how to notify administrators. Depending on what methods are enabled on the Alerts page, you may be able to choose **Email**, **Pop-up**, and **SNMP**.

To disable an alert, clear all check boxes to the right of the message summary.

All alerts are enabled, by default. If you have provided SMTP information for email notifications, 4 Web Security events cannot be disabled:

- ◆ A Websense Master Database download failed.
- ◆ The number of current users exceeds your subscription level.
- ◆ Your subscription expires in one month.
- ◆ Your subscription expires in one week.

There are also 3 optional alerts:

- ◆ The number of current users has reached 90% of your subscription level.
- ◆ The search engines supported by Search Filtering have been changed.
- ◆ The Websense Master Database has been updated.

In Websense Web Security Gateway and Gateway Anywhere environments, you have the option to enable the following additional system alerts:

- ◆ A domain controller is down.
- ◆ A scanning database download failed.
- ◆ Decryption and inspection of secure content has been disabled.
- ◆ Log space is critically low.

- ◆ Subscription information could not be retrieved.
- ◆ The connection limit is approaching, and connections will be dropped.
- ◆ Non-critical alerts have been received. (See [Content Gateway non-critical alerts, page 443](#), for information about conditions that can trigger this alert.)

When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring category usage alerts

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Flood control, page 333](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Adding category usage alerts, page 337](#)

Websense software can notify you when Internet activity for particular URL categories reaches a defined threshold. You can define alerts for permitted requests or for blocked requests to the category.

For example, you might want to be alerted each time 50 requests for sites in the Shopping category have been permitted to help decide whether to place restrictions on that category. Or, you might want to receive an alert each time 100 requests for sites in the Entertainment category have been blocked, to see whether users are adapting to a new Internet use policy.

On the Settings tab, use the **Alerts > Category Usage** page to view the alerts that have already been established, and to add or delete usage alert categories.

1. View the **Permitted Category Usage Alerts** and **Blocked Category Usage Alerts** lists to learn which categories are configured for alerts, the threshold for each, and the selected alert methods.
2. Click **Add** below the appropriate list to open the Add Category Usage Alerts page (see [Adding category usage alerts, page 337](#)) and configure additional URL categories for alerting.
3. Mark the check box for any categories you want to delete from its list, and then click **Delete** below the appropriate list.
4. When you are finished, click **OK** to cache your changes and return to the Category Usage page. Changes are not implemented until you click **Save All**.

Adding category usage alerts

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Configuring category usage alerts, page 336](#)

The **Add Category Usage Alerts** page appears when you click Add on the Category Usage page. Here, you can select new categories for usage alerts, establish the threshold for these alerts, and select the alert methods.

1. Mark the check box beside each category to be added with the same threshold and alert methods.



Note

You cannot add usage alerts for any category that is excluded from logging. See [Configuring how filtered requests are logged, page 355](#).

2. Set the **Threshold** by selecting the number of requests that cause an alert to be generated.
3. Mark the check box for each desired alert method (**Email, Pop-up, SNMP**) for these categories.
Only the alert methods that have been enabled on the Alerts page (see [Configuring general alert options, page 333](#)) are available for selection.
4. Click **OK** to cache your changes and return to the Category Usage page (see [Configuring category usage alerts, page 336](#)). Changes are not implemented until you click **Save All**.

Configuring protocol usage alerts

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Flood control, page 333](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Adding protocol usage alerts, page 338](#)

Websense software can notify you when Internet activity for a particular protocol reaches a defined threshold. You can define alerts for permitted or blocked requests for the selected protocol.

For example, you might want to be alerted each time 50 requests for a particular instant messaging protocol are permitted to help decide whether to place restrictions on that protocol. Or, you might want to receive an alert each time 100 requests for a particular peer-to-peer file sharing protocol have been blocked, to see whether users are adapting to a new Internet use policy.

On the Settings tab, use the **Alerts > Protocol Usage** page to view the alerts that have already been established, and to add or delete protocols for usage alerts.

1. View the **Permitted Protocol Usage Alerts** and **Blocked Protocol Usage Alerts** lists to learn which protocols are configured for alerts, the threshold for each, and the selected alert methods.
2. Click **Add** below the appropriate list to open the Add Protocol Usage Alerts page (see [Adding protocol usage alerts, page 338](#)) and configure additional protocols for alerting.
3. Select the check box for any protocols you want to delete, and then click **Delete** under the appropriate list.
4. When you are finished, click **OK** to cache your changes and return to the Protocol Usage page. Changes are not implemented until you click **Save All**.

Adding protocol usage alerts

Related topics:

- ◆ [Alerting, page 332](#)
- ◆ [Configuring general alert options, page 333](#)
- ◆ [Configuring protocol usage alerts, page 337](#)

Use the **Protocol Usage > Add Protocol Usage Alerts** page to select new protocols for usage alerts, establish the threshold for these alerts, and select the alert methods.

1. Mark the check box beside each protocol to be added with the same threshold and alert methods.



Notes

You cannot select a protocol for alerting unless it is configured for logging in one or more protocol filters.

Protocol alerts only reflect usage by clients governed by a protocol filter that logs the protocol.

2. Set the **Threshold** by selecting the number of requests that cause an alert to be generated.
3. Select each desired alert method (**Email, Pop-up, SNMP**) for these protocols. Only the alert methods that have been enabled on the Alerts page (see [Configuring general alert options, page 333](#)) are available for selection.

4. Click **OK** to cache changes and return to the Protocol Usage page (see [Configuring protocol usage alerts, page 337](#)). Changes are not implemented until you click **Save All**.

Reviewing current system status

Use the **Status > Alerts** page to find information about problems affecting the health of your Websense software, get troubleshooting help, and review the details of recent real-time updates to the Websense Master Database.

The **Active Alerts** list shows the status of monitored Websense software components.

- ◆ For detailed information about which components are monitored, click **What is monitored?** above the list of alert messages.
- ◆ To troubleshoot a problem, click the **Solutions** button next to the error or warning message.
- ◆ To hide an alert message, click **Advanced**. If your organization does not use Log Server, Network Agent, or User Service, or if you do not plan to enable WebCatcher, mark a check box to hide the associated alert. When you are finished, click **OK** to enact the change.

Click **Advanced** again to hide the advanced options.

The **Real-Time Database Updates** list provides information about emergency updates to the Websense Master Database, showing:

- ◆ When the update occurred
- ◆ The update type
- ◆ The new database version number
- ◆ The reason for the update
- ◆ The IP address of the Filtering Service instance that received the update

These supplemental updates occur in addition to regular, scheduled Master Database updates, and can be used, for example, to recategorize a site that has been temporarily miscategorized. Websense software checks for database updates every hour.

For Websense Web Security users, the Alerts page includes a third list: **Real-Time Security Updates**. This list has the same format as the Real-Time Database Updates list, but specifically shows security-related database updates.

Installing security updates as soon as they are created eliminates vulnerability to threats such as new phishing (identity fraud) scams, rogue applications, or malicious code infecting a mainstream Web site or application.

For more information about Real-Time Security Updates, see [Real-Time Security Updates™, page 35](#).

Use the **Print** button, above the page, to open a secondary window with a printable version of the Alerts area. Use browser options to print this page, which omits all the navigation options found in the main TRITON - Web Security window.

Backing up and restoring your Websense data

Related topics:

- ◆ [Scheduling backups](#), page 342
- ◆ [Running immediate backups](#), page 343
- ◆ [Maintaining the backup files](#), page 344
- ◆ [Restoring your Websense data](#), page 345
- ◆ [Discontinuing scheduled backups](#), page 346
- ◆ [Command reference](#), page 346

The Websense Backup Utility makes it easy to back up your Websense software settings and policy data, and to revert to a previous configuration. Data saved by the utility can also be used to import Websense configuration information after an upgrade.



Important

Make sure that all administrators log off of TRITON - Web Security before you back up or restore your configuration.

The Backup Utility saves:

- ◆ Global configuration information, including client and policy data, stored in the Policy Database.
- ◆ Local configuration information, such as Filtering Service and Log Server settings, stored by each Policy Server.
- ◆ Websense component initialization and configuration files.

The backup process works as follows:

1. You initiate an immediate backup (see [Running immediate backups](#), page 343) or define a backup schedule (see [Scheduling backups](#), page 342).
 - Manually launch a backup at any time.
 - Backup files are stored in a directory you specify when you run or schedule the backup.
2. The Backup Utility checks all Websense components on the machine, collects the data eligible for backup, and creates an archive file. The file name is given the format:

```
wsbackup_YYYY-mm-dd_hhmmss.tar.gz
```

Here, `YYYY-mm-dd_hhmmss` represents the date and time of the backup. **tar.gz** is a portable compressed file format.

Only root (Linux) and members of the Administrators group (Windows) can access the backup files.

Run the Websense Backup Utility on each machine that includes Websense components. The tool identifies and saves any of the following files that it finds on the current machine:

Path	File name
\Program Files\Websense\bin <i>or</i> /opt/Websense/bin	authserver.ini BrokerService.cfg config.xml eimserver.ini LogServer.ini netcache.conf securewispproxy.ini transid.ini upf.conf websense.ini WebUI.ini wsauthserver.ini wscitrix.ini WSE.ini wsedir.ini wsradius.ini wsufpserver.ini
bin/i18n	i18n.ini
bin/postgres/data	postgresql.conf pg_hba.conf
BlockPages/*/Custom	All custom block page settings
tomcat/conf/Catalina/Localhost	mng.xml
Windows\system32	isa_ignore.txt ignore.txt
/etc/wsLib	wsSquid.ini

Store Websense backup files in a safe and secure location. These files should be part of your organization's regular backup procedures.

To revert to an earlier configuration:

1. Retrieve the backup files from their storage site.
2. Copy each backup file to the Websense machine on which it was created.

3. Run the Backup Utility in restore mode.

**Important**

Always use the Backup Utility to restore a Websense software configuration. Do not extract the files from the archive using other extraction utilities.

If the backup file is corrupted, you will not be able to restore your settings.

During the restore process, any error messages or warnings are displayed on the machine where the restore is being run.

Scheduling backups

Related topics:

- ◆ [Running immediate backups, page 343](#)
- ◆ [Maintaining the backup files, page 344](#)
- ◆ [Restoring your Websense data, page 345](#)
- ◆ [Discontinuing scheduled backups, page 346](#)
- ◆ [Command reference, page 346](#)

Notify Websense administrators of the backup schedule, so that they can be sure to log out of TRITON - Web Security during the backup process.

To schedule backups:

- ◆ Windows:
 1. Open a command prompt and navigate to the Websense **bin** directory (see [Where is the Websense "bin" directory?, page 457](#)).

2. Enter the following command.

```
wsbackup -s -t "<m> <h> <day_of_month> <month>  
<day_of_week>" -d <directory>
```

- ◆ Linux:

1. Open a command shell and navigate to the **Websense** directory (/opt/Websense/, by default).

2. Enter the following command:

```
./WebsenseTools -b -s -t "<minute> <hour>  
<day_of_month> <month> <day_of_week>"  
-d <directory>
```

Note that the time information uses **crontab** format, and the quotation marks and spaces are required.

In place of the variables shown in the example, provide the following information:

Variable	Information
<m>	0 - 59 Specify the precise minute to start the backup.
<h>	0 - 23 Specify the general hour of the day to start the backup.
<day_of_month>	1 - 31 Specify the date to perform the backup. If you schedule a backup for days 29 - 31, the utility uses the standard substitution procedure for the operating system in months that do not include that date.
<month>	1 - 12 Specify the month to perform the backup.
<day_of_week>	0 - 6 Specify a day of the week. 0 represents Sunday.

Each field can take a number, an asterisk, or a list of parameters. Refer to any **crontab** reference for details.

Running immediate backups

Related topics:

- ◆ [Scheduling backups, page 342](#)
- ◆ [Maintaining the backup files, page 344](#)
- ◆ [Restoring your Websense data, page 345](#)
- ◆ [Discontinuing scheduled backups, page 346](#)
- ◆ [Command reference, page 346](#)

Before running the Backup Utility, make sure that all administrators are logged out of TRITON - Web Security.

To launch an immediate backup:

- ◆ Windows:
 1. Open a command prompt and navigate to the Websense **bin** directory (see [Where is the Websense "bin" directory?, page 457](#)).
 2. Enter the following command.


```
wsbackup -b -d <directory>
```
- ◆ Linux:
 1. Open a command shell and navigate to the **Websense** directory (/opt/Websense/, by default).

2. Enter the following command:

```
./WebsenseTools -b -b -d <directory>
```

Here, *directory* indicates the destination directory for the backup archive.



Warning

Do not store backup files in the Websense **bin** directory. This directory is deleted if you uninstall your Websense software.

When you initiate an immediate backup, any error messages and notifications are displayed on the console of the machine running the backup.

Maintaining the backup files

Related topics:

- ◆ [Scheduling backups, page 342](#)
- ◆ [Running immediate backups, page 343](#)
- ◆ [Restoring your Websense data, page 345](#)
- ◆ [Discontinuing scheduled backups, page 346](#)
- ◆ [Command reference, page 346](#)

When you perform a backup, a configuration file (**WebsenseBackup.cfg**) is created and stored with the backup archive. This configuration file specifies:

- ◆ How long to keep the backup archive in the backup directory
- ◆ The maximum amount of disk space that may be consumed by all backup files in the directory

Edit the **WebsenseBackup.cfg** file in any text editor to change either of these parameters:

Parameter	Value
KeepDays	Number of days archive files should remain in the backup directory. The default is 365.
KeepSize	Number of bytes allotted for backup files. The default is 10857600.

Any files older than the **KeepDays** value are deleted from the backup directory. If the amount of allotted disk space is exceeded, the oldest files are deleted from the backup directory to make room for newer files.

Restoring your Websense data

Related topics:

- ◆ [Scheduling backups](#), page 342
- ◆ [Running immediate backups](#), page 343
- ◆ [Maintaining the backup files](#), page 344
- ◆ [Discontinuing scheduled backups](#), page 346
- ◆ [Command reference](#), page 346

When you restore Websense configuration data, make sure that you are restoring data for the components that exist on the current machine. Also make sure that all administrators are logged off of TRITON - Web Security.

If you run the restore process on the Policy Broker machine, once the restore is complete, restart all Websense services in your deployment. This includes the services on and off the Policy Broker machine.

To initiate the restore process:

- ◆ Windows:
 1. Open a command prompt and navigate to the Websense **bin** directory (see [Where is the Websense “bin” directory?](#), page 457).
 2. Enter the following command:

```
wsbackup -r -f archive_file.tar.gz
```
- ◆ Linux:
 1. Open a command shell and navigate to the **Websense** directory (/opt/Websense/, by default).
 2. Enter the following command:

```
./WebsenseTools -b -r -f archive_file.tar.gz
```



Important

The restore process may take several minutes. Do not stop the process while restoration is underway.

During the restore process, the Backup Utility stops all Websense services. If the utility is unable to stop the services, it sends a message asking the user to manually stop them. Services must be stopped in the order described in [Stopping and starting Websense services](#), page 329.

The Backup Utility saves some files used for communication with third-party integration products. Because these files reside outside the Websense directory structure, you must restore them manually, by copying each file to the correct directory.

Files that must be restored manually include:

File name	Restore to
isa_ignore.txt	Windows\system32
ignore.txt	Windows\system32\bin
wsSquid.ini	/etc/wsLib

Discontinuing scheduled backups

Related topics:

- ◆ [Scheduling backups, page 342](#)
- ◆ [Running immediate backups, page 343](#)
- ◆ [Maintaining the backup files, page 344](#)
- ◆ [Restoring your Websense data, page 345](#)
- ◆ [Command reference, page 346](#)

To clear the backup schedule and stop running currently scheduled backups, open a command shell and navigate to the Websense bin directory (see [Where is the Websense "bin" directory?, page 457](#)). Enter the following command.:

```
wsbackup -u
```

Command reference

Related topics:

- ◆ [Scheduling backups, page 342](#)
- ◆ [Running immediate backups, page 343](#)
- ◆ [Maintaining the backup files, page 344](#)
- ◆ [Restoring your Websense data, page 345](#)
- ◆ [Discontinuing scheduled backups, page 346](#)

Only root (Linux) or a member of the Administrators group (Windows) can run the Backup Utility.

The **wsbackup** and **WebsenseTools -b** commands take the following options:

- ◆ `-b or --backup`
- ◆ `-d directory_path or --dir directory_path`
- ◆ `-f full_file_name or --file full_file_name`
- ◆ `-h or --help or -?`

- ◆ `-r or --restore`
- ◆ `-s or --schedule`
- ◆ `-t or --time`
- ◆ `-u or --unschedule`
- ◆ `-v or --verbose [0...3]`

15

Reporting Administration

Related topics:

- ◆ [Planning your configuration](#), page 349
- ◆ [Managing access to reporting tools](#), page 352
- ◆ [Basic configuration](#), page 352
- ◆ [Log Server Configuration utility](#), page 357
- ◆ [Log Database administration settings](#), page 361
- ◆ [Configuring investigative reports](#), page 368
- ◆ [Self-reporting](#), page 372

To use Websense presentation reports and investigative reports, you must install Websense Log Server (a Windows-only reporting component). You also must configure Websense software to log Internet filtering activity.

When logging is enabled, records of Internet activity are sent to Log Server, which processes them into a Log Database. The Log Database is created within a supported database engine: Microsoft SQL Server or SQL Server Express. See the *Deployment and Installation Center* for supported versions and information about installing reporting components.

When you generate a report, TRITON - Web Security uses filters that you define to display appropriate information from the Log Database in the format that you specify.

Planning your configuration

Depending on the volume of Internet traffic in your network, the Log Database can become very large. To help determine an effective logging and reporting strategy for your organization, consider these questions:

- ◆ When is the network traffic busiest?
Consider scheduling resource intensive database jobs and reporting jobs at times when the traffic volume is lower. This improves logging and reporting

performance during peak periods. See *Configuring Internet browse time options*, page 366, and *Configuring Log Database maintenance options*, page 364.

- ◆ How long should log data be kept to support historical reporting?
Consider automatically deleting partitions after they reach this age. This reduces the amount of disk space required for the Log Database. See *Configuring Log Database maintenance options*, page 364.
- ◆ How much detail is really needed?
Consider which logging options to activate: logging full URLs and hits increase the Log Database size. To decrease Log Database size, consider:
 - disabling full URL logging (see *Configuring how URLs are logged*, page 365)
 - logging visits instead of hits (see *Log Server Configuration utility*, page 357)
 - enabling consolidation (see *Log Server Configuration utility*, page 357)
 - enabling selective category logging (see *Configuring how filtered requests are logged*, page 355)

Successful reporting implementations are deployed on hardware that matches or exceeds the requirements for expected load and for historical data retention.

Configuring distributed logging

If you have a large or distributed environment that requires multiple Log Server instances, but you want all of your reporting data centralized, you have 2 options:

- ◆ Configure all Log Server instances to independently record their data in the same Log Database.
- ◆ Configure distributed Log Server instances to pass their data to a central Log Server, which then records all log records from all instances into the Log Database.

The first option does not require special configuration steps. You need only ensure that each Log Server instance points to the same database (both database engine IP address or host name and database instance name).

The second option requires more planning and configuration detail, as outlined in the sections that follow.

Centralized logging configuration

1. Identify or create a domain user account that will be used to run each Log Server service.
This ensures that permissions are consistent for all instances, and facilitates communication between distributed Log Server instances and the central instance.
2. Identify which Log Server instance will serve as the central Log Server.
All remote Log Server instances must be able to communicate with the central Log Server machine.
3. Create a shared folder on the central Log Server machine that all other Log Server instances will be able to access:

- a. Create the folder (for example, C:\Program Files\WebSense\bin\logs_cache\).
- b. Right-click the new folder and select **Properties**. On the **Sharing** tab, select **Share this folder** and provide the information requested.

Optionally, also restrict access to the folder to the domain user account assigned to all Log Server instances.

The shared folder is available within the network via its UNC file path (`\\<host_name>\<folder_name>`).

4. On the remote Log Server machines, create a mapped drive for the cache folder created in step 3:
 - a. Log on to each Log Server machine as the domain user assigned to all Log Server instances.
 - b. Open Windows Explorer and go to **Tools > Map Network Drive**.
 - c. Select a drive letter for the mapped drive, then browse to the shared folder created in step 3, then click **Finish**.
 - d. Make sure that you can copy a small text file from the remote Log Server machine to the shared drive.
5. Go to the central Log Server machine and stop the Websense Log Server service (see *Stopping and starting Websense services*, page 329), then navigate to the Websense **bin** directory (see *Where is the Websense "bin" directory?*, page 457) and open the **LogServer.ini** file in a text editor.
6. Search for the phrase "Centralized LogServer," then make the following changes:

```
[CacheFileWatcher]
Active=true
TimeInterval=180
FilePath=<path_to_folder_from_step_3>
```

Optionally, you can also edit the **TimeInterval** value to determine how frequently (in seconds) the central Log Server checks the cache directory for new files to process.

If you find that reporting on visits is too inconsistent when centralized logging is used, do to the way that log cache files are processed, you can also edit the **[Visits]** section of the file to change the **UsingVisits** value to **false**. (This can also be configured via the Log Server Configuration utility).

7. Save and close the file, then start the Websense Log Server service.
8. Go to a remote Log Server machine and stop the Websense Log Server service, then open the **LogServer.ini** file for that instance in a text editor.
9. Search for the phrase "Remote LogServer," then make the following changes:

```
[CacheLogging]
Active=true
...
CacheFileOutputPath=<UNC_path_to_mapped_drive>
```

Optionally, change the **TimeInterval** value to determine how often (in seconds) the remote Log Server closes the current cache file and creates a new one.

You can also edit the **MinFileSize** and **MaxFileSize** (in bytes) for each cache file. The default minimum is 1 MB; the default maximum is 5 MB.

Note that **CacheFileProcessingPath** shows where cache files are created on the local machine before being sent to the mapped drive for processing.

If you have turned off visits processing for the central Log Server service, also make the same change in the **[Visits]** section of the INI file for the remote Log Server instance.

10. Save and close the file, then start the Websense Log Server Service.
Repeat the process for each remote Log Server machine.

Managing access to reporting tools

When Log Server (a Windows-only component) is installed, a variety of reporting options are enabled in TRITON - Web Security.

In organizations that use only the admin account, everyone who uses TRITON - Web Security has access to all reporting options, including the Today and History charts, Toolbox options, presentation reports, investigative reports, and reporting settings.

In organizations that use delegated administration, access to reporting tools is controlled by members of the Super Administrator role. When creating a role, the Super Administrator designates whether administrators in that role have access to specific reporting options (see *Editing roles*, page 298).

When you install Log Server, the Log Server Configuration utility is also installed (see *Log Server Configuration utility*, page 357). This graphical tool, available from the Windows Start menu, makes it easy to change the way that Log Server connects and sends data to the Log Database. You must have access to the Log Server machine to access the utility.

Basic configuration

Related topics:

- ◆ *Configuring how filtered requests are logged*, page 355
- ◆ *Assigning categories to risk classes*, page 353
- ◆ *Configuring reporting preferences*, page 354
- ◆ *Log Server Configuration utility*, page 357
- ◆ *Log Database administration settings*, page 361

You can use a variety of configuration options to customize reporting for your environment.

- ◆ The Websense Master Database organizes categories into **risk classes**. Risk classes suggest possible types or levels of vulnerability posed by sites in those

categories. Use the Settings > General > Risk Classes page to customize risk classes for your organization. See [Assigning categories to risk classes, page 353](#).

- ◆ Use the Settings > Reporting > Preferences page to configure the email server used to distribute reports, activate self-reporting, and configure how long scheduled reports are stored on the TRITON - Web Security machine. See [Configuring reporting preferences, page 354](#).

Logging is the process of storing information about Websense filtering activities in a Log Database so that you can generate reports.

- ◆ Use the Settings > General > Logging page to enable logging, select the categories to be logged, and determine what user information is logged. See [Configuring how filtered requests are logged, page 355](#), for more information.
- ◆ Use the Log Server Configuration utility to manage the way the log records are processed and connections to the Log Database. See [Log Server Configuration utility, page 357](#).
- ◆ Use the Settings > Reporting > Log Database page to administer the Log Database, including Internet browse time controls, database partition options, and error logs. See [Log Database administration settings, page 361](#).

Assigning categories to risk classes

Related topics:

- ◆ [Risk classes, page 43](#)
- ◆ [Block Pages, page 91](#)
- ◆ [Use Reports to Evaluate Filtering, page 105](#)

The Websense Master Database organizes categories into **risk classes**. Risk classes suggest possible types or levels of vulnerability posed by sites in those categories.

Risk classes are used primarily in reporting. The Today and History pages offer charts where Internet activity is tracked by risk class, and you can generate presentation or investigative reports organized by risk class.

Unconditional Super Administrators can view or change which categories comprise each risk class on the **Settings > General > Risk Classes** page. For example, some businesses may consider user-posted video sites to fall under the risk classes of legal liability, network bandwidth loss, and productivity loss. However, if your company does market research on a certain demographic, you might consider these part of the Business Usage risk class.

Risk class information in Websense reports reflects the assignments you make on this page.

1. Select an entry in the **Risk Classes** list.
2. Review the **Categories** list to see which categories are currently included in that risk class.

A check mark shows that the category is currently assigned to the selected risk class. The blue W icon indicates categories that are included in the risk class by default.

3. Mark or clear entries in the category tree to include or exclude a category from the selected risk class. Categories can belong to more than one risk class.

Other choices include:

Option	Description
Select All	Selects all categories in the tree.
Clear All	Deselects all categories in the tree.
Restore Defaults	Resets the category choices for the selected risk class to those provided by the Websense software. A blue W icon indicates a default category.

4. Repeat this process for each risk class.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring reporting preferences

Related topics:

- ◆ [Self-reporting, page 372](#)
- ◆ [Scheduling presentation reports, page 118](#)
- ◆ [Scheduling investigative reports, page 146](#)

Use the **Settings > Reporting > Preferences** page to provide information used to send completed scheduled reports to selected recipients via email, activate self-reporting, and determine how long scheduled presentation reports are stored.

Also use this page to configure when Real-Time Monitor starts to capture data.

1. Enter the **Email address** to appear in the from field when scheduled reports are distributed via email.
2. Enter the **SMTP server IP or name** for the email server used to distribute scheduled reports via email.
3. Mark the **Allow self-reporting** check box to permit end users in your organization to access TRITON - Web Security and run investigative reports on their personal Internet activity. See [Self-reporting, page 372](#).
4. Use the **Store reports for** drop-down list to indicate how long scheduled reports are stored on the TRITON - Web Security machine (5 days, by default).

Note that as you increase the length of time that reports are stored, you affect the amount of disk space required on the TRITON - Web Security machine. The

TRITON - Web Security machine is not an appropriate location for a long-term reporting archive.

5. Use the **Warn administrators...** drop-down list to indicate how long before a report is deleted a warning is displayed on the Review Reports page (3 days, by default).

The warning is intended to give administrators time to archive important reports in an appropriate location before they are deleted from the TRITON - Web Security machine.

6. Under Real-Time Monitor, select a radio button to determine when Real-Time Monitor starts to capture user data:
 - Select **Capture data only when Real-Time Monitor is active** (default) to improve system performance and only see data for the period beginning when Real-Time Monitor is launched. This may mean that there is a slight delay (of a few seconds) before records start appearing on the screen.
 - Select **Always capture data** to have the Real-Time Monitor client continually process data into the RTM database, even when no one is viewing the data. This may have a noticeable effect on system performance.
7. Click **Save Now** to implement your changes.

Configuring how filtered requests are logged

Related topics:

- ◆ [Introducing the Log Database, page 360](#)
- ◆ [Log Server Configuration utility, page 357](#)

Use the **Settings > General > Logging** page to:

- ◆ Provide the IP address and port that Filtering Service uses to send log records to Log Server.
- ◆ (*Websense Web Security Gateway Anywhere*) Provide the port that Sync Service uses to send hybrid log records to Log Server.
- ◆ Specify what client-identifying information, if any, Filtering Service sends to Log Server for use in reporting.
- ◆ Determine which URL categories are logged for use in reporting and category usage alerting (see [Configuring category usage alerts, page 336](#)).

In an environment with multiple Policy Servers, configure the Logging page separately for each Policy Server instance. All Filtering Service instances associated with the active Policy Server send their log records to the Log Server identified on this page.

When working with multiple Policy Servers, note that:

- ◆ For reporting data to display properly in TRITON - Web Security, there must be a Log Server associated with the base Policy Server (the Policy Server instance

specified during TRITON - Web Security installation; noted on the Settings > General > Policy Server page).

- ◆ If the Log Server IP address and port are blank for any Policy Server, the Filtering Services associated with that Policy Server cannot log any traffic for reporting or alerts.
- ◆ Information about whether or not user names and IP addresses are logged is stored centrally, so the same settings are used throughout your deployment.
Likewise, any changes you make to how categories are logged is shared by all Filtering Service and Log Server instances.

If your environment includes both multiple Policy Servers and multiple Log Servers, make sure you log on to each Policy Server separately, and verify that it is communicating with the correct Log Server.

1. Enter the **IP address or name** of the Log Server machine.



Important

If Log Server is installed on a separate machine from Policy Server, this entry may default to localhost. If this happens, enter the correct IP address of the Log Server machine to enable display of charts on the Today and History pages, as well as other reporting features.

2. Enter the port that Filtering Service uses to send log records to Log Server.
3. (*Websense Web Security Gateway Anywhere*) Enter the port that Sync Service uses to send log records from the hybrid service to Log Server.
4. Click **Check Status** to determine whether TRITON - Web Security is able to communicate with Log Server using the specified location and port.
A message indicates whether the connection test passed. Update the IP address or host name and port, if needed, until the test is successful.
5. Specify how much user data is stored in log records and displayed in reports:
 - To log identifying information for machines accessing the Internet, mark **Log IP addresses**.
 - To log identifying information for users accessing the Internet, mark **Log user names**.



Note

If you do not log IP addresses or user names, there can be no user data in your reports. This is sometimes called **anonymous logging**.

6. Use the **Selective Category Logging** list to indicate any URL categories that should not be logged. Changes made here apply to all category filters in all active policies.

**Note**

If you disable logging for categories that have usage alerts set up (see [Configuring category usage alerts](#), page 336), no usage alerts can be sent.

Reports cannot include information on categories that are not logged.

- Expand parent categories as needed to change logging for subcategories.
 - Clear the check box next to a category name to stop logging the category. You must select or deselect each category separately. Selecting a parent category does not automatically select its subcategories. Use **Select All** and **Clear All** to assist with selections.
7. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Log Server Configuration utility

Related topics:

- ◆ [Stopping and starting Log Server](#), page 359

During installation, you configure certain aspects of Log Server operation, including how Log Server interacts with Websense filtering components. Use the Log Server Configuration utility to change these settings when needed, and configure other details about Log Server operation.

This utility is installed on the same machine as Log Server. Refer to the Log Server Configuration Help for detailed information.

1. From the Windows Start menu, select **Programs > Websense > Web Security > Web Security Log Server Configuration**.

The Log Server Configuration utility opens with the Connections tab selected.

2. Select a tab to display its options and make any changes.
 - Use the **Connection** tab to configure connections between Log Server and filtering components.
 - Use the **Database** tab to configure how Log Server works with the Log Database.

- Use the **Settings** tab to manage log cache file creation options, and to specify whether Log Server tracks the individual files that make up each Web site requested, or just the Web site.
- Use the **Consolidation** tab to enable consolidation of log records (reducing the size of the Log Database) and set consolidation preferences.



Note

With Websense Web Security Gateway Anywhere, the results of enabling consolidation may vary depending on whether users are filtered by on-premises components or the hybrid service.

Because reporting data is collected from the hybrid service at configured intervals, records for hybrid filtering user browsing sessions that overlap those intervals may be partially consolidated, while records for on-premises users will be fully consolidated.

- Use the **WebCatcher** tab to choose which types of URLs to send to Websense, Inc., for analysis, and to set the file size and processing time for sending URLs.
3. Click **Apply** to save the changes.
 4. Use the **Connection** tab to stop and restart Log Server for the changes to take effect.



Important

To avoid restarting Log Server multiple times, make and apply all of your changes before restarting Log Server.

Configuring WebCatcher

WebCatcher is an optional feature that collects uncategorized sites and security URLs, and submits them to Websense, Inc., where they are analyzed for potential security and liability risks, and for categorization. (Full URL logging is not required for WebCatcher processing.) Websense, Inc., reviews information and updates the Master Database with newly categorized URLs.

Choose the types of URLs to send, and set the file size and processing time on the **WebCatcher** tab of the Log Server Configuration utility.

**Note**

If you have multiple Log Server instances, enable WebCatcher for only one instance. Once WebCatcher is enabled for one instance, the WebCatcher tab of the Log Server Configuration utility is disabled for other Log Server instances.

The information sent to Websense, Inc., contains only URLs and does not include user information, as shown below. The IP address in this example is for the machine hosting the URL (target), not the requestor (source).

```
<URL HREF="http://www.ack.com/uncategorized/" CATEGORY="153"
IP_ADDR="200.102.53.105" NUM_HITS="1" />
```

WebCatcher data is sent to Websense, Inc., via HTTP Post. You may need to create roles or make other changes on your proxy server or firewall to permit the outgoing HTTP traffic. Refer to the proxy server or firewall documentation for instructions.

Use the WebCatcher tab of the Log Server Configuration utility (Start > Programs > Web Security) to specify whether to send URL data to Websense, Inc. If you enable WebCatcher, you can also determine whether uncategorized sites, security risk sites, or both are sent.

Stopping and starting Log Server

Related topics:

- ◆ [Log Server Configuration utility, page 357](#)

Log Server receives information from Filtering Service and saves it in the Log Database for use when generating reports. It runs as a Windows service, typically started during installation, and starts any time you restart the machine.

Changes you make in the Log Server Configuration utility take effect only after you stop and restart Log Server. This can be done easily through the Connection tab in the Log Server Configuration utility.

1. From the Windows Start menu, select **Programs > Websense > Web Security > Web Security Log Server Configuration**.
2. In the **Connections** tab, click **Stop**.
3. Wait several seconds, and then click **Start** to restart the Log Server service.

4. Click **OK** to close the Log Server Configuration utility.

**Note**

Websense software cannot log Internet access that occurs while the Log Server is stopped.

Introducing the Log Database

Related topics:

- ◆ [Database jobs, page 361](#)
- ◆ [Log Database administration settings, page 361](#)

The Log Database stores the records of Internet activity and the associated Websense filtering actions. Installation creates the Log Database with a catalog database and one database partition.

The **catalog database** provides a single connection point for the various Websense components that need to access the Log Database: Status pages, Log Server, presentation reports, and investigative reports. It contains supporting information for the database partitions, including the list of category names, risk class definitions, the mapping of users to groups, database jobs, and so forth. The catalog database also maintains a list of all the available database partitions.

Database partitions store the individual log records of Internet activity. New partitions are created based on size or date interval. See [Configuring database partition options, page 362](#), for more information.

When partitions are based on size, all incoming log records are inserted into the most recent active partition that satisfies the size rule. When the partition reaches the designated maximum size, a new partition is created for inserting new log records.

When the partitions are based on date, new partitions are created according to the established cycle. For example, if the rollover option is monthly, a new partition is created as soon as any records are received for the new month. Incoming log records are inserted into the appropriate partition based on date.

Database partitions provide flexibility and performance advantages. For example, you can generate reports from a single partition to limit the scope of data that must be analyzed to locate the requested information.

Database jobs

The following database jobs are installed along with the Log Database.



Important

If you are using Microsoft SQL Server 2005 or 2008 (not Express), the SQL Server Agent service must be running on the database engine machine. Make sure that this service is configured to start automatically when SQL Server or the machine is restarted.

- ◆ The Extract, Transform, and Load (ETL) job runs continuously, receiving data from Log Server, processing it, and then inserting it into the partition database. The ETL job must be running to process log records into the Log Database.
- ◆ The database maintenance job performs database maintenance tasks and preserves optimal performance. This job runs nightly, by default.
- ◆ The Internet browse time (IBT) job analyzes the data received and calculates browse time for each client. The IBT database job is resource intensive, affecting most database resources. This job runs nightly, by default.

Certain aspects of these database jobs can be configured on the Settings > Reporting > Log Database page. See [Log Database administration settings, page 361](#), for more information.

When configuring the start time for the maintenance job and the Internet browse time job, consider system resources and network traffic. These jobs are resource intensive, and can slow logging and reporting performance.

Log Database administration settings

Use the **Settings > Reporting > Log Database** page to manage:

- ◆ When, where, and how the Log Database creates new database partitions, and which partitions are used in creating reports ([Configuring database partition options, page 362](#))
- ◆ Whether log records include the full URL, including both the domain and the full path to the page or item (see [Configuring how URLs are logged, page 365](#))
- ◆ How Internet Browse Time is calculated (see [Configuring Internet browse time options, page 366](#))
- ◆ When and how maintenance jobs are run (see [Configuring Log Database maintenance options, page 364](#))

In addition, use the **Error Log Activity** section to review status and error messages recorded during the jobs run on the Log Database. Use the drop-down list next to the group title to select the maximum number of messages to display, or select **View None** to hide all error log entries.

The name of the active Log Database instance is displayed at the top of the page.

Configuring database partition options

Related topics:

- ◆ [Log Database administration settings, page 361](#)
- ◆ [Configuring Internet browse time options, page 366](#)
- ◆ [Configuring how URLs are logged, page 365](#)
- ◆ [Configuring Log Database maintenance options, page 364](#)

Use the **Database Rollover Options** section of the Settings > Reporting > Log Database page to specify when you want the Log Database to create a new database partition (roll over), where database partitions are stored, and how large partitions are. Also create new partitions manually, rather than waiting for the planned rollover, and review all database partitions available for reporting.

1. Next to **Roll over every**, indicate how often you want a new partition to be created.
 - For all supported database engines, you can enter a size limit for each partition. When the size limit is reached, a new partition is created.
The size limit can be set as follows:
 - *Microsoft SQL Server Express*: 100-8000 MB, default 5000.
 - *SQL Server Standard or Enterprise*: 100-1,000,000 MB, default 5000.
 - If you are using Microsoft SQL Server Standard or Enterprise, you can alternatively specify a partition rollover time interval (every 1-52 weeks, or every 1-12 months).



Note

If the rollover begins during a busy part of the day, performance may slow during the rollover process.

To avoid this possibility, some environments choose to set the automatic rollover to a long time period or large maximum size. Then, they perform regular manual rollovers to prevent the automatic rollover from occurring. See [Configuring Log Database maintenance options, page 364](#), for information on manual rollovers.

Keep in mind that extremely large individual partitions are not recommended. Reporting performance can slow if data is not divided into multiple, smaller partitions.

When a new partition database is created based reporting is automatically enabled for the partition.

2. Under Database Partition Creation, provide the following information:
 - a. Specify the **File Path** for creating both the **Data** and **Log** files for new database partitions.
 - b. Under **Init Size** set the initial file size for both the **Data** (*Express*: 1-5000 MB, default 100; *Standard or Enterprise*: 100-500,000 MB, default 2000) and **Log** (*Express*: 1-4000 MB, default 100; *Standard or Enterprise*, 1-250,000 MB, default 100) files that make up new database partitions.



Note

As a best practice, calculate the average partition size over a period of time. Then, update the initial size to that value. This approach minimizes the number of times the partition must be expanded, and frees resources to process data into the partitions.

- c. Under **Growth**, set the increment by which to increase the size of a partition's **Data** (*Express*: 1-1000 MB, default 100; *Standard or Enterprise*: 100-500,000 MB, default 500) and **Log** (*Express*: 1-1000 MB, default 100; *Standard or Enterprise*: 1-250,000 MB, default 100) files when additional space is required.
3. If you want to create a partition the next time the ETL job runs (see [Database jobs, page 361](#)), rather than waiting for the next automatic rollover, click **Manually Create Partition**. This process usually takes a few minutes.

To have the new partition use changes made on the Log Database page, click **OK** and **Save All** before click **Manually Create Partition**.

Click the Refresh link under the Available Partitions list periodically. The new partition is added to the list when the creation process is complete.

4. Use the **Available Partitions** list to review the partitions available for reporting. The list shows the dates covered, as well as the size and name of each partition. Mark the check box next to a partition name, and use the buttons below the list to determine whether the partition's data is used in or excluded from reports, or to delete the partition.
 - Click **Enable** to include a selected partition's data in reports. You must enable at least one partition for reporting.
 - Click **Disable** to exclude a selected partition's data from reports.

Together, the Enable and Disable options allow you to manage how much data is analyzed during report generation and speed report processing. For example, if you plan to generate a series of reports for June, deselect all partitions except those with dates in June.



Important

This selection affects scheduled reports as well as reports run interactively. To avoid generating reports with no data, make sure the relevant partitions are enabled when reports are scheduled to run.

- Click **Delete** to remove a partition that is no longer needed. The partition is actually deleted the next time the nightly database maintenance job runs.

**Warning**

Use this option with care. You cannot recover deleted partitions.

Deleting obsolete partitions minimizes the number of partitions in the Log Database, which improves database and reporting performance. Use this Delete option to delete individual partitions as needed. See [Configuring Log Database maintenance options, page 364](#), if you prefer to delete older partitions according to a schedule.

5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring Log Database maintenance options

Related topics:

- ◆ [Log Database administration settings, page 361](#)
- ◆ [Configuring database partition options, page 362](#)
- ◆ [Configuring Internet browse time options, page 366](#)
- ◆ [Configuring how URLs are logged, page 365](#)

Use the **Maintenance Configuration** section of the Settings > Reporting > Log Database page to control when the database maintenance job runs, whether and how often database partitions are automatically deleted, and how often tasks like reindexing partitions and deleting error log messages occur.

1. For **Maintenance start time**, select the time of day for running the database maintenance job.

The time and system resources required by this job vary depending on the tasks you select in this area. To minimize any impact on other activities and systems, it is best to run this job during a slow time on the network, different from the time set for the IBT job (see [Configuring Internet browse time options, page 366](#)).

2. Mark **Automatically delete partitions**, and then specify the number of days (from 2 to 365) after which partitions should be deleted.

**Warning**

After a partition has been deleted, the data cannot be recovered. See [Configuring database partition options, page 362](#), for an alternative way to delete partitions.

3. Check **Enable automatic reindexing**, and then select a day of the week to have this processing performed automatically each week.

Reindexing the database is important to maintain database integrity and to optimize reporting speed.



Important

It is best to perform this processing during a quiet time on the network. Reindexing database partitions is resource intensive and time-consuming. Reports should not be run during the process.

4. Check **Number of days before deleting failed batches** and then enter a number of days (from 0 to 90) after which to delete any failed batches.

If this option is not checked, failed batches are retained indefinitely for future processing.

If there is insufficient disk space or inadequate database permissions to insert log records into the database, the records are marked as a **failed batch**. Typically, these batches are successfully reprocessed and inserted into the database during the nightly database maintenance job.

However, this reprocessing cannot be successful if the disk space or permission problem has not been resolved. Additionally, if **Process the unprocessed batches** is not selected, failed batches are never reprocessed. They are deleted after the time specified here.

5. Check **Number of days before deleting error log**, and then enter a number of days (0 to 90) after which to delete database error records from the catalog database.

If this option is not checked, error logs are retained indefinitely.

6. Check **Process the unprocessed batches** to have the nightly database maintenance job reprocess any failed batches.

If this option is unchecked, failed batches are never reprocessed. They are deleted after the time specified above, if any.

7. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring how URLs are logged

Related topics:

- ◆ [Log Database administration settings, page 361](#)
- ◆ [Configuring database partition options, page 362](#)
- ◆ [Configuring Internet browse time options, page 366](#)
- ◆ [Configuring Log Database maintenance options, page 364](#)

Use the **Full URL Logging** section of the Settings > Reporting > Log Database page to determine how much of each requested URL is logged.

**Note**

Managing Log Database size is an important concern in high-volume networks. Disabling the Full URL Logging option is one way to control database size and growth.

1. Mark **Record full URL of each site requested** to log the entire URL, including the domain (www.domain.com) and the path to the particular page (/products/productA.html).

**Important**

Enable full URL logging if you plan to generate reports of scanning activity (see [Reporting on scanning activity](#), page 168). Otherwise, reports can display only the domain (www.domain.com) of the site categorized, even though individual pages within the site may fall into different categories, or contain different threats.

If this option is not checked, only domain names are logged. This choice results in a smaller database, but provides less detail.

Logging full URLs produces a larger Log Database, but provides greater detail.

If you activate full URL logging when consolidation is active, the consolidated record contains the full URL from the first record in the consolidation group. See [Log Server Configuration utility](#), page 357, for more information.

2. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring Internet browse time options

Related topics:

- ◆ [Log Database administration settings](#), page 361
- ◆ [Configuring database partition options](#), page 362
- ◆ [Configuring how URLs are logged](#), page 365
- ◆ [Configuring Log Database maintenance options](#), page 364

Internet browse time (IBT) reports give a view into the amount of time users spend on the Internet. A nightly database job calculates browse time for each client based on the new log records received that day. Set browse time options in the **Internet Browse Time Configuration** section of the Settings > Reporting > Log Database page.

1. Choose a **Job start time** for the IBT database job.

The time and system resources required by this job vary depending on the volume of data logged each day. It is best to run this job at a different time than the nightly maintenance job (see *Configuring Log Database maintenance options*, page 364), and to select a slow time on the network to minimize any impact on generating reports.

The IBT database job is resource intensive, affecting most database resources. If you enable this job, set the start time so that it does not interfere with the database system's ability to process scheduled reports and other important operations. Also, monitor the job to determine whether more robust hardware is needed to accommodate all processing needs.

2. For **Read time threshold**, set an average number of minutes for reading a specific Web site.

The read time threshold defines browse sessions for the purpose of Internet browse time reports. Opening a browser generates HTTP traffic. This represents the beginning of a browse session. The session is open as long as HTTP traffic is continually generated within the time set here. The browse session is considered closed once this amount of time passes with no HTTP traffic. A new browse session begins as soon as HTTP traffic is generated again.

**Note**

It is best to change the Read Time Threshold as seldom as possible, and to start a new database partition whenever you do make a change.

To avoid inconsistent data on the reports, generate IBT reports from database partitions that use the same Read Time Threshold value.

Be aware that some Web sites use an automatic refresh technique to update information frequently. One example is a news site that rotates a display of the latest news stories. This refresh generates new HTTP traffic. Therefore, when this kind of site is left open, new log records are generated each time the site refreshes. There is no gap in HTTP traffic, so the browser session is not closed.

3. Set a **Last read time** value to account for time spent reading the last Web site before the end of a browse session.

When the time gap of HTTP traffic is longer than the read time threshold, the session is ended and the value of the Last Read Time is added to the session time.

4. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring investigative reports

Related topics:

- ◆ [Database connection and report defaults, page 368](#)
- ◆ [Display and output options, page 370](#)

Investigative reports let you interactively delve into the information about your organization's Internet usage. See [Investigative reports, page 127](#).

The Options link on the main investigative reports page gives you the opportunity to modify which Log Database is used for reporting. It also lets you modify the default view of detail reports. See [Database connection and report defaults, page 368](#).

The **wse.ini** file lets you configure certain defaults for viewing summary and multi-level reports. It also gives you control over the default page size used when a report is output to PDF. See [Display and output options, page 370](#).

Database connection and report defaults

Related topics:

- ◆ [Configuring investigative reports, page 368](#)
- ◆ [Display and output options, page 370](#)
- ◆ [Summary reports, page 129](#)
- ◆ [Multi-level summary reports, page 133](#)

Use the **Investigative Reports > Options** page to connect to the desired Log Database, and to control defaults for investigative reports detail view.

Changes made to this page affect your reports. Other administrators, or even users logging on for self-reporting, can change these values for their own reporting activities.

1. Choose the Log Database to use for investigative reports.
 - Check **View the catalog database** to connect to the Log Database where Log Server is logging. Proceed to step 2.
 - To access a different Log Database:
 - a. Uncheck the **View the catalog database** option.

- b. Enter the following information to identify the desired Log Database.
(Investigative reports can be generated from a v6.3.x or v7 database.)

Field	Description
Server	Enter the machine name or IP address where the Log Database is located.
Database	Enter the name of the Log Database.
User ID	Enter the user ID for an account that has permission to access the database. Leave this blank if Log Server is configured to use a trusted connection to access the Log Database. If you are uncertain, enter sa . That is the default user or administrator account for Microsoft SQL Server.
Password	Enter the password for the specified account. Leave this blank for a trusted connection.

2. Select the following defaults for detail reports.

Field	Description
Select default Investigative Reports date range	Choose the date range for the initial summary report display.
Select the default detail report format	Choose Smart columns selection to display detail reports with the default columns set for the information being reported. Choose Custom columns selection to specify the exact columns for initial display on all detail reports. Use the Available Columns list to make your selections. Users can modify the columns displayed after generating the report.
Select report type	Choose whether to open detail reports initially showing: <ul style="list-style-type: none"> • Detail: each record appears on a separate row; time can be displayed. • Summary: combines into a single entry all records that share a common element. The specific element varies, according to the information reported. Typically, the right-most column before the measure shows the summarized element. Time cannot be displayed.
Available Columns / Current Report	Select a column name in the Available Columns list and click the appropriate arrow to move it to the Current Report list. Up to 7 columns can be on the Current Report list. After the Current Report list contains all the columns for initial detail reports, set the order of the columns. Select an entry in the list, and use the up and down arrow buttons to change its position.

3. Click **Save Options** to immediately save all changes.

Display and output options

Related topics:

- ◆ [Configuring investigative reports, page 368](#)
- ◆ [Database connection and report defaults, page 368](#)
- ◆ [Output to file, page 150](#)

You can make adjustments to the way certain report choices and report results are displayed in summary and multi-level investigative reports, and specify the default page size when reports are output to PDF format.

These investigative reports configuration options are set in the **wse.ini** file. The default location is:

```
C:\Program Files\WebSense\Web Security\webroot\Explorer\
wse.ini
```

The following table lists the parameters that affect display and output of investigative reports, what each controls, and its default value. (Do NOT modify any other settings in the wse.ini file.)

Parameter	Description
maxUsersMenu	The database must have fewer users than this value (by default, 5000) to show User as a report choice in the Internet Use by list.
maxGroupsMenu	The database must have fewer groups than this value (by default, 3000) to show Group as a report choice in the Internet Use by list. Note: There must be 2 or more groups for Group to appear in the Internet Use by list. There also must be 2 or more domains for Domain to appear in the Internet Use by list. There is no maximum value for domains.
maxUsersDrilldown	This works with the warnTooManyHits parameter to control when the User option displays in red. The red lettering indicates that selecting User will produce a very large report, which could be slow to generate. If there are more users than this value (by default, 5000), and more hits than the warnTooManyHits value, the User option displays red in various drop-down lists and values lists. If there are more users than this value, but fewer hits than the warnTooManyHits value, the User option displays in normal color, as the resulting report will be a more reasonable size.

Parameter	Description
maxGroupsDrilldown	The Group option displays in red during drill down if the proposed report includes more groups than this number (by default, 2000). The red lettering indicates that selecting Group will produce a very large report, which could be slow to generate.
warnTooManyHits	<p>This works with the maxUsersDrilldown parameter to control when the User option displays in red.</p> <p>If there are more users than the maxUsersDrilldown value, but fewer hits than this value (by default, 10000), the User option does <i>not</i> display in red.</p> <p>If there are more users than the maxUsersDrilldown value, and more hits than this value, the User option does display in red. The red lettering indicates that selecting User will produce a very large report, which could be slow to generate.</p>
hitsPerPage	This determines the maximum number of items (by default, 100) displayed per page. (This does not affect printed reports.)
maxOutputBufferSize	<p>This is the maximum amount of data (in bytes) that can be displayed on the main investigative reports page. If the requested data exceeds this limit (by default, 4000000, or, 4 million bytes), a message stating that some results are not shown appears in red at the end of the report.</p> <p>Larger values enable you to display larger amounts of data in one report, if this is an issue. However, if you encounter memory errors, consider decreasing this value.</p>
sendMulti	This option is disabled (0) by default. Set it to 1 (enabled) to divide very large, scheduled detail reports into multiple files of 10,000 rows each. The files that represent one report are zipped and sent to the email recipients. The report files can be extracted with most common file compression utilities.
maxSlices	This is the maximum number of distinct slices (by default, 6) in a pie chart, including an Other slice, which combines all values that do not have individual slices.
timelineCompressionThreshold	This option is used only for User Activity by Day or Month, when the Group Similar Hits/View All Hits option is available. The report collapses all hits with the same category that occur within the number of seconds set here (by default, 10).
PageSize	<p>Investigative report results can be output to Portable Document Format (PDF) for easy distribution or printing. The page size (by default, Letter) can be:</p> <ul style="list-style-type: none"> • A4 (8.27 X 11.69 inches) • Letter (8.5 X 11 inches)

Self-reporting

Related topics:

- ◆ [Configuring reporting preferences, page 354](#)
- ◆ [Accessing self-reporting, page 151](#)
- ◆ [Investigative reports, page 127](#)

Self-reporting is a feature you can enable to allow users to view investigative reports on their personal Internet activity. This allows them to see what kind of information is being gathered and monitored about them, which accommodates government regulations in many countries. In addition, viewing their own activity may encourage some users to alter their browsing habits so they meet the organization's Internet policy.

To enable self-reporting:

1. Go to **Settings > General > Directory Services**, and configure the directory service used to authenticate users who access TRITON - Web Security with their network credentials. This may have been done previously to enable filtering by user and group names. See [Directory services, page 65](#).

If your installation includes multiple Policy Servers, you must log on to each one and configure the Directory Services page with information for the appropriate directory service.

2. Go to the **Settings > Reporting > Preferences**, and mark the **Allow self-reporting** check box. See [Configuring reporting preferences, page 354](#).

After enabling the option, be sure to give users the information they need to run the reports:

- ◆ The URL for accessing the self-reporting interface. Remind users that they can save the URL as a favorite or bookmark for future use.

Read on for detailed information about the URL.

- ◆ Which Policy Server to select during logon.

In networks with only one Policy Server, this is not needed. If your network includes multiple Policy Servers, give users the IP address of the Policy Server configured to communicate with the directory service that authenticates their network logon.

- ◆ What user name and password to use during logon.

Self-reporting users must enter their network user name and password during logon.

The **URL** for accessing the self-reporting interface is:

```
https://<IP address>:9443/mng/login/pages/  
selfReportingLogin.jsf
```

Replace *<IP address>* with the IP address of the TRITON - Web Security machine.

Administrators and users can also access the self-reporting logon page by opening the TRITON - Web Security logon page and clicking the Self-Reporting link.

If your network includes **multiple Policy Servers**, you must inform users which one to choose during self-reporting logon.

16

Network Configuration

Related topics:

- ◆ [Hardware configuration, page 376](#)
- ◆ [Network Agent configuration, page 379](#)
- ◆ [Verifying Network Agent configuration, page 385](#)

When you run Websense Web Security or Websense Web Filter in standalone mode (not integrated with a proxy or firewall product), Websense Network Agent enables:

- ◆ Internet content filtering
- ◆ Network protocol and Internet application management
- ◆ Bandwidth management
- ◆ Logging of bytes transferred

In an integrated Websense software deployment, a gateway, firewall, or caching product may handle the task of routing user requests to Websense software for filtering, and routing block pages back to the client. In this environment, Network Agent may still be used to filter non-HTTP requests, provide enhanced logging detail, or both.

In addition, Websense Web Security Gateway can detect protocols that tunnel over HTTP (see [Tunneled protocol detection, page 160](#)) and provide some bandwidth management capabilities ([Using Bandwidth Optimizer to manage bandwidth, page 236](#)), independent of Network Agent.

Network Agent works by continually monitoring overall network usage, including bytes transferred over the network. The agent sends usage summaries to Websense software at predefined intervals. Each summary includes start time and end time, overall bytes used, and bytes used per protocol.

By default, Network Agent also provides bandwidth usage data to Policy Server, and filtering log data to Filtering Service.

Network Agent is typically configured to see all traffic in your network. The agent distinguishes between:

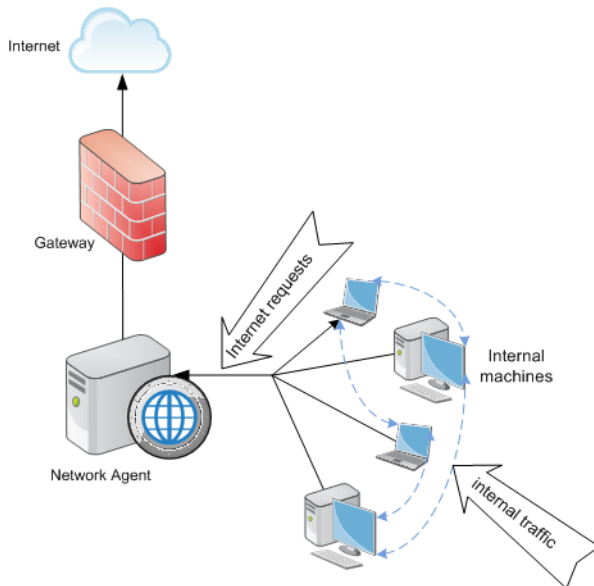
- ◆ Requests sent from internal machines to internal machines (hits to an intranet server, for example)

- ◆ Requests sent from internal machines to external machines such as Web servers (user Internet requests, for example)


The latter is the primary concern in monitoring employee Internet usage.

Hardware configuration

Each Network Agent instance monitors traffic from the machines you identify as belonging to your network. By default, it monitors traffic to external sites, while ignoring internal network traffic.



Network Agent can be configured to monitor traffic to internal machines that you specify (for example, internal Web servers).



You can customize which internal machines (network segments) are monitored by each Network Agent instance, or even by each network interface card (NIC) on a Network Agent machine.

Each Network Agent instance must:

- ◆ Be positioned appropriately in the network to detect traffic to and from all monitored machines.
- ◆ Have at least 1 NIC dedicated to monitoring traffic.

Network Agent can be installed on a machine with multiple NICs, and can use multiple NICs for both monitoring requests and sending block pages. If you add a new NIC to the Network Agent machine, restart the Network Agent service, and then configure the new NIC (see [Configuring NIC settings](#), page 382).



Note

To determine whether Network Agent can see traffic in a network segment, use any third-party packet analyzer.

More information about Network Agent placement and NIC requirements can be found in the *Deployment and Installation Center*.

For information about configuring Network Agent to monitor internal network requests, use specific NICs, and perform enhanced logging, see [Network Agent configuration](#), page 379.

Filtering IPv6 traffic

Use the **Settings > General > IPv6** page to configure how each Network Agent instance in your network handles any Internet Protocol Version 6 (IPv6) requests that it sees.

Each Network Agent instance visible to the current Policy Server is listed in the IPv6 Configuration table.

- ◆ To configure all instances to treat IPv6 traffic in the same way, click the **Permit All** or **Block All** button below the table.

If your network does not use IPv6, you may reduce Network Agent overhead and improve performance by selecting to block all IPv6 traffic.

- ◆ To configure the behavior of a single Network Agent instance, select the instance. The page section under the table is updated to show the current options for the selected instance.

Mark or clear the **Permit IPv6 traffic...** check box (under the bar showing the Network Agent instance IP address) to specify how the selected Network Agent responds to the IPv6 traffic it sees.

When any Network Agent instance blocks IPv6 traffic, you can specify exceptions. Exceptions can be IPv6 addresses or network ports. Traffic coming from exception IPv6 addresses, as well as IPv6 traffic on exception ports, is ignored (permitted), while all other IPv6 traffic is blocked.

To specify any exceptions for a selected Network Agent instance, use the list below the “Permit IPv6 traffic...” check box:

- ◆ To add an IPv6 address to the list, click **Add**, then enter the IPv6 address or range (see [Adding or editing IP addresses](#), page 385).
- ◆ To specify ports on which IPv6 traffic should be ignored, enter one or more comma-separated port numbers in the **Ports** field.
- ◆ If you want all Network Agent instances to share the IPv6 address and port exceptions that you have defined, click **Apply to All**.

When you are finished configuring IPv6 filtering, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Note that any changes made on this page are also reflected on the **Settings > Network Agent > Local Settings** page for each configured Network Agent instance.

Likewise, any changes made on the **Local Settings** page for an instance will also appear on the **IPv6** page.

Network Agent configuration

Related topics:

- ◆ [Hardware configuration, page 376](#)
- ◆ [Configuring global settings, page 379](#)
- ◆ [Configuring local settings, page 381](#)
- ◆ [Configuring NIC settings, page 382](#)
- ◆ [Adding or editing IP addresses, page 385](#)

After installing Network Agent, use TRITON - Web Security to configure its network monitoring behavior. Network Agent settings are divided into two main areas:

- ◆ **Global settings** affect all Network Agent instances. Use these settings to:
 - Identify the machines in your network.
 - List machines in your network that Network Agent should monitor for **incoming** requests (for example, internal Web servers).
 - Specify bandwidth calculation and protocol logging behavior.
- ◆ **Local settings** apply only to the selected Network Agent instance. Use these settings to:
 - Identify which Filtering Service instance is associated with each Network Agent.
 - Note proxies and caches used by the machines that this Network Agent monitors.
 - Configure how each network card (NIC) in the Network Agent machine is used (to monitor requests, send block pages, or both).

Network card settings also determine which segment of the network each Network Agent instance monitors.

Configuring global settings

Related topics:

- ◆ [Hardware configuration, page 376](#)
- ◆ [Configuring local settings, page 381](#)
- ◆ [Configuring NIC settings, page 382](#)
- ◆ [Adding or editing IP addresses, page 385](#)

Use the **Settings > Network Agent > Global** page to define basic monitoring and logging behavior for all instances of Network Agent.

The **Internal Network Definition** list identifies the IP addresses that are part of your network. By default, Network Agent does not monitor the traffic sent **between** these IP addresses (internal network communications).

Network Agent does not use this list to determine which IP addresses to monitor for Internet requests. That behavior is configured separately for each Network Agent NIC (see [Configuring NIC settings, page 382](#)). This list is used only to exclude internal traffic (like LAN and intranet connections) from monitoring.

An initial set of entries is provided by default. You can add more entries, or edit or delete existing entries.

The **Internal Traffic to Monitor** list includes any machines included within the Internal Network Definition for which you **do** want Network Agent to monitor traffic. This might include internal Web servers, for example, to help you to track internal connections.

Any request sent from anywhere in the network to the specified internal machines is monitored. By default, this list is blank.

- ◆ Click **Add** to add an IP address or range to the appropriate list. See [Adding or editing IP addresses, page 385](#), for more information.
- ◆ To edit an entry in the list, click the IP address or range. See [Adding or editing IP addresses, page 385](#), for more information.
- ◆ To remove an entry from the list, mark the check box next to an IP address or range, and then click **Delete**.

The **Additional Settings** options allow you to determine how often Network Agent calculates bandwidth usage, and whether and how often protocol traffic is logged:

Field	What to do
Bandwidth calculation interval	Enter a number between 1 and 300 to specify how frequently, in seconds, Network Agent should calculate bandwidth usage. An entry of 300, for example, indicates that Network Agent will calculate bandwidth every 5 minutes. The default is 10 seconds.
Log protocol traffic periodically	Mark this option to enable the Logging interval field.
Logging interval	Enter a number between 1 and 300 to specify how frequently, in minutes, Network Agent logs protocols. An entry of 60, for example, indicates that Network Agent will write to the log file every hour. The default is 1 minute.

When you are finished making changes, click **OK** to cache the changes. Changes are not implemented until you click **Save All**.

Configuring local settings

Related topics:

- ◆ [Hardware configuration, page 376](#)
- ◆ [Configuring global settings, page 379](#)
- ◆ [Configuring NIC settings, page 382](#)

Use the **Settings > Network Agent > Local Settings** page to configure filtering behavior, proxy information, and other settings for the selected instance of Network Agent. The IP address of the selected Network Agent instance appears in the title bar of the content pane.

Use the **Filtering Service Definition** settings to specify which Filtering Service is associated with the selected Network Agent instance, and how to respond to Internet requests if Filtering Service is not available.

Field	What to do
Filtering Service IP address	Select the Filtering Service that is associated with this Network Agent.
If Filtering Service is unavailable	Select Permit to permit all requests or select Block to block all requests until Filtering Service is available again. The default is Permit.

If your network is configured for IPv6 traffic, use the **IPv6 Configuration** section to specify whether to **Permit all IPv6 traffic**. If your network does not use IPv6, leave the checkbox blank to block IPv6 traffic.

If you choose to block IPv6 traffic, you can specify exceptions. Traffic coming from the IPv6 addresses and IPv6 traffic on the ports you specify will be ignored (permitted). All other IPv6 traffic is blocked.

- ◆ To add an IPv6 address to the list, click **Add**, then enter the IPv6 address or range (see [Adding or editing IP addresses, page 385](#)).
- ◆ To specify ports on which IPv6 traffic should be ignored, enter one or more comma-separated port numbers in the **Ports** field.

To ensure that user requests are monitored, filtered, and logged correctly, use the **Proxies and Caches** list to specify the IP address of any proxy or cache server that communicates with Network Agent.

- ◆ Click **Add** to add an IP address or range to the list (see [Adding or editing IP addresses, page 385](#)).
- ◆ To edit an entry in the list, click the IP address or range.
- ◆ To remove an entry from the list, mark the check box next to an IP address or range, and then click **Delete**.

Use the **Network Interface Cards** list to configure individual NICs. Click on a NIC in the **Name** column, and then see [Configuring NIC settings](#), page 382, for further instructions.

The **Advanced Network Agent Settings** options are used when:

- ◆ HTTP requests in your network are passed through a non-standard port.
By default the **Ports used for HTTP traffic** are **8080, 80** (when Websense software is integrated with a firewall, proxy, or cache) or **All** (in a standalone deployment).
- ◆ You want Network Agent to ignore traffic on specific ports.
Mark **Configure this Network Agent instance to ignore traffic on the following ports**, and then enter one or more ports.
If you have deployed Websense Content Gateway, this may be used to prevent double logging of HTTPS traffic.
- ◆ Websense Technical Support instructs you to change debugging options for troubleshooting purposes.
Debug Settings options should not be changed without direction from Technical Support.

Field	Description
Mode	<ul style="list-style-type: none"> ◆ None (default) ◆ General ◆ Error ◆ Detail ◆ Bandwidth
Output	<ul style="list-style-type: none"> ◆ File (default) ◆ Window
Port	55870 (default)

When you are finished making changes to your Network Agent settings, click **OK** to cache the changes. Changes are not implemented until you click **Save All**.

Configuring NIC settings

Related topics:

- ◆ [Hardware configuration](#), page 376
- ◆ [Network Agent configuration](#), page 379
- ◆ [Configuring monitoring settings for a NIC](#), page 384
- ◆ [Adding or editing IP addresses](#), page 385

Use the **Network Agent > Local Settings > NIC Configuration** page to specify how Network Agent uses each available network interface card (NIC) to monitor and manage network usage.

The **NIC Information** area provides the context for the changes that you make, showing the **IP address**, brief **NIC Description**, and card **Name**. Use this information to ensure that you are configuring the correct NIC.

Monitoring

In a multiple-NIC configuration, you can identify one NIC to monitor network traffic, and another NIC to serve block pages. At least one NIC must be used for monitoring, and more than one NIC can monitor traffic.

Use the **Monitoring** section to indicate whether or not to **Use this NIC to monitor traffic**.

- ◆ If this NIC is not used for monitoring, deselect the check box and then continue with the next section.
- ◆ If the NIC is used for monitoring, select the check box, and then click **Configure**. You are taken to the Configure Monitoring Behavior page. See [Configuring monitoring settings for a NIC, page 384](#), for instructions.

Other NIC options

In addition to configuring monitoring options, you can also determine other NIC behaviors:

1. Under Blocking, make sure that the appropriate NIC is listed in the **Blocking NIC** field. If you are configuring multiple NICs, the settings for each NIC should show the same value in this field. In other words, only one NIC is used for blocking.
2. If you are running Websense software in **Stand-Alone** mode, **Filter and log HTTP requests** is selected, and cannot be changed.
3. If you have integrated Websense software with a third-party device or application, use the **Integrations** options to indicate how this Network Agent should filter and log HTTP requests. Options that do not apply to your environment are disabled.
 - Select **Log HTTP requests** to improve accuracy in Websense reports.
 - Select **Filter all requests not sent over HTTP ports** to use Network Agent to filter only those HTTP requests not sent through the integration product.
4. Under Protocol Management, indicate whether Network Agent should use this NIC to filter non-HTTP protocols:
 - Check **Filter non-HTTP protocol requests** to activate the protocol management feature. This allows Websense software to filter Internet applications and data transfer methods, such as those used for instant messaging, streaming media, file sharing, Internet mail, and so on. See [Filtering categories and protocols, page 40](#), and [Working with protocols, page 229](#), for more information.

- Check **Measure bandwidth usage by protocol** to activate the Bandwidth Optimizer feature. Network Agent uses this NIC to track network bandwidth usage by each protocol or application. See [Using Bandwidth Optimizer to manage bandwidth, page 236](#), for more information.

Configuring monitoring settings for a NIC

Use the **Local Settings > NIC Configuration > Monitor List** page to specify which machines Network Agent monitors via the selected network interface card (NIC).

1. Under Monitor List, specify which requests Network Agent monitors:
 - **All:** Network Agent monitors requests from all machines it sees using the selected NIC. Typically, this includes all machines in the same network segment as the current Network Agent machine or NIC.
 - **None:** Network Agent does not monitor any requests.
 - **Specific:** Network Agent monitors only the network segments included in the Monitor List.
2. If you selected Specific, click **Add**, and then specify the IP addresses of the machines Network Agent should monitor. See [Adding or editing IP addresses, page 385](#), for more information.



Note

You cannot enter overlapping IP address ranges. If ranges overlap, network bandwidth measurements may not be accurate, and bandwidth-based filtering may not be applied correctly.

To remove an IP address or network range from the list, check the appropriate list item, and then click **Delete**.

3. Under Monitor List Exceptions, identify any internal machines Network Agent should exclude from monitoring.

For example, Network Agent could ignore requests made by CPM Server. This way, CPM Server requests will not clutter Websense log data or any of the status monitors output.

 - a. To identify a machine, click **Add**, and then enter its IP address.
 - b. Repeat the process to identify additional machines.
4. Click **OK** to cache your changes and return to the NIC Configuration page. Changes are not implemented until you click **Save All**.

Adding or editing IP addresses

Related topics:

- ◆ [Configuring global settings, page 379](#)
- ◆ [Configuring local settings, page 381](#)
- ◆ [Configuring NIC settings, page 382](#)

Use the **Add IP Addresses** or **Edit IP Addresses** page to make changes to any of the following Network Agent lists: Internal Network Definition, Internal Traffic to Monitor, IPv6 Configuration, Proxies and Caches, Monitor List, or Monitor List Exceptions.

- ◆ When you add or edit an IP address range, make sure that it does not overlap any existing entry (single IP address or range) in the list.
- ◆ When you add or edit a single IP address, make sure that it does not fall into a range that already appears in the list.

To add a new IP address or range:

1. Select the **IP address** or **IP address range** radio button.
2. Enter a valid IP address or range.
3. Click **OK** to return to the previous Network Agent Settings page. The new IP address or range appears in the appropriate table.

To return to the previous page without caching your changes, click **Cancel**.

4. Repeat this process for additional IP addresses, as needed.

When you edit an existing IP address or range, the Edit IP Addresses page displays the selected item with the correct radio button already selected. Make any necessary changes, and then click **OK** to return to the previous page.

When you are finished adding or editing IP addresses, click **OK** on the Network Agent Settings page. Changes are not implemented until you click **Save All**.

Verifying Network Agent configuration

After configuring Network Agent in TRITON - Web Security, you can use any third-party packet analyzer to ensure that computers on your network are visible to Websense software.

Make sure that the network card that you have configured as the Network Agent monitoring NIC can see traffic from IP addresses in all network segments that the Network Agent instance has been configured to monitor. (This configuration is done on the Local Settings > NIC Configuration > Monitor List page. See [Configuring monitoring settings for a NIC, page 384.](#))

If packets from some IP addresses are not visible to the monitoring NIC:

- ◆ Review the network configuration and NIC placement requirements (see [Hardware configuration, page 376](#)).
- ◆ Review the more detailed network configuration information in the *Deployment and Installation Center*.
- ◆ Verify that you have properly configured the monitoring NIC ([Configuring NIC settings, page 382](#)).

17

Troubleshooting

Use this section to find solutions to common issues before contacting Technical Support.

The Websense Support Web site features an extensive Knowledge Base and customer forums, available at websense.com/support/. Search for topics by keyword or phrase, or browse content by product and version.

Troubleshooting instructions are grouped into the following sections:

- ◆ *Installation and subscription issues*
- ◆ *Master Database issues*, page 389
- ◆ *Filtering issues*, page 395
- ◆ *Network Agent issues*, page 400
- ◆ *User configuration and identification issues*, page 402
- ◆ *Block message issues*, page 417
- ◆ *Log, status message, and alert issues*, page 420
- ◆ *Policy Server and Policy Database issues*, page 423
- ◆ *Delegated administration issues*, page 425
- ◆ *Log Server and Log Database issues*, page 426
- ◆ *Investigative report and presentation report issues*, page 434
- ◆ *Other reporting issues*, page 440
- ◆ *Interoperability issues*, page 442
- ◆ *Troubleshooting tips and tools*, page 457

Installation and subscription issues

- ◆ *There is a subscription problem*, page 388
- ◆ *Unable to verify the subscription key*, page 388
- ◆ *After upgrade, users are missing from TRITON - Web Security*, page 388

There is a subscription problem

A valid subscription key is needed to download the Websense Master Database and perform Internet filtering. When your subscription expires or is invalid, and when the Master Database has not been downloaded for more than 2 weeks, the Health Alerts summary on the Status > Today page displays a warning.

- ◆ Verify that you have entered your subscription key exactly as you received it. The key is case sensitive.
- ◆ Make sure that your subscription has not expired. See [Subscription key, page 390](#).
- ◆ Ensure that the Master Database has been downloaded successfully within the last 2 weeks. You can check download status in TRITON - Web Security: click **Database Download** on the Status > Today page.

See [The Master Database does not download, page 390](#), for help troubleshooting database download problems.

If you have entered the key correctly, but continue to receive a status error, or if your subscription has expired, contact Websense, Inc., or your authorized reseller.

When your subscription expires, TRITON - Web Security settings determine whether all users are given unfiltered Internet access or all Internet requests are blocked. See [Your subscription, page 31](#), for more information.

Unable to verify the subscription key

After you enter your subscription key, Filtering Service attempts to connect to the Websense database download server to both verify the key and download the Master Database.

If Filtering Service is unable to connect to the database download server, both subscription errors and database download errors will appear in TRITON - Web Security.

- ◆ If the database download server is down, the problem should resolve itself within a short period of time.
- ◆ If Filtering Service is unable to connect to the download server, see [Internet access, page 391](#), and [Verify firewall or proxy server settings, page 392](#), to make sure that Filtering Service and the network environment are properly configured to enable the connection.

After upgrade, users are missing from TRITON - Web Security

If you defined Active Directory as your directory service after upgrading your Websense software, user names may not appear in TRITON - Web Security. This occurs when user names include characters that are not part of the UTF-8 character set.

To support LDAP 3.0, the Websense installer changes the character set from MBCS to UTF-8 during upgrade. As a result, user names that include non-UTF-8 characters are not properly recognized.

To fix this problem, manually change the character set to MBCS:

1. Go to **Settings > General > Directory Services**.
2. Make sure that **Active Directory (Native Mode)** is selected under Directories, near the top of the page.
3. Click **Advanced Directory Settings**.
4. Under Character Set, click **MBCS**. You may have to scroll down to see this option.
5. Click **OK** to cache the change. Changes are not implemented until you click **Save All**.

Master Database issues

- ◆ [The initial filtering database is being used, page 389](#)
- ◆ [The Master Database is more than 1 week old, page 389](#)
- ◆ [The Master Database does not download, page 390](#)
- ◆ [Master Database download does not occur at the correct time, page 394](#)
- ◆ [Contacting Technical Support for database download issues, page 395](#)

The initial filtering database is being used

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content.

A partial version of the Master Database is installed with your Websense software on each Filtering Service machine. This partial database is used to enable basic filtering functionality from the time you enter your subscription key.

You must download the full database for complete filtering to occur. See [The Websense Master Database, page 34](#), for more information.

The process of downloading the full database may take a few minutes or more than 60 minutes, depending on factors such as Internet connection speed, bandwidth, available memory, and free disk space.

The Master Database is more than 1 week old

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content. Websense software downloads changes to the Master Database according to the schedule defined in TRITON - Web Security. By default, download is scheduled to occur once a day.

To manually initiate a database download:

1. Go to the **Status > Today** page, and then click **Database Download**.
2. Click **Update** next to the appropriate Filtering Service instance to start the database download, or click **Update All** to start the download on all Filtering Service machines.



Note

After downloading updates to the Master Database, CPU usage can be 90% or more for a short time while the database is loaded into local memory. It is a good idea to perform the download at off-peak times.

3. To continue working while the database is downloaded, click **Close**.

Click the **Database Download** button at any time to view download status.

If a new version of the Master Database adds or removes categories or protocols, administrators performing category- or protocol-related policy management tasks (like editing a category set) at the time of the download may receive errors. Although such updates are somewhat rare, as a best practice, try to avoid making changes to category- and protocol-related while a database is being updated.

The Master Database does not download

If you are unable to download the Websense Master Database successfully:

- ◆ Make sure that you have entered your subscription key correctly in TRITON - Web Security, and that the key has not expired (*Subscription key, page 390*).
- ◆ Verify that the Filtering Service machine is able to access the Internet (*Internet access, page 391*).
- ◆ Check firewall or proxy server settings to make sure that Filtering Service can connect to the Websense download server (*Verify firewall or proxy server settings, page 392*).
- ◆ Make sure that there is enough disk space (*Insufficient disk space on the Filtering Service machine, page 392*) and memory (*Insufficient memory on the Filtering Service machine, page 393*) on the download machine.
- ◆ Look for any application or appliance in the network, such as anti-virus software, that might prevent the download connection (*Restriction applications, page 394*).

Subscription key

To verify that the subscription key is entered correctly and has not expired:

1. Go to **Settings > General > Account**.
2. Compare the key that you received from Websense, Inc., or your reseller to the **Subscription key** field. The key must use the same capitalization as in your key document.

3. Check the date next to **Key expires**. If the date has passed, contact your reseller or Websense, Inc., to renew your subscription.
4. If you have made changes to the key in the Settings dialog box, click **OK** to activate the key and enable database download.

To manually initiate a database download, or to check the status of the most recent database download, click **Database Download** in the toolbar at the top of the Status > Today page.

Internet access

To download the Master Database, the Filtering Service machine sends an **HTTP post** command to the download servers at the following URLs:

download.websense.com
 ddsdom.websense.com
 ddsint.websense.com
 portal.websense.com
 my.websense.com

To verify that Filtering Service has the Internet access necessary to communicate with the download server:

1. Open a browser on the machine running Filtering Service.
2. Enter the following URL:

<http://download.websense.com/>

If the machine is able to open an HTTP connection to the site, a redirect page is displayed, and then the browser displays the Websense home page.

If this does not happen, ensure that the machine:

- Can communicate over port 80, or the port designated in your network for HTTP traffic
- Is configured to properly perform DNS lookups
- Is configured to use any necessary proxy servers (see *Verify firewall or proxy server settings*, page 392)

Also make sure that your gateway does not include any rules that block HTTP traffic from the Filtering Service machine.

3. Use one of the following methods to confirm that the machine can communicate with the download site:
 - From the command prompt, enter the following command:


```
ping download.websense.com
```

 Verify that the ping receives a reply from the download server.
 - Use telnet to connect to **download.websense.com 80**. If you see a cursor and no error message, you can connect to the download server.

Verify firewall or proxy server settings

If the Master Database is downloaded through a firewall or proxy server that requires authentication, ensure that a browser on the Filtering Service machine can load Web pages properly. If pages open normally, but the Master Database does not download, check the proxy server settings in the Web browser.

Microsoft Internet Explorer:

1. Select **Tools > Internet Options**.
2. Open the **Connections** tab.
3. Click **LAN Settings**. Proxy server configuration information appears under **Proxy server**.

Make a note of the proxy settings.

Mozilla Firefox:

1. Select **Tools > Options > Advanced**.
2. Select the **Network** tab.
3. Click **Settings**. The Connection Settings dialog box shows whether the browser is configured to connect to a proxy server.

Make a note of the proxy settings.

Next, make sure that Websense software is configured to use the same proxy server to perform the download.

1. Go to **Settings > General > Database Download**.
2. Verify that **Use proxy server or firewall** is selected, and that the correct server and port are listed.
3. Make sure that the **Authentication** settings are correct. Verify the user name and password, checking spelling and capitalization.

If Websense software must provide authentication information, the firewall or proxy server must be configured to accept clear text or basic authentication.

Information about enabling basic authentication is available from the Websense [Knowledge Base](#).

If a firewall restricts Internet access at the time Websense software normally downloads the database, or restricts the size of a file that can be transferred via HTTP, Websense software cannot download the database. To determine if the firewall is causing the download failure, search for a rule on the firewall that might be blocking the download, and change the download times in TRITON - Web Security (*Configuring database downloads*, page 35), if necessary.

Insufficient disk space on the Filtering Service machine

Filtering Service needs adequate space to download compressed Master Database updates to the Websense **bin** directory (see *Where is the Websense "bin" directory?*, page 457). It also needs space to decompress and load the database. As a general rule, Websense, Inc., recommends at least 4 GB of free disk space on the download drive.

A disk space warning indicates that free disk space on the Filtering Service machine has dipped below 4 GB.

On Windows systems, use Windows Explorer to check disk space:

1. Open **My Computer** in Windows Explorer (not Internet Explorer).
2. Select the drive on which Websense software is installed. By default, Websense software is located on the C drive.
3. Right-click and select **Properties** from the pop-up menu.
4. On the General tab, verify that at least 4 GB of free space is available. If there is insufficient free space on the drive, delete any unnecessary files to free up the required space.

On Linux systems, use the **df** command to verify the amount of available space in the file system in which Websense software is installed:

1. Open a terminal session.
2. At the prompt, enter:

```
df -h /opt
```

Websense software is usually installed in the `/opt/Websense/bin` directory. If it is installed elsewhere, use that path.

3. Make sure that at least 4 GB of free space is available. If there is insufficient free space on the drive, delete any unnecessary files to free up the required space.

If, after addressing any disk space issues, you are unable to download the Master Database:

1. Stop all Websense services on the Filtering Service machine (see [Stopping and starting Websense services](#), page 329).
2. Delete the **Websense.xfr** and **Websense** (no extension) files from the Websense **bin** directory.
3. Restart the Websense services.
4. Manually initiate a database download (go to the **Status > Today** page in TRITON - Web Security, and then click **Database Download**).

Insufficient memory on the Filtering Service machine

The memory required to run Websense software, download the Master Database, and apply Master Database updates varies, depending on the size of the network.

- ◆ In a small network, at least 2 GB of memory is recommended (Windows and Linux).
- ◆ Refer to the *Deployment and Installation Center* for complete system recommendations.

When free memory drops below 512 MB on the Filtering Service machine, a Health Alert message is generated. Buffer and cache space are not included in this calculation.

If the machine meets or exceeds the requirements in the *Deployment and Installation Center*, and Filtering Service is able to load the Master Database, the low memory condition is unlikely to cause problems.

If Filtering Service is unable to load the Master Database, however, you will need to free up memory on the machine, or add additional RAM.

To check the memory in a Windows system:

1. Open the Task Manager.
2. Select the **Performance** tab.
3. Check the total **Physical Memory** available.

You also can select **Control Panel > Administrative Tools > Performance** to capture information.

To check the memory in a Linux system:

1. Open a terminal session.
2. At the prompt, enter:

```
top
```
3. Compute the total memory available by adding **Mem: av** and **Swap: av**.

To address problems with insufficient memory, you can either upgrade the machine's RAM or move applications with high memory usage to another machine.

Restriction applications

Some restriction applications or appliances, such as virus scanners, size-limiting applications, or intrusion detection systems can interfere with database downloads. Ideally, configure Websense software to go straight to the last gateway so that it does not connect to these applications or appliances. Alternatively:

1. Disable the restrictions relating to the Filtering Service machine and to the Master Database download location.
 See the appliance or software documentation for instructions on changing the device's configuration.
2. Attempt to download the Master Database.

If this change has no effect, reconfigure the application or appliance to include the machine running Filtering Service.

Master Database download does not occur at the correct time

The system date and time may not be set correctly on the Filtering Service machine. Websense software uses the system clock to determine the proper time for downloading the Master Database.

If the download is not occurring at all, see [The Master Database does not download, page 390](#).

Contacting Technical Support for database download issues

If you are still experiencing Master Database download problems after completing the troubleshooting steps in this Help section, send the following information to Websense Technical Support:

1. The exact error message that appears in the Database Download dialog box
2. External IP addresses of the machines attempting to download the database
3. Your Websense subscription key
4. Date and time of the last attempt
5. Number of bytes transferred, if any
6. Open a command prompt and perform an **nslookup** on **download.websense.com**. If connection to the download server is made, send the IP addresses returned to Technical Support.
7. Open a command prompt and perform a **tracert** to **download.websense.com**. If connection to the download server is made, send the route trace to Technical Support.
8. A packet trace or packet capture performed on the Websense download server during an attempted download.
9. A packet trace or packet capture performed on the network gateway during the same attempted download.
10. The following files from the Websense **bin** directory (*Where is the Websense "bin" directory?*, page 457): **websense.ini**, **eimserver.ini** and **config.xml**.

Go to websense.com/support/ for Technical Support contact information.

Filtering issues

- ◆ *Filtering service is not running, page 395*
- ◆ *User Service is not available, page 396*
- ◆ *Sites are incorrectly categorized as Information Technology, page 397*
- ◆ *Keywords are not being blocked, page 398*
- ◆ *Custom or limited access filter URLs are not filtered as expected, page 398*
- ◆ *A user cannot access a protocol or application as expected, page 398*
- ◆ *An FTP request is not blocked as expected, page 399*
- ◆ *Websense software is not applying user or group policies, page 399*
- ◆ *Remote users are not filtered by correct policy, page 399*

Filtering service is not running

When Filtering Service is not running, Internet requests cannot be filtered and logged.

Filtering Service may stop running if:

- ◆ There is insufficient disk space on the Filtering Service machine (see *Insufficient disk space on the Filtering Service machine*, page 392).
- ◆ A Master Database download failed due to lack of disk space (see *The Master Database does not download*, page 390).
- ◆ The **websense.ini** file is missing or corrupted.
- ◆ You stop the service (after creating custom block pages, for example) and do not restart it.

Filtering Service may also appear to have stopped if you restarted multiple Websense services, and they were not started in the correct order. When you restart multiple services, remember to start the Policy Database, Policy Broker, and Policy Server before starting other Websense services.

To troubleshoot these problems:

- ◆ Verify that there is at least 3 GB of free disk space on the Filtering Service machine. You may need to remove unused files or add additional capacity.
- ◆ Navigate to the Websense **bin** directory (see *Where is the Websense “bin” directory?*, page 457), and confirm that you can open **websense.ini** in a text editor. If this file has been corrupted, replace it with a backup file.
- ◆ Check the Windows Event Viewer or **websense.log** file for error messages from Filtering Service (see *Troubleshooting tips and tools*, page 457).
- ◆ Log off of TRITON - Web Security, restart Policy Server, and then restart Filtering Service (see *Stopping and starting Websense services*, page 329).
Wait 1 minute before logging on to TRITON - Web Security again.

User Service is not available

When User Service is not running, or when Policy Server cannot communicate with User Service, Websense software cannot correctly apply user-based filtering policies.

User Service may appear to have stopped if you restarted Policy Server after restarting other Websense Services. To correct this issue:

1. Restart the Websense Policy Server service (see *Stopping and starting Websense services*, page 329).
2. Start or restart Websense User Service.
3. Close TRITON - Web Security.

Wait 1 minute before logging on to TRITON - Web Security again.

If the previous steps do not fix the problem:

- ◆ Check the Windows Event Viewer or **websense.log** file for error messages from User Service (see *Troubleshooting tips and tools*, page 457).
- ◆ Navigate to the Websense **bin** directory (*Where is the Websense “bin” directory?*, page 457), and make sure that you can open **websense.ini** in a text editor. If this file has been corrupted, replace it with a backup file.

High CPU usage on the Filtering Service machine

When the CPU on the Filtering Service machine is overloaded (whether by the volume of processing being performed by Filtering Service, or by demands from other software running on the Filtering Service machine), users may experience slow browsing, as requests for sites take longer to process.

During times of peak CPU usage (exceeding 95%), Filtering Service may be unable to process requests at all, leading to incorrect filtering.

To address this issue, start by using the Task Manager (Windows) or **top** command (Linux) to determine which processes on the machine are causing CPU usage to peak.

- ◆ Are there applications that could be run from another machine?
- ◆ Can you move Filtering Service to a dedicated machine?

If Filtering Service is using a high amount of processing time:

- ◆ Evaluate the amount of traffic being processed by Filtering Service. DNS lookups can require a fair amount of processing time; you may want to install an additional Filtering Service instance for load balancing.
- ◆ Evaluate your use of keywords and regular expressions. Are you using a large number of regular expressions or keywords, or using very complex regular expressions?

Reducing the number of keywords and regular expressions, or removing or simplifying complex regular expressions can improve Filtering Service performance.

Sites are incorrectly categorized as Information Technology

Internet Explorer versions 4.0 and later have the ability to accept searches from the Address bar. When this option is enabled, if a user enters only a domain name in the Address bar (**websense** instead of **http://www.websense.com**, for example), Internet Explorer considers the entry a search request, not a site request. It displays the most likely site the user is looking for, along with a list of closely matching sites.

As a result, Websense software permits, blocks, or limits the request based on the status of the Information Technology/Search Engines and Portals category in the active policy—not on the category of the requested site. For Websense software to filter based on the category of the requested site, you must turn off searching from the Address bar:

1. Go to **Tools > Internet Options**.
2. Go to the **Advanced** tab.
3. Under Search from the Address bar, select:
 - Internet Explorer 5, 6, and 7: **Do not search from the Address bar**
 - Internet Explorer 8: **Do not submit unknown addresses to your auto-search provider**
4. Click **OK**.

Keywords are not being blocked

There are 2 possible reasons for this problem: **Disable keyword blocking** is selected, or the site whose URL contains the keyword uses **post** to send data to your Web server.

To ensure that keyword blocking is enabled:

1. Go to **Settings > General > Filtering**.
2. Under General Filtering, check the **Keyword search options** list. If **Disable keyword blocking** is shown, select another option from the list. See [Configuring Websense filtering settings, page 57](#), for more information about the available options.
3. Click **OK** to cache the change. Changes are not implemented until you click **Save All**.

If a site uses **post** to send data to your Web server, Websense software does not recognize keyword filtering settings for that URL. Unless your integration product recognizes data sent via post, users can still access URLs containing blocked keywords.

To see whether a site uses a post command, view the site's source from within your browser. If the source code contains a string like `<method=post>`, then post is used to load that site.

Custom or limited access filter URLs are not filtered as expected

If an HTTPS URL in a limited access filter or custom URL list (recategorized or unfiltered) is not filtered as expected, an integration product may be transforming the URL into a format that Filtering Service cannot recognize.

Non-proxy integration products translate URLs from domain format into IP format. For example, the URL `https://<domain.com>` is read as `https://<IP_address>:443`. When this occurs, Filtering Service cannot match the URL received from the integration product with a custom URL or limited access filter, and does not filter the site appropriately.

To work around this problem, add both the IP addresses and URLs for sites you want to filter using custom URLs or limited access filters.

A user cannot access a protocol or application as expected

If your network includes Microsoft ISA Server, certain authentication method configurations may result in dropped connections to messaging applications.

If any method other than Anonymous Authentication is active, the proxy server attempts to identify data packets received when users request application connections. The proxy server fails to identify the data packet, and the connection is dropped. This may skew Websense protocol filtering activity.

An inability to access a protocol or Internet application might also occur if the port used by the application is blocked. This could occur if:

- ◆ The port is blocked by a firewall.
- ◆ A blocked custom protocol includes the port (as a single port or as part of a port range) in any of its identifiers.

An FTP request is not blocked as expected

When integrated with Check Point® firewalls, Websense software requires **folder view** to be enabled in the client's browser to recognize and filter FTP requests.

When folder view is not enabled, FTP requests sent to the FireWall-1 proxy are sent to Websense software with an "http:/" prefix. As a result, Websense software filters these requests as HTTP requests, rather than FTP requests.

Websense software is not applying user or group policies

Users may not be filtered by the user or group policy you assigned for a variety of reasons. Check the following topics, and search the [Knowledge Base](#) for additional information.

- ◆ *User Service is not available, page 396*
- ◆ *Remote users are not filtered by correct policy, page 399*
- ◆ *Users are incorrectly filtered by the Default policy, page 406*
- ◆ *Directory service connectivity and configuration, page 413*
- ◆ *Directory service configuration in TRITON - Web Security, page 413*
- ◆ *User identification and Windows Server 2008, page 414*
- ◆ *User Service on Linux, page 415*
- ◆ *Remote users are not being filtered correctly, page 417*

Remote users are not filtered by correct policy

If a remote user accesses the network by logging on using cached domain credentials (network logon information), Websense software applies the policy assigned to that user, or to the user's group or domain, if appropriate. If there is no policy assigned to the user, group, or domain, or if the user logs on to the computer with a local user account, Websense software applies the Default policy.

Occasionally, a user is not filtered by a user or group policy or the Default policy. This occurs when the user logs on to the remote computer with a local user account, and the last portion of the remote computer's Media Access Control (MAC) address overlaps with an in-network IP address to which a policy has been assigned. In this case, the policy assigned to that particular IP address is applied to the remote user.

Network Agent issues

- ◆ [Network Agent is not installed](#), page 400
- ◆ [Network Agent is not running](#), page 400
- ◆ [Network Agent is not monitoring any NICs](#), page 401
- ◆ [Network Agent can't communicate with Filtering Service](#), page 401

Network Agent is not installed

Network Agent is required to enable protocol filtering. With some integrations, Network Agent is also used to provide more accurate logging.

If you are running with an integration product, and do not require Network Agent protocol filtering or logging, you can hide the “No Network Agent is installed” status message. See [Reviewing current system status](#), page 339, for instructions.

For stand-alone installations, Network Agent must be installed for network traffic to be monitored and filtered. See the *Deployment and Installation Center* for installation instructions, and then see [Network Agent configuration](#), page 379.

Network Agent is not running

Network Agent is required to enable protocol filtering. With some integrations, Network Agent is also used to provide more accurate logging.

For stand-alone installations, Network Agent must be running to monitor and filter network traffic.

To troubleshoot this problem:

1. Check the Windows Services dialog box (see [The Windows Services dialog box](#), page 458) to see if the **Websense Network Agent** service has started.
2. Restart the **Websense Policy Broker** and **Websense Policy Server** services (see [Stopping and starting Websense services](#), page 329).
3. Start or restart the **Websense Network Agent** service.
4. Close TRITON - Web Security.
5. Wait 1 minute, and then log on to TRITON - Web Security again.

If that does not fix the problem:

- ◆ Check the **Windows Event Viewer** for error messages from Network Agent (see [The Windows Event Viewer](#), page 458).
- ◆ Check the **Websense.log** file for error messages from Network Agent (see [The Websense log file](#), page 458).

Network Agent is not monitoring any NICs

Network Agent must be associated with at least one network interface card (NIC) to monitor network traffic.

If you add or remove network cards from the Network Agent machine, you must update your Network Agent configuration.

1. In TRITON - Web Security, go to **Settings**.
2. In the left navigation pane, under Network Agent, select the IP address of the Network Agent machine.
3. Verify that all NICs for the selected machine are listed.
4. Verify that at least one NIC is set to monitor network traffic.

See [Network Agent configuration](#), page 379, for more information.

Network Agent can't communicate with Filtering Service

Network Agent must be able to communicate with Filtering Service to enforce your Internet usage policies.

- ◆ Did you change the IP address of Filtering Service machine or reinstall Filtering Service?

If so, see [Update Filtering Service IP address or UID information](#), page 401.

- ◆ Do you have more than 2 network interface cards (NICs) on the Network Agent machine?

If so, see [Network Configuration](#), page 375, to verify your Websense software settings.

- ◆ Have you reconfigured the switch connected to the Network Agent?

If so, refer to the *Deployment and Installation Center* to verify your hardware setup, and see [Network Agent configuration](#), page 379, to verify your Websense settings.

If none of these apply, see [Configuring local settings](#), page 381, for information about associating Network Agent and Filtering Service.

Update Filtering Service IP address or UID information

When Filtering Service has been uninstalled and reinstalled, Network Agent does not automatically update the internal identifier (UID) for the Filtering Service. TRITON - Web Security attempts to query Filtering Service using the old UID, which no longer exists.

Likewise, when you change the IP address of the Filtering Service machine, this change is not automatically registered.

To re-establish connection to the Filtering Service:

1. Open TRITON - Web Security.

A status message indicates that a Network Agent instance is unable to connect to Filtering Service.

2. Click **Settings** at the top of the left navigation pane.
3. In the left navigation pane, under Network Agent, select the IP address of the Network Agent machine.
4. At the top of the page, under Filtering Service Definition, expand the **Server IP address** list, and then select the IP address of the Filtering Service machine.
5. Click **OK** at the bottom of the page to cache the update. Changes are not implemented until you click **Save All**.

Insufficient memory on the Network Agent machine

Network Agent allocates the operation memory that it needs at startup. If there are severe memory constraints on the Network Agent machine, the agent will either:

- ◆ Fail to start
- ◆ Be unable to monitor traffic

In either case, filtering and logging based on information from Network Agent does not occur. As a result, users may be given access to sites or applications that would typically be blocked.

Use the Task Manager (Windows) or **top** command (Linux) to evaluate memory usage on the Network Agent machine. To solve the problem, you can:

- ◆ Upgrade the RAM on the machine.
- ◆ Move applications or components with high memory requirements to another machine.
- ◆ Simplify your Network Agent configuration to reduce memory needs.

High CPU usage on the Network Agent machine

When the CPU on the Network Agent machine is overloaded by demands from other software running on the machine, the agent may be unable to detect and log traffic. In a Stand-Alone environment, this can mean that all user requests for Web sites and Internet applications are permitted, even those that would be typically be blocked.

To address this issue, start by using the Task Manager (Windows) or **top** command (Linux) to determine which processes on the machine are causing CPU usage to peak.

- ◆ Are there applications that could be run from another machine?
- ◆ Can you move Network Agent to a dedicated machine?

User configuration and identification issues

- ◆ *User and group-based policies are not applied, page 403*

- ◆ [DC Agent has insufficient permissions](#), page 404
- ◆ [DC Agent unable to access required file](#), page 404
- ◆ [General DC Agent troubleshooting](#), page 405
- ◆ [Troubleshooting Logon Agent](#), page 406
- ◆ [Troubleshooting eDirectory Agent](#), page 409
- ◆ [Troubleshooting RADIUS Agent](#), page 411
- ◆ [I cannot add users and groups to TRITON - Web Security](#), page 412
- ◆ [User Service on Linux](#), page 415

User and group-based policies are not applied

If Websense software is using computer or network policies, or the **Default** policy, to filter Internet requests, even after you have assigned user or group-based policies, or if the wrong user or group-based policy is being applied, use the following steps to pinpoint the problem:

- ◆ If you are using Microsoft ISA Server, and changed its authentication method, ensure that the Web Proxy Service was restarted.
- ◆ If you are using nested groups in Windows Active Directory, policies assigned to a parent group are applied to users belonging to a sub-group, and not directly to the parent group. For information on user and group hierarchies, see your directory service documentation.
- ◆ The User Service cache may be outdated. User Service caches user name to IP address mappings for 3 hours. To clear and recreate the cache, go to the User Service Cache section of the Settings > General > Directory Services page in TRITON - Web Security, and then click **Clear Cache**.
- ◆ User Service may have been installed using the Guest account, equivalent to an anonymous user to the domain controller. If the domain controller has been set not to give the list of users and groups to an anonymous user, User Service is not allowed to download the list. See [Changing DC Agent, Logon Agent, and User Service permissions](#), page 415.
- ◆ If the user being filtered incorrectly is on a machine running Windows XP SP2, the problem could be due to the Windows Internet Connection Firewall (ICF), included and enabled by default in Windows XP SP2. For more information about the Windows ICF, see Microsoft Knowledge Base Article #320855.

For DC Agent or Logon Agent to get user logon information from a machine running Windows XP SP2:

1. On the Client machine, go to **Start > Settings > Control Panel > Security Center > Windows Firewall**.
2. Go to the **Exceptions** tab.
3. Check **File and Printer Sharing**.
4. Click **OK** to close the ICF dialog box, and then close any other open windows.

If none of these steps addresses your issue, check the following topics, or search the [Knowledge Base](#) for additional information.

- ◆ *Directory service connectivity and configuration*, page 413
- ◆ *Directory service configuration in TRITON - Web Security*, page 413
- ◆ *User identification and Windows Server 2008*, page 414

DC Agent has insufficient permissions

DC Agent may have been installed as a service using the Guest account, equivalent to an anonymous user to the domain controller.

At minimum, the Websense DC Agent service requires **domain admin** permissions to retrieve user logon information from the directory service. In some environments (typically very large enterprise networks), DC Agent requires **enterprise admin** permissions.

1. On the DC Agent machine, create a user account with a descriptive name, like **WsUserID**. This account exists only to provide a security context for DC Agent when it requests information from the directory service.
 - Assign the new account **domain admin** privileges in all domains.
 - Assign the same password to this account in all domains.
 - Set the password to never expire.

Make a note of the user name and password.
2. Open the Windows **Services** dialog box (Start > Administrative Tools > Services).
3. Scroll to the **Websense DC Agent** service, right-click the service name, and then select Stop.
4. Double-click the service name, and then select the **Log On** tab.
5. Select **This account**, and then enter the account name and password that you created for DC Agent. Some domains require that the account name be entered in the format domain\username.
6. Click **OK** to return to the Services dialog box.
7. Right-click the service name again, and then select **Start**.
8. Close the Services dialog box.

You may also need to assign User Service the same administrative privileges as DC Agent.

DC Agent unable to access required file

DC Agent works by identifying domain controllers in the network, and then querying those domain controllers for user logon sessions. By default, the agent automatically verifies existing domain controllers and detects new domains or domain controllers added to the network. It stores this information in a file called **dc_config.txt**, located in the Websense **bin** directory on the DC Agent machine.

An alert stating that DC Agent is unable to access this file can occur if:

- ◆ DC Agent is unable to open the file with read or write permissions.
 - Make sure that account used to run DC Agent has read and write permissions to the file and directory.
 - If the file is present, and not write protected, make sure that the file can be opened manually, and has not be corrupted.
- ◆ DC Agent is unable to create the file, because it cannot find any domain controller information.
 - If User Service is installed on a Linux machine, make sure that you have performed required WINS setup steps. For complete instructions, see [User Service on Linux, page 415](#).
 - If User Service is installed on a Windows Server 2008 machine, make sure that the service is running with **domain admin** credentials. See [Changing DC Agent, Logon Agent, and User Service permissions, page 415](#).
 - Make sure that NetBIOS for TCP/IP is enabled, and that the NetBIOS ports (137, 138, 139, and 445) are open between the DC Agent machine and the domain controller.
If User Service is running on Windows, make sure that the NetBIOS ports are also open between the User Service machine and the domain controller.
 - Make sure that the Computer Browser Service is running on any Windows 2008 Server machine that hosts DC Agent, User Service, or Active Directory. See [Turning on the Computer Browser service, page 414](#).
- ◆ DC Agent does not find any valid entries in the file.
 - Make sure that at least one domain controller entry in the file is enabled. If all entries are disabled, DC Agent has effectively been instructed to stop working.
 - Make sure that all entries in the file are in a valid format. For example:


```
[WEST_DOMAIN]
dcWEST1=on
dcWEST2=on
[EAST_DOMAIN]
dcEAST1=on
dcEAST2=off
```

General DC Agent troubleshooting

To troubleshoot user identification problems with DC Agent:

1. Check all network connections.
2. Check the Windows Event Viewer for error messages (see [The Windows Event Viewer, page 458](#)).
3. Check the Websense log file (Websense.log) for detailed error information (see [The Websense log file, page 458](#)).

Common causes for DC Agent user identification problems include:

- ◆ Network or Windows services are communicating with the domain controller in a way that makes DC Agent see the service as a new user, to whom no policy has been defined. See [Users are incorrectly filtered by the Default policy](#), page 406.
- ◆ The User Service cache is outdated. User Service caches user-name-to-IP-address mappings for 3 hours, by default. The cache is also updated each time you make changes and click **Save All** in TRITON - Web Security.

Users are incorrectly filtered by the Default policy

When some services or Microsoft Windows 200x machines contact the domain controller, the account name they use can cause Websense software to believe that an unidentified user is accessing the Internet from the filtered machine. Because no user or group-based policy has been assigned to this user, the computer or network policy, or the Default policy, is applied.

- ◆ Network services may require domain privileges to access data on the network, and use the domain user name under which they are running to contact the domain controller.

To address this issue, see [Configuring an agent to ignore certain user names](#), page 275.

- ◆ Windows 200x services contact the domain controller periodically with a user name made up of the computer name followed by a dollar sign (jdoe-computer\$). DC Agent interprets the service as a new user, to whom no policy has been assigned.

To address this issue, configure DC Agent to ignore any logon of the form **computer\$**.

1. On the DC Agent machine, navigate to the Websense **bin** directory ([Where is the Websense "bin" directory?](#), page 457).
2. Open the **transid.ini** file in a text editor.
3. Add the following entry to the file:


```
IgnoreDollarSign=true
```
4. Save and close the file.
5. Restart DC Agent (see [Stopping and starting Websense services](#), page 329).

Troubleshooting Logon Agent

If some users in your network are filtered by the **Default** policy because Logon Agent is not able to identify them:

- ◆ Make sure that Windows Group Policy Objects (GPO) are being applied correctly to these users' machines (see [Group Policy Objects](#), page 407).
- ◆ If User Service is installed on a Linux machine and you are using Windows Active Directory (Native Mode), check your directory service configuration (see [User Service on Linux](#), page 415).
- ◆ Verify that the client machine can communicate with the domain controller from which the logon script is being run (see [Domain controller visibility](#), page 407).

- ◆ Ensure that NetBIOS is enabled on the client machine (see [NetBIOS](#), page 407).
- ◆ Make sure that the user profile on the client machine has not become corrupt (see [User profile issues](#), page 408).

Group Policy Objects

After verifying that your environment meets the prerequisites described in the *Deployment and Installation Center* for your Websense software, make sure that Group Policy Objects are being applied correctly:

1. On the Active Directory machine, open the Windows Control Panel and go to **Administrative Tools > Active Directory Users and Computers**.
2. Right-click the domain entry, and then select **Properties**.
3. Click the **Group Policy** tab, and then select the domain policy from the Group Domain Policy Objects Links list.
4. Click **Edit**, and then expand the User Configuration node in the directory tree.
5. Expand the Windows Settings node, and then select **Scripts**.
6. In the right pane, double-click **Logon**, and then verify that **logon.bat** is listed in the Logon Properties dialog box.

This script is required by the client Logon Application.

- If **logon.bat** is not in the script, refer to the *Deployment and Installation Center*.
- If **logon.bat** does appear in the script, but Logon Agent is not working, use the additional troubleshooting steps in this section to verify that there is not a network connectivity problem, or refer to the Websense [Knowledge Base](#).

Domain controller visibility

To verify that the client machine can communicate with the domain controller:

1. Attempt to map a drive on the client machine to the domain controller's root shared drive. This is where the logon script normally runs, and where **LogonApp.exe** resides.
2. On the client machine, open a Windows command prompt and execute the following command:

```
net view /domain:<domain name>
```

If either of these tests fails, see your Windows operating system documentation for possible solutions. There is a network connectivity problem not related to Websense software.

NetBIOS

NetBIOS for TCP/IP must be enabled and the TCP/IP NetBIOS Helper service must be running for the Websense logon script to execute on the user's machine.

To make sure that NetBIOS for TCP/IP is enabled on the client machine.

1. Right-click **My Network Places**, and then select **Properties**.
2. Right-click **Local Area Connection**, and then select **Properties**.
3. Select **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**.
5. Select the **WINS** tab, and then verify that the correct NetBIOS option is set.
6. If you make a change, click **OK**, then click **OK** twice more to close the different Properties dialog boxes and save your changes.

If no change was needed, click **Cancel** to close each dialog box without making changes.

Use the Windows Services dialog box to verify that the **TCP/IP NetBIOS Helper** service is running on the client machine (see [The Windows Services dialog box](#), page 458). The TCP/IP NetBIOS Helper service runs on Windows 2000, Windows XP, Windows Server 2003, and Windows NT.

User profile issues

If the user profile on the client machine is corrupt, the Websense logon script (and Windows GPO settings) cannot run. This problem can be resolved by recreating the user profile.

When you recreate a user profile, the user's existing My Documents folder, Favorites, and other custom data and settings are not automatically transferred to the new profile. Do not delete the existing, corrupted profile until you have verified that the new profile has solved the problem and copied the user's existing data to the new profile.

To recreate the user profile:

1. Log on to the client machine as a local administrator.
2. Rename the directory that contains the user profile:
`C:\Documents and Settings\<user name>`
3. Restart the machine.
4. Log on to the machine as the filtered user. A new user profile is created automatically.
5. Check to make sure the user is filtered as expected.
6. Copy the custom data (such as the contents of the My Documents folder) from the old profile to the new one. Do not use the File and Settings Transfer Wizard, which may transfer the corruption to the new profile.

Troubleshooting eDirectory Agent

Related topics:

- ◆ [Enabling eDirectory Agent diagnostics](#), page 410
- ◆ [eDirectory Agent miscounts eDirectory Server connections](#), page 410
- ◆ [Running eDirectory Agent in console mode](#), page 411

A user may not be filtered properly if the user name is not being passed to eDirectory Agent. If a user does not log on to Novell eDirectory server, eDirectory Agent cannot detect the logon. This happens because:

- ◆ A user logs on to a domain that is not included in the default root context for eDirectory user logon sessions. This root context is specified during installation, and should match the root context specified for Novell eDirectory on the **Settings > Directory Services** page.
- ◆ A user tries to bypass a logon prompt to circumvent Websense filtering.
- ◆ A user does not have an account set up in eDirectory server.

If a user does not log on to eDirectory server, user-specific policies cannot be applied to that user. Instead, the **Default** policy takes effect. If there are shared workstations in your network where users log on anonymously, set up a filtering policy for those particular machines.

To determine whether eDirectory Agent is receiving a user name and identifying that user:

1. Activate eDirectory Agent logging, as described under [Enabling eDirectory Agent diagnostics](#), page 410.
2. Open the log file you have specified in a text editor.
3. Search for an entry corresponding to the user who is not being filtered properly.
4. An entry like the following indicates that eDirectory Agent has identified a user:

```
WsUserData::WsUserData()  
User: cn=Admin,o=novell (10.202.4.78)  
WsUserData::~~WsUserData()
```

In the example above, the user **Admin** logged on to eDirectory server, and was identified successfully.

5. If a user is being identified, but is still not being filtered as expected, check your policy configuration to verify that the appropriate policy is applied to that user, and that the user name in TRITON - Web Security corresponds to the user name in Novell eDirectory.

If the user is *not* being identified, verify that:

- The user has a Novell eDirectory account.

- The user is logging on to a domain that is included in the default root context for eDirectory user logons.
- The user is not bypassing a logon prompt.

Enabling eDirectory Agent diagnostics

eDirectory Agent has built-in diagnostic capabilities, but these are not activated by default. You can enable logging and debugging during installation, or at any other time.

1. Stop eDirectory Agent (see [Stopping and starting Websense services, page 329](#)).
2. On the eDirectory Agent machine, go to the eDirectory Agent installation directory.
3. Open the file **wsedir.ini** in a text editor.
4. Locate the **[eDirAgent]** section.
5. To enable logging and debugging, change the value of **DebugMode** to **On**:

```
DebugMode=On
```
6. To specify the log detail level, modify the following line:

```
DebugLevel=<N>
```

N can be a value from 0-3, where 3 indicates the most detail.
7. Modify the **LogFile** line to specify the name of the log output file:

```
LogFile=filename.txt
```

By default, log output is sent to the eDirectory Agent console. If you are running the agent in console mode (see [Running eDirectory Agent in console mode, page 411](#)), you can keep the default value.
8. Save and close the **wsedir.ini** file.
9. Start the eDirectory Agent service (see [Stopping and starting Websense services, page 329](#)).

eDirectory Agent miscounts eDirectory Server connections

If eDirectory Agent is monitoring more than 1000 users in your network, but shows only 1000 connections to the Novell eDirectory server, it may be due to a limitation of the Windows API that conveys information from the eDirectory server to the Websense eDirectory Agent. This occurs very rarely.

To work around this limitation, add a parameter to the **wsedir.ini** file that counts server connections accurately (Windows only):

1. Stop the Websense eDirectory Agent service (see [Stopping and starting Websense services, page 329](#)).
2. Go to the Websense **bin** directory ([Where is the Websense "bin" directory?, page 457](#)).
3. Open the **wsedir.ini** file in a text editor.
4. Insert a blank line, and then enter:

```
MaxConnNumber = <NNNN>
```

Here, <NNNN> is the maximum number of possible connections to the Novell eDirectory server. For example, if your network has 1,950 users, you might enter 2000 as the maximum number.

5. Save the file.
6. Restart eDirectory Agent.

Running eDirectory Agent in console mode

1. Do one of the following:
 - At the Windows command prompt (**Start > Run > cmd**), enter the command:

```
eDirectoryAgent.exe -c
```
 - At the Linux command shell, enter the command:

```
eDirectoryAgent -c
```
2. When you are ready to stop the agent, press **Enter**. It may take a few seconds for the agent to stop running.

Troubleshooting RADIUS Agent

RADIUS Agent has built-in diagnostic capabilities, but these are not activated by default. To activate RADIUS Agent logging and debugging:

1. Stop the RADIUS Agent service (see [Stopping and starting Websense services, page 329](#)).
2. Navigate to the Websense **bin** directory on the RADIUS Agent machine.
3. Open the **wsradius.ini** file in a text editor.
4. Locate the **[RADIUSAgent]** section.
5. To enable logging and debugging, change the value of **DebugMode** to **On**:

```
DebugMode=On
```
6. To specify the log detail level, modify the following line:

```
DebugLevel=<N>
```

N can be a value from 0-3, where 3 indicates the most detail.
7. Modify the **LogFile** line to indicate the name of the output file:

```
LogFile=filename.txt
```

By default, log output is sent to the RADIUS Agent console. If you are running the agent in console mode (see [Running RADIUS Agent in console mode, page 412](#)), you can optionally keep the default value.

8. Save and close the **wsradius.ini** file.
9. Start the RADIUS Agent service (see [Stopping and starting Websense services, page 329](#)).

If remote users are not being identified and filtered as expected, the likely cause is communication problems between RADIUS Agent and your RADIUS server. Check your RADIUS Agent logs for errors to determine the cause.

Running RADIUS Agent in console mode

To start RADIUS Agent in console mode (as an application), enter the following:

- ◆ At the Windows command prompt:

```
RadiusAgent.exe -c
```

- ◆ At the Linux shell prompt:

```
./RadiusAgent -c
```

To stop the agent at any time, press **Enter** again. It may take a couple of seconds for the agent to stop running.

RADIUS Agent accepts the following command-line parameters:



Note

On Linux, Websense, Inc., recommends using the script provided to start or stop Websense RADIUS Agent (**WsRADIUSAgent start|stop**), instead of the **-r** and **-s** parameters.

Parameter	Description
-i	Installs RADIUS Agent service/daemon.
-r	Runs RADIUS Agent service/daemon.
-s	Stops RADIUS Agent service/daemon.
-c	Runs RADIUS Agent as an application process instead of as a service or daemon. When in console mode, RADIUS Agent can be configured to send log output to the console or to a text file.
-v	Displays the version number of RADIUS Agent.
-? -h -help <no option>	Displays usage information on the command line. Lists and describes all possible command line parameters.

I cannot add users and groups to TRITON - Web Security

Being unable to see the list of users and groups when you attempt to add clients into TRITON - Web Security, can be caused by a variety of problems. Check the following topics, and check the [Knowledge Base](#) for additional information.

- ◆ [Directory service connectivity and configuration, page 413](#)
- ◆ [Directory service configuration in TRITON - Web Security, page 413](#)

- ◆ [User identification and Windows Server 2008](#), page 414

Directory service connectivity and configuration

Make sure that the Websense User Service machine and your directory server are running, and able to communicate over the network. The default ports used for directory service communication are:

139	NetBIOS communication: Active Directory
389	LDAP communication: Active Directory, Novell eDirectory, Oracle (formerly Sun Java) Directory Server
636	SSL port: Novell eDirectory, Oracle (formerly Sun Java) Directory Server
3268	Active Directory
3269	SSL port: Active Directory

In addition, consider the following:

- ◆ If you are using Windows Active Directory in **mixed** mode, and User Service runs on a Windows Server 2008 machine, the account used to run User Service must have administrative privileges on the directory. This may sometimes also be required in other versions of Windows Server.
To check or change the User Service account, see [Changing DC Agent, Logon Agent, and User Service permissions](#), page 415.
- ◆ If you are running Active Directory in **native** mode, set the User Service to run as the Local System account. No account should be assigned to the actual service.
User Service connects to the directory with the administrator user name and password configured on the Settings > General > Directory Services > Add Global Catalog Server page in TRITON - Web Security.
- ◆ If you are running User Service on a Linux machine and communicating with a Windows-based directory service, see [User Service on Linux](#), page 415, for additional instructions.
- ◆ Determine whether a firewall is blocking communication between TRITON - Web Security and User Service on port 55815. If so, open the blocked port.

Directory service configuration in TRITON - Web Security

If you encounter problems adding users and groups in TRITON - Web Security, make sure that you have provided complete and accurate configuration for your directory service.

1. Go to Settings > General > Directory Services.
2. Select the directory service used by your organization.
3. Verify the configuration. See [Directory services](#), page 65, and its sub-topics for details.

If Websense User Service is installed on a Linux machine, and is configured to communicate with Active Directory, see [User Service on Linux](#), page 415, for additional configuration requirements.

User identification and Windows Server 2008

You may encounter problems adding users and groups in TRITON - Web Security if you install one or more of the following components on Windows Server 2008:

- ◆ Websense User Service
- ◆ Windows Active Directory

If your network uses Active Directory in mixed mode, the Windows Computer Browser service must be running on the machine where User Service is installed, and also on the machine running Active Directory 2008. This service is turned on by default in earlier versions of Windows. However, it is disabled by default on Windows Server 2008.

In addition, when User Service is installed on Windows Server 2008, and you are using Active Directory in mixed mode, you must configure User Service with domain rights to access information from Active Directory.

If you are running User Service on Linux and using Active Directory 2008, additional configuration is required. See [User Service on Linux](#), page 415.

To enable the Computer Browser service on a relevant machine, see [Turning on the Computer Browser service](#), page 414.

To configure User Service with rights to access directory information, see [Changing DC Agent, Logon Agent, and User Service permissions](#), page 415.

Turning on the Computer Browser service

Websense Setup offers the option to turn on the Computer Browser service during installation of the following components on Windows Server 2008.

- ◆ Websense User Service
- ◆ Websense DC Agent
- ◆ Websense Logon Agent

If you chose not to have it started, or the installer was not successful, you must turn on the service manually.

Perform the following procedure on each machine running an affected component:

1. Make sure that Windows Network File Sharing is enabled.
 - a. Go to **Start > Network** and click **Network and Sharing Center**.
 - b. Click **Advanced Sharing Settings**, then select **Turn on file and print sharing**.
2. Go to **Control Panel > Administrative Tools > Services**.
3. Double-click **Computer Browser** to open the Properties dialog box.

4. Set the Startup type to **Manual**.
5. Click **Start**.
6. Change the Startup type to **Automatic**. This ensures that the service is started automatically every time the machine is restarted.
7. Click **OK** to save your changes and close the Services dialog box.
8. Repeat these steps on each machine running Windows Server 2008 and an affected component.

Changing DC Agent, Logon Agent, and User Service permissions

Sometimes, DC Agent, Logon Agent, or User Service needs to run as an account that has permission to access the directory service.

1. On the machine running the domain controller, create a user account such as **Websense**. You can use an existing account, but a Websense account is preferable so the password can be set not to expire. No special privileges are required.
Set the password never to expire. This account only provides a security context for accessing directory objects.
Make note of the user name and password you establish for this account, as they must be entered in step 6 and 7.
2. On the machine running an affected component, go to **Start > Programs > Administrative Tools > Services**.
3. Select the appropriate Websense service entry, listed below, and then click **Stop**.
 - Websense DC Agent
 - Websense Logon Agent
 - Websense User Service
4. Double-click the Websense service entry.
5. On the **Log On** tab, select the **This account** option.
6. Enter the user name of the Websense account created in step 1. For example: **DomainName\websense**.
7. Enter and confirm the Windows password for this account.
8. Click **OK** to close the dialog box.
9. Select the Websense service entry in the Services dialog box, and then click **Start**.
10. Repeat this procedure for each instance of Websense DC Agent, Logon Agent, and User Service in the network.

User Service on Linux

If you plan to apply filtering policies to individual users and groups in your network, special configuration steps are required to assure that the Websense software can identify users successfully in networks that run Websense User Service on a Linux server and one of the following:

- ◆ Use Active Directory in mixed mode

- ◆ Plan to use Websense Logon Agent to transparently identify users via Active Directory in native mode
- ◆ Use DC Agent to transparently identify users

In these environments, Websense software must be configured to communicate with a Windows Internet Name Server (WINS) to resolve domain names to domain controller IP addresses. The precise steps vary, depending on your environment.

If your network uses Windows Active Directory in mixed mode:

1. In TRITON - Web Security, go to the **Settings > General > Directory Services** page.
2. Select **Windows Active Directory (Mixed Mode)**. This is the default option.
3. Enter the name and password for the administrative user.
4. Enter the **Domain** name.

If your organization uses multiple domains, enter the name of a domain that is trusted by all domains that authenticate your users.

5. Enter the IP address of a Windows Internet Name Server (WINS) that can resolve the domain name entered above to a domain controller IP address.
6. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

If your network uses Active Directory (Native Mode), and you need to configure WINS settings:

1. Go to the **Settings > General > Directory Services** page.
2. Provide administrative credentials and identify the Windows Internet Name Server (WINS), as follows.
 - a. Select **Windows Active Directory (Mixed Mode)**, which is the default.
 - b. Enter the name and password for the administrative user.
 - c. Enter the **Domain** name.

If your organization uses multiple domains, enter the name of a domain that is trusted by all domains that authenticate your users.
 - d. Enter the IP address of a Windows Internet Name Server (WINS) that can resolve the domain name entered above to a domain controller IP address.
 - e. Click **OK** to cache your changes.
 - f. Click **Save All** to implement these changes.
3. On the Directory Service page, select **Active Directory (Native Mode)**.
4. Configure the global catalog servers and other settings for your directory service. See *Windows Active Directory (Native Mode)*, page 66, for assistance.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Remote users are not prompted for manual authentication

If you have configured remote users to manually authenticate when accessing the Internet, there may be some occasions when individual users are not prompted for the authentication. This can occur in situations where some in-network IP addresses have been configured to bypass manual authentication.

When a remote user accesses the network, Websense software reads the last portion of the computer's Media Access Control (MAC) address. If this matches an in-network IP address that has been configured to bypass manual authentication, the remote user will not be prompted to authenticate manually when accessing the Internet.

One solution is to reconfigure the in-network IP address to use manual authentication. An alternative solution is to disable the manual authentication requirement for the affected remote user.

Remote users are not being filtered correctly

If remote users are not being filtered, or are not being filtered by particular policies assigned to them, check the RADIUS Agent logs for the message **Error receiving from server: 10060** (Windows) or **Error receiving from server: 0** (Linux).

This usually occurs when the RADIUS server does not recognize RADIUS Agent as a client (source of RADIUS requests). Ensure that your RADIUS server is configured properly (see [Configuring the RADIUS environment](#), page 262).

You can use RADIUS Agent's built-in diagnostic tool to troubleshoot filtering problems (see [Troubleshooting RADIUS Agent](#), page 411).

If you have installed remote filtering software (see [Filter Users Off Site](#), page 201), off-site users cannot be filtered if the Remote Filtering Client cannot communicate with the Remote Filtering Server within the network.

For instructions on configuring remote filtering software, see the [Remote Filtering Software](#) technical paper.

Block message issues

- ◆ [No block page appears for a blocked file type](#), page 418
- ◆ [Users receive a browser error instead of a block page](#), page 418
- ◆ [A blank white page appears instead of a block page](#), page 419
- ◆ [Protocol block messages don't appear as expected](#), page 419
- ◆ [A protocol block message appears instead of a block page](#), page 419

No block page appears for a blocked file type

When file type blocking is used, the block message may not always be visible to the user. For example, when a downloadable file is contained within an internal frame (iframe) on a permitted site, the block message sent to that frame is not visible because the frame size is zero.

This is only a display problem; users cannot access or download the blocked file.

Users receive a browser error instead of a block page

If users receive an error message instead of a block page, the 2 most likely causes are:

- ◆ The user's browser is configured to use an external proxy. In most browsers, there is a setting that enables use of an external proxy. Verify that the browser is not set to use an external proxy.
- ◆ There is a problem identifying or communicating with the Filtering Service machine.

If the user's browser settings are correct, make sure that the IP address of the Filtering Service machine is listed correctly in the **eimserver.ini** file.

1. Stop **Websense Filtering Service** (see [Stopping and starting Websense services](#), page 329).
2. Navigate to the Websense **bin** directory ([Where is the Websense "bin" directory?](#), page 457).
3. Open the **eimserver.ini** file in a text editor.
4. Under **[WebsenseServer]**, add a blank line, and enter the following:


```
BlockMsgServerName = <Filtering Service IP address>
```

For example, if the Filtering Service IP address is 10.201.72.15, enter:

```
BlockMsgServerName = 10.201.72.15
```
5. Save and close the file.
6. Restart Filtering Service.

If the Filtering Service machine has more than one NIC, and the block page still does not display correctly after editing the **eimserver.ini** file, try the IP addresses of the other NICs in the **BlockMsgServerName** parameter.

If the block page still does not appear, make sure that users have read access to the files in the Websense block page directories:

- ◆ Websense\BlockPages\en\Default
- ◆ Websense\BlockPages\en\Custom

If the block page problem persists, search the Websense [Knowledge Base](#) for additional troubleshooting hints.

A blank white page appears instead of a block page

When advertisements are blocked, or when a browser does not correctly detect the encoding associated with a block page, users may receive a blank white page instead of a block page. The reasons for this behavior are as follows:

- ◆ When the Advertisements category is blocked, Websense software sometimes interprets a request for a graphic file as an advertisement request, and displays a blank image instead of a block message (the normal method for blocking advertisements). If the requested URL ends in .gif or similar, have the user reenter the URL, leaving off the *.gif portion. See [Blocking graphical advertisements, page 92](#).
- ◆ Some older browsers may not detect the encoding of block pages. To enable proper character detection, configure your browser to display the appropriate character set (UTF-8 for French, German, Italian, Spanish, Brazilian Portuguese, Simplified Chinese, Traditional Chinese, or Korean; and Shift_JIS for Japanese). See your browser's documentation for instructions, or upgrade the browser to a newer version.

Protocol block messages don't appear as expected

Protocol block messages may not appear, or appear only after a delay, for any of the following reasons:

- ◆ User Service must be installed on a Windows machine in order for protocol block messages to display properly. For more information, see the *Deployment and Installation Center*.
- ◆ Protocol block messages may not reach client machines if Network Agent is installed on a machine with multiple network interface cards (NICs), and a NIC is monitoring a different network segment from Filtering Service. Ensure that the Filtering Service machine has NetBIOS and Server Message Block protocol access to client machines, and that port 15871 is not blocked.
- ◆ A protocol block message may be slightly delayed, or appear on an internal machine where the requested protocol data originated (instead of on the client machine), when Network Agent is configured to monitor requests **sent to** internal machines.
- ◆ If the filtered client or the Websense filtering machine is running Windows 200x, the Windows **Messenger** service must be running for the protocol block message to display. Use the Windows Services dialog box on the client or server machine to see if the Messenger service is running (see [The Windows Services dialog box, page 458](#)). Even though the block message does not appear, protocol requests are still blocked.

A protocol block message appears instead of a block page

If your integration product does not send HTTPS information to Websense software, or if Websense software is running in stand-alone mode, Network Agent may interpret an HTTPS site request that is blocked via category settings as a protocol request. As a

result, a protocol block message is displayed. The HTTPS request is also logged as a protocol request.

Log, status message, and alert issues

- ◆ [Where do I find error messages for Websense components?](#), page 420
- ◆ [Websense Health alerts](#), page 421
- ◆ [Two log records are generated for a single request](#), page 422
- ◆ [Usage Monitor is not available](#), page 423
- ◆ [Usage Monitor is not running](#), page 423

Where do I find error messages for Websense components?

When there are errors or warnings related to core Websense components, short alert messages are displayed in the **Health Alert Summary** list at the top of the **Status > Today** page in TRITON - Web Security (see [Websense Health alerts](#), page 421).

- ◆ Click an alert message to see more detailed information on the **Status > Alerts** page.
- ◆ Click **Solutions** next to a message on the Status > Alerts page for troubleshooting assistance.

Errors, warnings, and messages from Websense software components, as well as database download status messages, are recorded in the **websense.log** file in the Websense **bin** directory (see [The Websense log file](#), page 458).

For Websense software components installed on Windows machines, you can also check the Windows Event viewer. See [The Windows Event Viewer](#), page 458.

Websense Health alerts

The Websense Health Alert Summary lists any potential concerns encountered by monitored components of your Websense software. These include:

- User Service is not available
- There is a subscription problem
- No subscription key has been entered
- The subscription key is about to expire
- The initial filtering database is in use
- The Master Database is downloading for the first time
- The Master Database is being updated
- The Master Database is more than 1 week old
- The Master Database did not download successfully
- WebCatcher is not enabled
- A Log Server is not running
- The Log Database is not available
- Presentation reports scheduler is not connected to the Log Database
- One or more presentation report jobs failed
- Low disk space on the TRITON - Web Security machine
- There is no Log Server configured for a Policy Server
- Low disk space on the Log Server machine
- Network Agent is not running
- No monitoring NIC has been configured for a Network Agent
- No Filtering Service has been configured for a Network Agent
- Low memory on the Network Agent machine
- High CPU usage on the Network Agent machine
- There is no Network Agent configured for a Policy Server
- A Filtering Service is not running
- Low disk space on the Filtering Service machine
- Low memory on the Filtering Service machine
- High CPU usage on the Filtering Service machine
- A DC Agent instance has insufficient permissions
- A DC Agent instance is unable to access a required file
- Usage Monitor is not available
- Usage Monitor is not running

If you have subscribed to Websense Web Security Gateway or Gateway Anywhere, Websense software monitors Content Gateway to provide alerts about the following conditions:

- Content Gateway is not running
- Content Gateway is not available

If you have subscribed to Websense Web Security Gateway Anywhere, or if your subscription includes both Web and data security components, Websense software monitors interoperability components to provide alerts about the following conditions:

- A Sync Service is not running.
- There is no Sync Service associated with a Policy Server instance.
- On-premises components are unable to connect to the hybrid service.
- Disk space is low on the partition hosting Sync Service.
- 24 hours since Sync Service downloaded log files from the hybrid service.
- Missing information required to activate hybrid filtering.
- A Directory Agent is not running.
- There is no Directory Agent associated with a Policy Server instance.
- Alerts were received from the hybrid service.
- 24 hours since Sync Service sent log files to Log Server.

The icon next to the alert message indicates the potential impact of the related condition.



The message is informational, and does not reflect a problem with your installation (for example, WebCatcher is not enabled, or Filtering Service is downloading a Master Database update).



The alert condition has the potential to cause a problem, but will not immediately prevent filtering or reporting (for example, the Master Database is more than 1 week old, or the subscription key is about to expire).



A Websense software component is not functioning (has not been configured or is not running), which may impair filtering or reporting, or your subscription has expired.

Click an alert message in the Health Alerts Summary to go to the Alerts page, which provides additional information about current alert conditions. Click **Learn More** (for informational alerts) or **Solutions** (for errors or warnings) for details and troubleshooting tips.

If a Health Alert indicates that messages were received from the hybrid service, check the Hybrid Filtering Alerts table on the Status > Alerts page for details.

In some cases, if you are receiving error or status messages about a component that you are not using, or that you have disabled, you can choose to hide the alert messages. See [Reviewing current system status](#), page 339, for more information.

Two log records are generated for a single request

When Windows QoS Packet Scheduler is installed on the same machine as Network Agent, 2 requests are logged for each single HTTP or protocol request made from the Network Agent machine. (This duplication does not occur with requests made by client machines within your network.)

To fix the problem, disable Windows QoS Packet Scheduler on the Network Agent machine.

This problem does not occur if you use Network Agent for all logging. See [Configuring NIC settings, page 382](#), for details.

Usage Monitor is not available

In order to enable category and protocol usage alerting and Real-Time Monitor, Websense Usage Monitor must be installed. Typically, one Usage Monitor instance is installed for each Policy Server in your network. Usage Monitor may be installed on the Policy Server machine.

When installing Usage Monitor, make sure that it can communicate with:

- ◆ Policy Server on ports 55806 and 40000
- ◆ Policy Broker on port 55880
- ◆ Filtering Service and Real-Time Monitor on port 55809

Usage Monitor should also be able to receive information from Policy Server and Filtering Service on its listening port: 55813.

Usage Monitor is not running

When Websense Usage Monitor is stopped:

- ◆ Category and protocol access information cannot be collected for alerting purposes.
- ◆ Category and protocol usage alerts cannot be generated.
- ◆ Real-Time Monitor does not receive Internet activity data.

To start Usage Monitor:

- ◆ (*Windows*) Open the Windows Services dialog box, scroll to Websense Usage Monitor, right-click the service, and select **Start**.
- ◆ (*Linux*) Use the `/opt/Websense/WebsenseDaemonControl` command to start Usage Monitor.

If Usage Monitor will not start, check the Windows Event Viewer or **websense.log** file for error information from the service.

Policy Server and Policy Database issues

- ◆ [I forgot my password, page 424](#)
- ◆ [I cannot log on to Policy Server, page 424](#)
- ◆ [The Websense Policy Database service fails to start, page 424](#)
- ◆ [Policy Server stops unexpectedly, page 425](#)

I forgot my password

If you are a Super Administrator or delegated administrator using a local account to log on to the TRITON Unified Security Center, any Global Security Administrator can reset the password. Global Super Administrators can manage accounts and passwords on the TRITON Settings > Administrators page.

If a Global Super Administrator is not available, administrators using local accounts can request a new password via the **Forgot my password** link on the TRITON logon page.

- ◆ A temporary password will be sent to the email address associated with your administrator account.
- ◆ The temporary password is valid for only 30 minutes. If more than 30 minutes elapses before you attempt to log on with the temporary password, you must request a new password again.
- ◆ You will be prompted to enter a new password after you have logged on using the temporary password.

I cannot log on to Policy Server

Verify that the selected Policy Server IP address is correct. If the address of the Policy Server machine has changed since the Policy Server was added to TRITON - Web Security, you will need to log on to a different Policy Server, remove the old IP address from TRITON - Web Security, and then add the new Policy Server IP address. See [Reviewing Policy Server connections](#), page 320.

If TRITON - Web Security has stopped suddenly, or has been stopped via the kill (Linux) or End Task (Windows) commands, wait a few minutes before logging on again. Websense software detects and closes the terminated session within 3 minutes.

The Websense Policy Database service fails to start

The Websense Policy Database runs as a special account: **WebsenseDBUser**. If this account experiences logon problems, the Policy Database is unable to start.

To address this issue, change the WebsenseDBUser password.

1. Log on to the Policy Database machine as a local administrator.
2. Go to **Start > Programs > Administrative Tools > Computer Management**.
3. In the navigation pane, under System Tools, expand **Local Users and Groups**, and then select **Users**. User information is displayed in the content pane.
4. Right-click **WebsenseDBUser** and select **Set Password**.
5. Enter and confirm the new password for this user account, and then click **OK**.
6. Close the Computer Management dialog box.
7. Go to **Start > Programs > Administrative Tools > Services**.
8. Right-click **Websense Policy Database** and select **Properties**.

9. On the Log On tab of the Properties dialog box, enter the new WebsenseDBUser password information, and then click **OK**.
10. Right-click Websense Policy Database again, and then select **Start**.
When the service has started, close the Services dialog box.

Policy Server stops unexpectedly

If the hard disk on the Policy Server machine runs out of free space, the Websense Policy Server service or daemon stops. Even if the lack of disk space is the result of a transient condition (another application creates large temporary files, for example, and then removes them), Policy Server does not restart automatically.

- ◆ If Filtering Service or Network Agent is installed on the Policy Server machine, a Health Alert message in TRITON - Web Security will provide a warning that disk space is getting low.
- ◆ When Policy Server stops, a Health Alert message is displayed in TRITON - Web Security.

Manually restart Policy Server to address the immediate issue. Then, determine which application is sometimes filling up all available disk space on the machine. You can then decide whether the best solution is to move the application to another machine or to add disk space to the Policy Server machine.

Delegated administration issues

- ◆ [Managed clients cannot be deleted from role, page 425](#)
- ◆ [Logon error says someone else is logged on at my machine, page 426](#)
- ◆ [Some users cannot access a site in the Unfiltered URLs list, page 426](#)
- ◆ [Recategorized sites are filtered according to the wrong category, page 426](#)
- ◆ [I cannot create a custom protocol, page 426](#)

Managed clients cannot be deleted from role

Clients cannot be deleted directly from the managed clients list on the Delegated Administration >Edit Role page if:

- ◆ the administrator has applied a policy to the client
- ◆ the administrator has applied a policy to one or more members of a network, group, domain, or organizational unit

There may also be problems if, during TRITON - Web Security logon, the Super Administrator chooses a different Policy Server than the one that communicates with the directory service containing the clients to be deleted. In this situation, the current Policy Server and directory service do not recognize the clients.

For assistance deleting managed clients, see [Delete managed clients, page 306](#).

Logon error says someone else is logged on at my machine

When you attempt to log on to TRITON - Web Security you may sometimes receive the error “Logon failed. The role *<role name>* has been in use by *<user name>*, since *<date, time>*, on computer 127.0.0.1.” The IP address 127.0.0.1 is also called the loopback address, and typically indicates the local machine.

This message means that someone is logged on at the TRITON - Web Security installation machine, in the same role you are requesting. You can select a different role (if you administer multiple roles), log on for reporting only, or wait until the other administrator logs off.

Some users cannot access a site in the Unfiltered URLs list

Unfiltered URLs affect only the clients managed by role in which the URLs are added. For example, if a Super Administrator adds unfiltered URLs, clients managed by delegated administration roles are not granted access to those sites.

To make the site available to clients in other roles, the Super Administrator can switch to each role and add the relevant sites to that role’s unfiltered URLs list.

Recategorized sites are filtered according to the wrong category

Recategorized URLs affect only the clients managed by role in which the URLs are added. For example, when a Super Administrator recategorizes URLs, clients managed by delegated administration roles continue to be filtered according to the Master Database category for those sites.

To apply the recategorization to clients in other roles, the Super Administrator can switch to each role and recategorize the sites for that role.

I cannot create a custom protocol

Only Super Administrators are able to create custom protocols. However, delegated administrators can set filtering actions for custom protocols.

When Super Administrators create custom protocols, they should set the appropriate default action for most clients. Then, inform delegated administrators of the new protocol so they can update the filters for their role, as appropriate.

Log Server and Log Database issues

- ◆ [Log Server is not running, page 427](#)
- ◆ [Low disk space on the Log Server machine, page 427](#)
- ◆ [No Log Server is installed for a Policy Server, page 428](#)
- ◆ [Log Database was not created, page 429](#)

- ◆ [Log Database is not available](#), page 430
- ◆ [Log Database size](#), page 431
- ◆ [Configure Log Server to use a database account](#), page 431
- ◆ [Log Server is not recording data in the Log Database](#), page 431
- ◆ [Updating the Log Server connection account or password](#), page 432
- ◆ [Configuring user permissions for Microsoft SQL Server](#), page 432
- ◆ [Log Server cannot connect to the directory service](#), page 433
- ◆ [Wrong reporting page displayed](#), page 434

Log Server is not running

If Log Server is not running, or if other Websense components are unable to communicate with Log Server, Internet usage information is not stored, and you may not be able to generate Internet usage reports.

Log Server may be unavailable if:

- ◆ There is insufficient disk space on the Log Server machine.
- ◆ You changed the Microsoft SQL Server password without updating the ODBC or Log Server configuration.
- ◆ It has been more than 14 days since the Master Database was downloaded successfully.
- ◆ The **logserver.ini** file is missing or corrupted.
- ◆ You stopped Log Server to avoid logging Internet usage information.

To troubleshoot the problem:

- ◆ Verify the amount of free disk space, and remove extraneous files, as needed.
- ◆ If you believe that a password change is the source of the problem, see [Updating the Log Server connection account or password](#), page 432.
- ◆ Navigate to the Websense **bin** directory ([Where is the Websense “bin” directory?](#), page 457) and make sure that you can open **logserver.ini** in a text editor. If this file has been corrupted, replace it with a backup file.
- ◆ Check the Windows Services dialog box to verify that Log Server has started, and restart the service if necessary (see [Stopping and starting Websense services](#), page 329).
- ◆ Check the Windows Event Viewer and **websense.log** file for error messages from Log Server (see [Troubleshooting tips and tools](#), page 457).

Low disk space on the Log Server machine

Websense Log Server stores Internet filtering records in temporary log cache files or BCP (bulk copy program) files on the Log Server machine until they can be processed into the Log Database.

Websense software watches the space available for both log cache files and BCP files are stored. By default:

- ◆ Log cache files are stored in the **C:\Program Files\WebSense\bin\Cache** directory.
- ◆ BCP files are stored in the **C:\Program Files\WebSense\bin\Cache\BCP** directory.

The log cache file and BCP file location can be changed on the Database tab of the Log Server Configuration utility. See [Configuring WebCatcher](#), page 358.



Note

If you have multiple Log Servers that forward their data to a primary Log Server, disk space is tracked for the primary Log Server only.

A Health Alert message is displayed on the Status > Today page if the space available at either of these locations drops too low. If there is insufficient disk space, logging stops.

- ◆ A warning message appears when the free disk space falls below 10% on the drive where log cache files and BCP files are stored. Although logging continues, you should clear disk space on the machine as soon as possible to avoid loss of log data.
- ◆ An error message appears when there is less than 4 MB of free disk space on the drive where log cache files and BCP files are stored.

When disk space dips below 4 MB, logging may become intermittent or stop completely. To minimize loss of log data, clear disk space on the Log Server machine as soon as possible after the error message appears.

No Log Server is installed for a Policy Server

Websense Log Server collects Internet usage information and stores it in the Log Database for use in investigative reports, presentation reports, and the charts and summaries on the Today and History pages in TRITON - Web Security.

Log Server must be installed for reporting to occur.

You may see this message if:

- ◆ Log Server is installed on a different machine than Policy Server, and the Log Server IP address is incorrectly set to localhost in TRITON - Web Security.
- ◆ You are not using Websense reporting tools.
- ◆ Log Server is associated with a different Policy Server instance.

To verify that the correct Log Server IP address is set in TRITON - Web Security:

1. Select the **Settings** tab of the left navigation pane, and then go to **General > Logging**.

2. Enter the IP address of the Log Server machine in the **Log Server IP address or name** field.
3. Click **OK** to cache your change, and then click **Save All**.

If you are not using Websense reporting tools, or if Log Server is associated with a different Policy Server instance, you can hide the alert message in TRITON - Web Security.

1. On the Main tab of the left navigation pane, go to **Status > Alerts**.
2. Under Active Alerts, click **Advanced**.
3. Mark **Hide this alert** for the “No Log Server installed” message.
4. Click **Save Now**. The change is implemented immediately.

Log Database was not created

If the installer cannot create the Log Database, make sure that:

- ◆ The account used to log on for installation has inadequate SQL Server permissions to create a database. The required permissions depend on the version of Microsoft SQL Server:
 - SQL Server 2008 or 2005 (Standard or Enterprise):
 - **dbcreator** server role membership
 - **SQLServerAgentReader** permissions required
 - SQL Server Express: **sysadmin** permissions required

Update the logon account or log on with an account that already has the required permissions (see [Configuring user permissions for Microsoft SQL Server](#), page 432), then run the installer again.
- ◆ A file or files exist with the default Log Database names (wslogdb70 and wslogdb70_1), but the files are not properly connected to the database engine and cannot be used by the Websense installer.

Remove or rename the existing files, and then run the installer again.
- ◆ The account used to run the installer has inadequate permissions on the drive where the database is being installed.

Update the logon account to have read and write permissions for the installation location, or log on with a different account that already has these permissions. Then, run the installer again.
- ◆ There is insufficient disk space available to create and maintain the Log Database at the specified location.

Clear enough space on the selected disk to install and maintain the Log Database. Then, run the installer again. Alternatively, choose another location.

Log Database is not available

The Websense Log Database stores Internet usage information for use in presentation reports, investigative reports, and the charts and summaries on the Today and History pages in TRITON - Web Security.

If Websense software is unable to connect to the Log Database, first verify that the database engine (Microsoft SQL Server or Microsoft SQL Server Express) is running on the Log Database machine.

1. Open the Windows Services dialog box (see [The Windows Services dialog box, page 458](#)) and verify that the **MSSQLSERVER** service is running.
If you are running Microsoft SQL Server Standard or Enterprise (not Express), also make sure that the **SQLSERVERAGENT** service is running.
2. If a service has stopped, right-click the service name and click **Start**.
If the service does not restart, check the Windows Event Viewer (see [The Windows Event Viewer, page 458](#)) for Microsoft SQL Server errors and warnings.
3. If you are running Microsoft SQL Server Standard or Enterprise (not Express), double-click the SQLSERVERAGENT service to open a Properties dialog box, and verify that the **Startup type** is set to **Automatic**. This ensures that SQL Server Agent restarts each time Microsoft SQL Server, or the database engine machine, is restarted.
If the Startup type is Manual or Disabled, change it to **Automatic**, and then click **OK**.

If the database engine and (if applicable) SQL Server Agent services are running:

- ◆ Use the Windows Services dialog box to make sure that the **Websense Log Server** service is running.
- ◆ If Log Server and the Log Database are on different machines, make sure that both machines are running, and that the network connection between the machines is not impaired.
- ◆ Make sure that there is enough disk space on the Log Database machine, and that the Log Database has a sufficient quantity of allocated disk space (see [Log Server is not recording data in the Log Database, page 431](#)).
- ◆ Make sure that the SQL Server password has not been changed. If the password changes, you must update the password information Log Server uses to connect to the database. See [Updating the Log Server connection account or password, page 432](#).
- ◆ Make sure that there are no network interruptions that are preventing TRITON - Web Security from communicating with the Log Database.

After making sure that the database engine and related services are running, and that any network problems have been resolved, use the Windows Services dialog box to restart the Websense TRITON - Web Security service. This ensures that presentation reports scheduler can save job definitions (see [No Log Server is installed for a Policy Server, page 428](#)).

Log Database size

Log Database size is always a concern. If you have been successfully generating Websense reports and notice the reports are now taking much longer to display, or you begin receiving timeout messages from your Web browser, consider disabling some database partitions.

1. In TRITON - Web Security, go to **Settings > Reporting > Log Database**.
2. Locate the **Available Partitions** section of the page.
3. Clear the **Enable** check box for any partitions that are not required for current reporting operations.
4. Click **Save Now** to implement the change.

Configure Log Server to use a database account

When you install TRITON - Web Security on a Linux machine, reporting components are not available when Log Server uses a Windows Trusted Connection to connect to the Log Database.

To display all reporting components in TRITON - Web Security on Linux, configure Log Server to use a database account (like **sa**) to connect to the Log Database. To do this:

1. On the Log Server machine, open the Log Server Configuration Utility (Start > Programs > Websense > Log Server Configuration).
2. Select the Database tab, and then click **Connection**.
3. Select the Machine Data Source tab, and then double-click the data source name of the Log Database (**wslogdb70**, by default).
4. Deselect the **Use Trusted Connection** option, if selected, and then enter the user name (like **sa**) and password of a SQL Server account with create, read, and write permissions. See [Configuring user permissions for Microsoft SQL Server](#), page 432, for more information.
5. Click **OK** to close the SQL Server Login dialog box.
6. Select the Connection tab, and then stop and restart Log Server.
7. When Log Server has started, click **OK** to exit the utility.

Log Server is not recording data in the Log Database

Usually, when Log Server is unable to write data to the Log Database, the database has run out of allocated disk space. This can occur either when the disk drive is full, or in the case of Microsoft SQL Server, if there is a maximum size set for how large the database can grow.

If the disk drive that houses the Log Database is full, you must add disk space to the machine to restore logging.

If your SQL Server Database Administrator has set a maximum size for how large an individual database within Microsoft SQL Server can grow, do one of the following:

- ◆ Contact your SQL Server Database Administrator to increase the maximum.
- ◆ Find out the maximum size, and go to **Settings > Reporting > Log Database** to configure the Log Database to roll over when it reaches approximate 90% of the maximum size. See *Configuring database partition options*, page 362.

If your IT department has established a maximum amount of disk space for SQL Server operations, contact them for assistance.

Updating the Log Server connection account or password

To change the account or password that Log Server uses to connect to the Log Database:

1. On the Log Server machine, go to **Start > Programs > Websense > Utilities > Log Server Configuration**. The Log Server Configuration utility opens.
2. Click the **Database** tab, and then verify that correct database (by default, **wslogdb70**) appears in the ODBC Data Source Name (DSN) field.
3. Click **Connection**. The Select Data Source dialog box opens.
4. Click the **Machine Data Source** tab, and then double-click **wslogdb70** (or your Log Database name). The SQL Server Login dialog box opens.
5. If you have installed TRITON - Web Security on a Linux machine, make sure that the **Use Trusted Connection** check box is not selected.
6. Verify that a valid account name (like **sa**) appears in the **Login ID** field.
7. Enter the current password for the connection account.
8. Click **OK**, and then, in the Log Server Configuration dialog box, click **Apply**.
9. Click the **Connection** tab, and then stop and restart Log Server.
10. When Log Server is running again, click **OK** to close the utility.

Configuring user permissions for Microsoft SQL Server

Microsoft SQL Server Standard and Enterprise editions define SQL Server Agent roles that govern accessibility of the job framework. The SQL Server Agent jobs are stored in the SQL Server **msdb** database.

To install Websense Log Server successfully, the user account that owns the Websense database must:

1. Have membership in one of the following roles in the msdb database:
 - SQLAgentUserRole
 - SQLAgentReader Role
 - SQLAgentOperator Role
2. Have membership in the **db_datareader** role.
3. Be a member of the **dbcreator** fixed server role.

Use Microsoft SQL Server Management Studio to grant the database user account the necessary permissions to successfully install Log Server.

1. On the SQL Server machine, go to **Start > Programs > Microsoft SQL Server 2005 or 2008 > Microsoft SQL Server Management Studio**.
2. Select the **Object Explorer** tree, and then go to select **Security > Logins**.
3. Select the login account to be used during the installation.
4. Right-click the login account and select **Properties** for this user.
5. Select **User Mapping** and do the following:
 - a. Select **msdb** in database mapping.
 - b. Grant membership to one of these roles:
 - SQLAgentUserRole
 - SQLAgentReader Role
 - SQLAgentOperator Role
 - c. Grant membership to the **db_datareader** role.
 - d. Click **OK** to save your changes.
6. Select **Server Roles**, and then select **dbcreator**. The dbcreator role is created.
7. Click **OK** to save your changes.

Log Server cannot connect to the directory service

If either of errors listed below occurs, Log Server is unable to access the directory service, which is necessary for updating user-to-group mappings for reports. These errors appear in the Windows Event Viewer (see [The Windows Event Viewer, page 458](#)).

- ◆ EVENT ID:4096 - Unable to initialize the Directory Service. Websense Server may be down or unreachable.
- ◆ EVENT ID:4096 - Could not connect to the directory service. The groups for this user will not be resolved at this time. Please verify that this process can access the directory service.

The most common cause is that Log Server and User Service are on different sides of a firewall that is limiting access. To resolve this problem, configure the firewall to permit access over port 55815.

The default ports used for directory service communication are:

139	NetBIOS communication: Active Directory
389	LDAP communication: Active Directory, Novell eDirectory, Oracle (formerly Sun Java) Directory Server
636	SSL port: Novell eDirectory, Oracle (formerly Sun Java) Directory Server
3268	Active Directory
3269	SSL port: Active Directory

Wrong reporting page displayed

If you have deployed a V-Series appliance, the time zone settings on the TRITON - Web Security and Log Server machines must match the time zone on the appliance.

When the time zone settings are out of sync, the wrong page is displayed when administrators attempt to open the Reporting > Investigative Reports page or the Settings > Reporting > Log Database page. A login page or a “login failed” message is displayed instead of the expected functionality.

To resolve this issue, update the time zone on the TRITON - Web Security and Log Server machines to match the time zone on the appliance, then restart the off-box services.

Investigative report and presentation report issues

- ◆ [No Log Server is installed for a Policy Server, page 428](#)
- ◆ [Inadequate disk space to generate presentation reports, page 435](#)
- ◆ [Scheduled jobs in presentation reports failed, page 435](#)
- ◆ [Wrong reporting page displayed, page 434](#)
- ◆ [Bandwidth is larger than expected, page 436](#)
- ◆ [Some protocol requests are not being logged, page 436](#)
- ◆ [All reports are empty, page 436](#)
- ◆ [Microsoft Excel output is missing some report data, page 438](#)
- ◆ [Saving presentation reports output to HTML, page 438](#)
- ◆ [Error generating presentation report, or report does not display, page 439](#)
- ◆ [Investigative reports search issues, page 439](#)
- ◆ [General investigative reports issues, page 439](#)

Presentation Reports Scheduler not connected to Log Database

When a health alert warns that Presentation Reports Scheduler is disconnected from the Log Database, do **not** create any scheduled jobs in presentation reports until you resolve the problem.

Any scheduled jobs that you create in presentation reports while this connection is broken are only stored temporarily; they cannot be written to the Log Database and saved permanently. As a result, the job definitions are lost when the TRITON machine has to be restarted, or any other time the Websense TRITON - Web Security service is restarted.

Make sure that the database engine is running and any network problems have been resolved. Then, restart the Websense TRITON - Web Security service.

1. On the TRITON machine, open the Windows Services dialog box.

2. Select **Websense TRITON - Web Security** in the services list.
3. Click the Restart button in the toolbar.
4. Close the Services dialog box after the service has started.

Inadequate disk space to generate presentation reports

By default, to generate presentation reports, Websense software uses space in the following folder on the TRITON machine:

```
C:\Program Files\Websense\ReportingOutput
```

If the space available at this location falls below 1 GB, a warning message appears in the Health Alert Summary on the Status > Today page.

When this message appears, clear disk space on the appropriate disk of the TRITON machine to avoid problems generating presentation reports or other system performance problems.

Scheduled jobs in presentation reports failed

If one or more scheduled jobs cannot run successfully in presentation reports, the Health Alert Summary on the Status > Today page displays a warning message.

Scheduled jobs may fail for a variety of reasons, such as:

- ◆ Email server information has not been configured on the Settings > Reporting > Preferences page. See [Configuring reporting preferences, page 354](#), for instructions.
- ◆ There is insufficient disk space on the TRITON machine to generate presentation reports. See [Inadequate disk space to generate presentation reports, page 435](#), for more information.
- ◆ Connectivity with the Log Database has been lost. See [No Log Server is installed for a Policy Server, page 428](#), for more information.
- ◆ The configured email server is not running. Work with your system administrator to resolve the problem.

To find out which job has failed, go to the **Presentation Reports > Job Queue** page.

- ◆ If known problems have been resolved, mark the check box for the failed job, and then click **Run Now** to try the job again.
- ◆ Click the **Details** link for the failed job to display the Job History page, which gives information about recent attempts to run the selected job.

Data on Internet browse time reports is skewed

Be aware that consolidation may skew the data for Internet browse time reports. These reports show the time users spend accessing the Internet and can include details about the time spent at each site. Internet browse time is calculated using a special

algorithm, and enabling consolidation may skew the accuracy of the calculations for these reports.

Bandwidth is larger than expected

Many, but not all, Websense integrations provide bandwidth information. If your integration does not provide bandwidth information, you can configure Network Agent to perform logging so that bandwidth data is included.

When a user requests a permitted file download, the integration product or Network Agent sends the full file size, which Websense software logs as bytes received.

If the user subsequently cancels the actual download, or the file does not download completely, the bytes received value in the Log Database still represents the full file size. In these circumstances, the reported bytes received will be larger than the actual number of bytes received.

This also affects reported bandwidth values, which represent a combination of bytes received and bytes sent.

Some protocol requests are not being logged

A few protocols, such as those used by ICQ and AOL, prompt users to log into a server using one IP address, and then send a different identifying IP address and port number to the client for messaging purposes. In this case, all messages sent and received may not be monitored and logged by the Websense Network Agent, because the messaging server is not known at the time messages are exchanged.

As a result, the number of requests logged may not match the number of requests actually sent. This affects the accuracy of reports produced by Websense reporting tools.

All reports are empty

If there is no data for any of your reports, make sure that:

- ◆ The active database partitions include information for the dates included in the reports. See [Database partitions, page 436](#).
- ◆ The SQL Server Agent job is active on the Microsoft SQL Server machine. (This service is not used with SQL Server Express.) See [SQL Server Agent job, page 437](#).
- ◆ Log Server is correctly set up to receive log information from Filtering Service. See [Log Server configuration, page 437](#).

Database partitions

Websense log records are stored in partitions within the database. New partitions may be created based on size or date, depending on your database engine and configuration.

You can activate or deactivate individual partitions in TRITON - Web Security. If you attempt to generate report based on information stored in deactivated partitions, no information is found and the report is empty.

To make sure the appropriate database partitions are active:

1. Go to **Settings > Reporting > Log Database**.
2. Scroll down to the **Available Partitions** section.
3. Mark the **Enable** check box for each partition that contains data to be included on the reports.
4. Click **Save Now** to implement the change.

SQL Server Agent job

If you are using a Standard or Enterprise edition of Microsoft SQL Server, it is possible that the SQL Server Agent database job has been disabled. This job must be running for the log records to be processed into the database by the ETL database job.

1. Go to **Start > Administrative Tools > Services**.
2. Make sure that both the MSSQLSERVER and SQLSERVERAGENT services are started.
3. Make sure that the SQLSERVERAGENT service is configured for **Automatic** startup. (Double-click the service name in the Services list to open a Properties dialog box that includes **Startup type** information.)

This ensures that SQL Server Agent restarts automatically any time SQL Server or the host machine is restarted.

If you do not have access to the SQL Server machine, ask your Database Administrator to make sure the SQL Server Agent job is running, and configured for automatic startup.

Log Server configuration

Configuration settings must be correct in both TRITON - Web Security and Log Server to make sure that Log Server receives log information from Filtering Service. Otherwise, log data is never processed into the Log Database.

First, verify that TRITON - Web Security is connecting to the Log Server successfully.

1. Log on to TRITON - Web Security with unconditional Super Administrator permissions.
2. Go to **Settings > General > Logging**.
3. Enter the IP address or host name for the Log Server machine.
4. Enter the port that Log Server is listening on (the default is 55805).
5. Click **Check Status** to determine whether TRITON - Web Security is able to communicate with the specified Log Server.

A message indicates whether the connection test passed. Update the IP address or machine name and port, if needed, until the test is successful.

6. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Next, verify the settings in the Log Server Configuration utility.

1. On the machine where Log Server is running, go to **Start > Programs > Websense > Utilities > Log Server Configuration**.
2. On the **Connections** tab, verify that the Port matches the value entered in TRITON - Web Security.
3. Click **OK** to save any changes.
4. Use the button on the **Connections** tab to stop and then start Log Server.
5. Click **Quit** to close the Log Server Configuration utility.

Microsoft Excel output is missing some report data

The largest number of rows that can be opened in a Microsoft Excel worksheet is 65,536. If you export a report with more than 65,536 records to Microsoft Excel format, the 65,537th and all following records are not available in the worksheet.

To assure access to all information in the exported report, do one of the following:

- For presentation reports, edit the report filter to define a smaller report, perhaps by setting a shorter date range, selecting fewer users and groups, or selecting fewer actions.
- For investigative reports, drill down into the data to define a smaller report.
- Select a different export format.

Saving presentation reports output to HTML

If you generate a report directly from the Reporting > Presentation Reports page, you can choose from 3 display formats: HTML, PDF, and XLS. If you choose the HTML display format, you can view the report in the TRITON - Web Security window.

Printing and saving presentation reports from the browser are not recommended. The printed output includes the entire browser window, and opening a saved file launches TRITON - Web Security.

To print or save reports more effectively, choose PDF or XLS as the output format. You can open these file types immediately if the viewing software (Adobe Reader or Microsoft Excel) is installed on the local machine. You also can save the file to disk (the only option if the proper viewing software is not available).

After you open a report in Adobe Reader or Microsoft Excel, use that program's print and save options to produce the desired final output.

Error generating presentation report, or report does not display

Presentations reports offers 2 options for running a report immediately: schedule the report to run in the background (default) or run the report without scheduling (if you deselect the default option).

If you run the report without scheduling (in the foreground), a pop-up window indicates the status of the report, and then the report is displayed. In some cases, instead of displaying a completed report:

- ◆ The message “error generating report” is displayed.
- ◆ A “report complete” message is displayed, but no report is shown.

This typically occurs after an administrator closes the pop-up window displaying the “generating report” and “report complete” messages the browser’s close (X) button. To avoid this issue, instead use the embedded **Close** button to close the window.

If you encounter this issue, navigate away from the Presentation Reports page in TRITON - Web Security, and then run the report again. If that does not work, log out of TRITON - Web Security and log back on before running the report again.

If the problem persists, you can:

- ◆ Use the **Schedule the report to run in the background** option and open reports from the Review Reports page.
- ◆ Use Firefox, rather than Internet Explorer, when generating reports.

Investigative reports search issues

The Search fields above the bar chart on the main Investigative Reports page allow searches for a specific term or text string in the selected chart element. There are two potential concerns related to searching investigative reports: extended ASCII characters and search pattern matching.

- ◆ If you are using Mozilla Firefox on a Linux machine to access TRITON - Web Security, you cannot enter extended ASCII characters in the Search fields. This is a known limitation of Firefox on Linux.

If you need to search an investigative report for a text string that includes extended ASCII characters, access TRITON - Web Security from a Windows machine, using any supported browser.

- ◆ Sometimes, investigative reports is unable to find URLs associated with a pattern entered in the Search fields on the main investigative reports page. If this occurs, and you are reasonably certain that the pattern exists within the URLs reported, try entering a different pattern that would also find the URLs of interest.

General investigative reports issues

- ◆ Some queries take a very long time. You may see a blank screen or get a message saying that your query has timed out. This can happen for the following reasons:

- Web server times out
- Microsoft SQL Server times out
- Proxy or caching server times out

You may need to manually increase the timeout limit for these components.

- ◆ If users are not in any group, they will not show up in a domain either. Both Group and Domain choices will be inactive.
- ◆ Even if the Log Server is logging visits instead of hits, investigative reports label this information as **Hits**.

Other reporting issues

- ◆ [Low memory on the Real-Time Monitor machine, page 440](#)
- ◆ [Real-Time Monitor is not running, page 440](#)
- ◆ [Real-Time Monitor is not responding, page 441](#)
- ◆ [Cannot access certain reporting features, page 442](#)
- ◆ [No charts appear on Today or History pages, page 442](#)

Low memory on the Real-Time Monitor machine

This alert is displayed when available memory on the Real-Time Monitor machine is at 15% or less of total memory. Low memory can prevent Real-Time Monitor from receiving, displaying, and storing some or all filtering records.

This can result in gaps in the data displayed in the monitor, or prevent the monitor's server and database components from running at all.

Use the Windows Task Manager to evaluate memory usage on the Real-Time Monitor machine. To solve the problem, you can:

- ◆ Upgrade the RAM on the machine.
- ◆ Move applications or components with high memory requirements to another machine.

Move Real-Time Monitor to a machine with more available memory.

Real-Time Monitor is not running

This alert is displayed when the Websense RTM Server service is stopped.

Use the Windows Services dialog box to verify that all 3 Real-Time Monitor services are started, and to start any of the following services that have stopped:

- ◆ Websense RTM Database
- ◆ Websense RTM Server
- ◆ Websense RTM Client

If any service will not start:

- ◆ Check the Windows Event Viewer for any errors or warnings from Websense RTM Server.
- ◆ Check the **WebsenseRTMMemoryOutput0.log** file (located, by default, in the C:\Program Files (x86)\Websense\Web Security\rtm\logs directory) for information about Real-Time monitor memory usage.
- ◆ Make sure that there are sufficient resources (memory, hard disk, and CPU) available for the services to run.

If the service is running, but the alert continues to appear, this may indicate that Real-Time Monitor was unable to register with Policy Server. Verify that the Policy Server associated with this Real-Time Monitor instance is running, and that the Real-Time Monitor machine can communicate with the Policy Server machine on port 55836 (encrypted communication) or 55856 (non-encrypted communication).

If the services do start, make sure that they are configured for **Automatic** (not Manual) startup.

Real-Time Monitor is not responding

If Real-Time Monitor is installed on a different machine from the TRITON Unified Security Center, use the ping command to make sure that the 2 machines can communicate across the network. Also verify that the TRITON machine can communicate with Real-Time Monitor on port 9445 (for user interface display).

In addition, Real-Time Monitor must be able to communicate with:

- ◆ Usage Monitor on port 55835
- ◆ Policy Server on port 55836 (encrypted communication) or 55856 (non-encrypted communication)

If there is not a network communication problem, Real-Time Monitor may be experiencing resource constraints.

- ◆ Check memory, CPU usage, and available disk space on the Real-Time Monitor machine.

Note that the RTM Database can hold a maximum of 10,000 records, which should help to limit its impact on available disk space.

- ◆ The database may be receiving too many requests, or be unable to accept additional connections.

If the Windows Event Viewer shows Websense RTM Database errors, you can restart the service to address the problem.

Note that when the database is restarted, all records are cleared, so older data is lost. Data that is not available for display in Real-Time Monitor is still stored in the Log Database, and can be seen in investigative and presentation reports.

Cannot access certain reporting features

If your Web browser has pop-up blocking at a very strict setting, it may block certain reporting features. To use those features, you must decrease the blocking level or disable pop-up blocking entirely.

No charts appear on Today or History pages

In organizations that use delegated administration, review the reporting permissions for the delegated administrator's role. If **View reports on Today and History pages** is not selected, these chart do not appear for delegated administrators in that role.

In environments that use multiple Policy Servers, TRITON - Web Security only shows reporting data when connected to the Policy Server that is also configured to communicate with Log Server. You must log on to that Policy Server to view charts on the Today and History pages, or to access other reporting features.

If you have multiple Log Server instances, there are special considerations for also deploying multiple TRITON - Web Security instances. In these distributed logging environments, it is important that only one TRITON - Web Security instance be used for reporting. Administrators connecting to the reporting instance of TRITON - Web Security will see all reporting features (including the Today and History page charts). Administrators connecting to other TRITON - Web Security instances will not see reporting features.

Interoperability issues

- ◆ [Content Gateway is not running, page 443](#)
- ◆ [Content Gateway is not available, page 443](#)
- ◆ [Content Gateway non-critical alerts, page 443](#)
- ◆ [Administrator unable to access other TRITON modules, page 446](#)
- ◆ [Sync Service is not available, page 446](#)
- ◆ [Sync Service has been unable to download log files, page 447](#)
- ◆ [Sync Service has been unable to send data to Log Server, page 448](#)
- ◆ [Hybrid filtering data does not appear in reports, page 448](#)
- ◆ [Disk space is low on the Sync Service machine, page 449](#)
- ◆ [The Sync Service configuration file, page 449](#)
- ◆ [Directory Agent is not running, page 450](#)
- ◆ [Directory Agent cannot connect to the domain controller, page 450](#)
- ◆ [Directory Agent does not support this directory service, page 452](#)
- ◆ [The Directory Agent configuration file, page 452](#)
- ◆ [Directory Agent command-line parameters, page 454](#)
- ◆ [Alerts were received from the hybrid service, page 454](#)

- ◆ [Unable to connect to the hybrid service, page 455](#)
- ◆ [Hybrid service unable to authenticate connection, page 455](#)
- ◆ [Missing key hybrid configuration information, page 456](#)
- ◆ [User access no longer filtered following upgrade, page 456](#)

Content Gateway is not running

When a Content Gateway instance registers with Policy Server, that connection is tracked in TRITON - Web Security. Information about the Content Gateway instance appears on the Settings > General > Content Gateway Access page, and in the Filtering Service Summary on the Status > Today page.

If the registered instance stops, or is removed, a health alert message is displayed in TRITON - Web Security.

- ◆ If the instance has stopped unexpectedly, check the **syslog** file on the Content Gateway machine for information about what caused the failure.
- ◆ If you have relocated Content Gateway to another IP address or physical machine, or if you have removed an instance that was not needed, you can manually remove the instance from the Settings > General > Content Gateway Access page to stop the health alert from being displayed.

Content Gateway is not available

When a Content Gateway instance registers with Policy Server, that connection is tracked in TRITON - Web Security. Information about the Content Gateway instance appears on the Settings > General > Content Gateway Access page, and in the Filtering Service Summary on the Status > Today page.

If Web Security components can no longer communicate with the registered instance, a health alert message is displayed in TRITON - Web Security.

- ◆ Make sure that the Content Gateway machine is up, and that Content Gateway is running.
- ◆ This alert may indicate a network problem. Verify that Content Gateway can communicate with the Policy Server (ports 55806 and 55880) and Filtering Service (port 15868) machines.

Content Gateway non-critical alerts

When you receive notification that non-critical alerts have been received from a Content Gateway instance, any of the following errors or conditions may have occurred. To determine which error occurred, check the Content Gateway Manager associated with the affected Content Gateway instance.

Use the table below to get an overview of the error condition. More detailed information can be found in the system, error, and event log files on the Content Gateway machine.

Alert	Description
Content Gateway process reset	<p>A problem that caused Content Gateway to restart.</p> <p>See the Content Gateway syslog file for information about what caused the reset.</p>
Cache configuration issue	<p>Content Gateway was unable to configure a cache.</p> <p>See “Configuring the Cache” in the Content Gateway Manager Help for more information.</p>
Unable to create cache partition	<p>An error occurred during cache configuration.</p> <p>See “Configuring the Cache” in the Content Gateway Manager Help.</p>
Unable to initialize cache	<p>A cache failure occurred.</p> <p>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.</p> <p>See “Configuring the Cache” in the Content Gateway Manager Help.</p>
Unable to open configuration file	<p>There is a problem in a configuration file.</p> <ul style="list-style-type: none"> • Check the system log for information about which file is affected. • Permissions to the file or directory may have changed. • If the file was edited outside Content Gateway Manager, there may be invalid syntax or other problems preventing the file from being read.
Invalid fields in configuration file	<p>One or more parameters or parameter values in a configuration file is incorrect.</p> <p>Check the system log for information about which file is affected.</p>
Unable to update configuration file	<p>There is a problem preventing a configuration file from being saved.</p> <p>Check the system log for information about which file is affected.</p>
Clustering peer operating system mismatch	<p>The nodes in a cluster must be homogeneous, with the same:</p> <ul style="list-style-type: none"> • Hardware platform • Operating system version

Alert	Description
Could not enable virtual IP addressing	<p>Content Gateway attempted to enable virtual IP address failover, but failed.</p> <p>This often occurs when the designated virtual IP address is already in use in the network.</p> <p>Like all IP addresses, virtual IP addresses must be pre-reserved before they can be assigned to Content Gateway.</p>
Connection throttle too high	<p>A connection throttle event occurs when client or origin server connections reach 90% of half the configured connection limit (45000 by default).</p> <p>When you raise the connection throttle limit, the system must have adequate memory to handle the client connections required. A system with limited RAM might need a throttle limit lower than the default value.</p>
Host database disabled	<p>The host database stores the Domain Name Server (DNS) entries of origin servers to which the proxy connects. It tracks:</p> <ul style="list-style-type: none"> • 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)
Logging configuration error	<p>Content Gateway can be configured to log transactions, errors, or both to a location that you specify.</p> <p>See “Working with Log Files” in the Content Gateway Manager Help for information about logging.</p>
Unable to open Content Gateway Manager	<p>Content Gateway is unable to set up a socket to handle management API calls to start the Web interface.</p>
ICMP echo failed for a default gateway	<p>A Content Gateway node failed to contact its default gateway while assigning virtual IP addresses for a cluster. The node will shut down.</p>

Alert	Description
HTTP origin server is congested	<p>When Content Gateway is deployed as a Web proxy cache, user requests for Web content pass through Content Gateway on the way to the destination Web server (origin server).</p> <p>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.</p> <p>If the origin server is congested (unable to accept additional connections), and does not respond to the revalidation query, the proxy does not perform any validation; it serves the stale object from the cache.</p>
Congestion alleviated on the HTTP origin server	<p>An origin server that previously denied connection attempts is now accepting requests again.</p>
Content scanning skipped	<p>Content Gateway did not scan content for a requested site that would have ordinarily be scanned.</p> <p>This may occur when Content Gateway is experiencing too many connections, or inadequate system resources (CPU and memory).</p>
WCCP configuration error	<p>See the “WCCP Configuration” section of the Content Gateway Manager Help for configuration parameter details.</p>

Administrator unable to access other TRITON modules

If you receive an error when you click **Data Security** or **Email Security** in TRITON - Web Security, the local or network account that you use to log on to TRITON may not have been granted Data Security or Email Security access permissions.

A Global Security Administrator must give an administrator access to each module on the TRITON Settings > Administrators page before an administrator can switch between TRITON modules.

The default TRITON Unified Security Center administrator account, **admin**, has full access to all installed modules.

See the TRITON Unified Security Center Help (which can be opened from the Help menu when you are on any TRITON Settings page) for more information.

Sync Service is not available

In Websense Web Security Gateway Anywhere deployments, Websense Sync Service is responsible for communication between the on-premises and hybrid services. Sync Service:

- ◆ Sends policy configuration data to the hybrid service
- ◆ Sends user information collected by Directory Agent to the hybrid service
- ◆ Receives reporting log records from the hybrid service

If you have not yet activated hybrid filtering, or if you have attempted to activate hybrid filtering, but have not been able to do so, note that your local Websense software components must be able to communicate with Sync Service before the connection to the hybrid service can be created.

To troubleshoot this issue, make sure that:

- ◆ Sync Service is running.
- ◆ Sync Service is successfully binding to the correct IP address and port.
 - The IP address and port that Sync Service is attempting to use are listed in the **syncservice.ini** file, located in the Websense **bin** directory on the Sync Service machine.
 - The IP address and port shown on the Settings > Hybrid Configuration > Shared User Data page in TRITON - Web Security must match those listed in the **syncservice.ini** file. If you update the configuration file, also manually update the Settings page.
 - The IP address and port in the **syncservice.ini** file must match the Sync Service IP address and port values in the **das.ini** file (located in the Websense **bin** directory on the Directory Agent machine).

Verify that no other service on the Sync Service machine is binding to the IP address and port that Sync Service is attempting to use. If you suspect that Sync Service is unable to bind to the correct IP address and port, stop the service, open a command prompt, and try to start the service in console mode:

```
syncservice -c
```

In console mode, Sync Service displays the IP address and port that it is using, or displays an error, if it is unable to bind to the IP address and port.

- ◆ The Sync Service machine can communicate with the Policy Broker machine on port 55880.
- ◆ The Sync Service machine can connect to the Policy Server machine on ports 55806 and 40000, and receive data from Policy Server on ports 55830 and 55831.
- ◆ The TRITON - Web Security machine can create an HTTP connection to the Sync Service machine on port 55832.

Also check the Windows Event Viewer or **websense.log** file for errors from Sync Service.

Sync Service has been unable to download log files

Sync Service attempts to connect to the hybrid service to download reporting log files at an interval that you configure (see [Schedule communication with hybrid filtering, page 196](#)). If Sync Service is unable to make the connection, or if Sync Service is unable to retrieve the log files after connecting, the following problems may occur:

- ◆ The hybrid service stores log files for only 14 days. After that period, the files are deleted, and cannot be recovered. When this occurs, your organization is no longer able to report on hybrid filtering activity recorded in those logs.
- ◆ Depending on the volume of Internet activity that your organization sends through the hybrid service, reporting log files may grow quickly. If Sync Service is unable to download log files for a day or more, the bandwidth required to download the files and the disk space required to temporarily store them may be substantial.

To address this issue, check the Status > Today > Hybrid Service Status page to verify that Sync Service is able to connect to hybrid filtering. See [Unable to connect to the hybrid service](#), page 455, for more troubleshooting steps.

If Sync Service is connecting to the hybrid service, but cannot retrieve log records, check the Status > Alerts page for information from the hybrid service. Also check the administrative email address associated with your hybrid filtering account.

Sync Service has been unable to send data to Log Server

After Sync Service downloads reporting log files from the hybrid service, it passes the files to Log Server so that they can be processed into the Log Database and included in reports. If Sync Service cannot pass the data to Log Server, log files may accumulate on the Sync Service machine, consuming potentially large amounts of disk space.

- ◆ Use the **telnet** command to verify that it is possible for the Sync Service machine to connect to the Log Server machine on port **55885**.
- ◆ Make sure that Log Server is running, and that no Log Server errors appear on the Status > Alerts page.

Hybrid filtering data does not appear in reports

If Internet activity information for users filtered by the hybrid service does not appear in reports, first make sure that:

- ◆ A hybrid logging port is configured on the Settings > General > Logging page. See [Configuring how filtered requests are logged](#), page 355.
- ◆ The **Have the hybrid service collect reporting data for the clients it filters** check box is selected on the Settings > Hybrid Configuration > Scheduling page. See [Schedule communication with hybrid filtering](#), page 196.
- ◆ The Status > Today > Hybrid Service Status page shows that Sync Service has successfully connected to the hybrid service, and retrieved log records. See [Monitor communication with the hybrid service](#), page 197.
- ◆ No health alerts appear on the Status > Today page indicating Sync Service communication problems or Log Server errors. See [Sync Service has been unable to send data to Log Server](#), page 448.

If your deployment uses distributed logging, in which multiple, remote Log Servers send data to a centralized Log Server instance, also make sure that Sync Service is

configured to communicate with the central Log Server. Hybrid logging data cannot be passed to the central Log Server by remote Log Server instances.

Disk space is low on the Sync Service machine

If Sync Service is unable to pass reporting log files collected by the hybrid service to Log Server in a timely manner, log files may accumulate on the Sync Service machine, consuming large amounts of disk space. To avoid this issue:

- ◆ Make sure that Sync Service is collecting reporting log data from the hybrid service at appropriate intervals. The more Internet activity your organization sends through the hybrid service, the more frequently log files should be downloaded to avoid large backlogs.
- ◆ Make sure that the Sync Service machine is able to connect to the Log Server machine on port **55885**.
- ◆ Allocate sufficient resources on the Sync Service machine for the volume of reporting data being processed.

The Sync Service configuration file

Use the **syncservice.ini** file to configure aspects of Sync Service behavior that cannot be configured in TRITON - Web Security.

The **syncservice.ini** file is located in the Websense **bin** directory (*Where is the Websense "bin" directory?*, page 457).

- ◆ Use a text editor to edit the file.
- ◆ When you are finished making changes, save and close the file, and then restart Sync Service. Changes do not take effect until the service has restarted.

The file contains the following information:

- ◆ **SyncServiceHTTPAddress**: The IP address that Sync Service binds to for communication with Directory Agent and TRITON - Web Security. Must match the Sync Service IP address on the Settings > Hybrid Configuration > Shared User Data page in TRITON - Web Security.
- ◆ **SyncServiceHTTPPort**: The port that Sync Service listens on for communication from Directory Agent and TRITON - Web Security (default 55832). Must match the Sync Service port displayed on the Settings > Hybrid Configuration > Shared User Data page in TRITON - Web Security.
- ◆ **UseSyncServiceProxy**: Indicates whether Sync Service goes through a proxy to connect to the hybrid service. Values are **true** or **false**.
 - **SyncServiceProxyAddress**: The IP address of the proxy through which Sync Service connects to the hybrid service.
 - **SyncServiceProxyPort**: The port of the proxy through which Sync Service connects to the hosted service.
 - **SyncServiceProxyUsername**: The user name (if required) that Sync Service uses to connect to the proxy in order to contact the hybrid service.

- **SyncServiceProxyPassword:** The password (if required) that Sync Service uses to connect to the proxy in order to contact the hybrid service.

Directory Agent is not running

In Websense Web Security Gateway Anywhere deployments, Websense Directory Agent gathers user information from your directory service and sends it to the hybrid service for use in applying filtering policies.

When Directory Agent is not available, the hybrid service's user data may become outdated.

Make sure that Directory Agent is installed, and that the service or daemon is running.

- ◆ Windows: Use the Windows Services dialog box to start the service.
- ◆ Linux: Use the `/opt/Websense/WebsenseDaemonControl` command to start the daemon.

If Directory Agent is running, but the alert message continues to appear, verify that:

- ◆ The Directory Agent machine can communicate with the Policy Server machine (ports 40000 and 55806).
- ◆ The Directory Agent machine can communicate with the Sync Service machine (port 55832).
- ◆ The firewall permits communication on the Directory Agent port (55900).

If the service starts, but does not continue to run:

- ◆ Check the Event Viewer (Windows) or `websense.log` file (Linux) for errors.
- ◆ Navigate to the Websense **bin** directory ([Where is the Websense "bin" directory?](#), page 457) and verify that the **das.ini** file exists, and that it has not been corrupted or truncated.
- ◆ Make sure that there is enough disk space on the Directory Agent machine to store a full snapshot of your directory. For example, a snapshot of a 200,000 user directory requires about 100 MB of disk space.
- ◆ Make sure that there is enough available memory for Directory Agent to compare its current snapshot with the previous one. For example, comparing snapshots of a 200,000 user directory requires about 100 MB of memory.

Directory Agent cannot connect to the domain controller

Directory Agent must be able to connect to the domain controller to gather user information from the directory service. If there are communication problems between the Directory Agent machine and the domain controller, the hybrid service's user data may become outdated, leading to incorrect filtering.

To troubleshoot this problem:

- ◆ Make sure that the Directory Agent machine is bound to the domain, and that the firewall permits communication on the directory service port.

Port	Used for:
139	NetBIOS communication: Active Directory
389	LDAP communication: Active Directory, Novell eDirectory, Oracle (formerly Sun Java) Directory Server
636	SSL port: Novell eDirectory, Oracle (formerly Sun Java) Directory Server
3268	Active Directory
3269	SSL port: Active Directory

- ◆ Go to the Settings > General > Directory Services page and verify that your directory service configuration has not changed since you last updated your Directory Agent settings.
- ◆ Go to the Settings > Hybrid Configuration > Shared User Data page and verify that Directory Agent is attempting to search a valid context (path) for user and group information. To do this:
 - If you are using Windows Active Directory, click a directory server name or IP address, and then click Test Context. Repeat this process for each global catalog server.
 - If you are using Oracle (formerly Sun Java) Directory Server or Novell eDirectory, click Test Context.
- ◆ On the Shared User Data page, also make sure that the context is not only valid, but appropriate. The context should be limited to include only those users and groups filtered by the hybrid service.
- ◆ Still on the Shared User Data page, make sure that the Directory Search option is set correctly, so that Directory Agent is searching only the relevant portion of your directory service.
- ◆ Verify that it is possible to connect to the directory service IP address and port from the Directory Agent machine.

Directory Agent communication issues

If Directory Agent is prevented from communicating with directory service to gather user information, or if Directory Agent cannot connect to Sync Service, updated user and group information cannot be sent to the hybrid service.

Communication problems can occur if:

- ◆ There is problem in the network.

- ◆ The ports used for directory service (see table) or Sync Service (55832) communication are blocked between the Directory Agent machine and the target machine.

Port	Used for:
139	NetBIOS communication: Active Directory
389	LDAP communication: Active Directory, Novell eDirectory, Oracle (formerly Sun Java) Directory Server
636	SSL port: Novell eDirectory, Oracle (formerly Sun Java) Directory Server
3268	Active Directory
3269	SSL port: Active Directory

- ◆ Directory Agent is using incorrect credentials, or the target service is unable to authenticate the connection.
- ◆ A service is not available, because of a service restart or a machine reboot, for example.

To determine what is causing the communication problem, consult the Windows Event Viewer or **websense.log** file for detailed information.

Directory Agent does not support this directory service

Directory Agent is only able to retrieve user and group information from LDAP-based directory services. Windows Active Directory (Mixed Mode) is not supported. The supported directory services include:

- ◆ Windows Active Directory (Native Mode)
- ◆ Sun Java System Directory
- ◆ Novell eDirectory
- ◆ If you are not using a supported directory service, hybrid filtering can still be applied to filtered locations. User and group-based filtering, however, cannot be performed.

The Directory Agent configuration file

Use the **das.ini** file to configure aspects of Directory Agent behavior that cannot be configured in TRITON - Web Security. These include the maximum memory the agent can use, the maximum threads it can create, the directory where it should store user information snapshots, and more.

The **das.ini** file is located in the Websense **bin** directory (*Where is the Websense “bin” directory?*, page 457).

- ◆ Use a text editor to edit the file.
- ◆ For parameters that can take multiple values, use the pipe symbol (“|”) to separate entries.

- ◆ When you are finished making changes, save and close the file, and then restart the Directory Agent service or daemon. Changes do not take effect until the service has restarted.

Key values that can be configured in the file include:

- ◆ The maximum amount of memory that Directory Agent can use, in megabytes (MB). If Directory Agent is configured to collect a very large number of directory entries (more than 200,000 user or group definitions), you may need to increase this number.

```
MaxMemory=100
```

- ◆ The full directory path showing where Directory Agent stores directory service snapshots (complete views of the directory, used to determine what has changed between one query and the next).

```
SnapshotDir=./snapshots/
```

This relative path translates to C:\Program Files\WebSense\bin\snapshots (Windows) or /opt/WebSense/bin/snapshots/ (Linux).

- ◆ The full directory path showing where Directory Agent stores the LDIF files that Sync Service sends to hybrid filtering.

```
DiffDir=./diffs/
```

- ◆ The regular expression Directory Agent uses to validate email addresses in LDAP records. Records whose email addresses do not match the pattern are dropped. For example, `[a-z0-9!#$%&'*/+=?^_`{|}~-]+(?:\.[a-z0-9!#$%&'*/+=?^_`{|}~-]+)*@(?:[a-z0-9](?:[a-z0-9-]*[a-z0-9])?\.)+[a-z0-9](?:[a-z0-9-]*[a-z0-9])?`

Leave the parameter blank (default) if you do not want Directory Agent to perform email address validation.

```
EmailValidateRegex=
```

- ◆ The number of times Directory Agent retries after a failed attempt to connect to Sync Service. Takes an integer value between 1 and 65535.

```
SyncServiceRetryCount=5
```

- ◆ The number of seconds Directory Agent waits between retry attempts when establishing a connection to Sync Service. Takes an integer value between 1 and 65535.

```
SyncServiceRetryDelay=60
```

- ◆ The number of times Directory Agent retries after a failed attempt to connect to the directory service. Takes an integer value between 1 and 65535.

```
DirServiceRetryCount=5
```

- ◆ The number of seconds Directory Agent waits between retry attempts when establishing a connection to the directory service. Takes an integer value between 1 and 65535.

```
DirServiceRetryDelay=60
```

- ◆ The number of seconds the Directory Agent backup subsystem waits between attempts to reconnect to Sync Service. The backup subsystem is responsible for verifying that user data is successfully received by Sync Service and sent to the

hybrid service. In the event of a failure, the backup subsystem makes sure that the LDIF file that could not be sent is preserved for a later retry attempt.

Takes an integer value between 1 and 65535.

`BackupPollPeriod=60`

- ◆ The number of times the Directory Agent backup subsystem attempts to reconnect to Sync Service to determine the status of the last transaction. Takes an integer value between 1 and 65535.

`BackupRetryCount=60`

Directory Agent command-line parameters

Directory Agent has a command-line interface that you can use to install, uninstall, start, and stop the agent if necessary. You can also print version and usage information about the agent.

To start Directory Agent in console mode (as an application), open a command prompt and navigate to the Websense **bin** directory ([Where is the Websense “bin” directory?](#), [page 457](#)) and enter the following:

```
DAS.exe -c
```

Directory Agent accepts the following command-line parameters. Note that some parameters can only be used in Microsoft Windows environments.

Parameter	Description
-i	Installs Directory Agent service. Registers itself with the operating system (Windows only).
-u	Uninstalls Directory Agent service. (Windows only).
-c	Runs Directory Agent in console mode.
-r	Runs Directory Agent as daemon or service.
-s	Stops the Directory Agent service. (Windows only).
-v	Prints version information about the Directory Agent service.
-h -? -help <no option>	Prints usage information about the Directory Agent service.

Alerts were received from the hybrid service

When the hybrid service encounters a problem that could affect your organization, it sends an alert to your installation of Sync Service. Alerts are sent for issues that affect either hybrid filtering as a whole, or that are specific to your account. When the alert is received:

- ◆ A general alert is displayed under Health Alerts on the Status > Today page in TRITON - Web Security.

- ◆ A more specific alert is shown on the Status > Alerts page under Hybrid Filtering Alerts.

If there are steps that you can take to correct the problem (for example, prompting Directory Agent to re-send user information, or clicking Save All to prompt Sync Service to re-send policy information), that information is included in the detailed alert message on the Status > Alerts page.

In many cases, alerts from hybrid filtering are informational, making sure that you are aware that a temporary issue may be preventing user or policy information from being received, or reporting data from being sent. No action on your part is required to address such issues.

When the condition causing the problem has been resolved, both the Health Alert on the Status > Today page and the alerts on the Status > Alerts page are cleared.

Unable to connect to the hybrid service

The on-premises and hybrid portions of your Websense Web Security Gateway Anywhere solution must communicate regularly to ensure consistent filtering and accurate reporting.

Sync Service may be prevented from accessing the hybrid service due to network problems, either affecting Internet or internal network connections.

- ◆ Use a browser or the **ping** utility to verify that the Sync Service machine can connect to the Internet.
- ◆ Make sure that an HTTPS connection to the Internet can be established from the Sync Service machine. Sync Service uses port 443 to connect to the hybrid service.
- ◆ Make sure that Sync Service can communicate with other on-premises components in the network via ports 55830 and 55831.

Also verify that there is not a problem preventing the hybrid service from accepting the Sync Service connection.

- ◆ Check the Hybrid Filtering Alerts table on the Status > Alerts page for information from the hybrid service.
- ◆ Make sure that administrators have been monitoring the email account provided as a contact address on the Settings > General > Account page for messages from Websense Technical Support.

Hybrid service unable to authenticate connection

In hybrid filtering environments, Sync Service provides an account identifier each time it connects to the hybrid service to send or retrieve information. This identifier is unique to your organization, and updated each time the **admin** password changes.

Under rare circumstances, possibly involving a serious problem with the Policy Database, the connection between your on-premises software and the hybrid service may be lost. In these cases, you must request a security token, used to generate a new

identifier for your hybrid filtering account. The security token is sent to the **contact email address** specified on the Settings > General > Account page in TRITON - Web Security.

To request a new token:

1. Click the **Get Token** button that appears next to the “unable to authenticate connection” alert on the Status > Alerts page in TRITON - Web Security.
2. Verify that you receive a success message stating that the request has been sent to the hybrid service.
3. Monitor the administrative email account associated with your hybrid filtering account. It may take some time for the request for a new security token to be processed.
4. When you receive an email message from the hybrid service, go to the Settings > General > Account page in TRITON - Web Security.
5. Scroll down to the Hybrid Filtering section of the page and enter the **Security token** provided in the email message,
6. Click **Connect**.

The temporary token is verified and used to resume communication between Sync Service and the hybrid service.

Missing key hybrid configuration information

In hybrid filtering environments, Sync Service provides an account identifier each time it connects to the hybrid service to send or retrieve information. This identifier is unique to your organization, and updated each time the **admin** password changes.

Under rare circumstances, possibly involving a serious problem with the Policy Database, the connection between your on-premises software and the hybrid service may be lost. In these cases, you must request a security token, used to generate a new identifier for your hybrid filtering account. The security token is sent to the **contact email address** specified on the Settings > General > Account page.

If you receive the alert message, “Missing configuration information; connection to hybrid filtering lost,” either no contact email address has been provided, or the contact email address is no longer valid.

In this case, in order to maximize the security of your organization’s private data, you must contact [Websense Technical Support](#) directly to update your hybrid filtering account.

User access no longer filtered following upgrade

Due to changes in off-site user configuration for Websense Web Security Gateway Anywhere version 7.6, following upgrade from version 7.5:

- ◆ The previous method of configuring hybrid filtering for off-site users (by entering a host name that can only be resolved outside the network) is no longer available.

- ◆ All filtered locations are configured to use the hybrid proxy
- ◆ The **Enable hybrid filtering of off-site users** check box on the **Settings > Hybrid Configuration > User Access** page, which is new for version 7.6, is not selected.

As a result, all users from unknown IP addresses will not be filtered. This includes off-site users as well as any on-premises users whose browsers are configured to use the PAC file supplied by the hybrid service.

If you had previously enabled filtering for off-site users using a transparent proxy for in-network machines, this issue affects your off-site users only and you receive the alert message “Off-site user access no longer filtered following upgrade.” In this case, go to the **Settings > Hybrid Configuration > User Access** page and mark **Enable hybrid filtering of off-site users**.

If you had previously enabled filtering for off-site users using an explicit proxy for in-network machines, this issue affects both off-site and on-site users and you receive the alert message “Off-site and on-site user access no longer filtered following upgrade.” Go to the **Settings > Hybrid Configuration > User Access** page and mark **Enable hybrid filtering of off-site users**. You can also set up your filtered locations to use a transparent, explicit, or hybrid proxy.

Troubleshooting tips and tools

- ◆ [The Windows Services dialog box, page 458](#)
- ◆ [The Windows Event Viewer, page 458](#)
- ◆ [The Websense log file, page 458](#)

Where is the Websense “bin” directory?

Many Websense Web Security executable and configuration files are installed in the Websense **bin** directory. When a troubleshooting step prompts you to navigate to this directory, the location depends on the operating system and installed components.

Linux platforms

For all supported Linux operating systems, the default path is:

```
/opt/Websense/bin/
```

Windows platforms

On Windows systems that meet either of the following descriptions:

- ◆ No TRITON console components, and only **upgraded** policy, filtering, or user identification components
- ◆ Upgraded TRITON - Web Security (formerly Websense Manager) components
 - No Data Security or Email Security modules installed

- Does not include a fresh installation of TRITON - Web Security that has imported configuration from a previous version

The default path is:

```
C:\Program Files\WebSense\bin
```

On Windows systems that have a new installation of Websense Web Security components, or that have multiple TRITON modules installed (Web Security with Data Security and/or Email Security), the default path is either:

```
C:\Program Files\WebSense\Web Security\bin
```

```
C:\Program Files (x86)\WebSense\Web Security\bin
```

The Windows Services dialog box

On Microsoft Windows machines, Filtering Service, Network Agent, Policy Server, User Service, and all Websense transparent identification agents run as services. You can use the Windows Services dialog box to check the status of these services.

1. In the Windows Control Panel, open the **Administrative Tools** folder.
2. Double-click **Services**.
3. Scroll through the list of services to find the service you are troubleshooting.
The service entry includes the service name, a brief service description, the service status (started or stopped), how the service starts, and what account the service uses to perform its tasks.
4. Double-click a service name to open a properties dialog box with more detailed information about the service.

The Windows Event Viewer

The Windows Event Viewer records error messages about Windows events, including service activities. These messages can help you identify network or service errors that may be causing Internet filtering or user identification problems.

1. In the Windows Control panel, open the **Administrative Tools** folder.
2. Double-click **Event Viewer**.
3. In the Event Viewer, click **Application** for a list of error messages, warnings, and informational messages.
4. Scroll through the list to identify errors or warnings from Websense services.

The Websense log file

Websense software writes error messages to the **websense.log** file, located in the Websense **bin** directory (see [Where is the Websense “bin” directory?](#), page 457).

The information in this file is comparable to that found in the Windows Event Viewer. In Windows environments, the Event Viewer presents messages in a more user-

friendly format. The **websense.log** file, however, is available on Linux systems, and can be sent to Websense Technical Support if you need help troubleshooting a problem.

Index

A

- accessing TRITON - Web Security, 19, 295
- account information
 - configuring, 32
 - hybrid filtering, 174
- account override, 74
 - in multiple Policy Server environment, 322
- actions, 46
 - Block, 46
 - block based on bandwidth, 47
 - Block File Types, 47
 - Block Keywords, 47
 - Confirm, 47
 - Permit, 46
 - Quota, 47
 - selecting for presentation reports, 113
- active content
 - removing, 164
- Active Directory
 - hybrid configuration, 189
 - Native Mode, 66
- ActiveX content
 - removing, 164
- Add
 - category filter, 50
 - custom LDAP groups, 70
 - keywords, 223
 - limited access filter, 213
 - policies, 82
 - protocol filter, 53
- adding
 - Always Scan or Never Scan list entries, 166
 - category filters, 50
 - clients, 71
 - email domains (hybrid), 184
 - explicit proxy, 179
 - file types, 240
 - filtered locations, 176
 - limited access filters, 213
 - policies, 82
 - protocol filters, 53
 - root contexts for hybrid service, 193
 - to Websense-defined protocols, 235
 - unfiltered destinations, 181
- admin, 19
 - deleting, 286
 - user, 286
- administrative roles, 286
- administrator accounts
 - viewing, 297
- administrators, 286
 - accessing TRITON - Web Security, 312
 - adding to role, 298, 301
 - concurrent access to same role, 290
 - conditional policy permissions, 288
 - deleting from role, 298
 - Filter Lock, effects of, 292
 - in multiple roles, 290, 301
 - notifying of responsibilities, 295
 - overview, 287
 - permissions, 288
 - permissions, setting, 299, 302
 - reporting, 308
 - reporting permissions, 288, 300
 - Super Administrator, 288
 - tasks for delegated, 307
 - tasks for Super Administrator, 291
 - tracking changes made, 328
 - unable to switch modules, 446
 - unconditional policy permissions, 288
 - viewing role definition, 308
- advertisements, blocking, 92
- alerts, 339
 - category usage, 332
 - category usage, adding, 337
 - category usage, configuring, 336
 - configuring limits, 333

- configuring methods, 333
- email, 334
- Health Summary, 24
- hybrid filtering, 454
- methods for sending, 332
- pop-up, 334
- preventing excessive, 333
- protocol usage, 332
- protocol usage, adding, 338
- protocol usage, configuring, 337
- real-time database updates, 339
- Real-Time Security Updates, 339
- SNMP, 334
- system, 332
- system, configuring, 335
- Websense health, 339
- alternate block messages, 101
- Always Scan list
 - adding sites, 166
 - deleting entries, 167
- anonymous logging, 356
- applets
 - quota time, 48
- application scanning, 162
- Apply Policy to Clients, 85
- Apply to Clients, 83
- ASCII characters, extended
 - searching investigative reports, 439
- audit log, 328
- auditor permissions, 290
- authentication
 - selective, 249
- Authentication Service
 - defined, 280
 - deploying, 280
- B**
- backing up Websense data, 340
- backup files
 - naming, 340
 - storing, 344
- backup utility, 340
 - command reference, 346
 - configuration file, 344
 - running, 343
 - scheduling backups, 342
 - stopping scheduled backups, 346
- bandwidth
 - larger than expected, 436
 - logged for blocked requests, 130
 - managed by Content Gateway, 237
 - managed by Network Agent, 236
 - managing, 236
 - setting limits, 237
 - used by categories, 236
 - used by protocols, 236
- Bandwidth category, 42
- bandwidth logged, blocked requests, 138
- bandwidth savings
 - History page, 28, 30
- banner, 21
- bar chart, 131
- bin directory, 457
- blank block page, 92
- blank page instead of block page, 419
- Block, 46
 - based on bandwidth, 47
 - File Types, 47
 - Keywords, 47
- Block All filter, 55
 - and filtering precedence, 87
- block messages
 - changing frame size, 98
 - creating alternate, 101
 - creating custom, 97
 - customizing, 95
 - for file types, 239
 - protocol, 94
- block pages, 91
 - account override, 74
 - advertisements, 92, 419
 - blank white page, 419
 - changing logo, 98
 - content variables, 99
 - Continue button, 47
 - customizing for hybrid service, 185
 - error message appears, 418
 - hybrid filtering, 184

- not displayed for file type block, 418
- partial, 93
- password override, 73
- reverting to default, 100
- source files, 95
- Use Quota Time button, 47
- blocked advertisements, 92
- blocked and locked, 293
 - categories, 293
 - file types, 294
 - keywords, 293
 - protocols, 294
- blocked requests
 - bandwidth logged, 130
- blocked requests, bandwidth logged, 138
- blocking
 - based on keyword, 223
 - embedded content, 93
 - file types, 238
 - posts to some sites, 228
 - protocols, 230
- blocking NIC, 383
- BlockMessageBoardPosts, 228
- browse session, 367
- browse time
 - Internet (IBT), 106, 366
- bypassing decryption, 171
- C**
- cache file
 - logging, 358
- cached changes, 22
- catalog
 - database, 360
 - report, 107
- categories
 - added to Master Database, 42
 - adding custom, 221
 - Bandwidth, 42
 - bandwidth usage, 236
 - custom, 218
 - defined, 34, 40
 - editing custom, 218
 - Extended Protection, 43
 - list of all, 41
 - locking for all roles, 293
 - logging, 355
 - Productivity, 42
 - renaming custom, 220
 - Security, 43
 - selecting for presentation reports, 111
 - Special Events, 42
- categorizing content, 158
- Category Bypass settings, 171
- category filters, 49
 - adding, 83
 - creating, 50
 - defined, 39
 - duplicating, 50
 - editing, 51
 - renaming, 51
 - templates, 50, 56
- category management, 217
- category map
 - User Activity Detail report, 140
- category usage alerts
 - adding, 337
 - and logging, 355
 - configuring, 336
 - deleting, 336
- changes
 - caching, 23
 - reviewing, 23
 - saving, 23
- changing roles, 289
- changing URL category, 226
- character set
 - MBCS, 388
- character sets
 - used with LDAP, 69
- charts
 - choosing for Today page, 26
 - Current Filtering Load, 25
 - Filtering Service Summary, 26
 - History page, 28
 - Today page, 24
 - Today's value, 24
- Check Point

- FTP request not blocked, 399
- Check Policy
 - Find User, 244
- Check Policy tool, 242
- CLI, Directory Agent, 454
- clients, 61
 - adding, 71
 - applying policies, 61
 - assigning policies, 83, 85
 - computers, 61, 63
 - custom LDAP groups, 63
 - Directory, 61
 - editing, 72
 - groups, 64
 - managing, 62
 - move to role, 75
 - networks, 61, 63
 - selecting for presentation reports, 110
 - users, 64
- clients, managed, 286
 - adding in roles, 307
 - applying policies, 311
 - assigning to roles, 299, 302, 309
 - deleting from roles, 299, 306
 - in multiple roles, 302, 309
 - moving to role, 75
 - overlapping roles, 304
- columns
 - for detail investigative reports, 136
- combining
 - hybrid and on-premises filtering, 173
- commands
 - WebsenseAdmin, 330
 - WebsenseDaemonControl, 330
- components, 314
 - DC Agent, 318
 - Directory Agent, 319
 - eDirectory Agent, 318
 - Filtering Plug-In, 319
 - Filtering Service, 315
 - Linking Service, 319
 - Log Database, 317
 - Log Server, 317
 - Logon Agent, 318
 - Master Database, 316
 - Network Agent, 315
 - Policy Broker, 315
 - Policy Database, 315
 - Policy Server, 315
 - RADIUS Agent, 318
 - Real-Time Monitor, 317
 - Remote Filtering Client, 202, 317
 - Remote Filtering Server, 202, 317
 - Sync Service, 319
 - TRITON - Web Security, 316
 - Usage Monitor, 316
 - User Service, 318
 - Websense Content Gateway, 316
 - Websense Security Gateway, 316
- Computer Browser service
 - enabling, 414
- computers
 - clients, 61
- conditional policy permissions, 288
- conditional Super Administrator, 288
- Configuration utility
 - accessing, 357
 - Log Server, 357
- configuring Log Server, 352
- Confirm, 47
 - in multiple Policy Server environment, 322
- console mode
 - eDirectory Agent, 411
- consolidation
 - and full URL logging, 366
 - and Internet browse time, 435
 - log records, 350
- contacting technical support, 31
- content
 - blocking embedded, 93
 - categorization, 158
 - threat scanning, 161
- Content Gateway, 316, 326
 - Access page, 327
 - accessing from TRITON console, 327
 - non-critical alerts, 443
 - not available, 443
 - not running, 443

- subscription key, 156
 - system alerts, 335
 - content stripping, 164
 - Continue button, 47
 - copy
 - presentation reports, 108
 - Copy to Role, 216
 - filters, 50
 - policies, 81
 - copying
 - category filters, 50
 - limited access filters, 50
 - presentation reports, 108
 - protocol filters, 50
 - creating
 - category filters, 83
 - limited access filters, 83
 - policies, 82
 - protocol filters, 84
 - Current Filtering Load chart, 25
 - custom block messages, 97
 - custom categories, 218
 - adding, 221
 - creating, 217
 - editing, 218
 - renaming, 220
 - custom LDAP groups, 69
 - adding, 70
 - editing, 70
 - custom logo
 - block pages, 98
 - hybrid block pages, 185
 - presentation reports, 110, 114
 - custom protocols, 229
 - creating, 234
 - editing, 232
 - identifiers, 232
 - renaming, 233
 - unable to create, 426
 - custom URLs
 - defined, 224
 - filtering precedence, 225
 - not filtered correctly, 398
 - customer support, 37
 - customize
 - block messages, 95
 - History page, 29, 30
 - hybrid filtering block pages, 185
 - proxy auto-configuration (PAC) file, 188
 - Today page, 26
- ## D
- das.ini, 452
 - database
 - catalog, 360
 - for scanning, 168
 - Log Database, 360
 - Log Database jobs, 361
 - Log Database partitions, 360
 - maintenance job, 364
 - Master Database, 34
 - Policy Database, 319
 - real-time database updates, 35
 - Real-Time Security Updates, 35
 - database account
 - used by Log Server, 431
 - database download, 34
 - configuring, 35
 - disk space requirements, 392
 - memory requirements, 393
 - Real-Time Security Updates, 35
 - real-time updates, 35
 - restriction application problems, 394
 - resuming, 326
 - scanning, 168
 - status, 325
 - subscription problems, 390
 - troubleshooting, 390
 - verify Internet access, 391
 - via proxy, 36
 - database engines
 - supported, 349
 - database jobs
 - ETL, 361
 - Internet browse time (IBT), 361
 - maintenance, 361
 - SQL Server Agent, 437
 - database partitions

- creating, 363
- deleting, 364
- enabling or disabling, 436
- rollover options, 362
- selecting for reports, 363
- database updates, 34
 - real-time, 35, 339
 - Real-Time Security, 35, 339
 - scanning, 168
- date range
 - investigative reports scheduled job, 148
 - presentation reports scheduled job, 122
- DC Agent, 255, 318
 - configuring, 256
 - dc_config.txt file, 404
 - insufficient permissions, 404
 - missing required file, 404
 - permissions, 415
 - troubleshooting, 405
- Default policy, 80
 - applied incorrectly, 406
 - incorrectly applied, 406
- default user, 286
 - deleting, 286
- delegated administration
 - adding administrators, 301
 - adding roles, 296, 297
 - applying policies, 307
 - auditors, 290
 - deleting clients from roles, 306
 - deleting roles, 296
 - deleting roles, effects of, 305
 - editing roles, 298
 - notifying administrators, 295
 - overview, 285
 - permissions, 288
 - policy permissions, 288
 - preparation, 291
 - reporting access, 352
 - reporting permissions, 288
 - role conflicts, 304
 - using, 296
- deleting entries from the Always Scan or Never Scan lists, 168
- deleting managed clients, 425
- detail view
 - columns, 136
 - configuring defaults, 369
 - investigative reports, 134
 - modifying, 135
- diagnostics
 - eDirectory Agent, 410
- Directory Agent, 189, 319
 - adding contexts, 193
 - and off-site users, 189, 207
 - communication issues, 451
 - configuration file, 452
 - directory service not supported, 452
 - exclude contexts, 194
 - not running, 450
 - status information, 197
 - unable to connect to domain controller, 450
- Directory Agent command-line interface, 454
- directory clients, 61
- directory service
 - configuration problems, 412
- directory service settings
 - troubleshooting, 412
- directory services
 - configuring, 65
 - configuring for TRITON - Web Security logon, 312
 - Log Server connecting to, 433
 - searching, 72
 - supported for hybrid filtering, 189
 - Windows Active Directory (Mixed Mode), 66
- directory settings
 - advanced, 68
- disk space
 - database download requirements, 392
 - Log Database requirements, 350
 - presentation reports usage, 119
- display options
 - investigative reports, 370
- domain controller
 - testing for visibility, 407
- drill down, investigative reports, 129
- dynamic content
 - categorizing, 158

E

- eDirectory, 67
- eDirectory Agent, 266, 318
 - configuring, 267
 - console mode, 411
 - diagnostics, 410
 - troubleshooting, 409
- eDirectory server replicas
 - configuring, 269
- Edit
 - category filter, 51
 - custom LDAP group, 70
- Edit Categories button, 217
- Edit Protocols button, 217
- editing
 - category filters, 51
 - client settings, 72
 - email domains (hybrid), 185
 - explicit proxy, 180
 - filtered locations, 177
 - limited access filters, 214
 - policies, 83
 - protocol filters, 53
 - root contexts for hybrid service, 194
 - unfiltered destinations, 182
- eimserver.ini
 - BlockMessageBoardPosts parameter, 229
 - SecurityCategoryOverride parameter, 228
- email
 - report distribution, 354
- email address
 - hybrid filtering contact, 174
- email alerts, 334
- email message
 - customizing for investigative reports, 147
 - customizing for presentation reports, 123
- empty block page, 92
- error log
 - deleting for Log Database, 365
 - Event Viewer, 458
 - viewing for Log Database, 361
 - Websense.log, 458
- error messages
 - finding, 420
 - health alerts, 421
 - estimates
 - bandwidth savings, 30
 - time savings, 30
- ETL job, 361
- evaluating filtering policies, 105
- Event Viewer, 458
- Example - Standard User policy, 79
- examples
 - category and protocol filters, 55
 - policies, 79
- Excel format
 - audit log, 328
 - hybrid service authentication reports, 199
 - investigative reports, 128, 148
 - presentation reports, 117, 119, 120, 123
 - reports incomplete, 438
- explicit proxies
 - adding, 179
 - editing, 180
 - hybrid filtering, 179
- extended ASCII characters
 - in DC Agent machine name, 257
 - in eDirectory Agent machine name, 267
 - in Logon Agent machine name, 259
 - in RADIUS Agent machine name, 263
 - searching investigative reports, 439
- Extended Protection, 43
- Extract, Transform, and Load (ETL) job, 361

F

- fail closed
 - remote filtering software, 203
 - timeout, 203, 205
- fail open
 - remote filtering software, 203
- failed batches, 365
- Favorites
 - investigative reports, 128, 144, 145, 146
 - presentation reports, 108, 114, 116
- file extensions
 - adding to file type, 240
 - adding to predefined file type, 240
 - filtering by, 238
 - for scanning, 163
 - in predefined file types, 239

- file name
 - scheduled presentation report, 119
 - file scanning
 - file extensions, 163
 - scanning timeout, 164
 - setting maximum size, 164
 - file types, 217
 - adding, 240
 - blocking, 47
 - editing, 240
 - locking for roles, 294
 - filter components, 217
 - Filter Lock
 - configuring, 292
 - creating, 288, 293
 - effect on roles, 311
 - locking categories, 293
 - locking file types, 294
 - locking keywords, 293
 - locking protocols, 294
 - logging protocols, 294
 - filter templates, 56
 - filtered locations
 - adding, 176
 - defined, 175
 - editing, 177
 - explicit proxies, 179
 - filtering
 - actions, 46
 - combining solutions, 173
 - component overview, 313
 - diagram, 86
 - file types, 238
 - order, 85
 - precedence, 86
 - precedence, custom URLs, 225
 - process, 86
 - protocols, 230
 - remote or roaming users, 201
 - search images, 48
 - toolbox, 242
 - with keywords, 222
 - filtering methods
 - combining, 175, 207
 - Filtering Plug-In, 319
 - Filtering Service, 315, 324
 - Content Gateway connections, 325
 - database downloads, 325
 - Details page, 325
 - finding version information, 325
 - high CPU alert, 397
 - IP address change, 401
 - not running, 395
 - Summary chart, 26
 - updating UID, 401
 - filtering settings
 - configuring, 57
 - filters, 49
 - category, 39, 49
 - copy to role, 216
 - creating for role, 310
 - determining usage, 84
 - editing active, 84
 - editing for role, 310
 - limited access, 49, 212
 - presentation reports, 108
 - protocol, 39, 49
 - restoring defaults, 56
 - firewall settings
 - database download, 392
 - flood control, alerts, 333
 - full URL logging, 350, 366
- ## G
- global catalog, 66
 - group search filters, 195
 - groups, 64
 - custom LDAP, 63
- ## H
- handling encrypted traffic, 171
 - health alerts, 339
 - described, 421
 - solutions, 422
 - Summary, 24
 - heartbeat
 - changing interval, 206
 - remote filtering software, 203

- hiding user names
 - investigative reports, 131
 - History page, 27
 - charts, 28
 - customizing, 29, 30
 - no charts appear, 442
 - hits
 - logging, 350
 - HTML format
 - presentation reports, 119
 - saving presentation reports, 438
 - HTML format, presentation reports, 117
 - HTTP Post, 359
 - HTTP tunneling, 160
 - exceptions, 160
 - HTTPS sites
 - custom filtering, 398
 - hybrid block pages
 - customizing, 185
 - logo, 185
 - text, 186
 - hybrid filtering, 173
 - account, 174
 - Active Directory configuration, 189, 191
 - Active Directory root context, 191
 - alerts, 454
 - Authentication Service, 280
 - connection lost, 455
 - contact email address, 174
 - custom block pages, 184, 185
 - exclude contexts, 194
 - filtered locations, 175
 - missing configuration information, 456
 - missing from reports, 448
 - Novell eDirectory configuration, 193
 - Novell eDirectory root context, 193
 - NTLM identification, 277
 - off-site user passwords, 189
 - off-site users, 175, 184, 206
 - Oracle Directory Server configuration, 192
 - Oracle Directory Server root context, 192
 - PAC file, 183
 - registering domains, 184
 - scheduling policy, user, and log record synchronization, 196
 - self-registration, 184, 207
 - SSL certificate, 186
 - status, 197
 - supported directory services, 189
 - transparent identification, 276
 - unable to authenticate, 455
 - unfiltered destinations, 181
 - user access, 183
 - user and group search filters, 195
 - user identification, 276
 - Web Endpoint, 277
 - hybrid log records
 - collecting, 196
 - not downloaded, 447
 - not sent to Log Server, 448
 - hybrid user identification
 - Active Directory, 189
 - Novell eDirectory, 193
 - Oracle Directory Server, 192
- ## I
- identifiers
 - protocol, 232
 - Information Technology
 - sites miscategorized, 397
 - initial database, 34
 - initial filtering database, 389
 - installation directories, 331
 - Internet browse time (IBT)
 - and consolidation, 435
 - configuration, 366
 - database job, 106
 - explained, 106
 - read time, 367
 - reports, 366
 - Investigate User tool, 243
 - investigative reporting
 - permissions, 289
 - role type, 286
 - investigative reports, 127, 349
 - accessing, 28
 - anonymous, 131
 - bar chart, 131

- choosing a Log Database, 368
- configuring, 368
- customing email, 147
- default settings, 369
- detail view, 134, 135, 136
- display options, 370
- Excel format, 128, 148, 150
- Favorites, 128, 144, 145
- general issues, 439
- hiding user names, 131
- job queue, 128, 148
- multi-level summary, 133
- options, 128
- outliers, 128, 149
- output options, 370
- overview, 105
- PDF format, 128, 148, 150
- pie chart, 131
- printing, 150
- red lettering, 130
- saving Favorites, 144
- scheduled jobs, 128, 146
- search patterns, 439
- searching, 131, 439
- self-reporting, 151, 372
- setting schedule for, 147
- standard, 128, 142
- summary, 129
- User Activity, 128
- User Activity Detail by Day, 139
- User Activity Detail by Month, 140
- XLS format, 150

IP address change

- Policy Server, 322

IP addresses

- filtering by, 63

J

JavaScript content

- removing, 164

job queue

- investigative reports, 128, 148
- presentation reports, 109

jobs

- ETL, 361
- IBT, 361
- Log Database, 361
- Log Database maintenance, 361
- scheduled investigative reports, 146, 148
- scheduled presentation reports, 119, 123
- SQL Server Agent, 437

K

- key, 31
- keyword blocking
 - troubleshooting, 398
- keywords, 217, 222
 - blocking, 47
 - defining, 223
 - locking for roles, 293
 - not being blocked, 398

L

LDAP

- character sets, 69
- custom groups, 69

limited access filters, 49, 212

- adding, 83
- creating, 213
- filtering precedence, 212
- regular expressions, 215
- renaming, 214
- URLs not permitted, 398

limiting access time, 48

Linking Service, 319

locating product information, 31

log

- audit, 328
- Remote Filtering Client, 203

log cache file, 358

Log Database, 317, 349, 353

- active, 362
- administering, 353, 361
- catalog database, 360
- connect for investigative reports, 368
- consolidation, 358
- creating partitions, 362
- database partitions, 360

- deleting errors, 365
- disk space requirements, 350
- IBT job, 106, 361
- introducing, 360
- jobs, 361
- maintenance configuration, 364
- maintenance job, 361, 364
- not available, 430
- not created, 429
- not recording data, 431
- reindexing, 365
- selecting partitions for reports, 363
- settings, 361
- size, 431
- viewing error log, 361
- log file, 458
 - Remote Filtering Client, 205
- log records, 168
 - for scanning activity, 170
 - Sync Service unable to download, 447
- Log Server, 317, 349
 - configuration, 437
 - Configuration utility, 352, 353, 357
 - connecting to directory service, 433
 - database account, 431
 - low disk space, 427
 - not installed, 428
 - not running, 427
 - starting, 358, 359
 - stopping, 358, 359
 - updating connection, 432
- logging
 - anonymous, 356
 - categories, 355
 - configuring, 355
 - multiple Policy Servers, 355
 - defined, 353
 - full URLs, 366
 - scanning options, 168
 - scanning options compare with filtering, 170
 - selective category, 350, 357
 - strategy, 349
 - user information, 355
- logging on, 19

- logging protocols
 - for all roles, 294
- logo
 - changing on block page, 98
 - changing on hybrid filtering block page, 185
 - presentation reports, 110
- logo, presentation reports, 114
- Logon Agent, 258, 318
 - configuring, 259
 - permissions, 415
 - troubleshooting, 406
- logon error, 426
- logon script
 - domain controller visibility issues, 407
 - enabling NetBIOS, 407
 - user profile issues, 408

M

- Main tab, 22
- maintenance job
 - configuring, 364
 - Log Database, 361, 364
- managed clients, 286
 - adding in roles, 307
 - assigning to role, 299, 302
 - deleting from roles, 299, 306
 - moving to roles, 75
- manual authentication, 247
 - enabling, 249
- Master Database, 34, 316
 - categories, 40
 - download problems, 390
 - download schedule, 35
 - download servers, 391
 - download status, 325
 - downloading, 34
 - enhancing, 358
 - initial, 389
 - more than 1 week old, 389
 - protocols, 41
 - Real-Time Security Updates, 35
 - real-time updates, 35
 - resuming download, 326
 - updating, 389

- maximum size for file scanning, 164
- memory requirements
 - database download, 393
- Microsoft Excel
 - incomplete reports, 438
- Microsoft ISA Server
 - authentication methods, 398
- Microsoft SQL Server, 349
- missing users
 - after upgrade, 388
- Mixed Mode
 - Active Directory, 66
- mobile users
 - filtering, 201
- monitor Internet activity, 152
- monitoring NIC, 383
- move to role, 75
 - clients, 75
- moving sites to another category, 226
- multiple group policies, 86
- multiple policies
 - filtering precedence, 61
- multiple Policy Servers, 322
- multiple roles, permissions, 290
- MyWebsense portal, 31

N

- Native Mode
 - Active Directory, 66
- navigating TRITON - Web Security, 21
- NetBIOS
 - enabling, 405, 407
- Network Agent, 315, 375
 - and remote filtering software, 202
 - blocking NIC, 383
 - communication with Filtering Service, 401
 - global settings, 379
 - hardware configuration, 376
 - high CPU use, 402
 - insufficient memory, 402
 - local settings, 381
 - monitoring NIC, 383
 - more than 2 NICs, 401
 - NIC configuration, 383
 - not installed, 400

- not monitoring, 401
- not running, 400
- protocol management, 383
- network configuration, 376
- network credentials
 - accessing TRITON - Web Security, 312
- networks
 - clients, 61
- Never Scan list, 158
 - adding sites, 166
 - deleting entries, 167
- NIC configuration, 377
 - blocking, 383
 - monitoring, 383
 - settings, 383
- Novell eDirectory, 67
 - hybrid configuration, 193
- NTLM identification
 - hybrid filtering, 277

O

- off-site users
 - configuring remote filtering software, 202
 - enable hybrid filtering for, 184
 - filtering options, 201
 - identifying (hybrid), 189, 207
 - not filtered by hybrid service, 456
 - self-registration (hybrid), 207
- optimizing search results, 195
- options, investigative reports, 128
- Oracle Directory Server
 - hybrid configuration, 192
- order
 - filtering, 86
- outliers reports, 128, 149
- output options
 - investigative reports, 370
- override action
 - categories, 220
 - protocols, 233

P

- PAC file. See proxy auto-configuration (PAC) file.
- partitions
 - creating, 363

- deleting, 350, 364
- Log Database, 360
- rollover options, 362
- selecting for reports, 363
- password
 - resetting, 424
- password override, 73
 - in multiple Policy Server environment, 322
- patches, 31
- PDF format
 - hybrid service authentication reports, 199
 - investigative reports, 128, 148, 150
 - presentation reports, 117, 119, 123
- permissions
 - administrator, 288
 - auditor, 290
 - conditional policy, 288
 - DC Agent, 404, 415
 - installation drive, 429
 - investigative reporting, 289
 - Logon Agent, 415
 - multiple roles, 290
 - policy, 288
 - policy management and reporting, 289
 - releasing policy, 295
 - reporting, 288, 311
 - setting, 299, 300, 302
 - SQL Server, 429, 432
 - unconditional policy, 288
 - User Service, 415
- Permit, 46
- Permit All filter
 - and filtering precedence, 87
- Permit All filters, 55
- permitting URLs for all users, 225
- permitting URLs for all users (hybrid), 181
- pie chart, 131
- policies
 - adding, 81, 82
 - applying to clients, 83, 85
 - applying to managed clients, 307, 311
 - applying to users and groups, 64
 - copy to role, 216
 - copying to roles, 81, 305
 - creating for role, 310
 - Default, 80
 - defined, 39, 79
 - descriptions, 82
 - determining applicable, 85
 - editing, 81, 83
 - editing for role, 310
 - enforcing, 85
 - Example - Standard User, 79
 - filtering precedence, 86
 - multiple group, 86
 - printing to file, 81
 - renaming, 83
 - troubleshooting for directory clients, 399
 - troubleshooting for remote clients, 399
 - Unrestricted, 79
 - viewing, 81
- Policy Broker, 315
 - and the Policy Database, 319
- policy configuration
 - restoring defaults, 56
- Policy Database, 315, 319
 - does not start, 424
- policy definition
 - schedule, 83
- policy management and reporting
 - permissions, 289
 - role type, 286
- policy permissions, 288
 - conditional, 288
 - releasing, 295
 - unconditional, 288
- Policy Server, 315, 320
 - and the Policy Database, 319
 - and TRITON - Web Security, 320
 - changing IP address, 322
 - managing TRITON - Web Security
 - connections, 321
 - multiple instances, 322
 - multiple instances, configuring logging, 355
 - removing from TRITON - Web Security, 321
 - stops unexpectedly, 425
- pop-up alerts, 334
- pop-up blocking

- reporting access, 442
- precedence
 - delegated administration role, 304
 - filtering, 86
 - filtering policy, 61
- preferences, reporting, 354
- presentation reports, 349
 - confirming report filter, 115
 - copying, 108
 - custom, 108
 - custom logo, 110, 114
 - disk space usage, 119
 - do not display, 439
 - errors, 439
 - Excel format, 117, 118, 123
 - Favorites, 108, 114, 116
 - file name, 119
 - HTML format, 117, 119
 - job history, 125
 - job queue, 109, 123
 - low disk space, 435
 - output format, 123
 - overview, 105
 - PDF format, 117, 119, 123
 - printing, 118
 - report catalog, 107
 - report catalog name, 113
 - report filter, 108, 109
 - retaining, 119
 - Review Reports, 109
 - Review Reports page, 125
 - running, 117
 - saving, 118
 - scheduled job fails, 435
 - scheduling, 108, 118, 120
 - setting date range for job, 122
 - viewing scheduled, 125
 - XLS format, 117, 119
- Presentation Reports Scheduler
 - not connected, 434
- Print Policies To File, 81
- printing
 - History page, 29
 - investigative reports, 150
 - presentation reports, 118
 - Today page, 26, 339
- priority, role, 297, 304
- Privacy Category, 171
- private IP addresses
 - and hybrid filtering, 175
- Productivity category, 42
- protocol
 - definitions, 229
 - management, 217
- protocol block message
 - instead of block page, 419
- protocol block messages, 94
 - not displayed, 419
- protocol detection, 160
- protocol filters, 49
 - adding, 84
 - creating, 53
 - defined, 39
 - editing, 53
 - renaming, 53
 - templates, 53, 56
- protocol identifiers, 232
 - IP addresses, 232
 - ports, 232
- protocol management
 - Network Agent, 383
- protocol usage alerts
 - adding, 338
 - configuring, 337
- protocols
 - added to master database, 42
 - bandwidth usage, 236
 - collecting usage information, 33
 - creating new, 231
 - defined, 34, 41
 - defining custom, 217
 - definitions, 229
 - filtering, 53, 230
 - list of all, 41
 - locking for all roles, 293, 294
 - logging for all roles, 294
 - modifying Websense-defined, 235
 - not logged, 436

- renaming custom, 233
- Security Protocol Groups, 46
- selecting for investigative reports, 137
- selecting for presentation reports, 112
- TCP and UDP support, 54
- proxy auto-configuration (PAC) file, 183
 - browser configuration, 187
 - customized fragment, 188
 - customizing, 188
 - default, 188
 - defined, 187
 - status information, 197
 - unfiltered destinations, 181
- proxy server
 - database download configuration, 36
- proxy settings
 - database download, 392
 - verifying, 392

Q

- Quick Start tutorials, 19
 - launching, 19
- Quota, 47
- quota time, 48
 - applets, 48
 - applying to clients, 48
 - in multiple Policy Server environment, 322
 - sessions, 48

R

- RADIUS Agent, 261, 318
 - command-line parameters, 412
 - configuring, 263
 - console mode, 412
 - diagnostics, 411
- read time, 367
- read time threshold, 367
- real-time database updates, 35, 339
- Real-Time Monitor, 152, 317
 - low memory, 440
 - multiple Policy Servers, 154
 - no data, 423
 - not responding, 441
 - not running, 440

- overview, 106
- starting and stopping, 330
- real-time options
 - file scanning, 162
- real-time scanning, see scanning, 155
- Real-Time Security Updates, 35, 339
- recategorized URLs, 225
 - adding, 226
 - editing, 226
 - explained, 217
 - not applied, 426
- red lettering, investigative reports, 130
- regular expressions, 217, 241
 - and unfiltered URLs, 226
 - in a limited access filter, 215
 - recategorizing URLs, 219
- reindexing the Log Database, 365
- release policy permissions, 295
- Remote Filtering Client, 202, 317
 - log file, 203, 205
- Remote Filtering Server, 202, 317
- Remote Filtering settings, 204
 - Block all requests, 204
 - timeout interval, 205
- remote filtering software, 202
 - and Network Agent, 202
 - bandwidth-based filtering, 202
 - changing heartbeat interval, 206
 - communication, 203
 - fail closed, 203
 - fail open, 203
 - heartbeat, 203
 - ignoring FTP traffic, 205
 - ignoring HTTPS traffic, 205
 - inside the network, 203
 - outside the network, 203
 - supported protocols, 202
 - timeout interval, 203
- remote users
 - filtered incorrectly, 417
 - filtering, 201
 - manual authentication issues, 417
- removing
 - active content, 164

- Always Scan or Never Scan list entries, 167
- Policy Server instances from TRITON - Web Security, 321
- VB Script content, 164
- rename
 - category, 220
 - category filters, 51
 - custom protocol, 233
 - limited access filters, 214
 - policies, 83
 - protocol filters, 53
- report catalog, 107
 - name, 113
- report filter, presentation reports, 108, 109
 - confirming, 115
 - selecting actions, 113
 - selecting categories, 111
 - selecting clients, 110
 - selecting protocols, 112
 - selecting risk classes, 111
- report title, presentation reports, 114
- reporting
 - access, 352
 - administrator, 308
 - components, 349
 - configuring email server, 354
 - configuring self-reporting, 372
 - missing hybrid filtering data, 448
 - overview, 105
 - permissions, 288, 300, 311
 - pop-up blocking, 442
 - preferences, 354
 - Real-Time Monitor, 106, 152
 - retrieve hybrid data, 196
 - scanning information, 170
 - scanning options, 168
 - self-reporting, 303
 - setting permissions, 300
 - strategy, 349
 - timeout, 431
- reports
 - configuring investigative, 368
 - editing, 108
 - email distribution, 354
 - empty, 436
 - incomplete, 438
 - investigative, 105
 - presentation, 105
 - retaining, 119
 - Review Reports page, 125
 - User Activity Detail by Day, 139
 - User Activity Detail by Month, 140
 - using, 105
 - viewing scheduled presentation, 125
- reputation filtering, 43
- restore utility, 340
 - command reference, 346
 - running, 345
- restoring Websense data, 340, 345
- Review Reports, 125
 - presentation reports, 109
- risk classes, 44, 352, 353
 - assigning categories, 353
 - Business Usage, 44
 - in reporting, 353
 - Legal Liability, 44
 - Network Bandwidth Loss, 44
 - Productivity Loss, 44, 45
 - Security Risk, 45
 - selecting for investigative reports, 137
 - selecting for presentation reports, 111
- roaming users
 - filtering, 201
- roles
 - adding, 296, 297
 - adding administrators, 298, 301
 - adding managed clients, 299, 302, 307, 309
 - administrative, 286
 - administrators in multiple, 301
 - applying policies, 307, 311
 - clients in multiple, 304
 - creating filters, 310
 - creating policies, 310
 - deleting, 296
 - deleting administrators, 298
 - deleting clients, 299
 - deleting Super Administrator, 286, 306
 - deleting, effects of, 305

- editing, 298
 - editing filters, 310
 - editing policies, 310
 - Filter Lock, effects of, 292
 - investigative reporting, 286
 - locking categories, 293
 - locking protocols, 294
 - names, 296
 - overlapping clients, 309
 - policy management and reporting, 286
 - priority, 297, 304
 - Super Administrator, 286
 - switching, 289
 - viewing definition, 308
 - rollover options, database partitions, 362
 - running TRITON - Web Security, 19
- S**
- samples
 - category and protocol filters, 55
 - policies, 79
 - Save All, 22
 - saving presentation reports, 118
 - scanning
 - database updates, 168
 - exceptions, 156
 - how to enable, 156
 - log records, 170
 - overview, 155
 - settings, 157
 - subscription key, 156
 - scanning applications, 162
 - scanning files, 162
 - scanning options, 161, 168
 - categorizing content, 158
 - reporting, 168
 - saving changes, 159, 160, 161, 163, 167, 172
 - stripping content, 164
 - schedule
 - policy definition, 83
 - scheduled jobs
 - activating, 125
 - customizing email, 123, 147
 - date range, 122, 148
 - deactivating, 125
 - deleting, 124
 - failed for presentation reports, 435
 - investigative reports, 128, 146
 - job history, 125
 - lost, 434
 - output format, 123
 - presentation reports, 119, 121, 123
 - report file name, 119
 - schedule, 120, 147
 - scheduled jobs list
 - investigative reports, 148
 - presentation reports, 108
 - Scheduler, presentation reports, 118
 - scheduling
 - backups, 342
 - hybrid directory synchronization, 196
 - hybrid log record synchronization, 196
 - hybrid policy synchronization, 196
 - routing traffic through proxy or firewall, 197
 - stopping scheduled backups, 346
 - scheduling hybrid communication, 196
 - search filtering, 48
 - search filters
 - hybrid filtering, 195
 - search pattern
 - investigative reports, 439
 - searching
 - directory clients, 72
 - from address bar, 397
 - investigative reports, 131, 439
 - securewispproxy.ini file, 205, 206
 - Security category, 43
 - Security Gateway, 316
 - Security Protocol Groups, 46
 - Security Risk
 - filtering sites, 227
 - security threats
 - scanning for, 161
 - SecurityCategoryOverride, 227
 - selective authentication, 249
 - selective category logging, 350, 357
 - self-registration, 207

- adding domains, 184
- editing domains, 185
- self-reporting, 151, 303
 - configuring, 372
 - enabling, 354
 - notifying users, 372
- services
 - stopping and starting, 329
- Services dialog box, 458
- session
 - timeout, 27
- session timeout, 19
- session, browse, 367
- setting scanning options, 157
- settings
 - Account, 32
 - Alerts, 333
 - Database Download, 35
 - Directory Services, 65
 - Filtered Locations, 175
 - Filtering, 57
 - Hybrid User Identification, 276
 - Log Database, 361
 - Network Agent, 379
 - Policy Server, 320
 - Remote Filtering, 204
 - Scheduling, 196
 - Shared User Data, 189, 192, 193
 - User Identification, 248
- Settings tab, 22
- SNMP alerts, 334
- Special Events, 42
- SQL Server
 - database account, 431
 - permissions, 429
- SQL Server Agent
 - job, 437
- SQL Server Agent job, 432
- SSL certificate
 - hybrid filtering, 186
- SSL decryption bypass, 171
 - overview, 156
 - Privacy Category, 171
- standard reports, investigative, 128, 142
- starting
 - Linux daemons, 330
 - Log Server, 358, 359
 - Websense services, 329
- Status
 - Alerts, 339
 - Audit Log, 328
 - History, 27
 - Today, 24
- status
 - hybrid filtering, 197
 - Hybrid Service, 197
- Status Monitor mode, 27
- stopping
 - Linux daemons, 330
 - Log Server, 358, 359
 - Websense services, 329
- storing backup files, 344
- stripping active content, 164
- subscription key, 31
 - entering, 32
 - invalid or expired, 388
 - not verified, 388
 - verifying, 390
- subscriptions, 31
 - exceeded, 31
 - expired, 31
 - MyWebsense portal, 31
- summary reports
 - investigative reports, 129
 - multi-level, 133
- Sun Java System Directory, 67
- Super Administrator
 - adding clients to role, 305
 - admin, 19
 - conditional, 288
 - deleting role, 286, 306
 - Filter Lock, effects of, 292
 - moving clients from role, 75
 - permissions, 288
 - role, 286
 - switching roles, 289
 - unconditional, 288
- switching roles, 289

- Sync Service, 189, 319
 - configuration file, 449
 - configuring, 196
 - low disk space, 449
 - not available, 446
 - status information, 197
 - unable to connect to hybrid service, 455
 - unable to connect to Log Server, 448
 - unable to download log records, 447
- syncservice.ini, 449
 - BlockMessageBoardPosts parameter, 229
 - SecurityCategoryOverride parameter, 228
- system alerts, 332
 - configuring, 335
 - Content Gateway, 335
 - Web Security, 335
- system status
 - monitoring, 27
- T**
- TCP and UDP support, 54
- technical support, 37
- templates, 56
 - category filter, 50, 56
 - protocol filter, 53, 56
- Test Filtering
 - Find User, 244
- Test Filtering tool, 243
- threat scanning, 161
- threats
 - in files, 162
 - in Web pages, 161
- time savings
 - History page, 28, 30
- time-based Internet access, 48
- timeout
 - disable for TRITON - Web Security, 27
 - reporting, 431
- title, presentation reports, 114
- Today page, 24
 - charts, 24
 - customizing, 26
 - Database Download, 25
 - Health Alert Summary, 24
 - monitoring, 27
 - no charts appear, 442
 - printing, 26
 - Status Monitor, 26
- Today's Value chart, 24
- Toolbox, 242
- tools
 - Check Policy, 242
 - Find User option, 244
 - Investigate User, 243
 - Test Filtering, 243
 - URL Access, 243
 - URL Category, 242
- tracking
 - Internet activity, 332
 - system changes, 328
- transparent identification
 - hybrid filtering, 276
- transparent user identification, 246
 - agents, 246
 - configuring, 248
 - DC Agent, 255
 - eDirectory Agent, 266
 - Logon Agent, 258
 - RADIUS Agent, 261
- Trap server
 - SNMP alert configuration, 334
- TRITON - Web Security, 19, 316
 - accessing with network account, 312
 - concurrent access by administrators, 290
 - disable timeout, 27
 - logging on, 19
 - managing Policy Server connections, 321
 - navigation, 21
 - Policy Server connections, 320
 - session timeouts, 19
- TRITON banner, 21
- TRITON console, 19
 - unable to switch modules, 446
- troubleshooting tools
 - Event Viewer, 458
 - Services dialog box, 458
 - websense.log, 458
- tunneled protocol detection, 160

tutorials

- Quick Start, 19

U

unable to

- add user and groups, 412

unblocking URLs, 225

unblocking URLs (hybrid), 181

unconditional Super Administrator, 288

unfiltered destinations

- adding, 181

- defined, 181

- editing, 182

- PAC file, 181

- syntax, 182

- webmail, 181

unfiltered URLs, 217, 225

- defining, 225

- for hybrid filtering, 181

- not applied, 426

Unrestricted policy, 79

updates

- Master Database, 389

updating the scanning database, 168

upgrade

- hybrid users not filtered, 456

- missing users, 388

URL Access tool, 243

URL Category tool, 242

usage alerts, 332

- category, adding, 337

- category, configuring, 336

- logging categories, 355

- not generated, 423

- protocol, adding, 338

- protocol, configuring, 337

Usage Monitor, 316

- not available, 423

- not running, 423

Use custom filters, 68

Use more restrictive blocking, 212

- with limited access filters, 212

use quota time, 48

- block page button, 47

user access to hybrid filtering, 183

user accounts

- admin, 286

User Activity Detail by Day report, 139

- category map, 140

User Activity Detail by Month report, 140

User by Day/Month reports, 128, 138

user identification

- Authentication Service, 280

- hybrid filtering, 276

- hybrid user reports, 198

- manual, 247

- remote users, 246

- transparent, 246

- troubleshooting, 403

- Web Endpoint, 277

User Identification page, 248

user information, logging, 355

user profile

- logon script issues, 408

user search, 72

user search filters, 195

User Service, 64, 318

- not available, 396

- on Linux, 415

- permissions, 415

- ports for directory communication, 413

- WINS settings, 416

users, 64

- identifying, 245

- manual authentication, 247

- transparent identification, 246

utilities

- Log Server Configuration, 357

V

View Pending Changes, 23

visits

- logging, 350

W

Web Endpoint

- defined, 277

- deploying, 278

Web Security

- installation directories, 331
 - system alerts, 335
 - Websense bin directory, 457
 - Websense configuration information, 319
 - Websense Manager, 19
 - Websense Master Database, 34
 - Websense proxy, 326
 - Websense software
 - components, 314
 - Websense status, 339
 - Alerts, 339
 - Audit Log, 328
 - History, 27
 - Today, 24
 - Websense Web Security Gateway subscription
 - key, 156
 - websense.log, 458
 - WebsenseAdmin command, 330
 - WebsenseDaemonControl command, 330
 - Windows
 - Event Viewer, 458
 - Services dialog box, 458
 - Windows Active Directory (Mixed Mode), 66
 - Windows Active Directory (Native Mode), 66
 - WINS
 - configuring User Service settings, 416
 - enabling, 416
- X**
- XLS format
 - audit log, 328
 - investigative reports, 128, 150
 - presentation reports, 117, 119

