



Administrator Help

Websense® TRITON® AP-WEB

v8.1.x

©1996–2015, Websense Inc.

All rights reserved.

10900 Stonelake Blvd, 3rd Floor, Austin, TX 78759, 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 and TRITON are registered trademarks of Websense, Inc., in the United States and certain international markets. Websense has numerous other unregistered trademarks in the United States and internationally. All other trademarks are the property of their respective owners.

Microsoft, Windows, 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	1
	Overview	2
	Working in the TRITON Manager	2
	Reviewing, saving, and discarding changes	6
	Your subscription	7
	Managing your account through the MyWebsense Portal	8
	Configuring your account information	8
	The Websense Master Database	11
	Configuring database downloads	12
	What is WebCatcher?	14
Topic 2	The Status Dashboards	17
	Threats dashboard	19
	Investigate threat event details	21
	How severity is assigned to suspicious activity	23
	Reviewing threat incident details	23
	Reviewing threat-related forensic data	25
	Risks dashboard	26
	Usage dashboard	27
	System dashboard	27
	Adding elements to a dashboard	28
	Status Monitor mode	30
Topic 3	Internet Usage Filters	31
	Managing access to categories and protocols	32
	When a category or protocol is blocked	33
	New Master Database categories and protocols	34
	Special categories	34
	Risk classes	36
	Security protocol groups	39
	Actions	40
	Using quota time to limit Internet access	41
	Search filtering	41
	Working with filters	42
	Creating a category filter	43
	Editing a category filter	44

	Creating a protocol filter	46
	Editing a protocol filter	47
	Websense-defined category and protocol filters	48
	Category and protocol filter templates	49
	Configuring filtering settings	50
Topic 4	Web Protection Clients.	55
	Working with clients	56
	Working with computers and networks	57
	Working with users and groups	58
	Connecting web protection software to a directory service	59
	Connecting to Windows Active Directory (Mixed Mode)	60
	Connecting to Windows Active Directory (Native Mode)	61
	Connecting to Novell eDirectory or Oracle Directory Server.	62
	Advanced directory settings	63
	Working with custom LDAP groups	65
	Adding or editing a custom LDAP group	66
	Adding a client	66
	Searching the directory service from the TRITON Manager	67
	Changing client settings	68
	Password override	69
	Account override	70
	Moving clients to roles	71
	Working with hybrid service clients.	72
Topic 5	Web Protection Policies	73
	The Default policy	74
	Working with policies	74
	Creating a policy	75
	Editing a policy	76
	Assigning a policy to clients	78
	Enforcement order	79
	Prioritizing group and domain policies	80
	Responding to a URL request	81
Topic 6	Content Gateway Analysis.	87
	Configuring Content Gateway analysis	89
	Configuring content categorization.	90
	Configuring tunneled protocol detection.	92
	Configuring content security.	93
	Configuring file analysis	94
	Content Gateway outbound security analysis	99
	Content Gateway advanced analysis options	100

	Configuring exceptions to Content Gateway analysis	103
	Data files used with Content Gateway analysis	105
	Reporting on advanced analysis activity	105
	How analysis activity is logged.	107
	SSL decryption bypass	108
Topic 7	Use Reports to Evaluate Internet Activity	111
	What is Internet browse time?	112
	Presentation reports	113
	Creating a new presentation report	115
	Defining the presentation report filter.	117
	Selecting clients for a presentation report	118
	Selecting categories for a presentation report	119
	Selecting protocols for a presentation report	120
	Selecting actions for a report	120
	Setting presentation report options	121
	Customizing the presentation report logo	122
	Confirming presentation report filter definitions.	123
	Working with presentation report Favorites	123
	Running a presentation report	124
	Scheduling presentation reports	125
	Setting the presentation reports schedule.	127
	Selecting presentation reports to schedule.	128
	Setting the date range for a scheduled presentation report	129
	Selecting output options for scheduled presentation reports	130
	Viewing the presentation reports scheduled jobs list	131
	Viewing the presentation reports scheduled job history	132
	Reviewing scheduled presentation reports	133
	Investigative reports	134
	Summary investigative reports	136
	Using search to generate a summary investigative report.	139
	Anonymizing investigative reports	140
	Multi-level summary investigative reports	141
	Flexible detail investigative reports	142
	Columns for flexible detail investigative reports.	144
	User Activity Detail investigative reports.	146
	Standard investigative reports	148
	Favorite investigative reports	149
	Scheduling investigative reports	150
	Managing scheduled investigative reports jobs	153
	Outliers investigative reports	153
	Output options for investigative reports	155
	Accessing self-reporting.	156
	Application reporting	156

	How is user agent data collected?	159
	Browser use details	160
	Platform use details	161
	Real-Time Monitor	163
	Real-Time Monitor in Multiple Policy Server Deployments	165
Topic 8	Exceptions to Web Protection Policies	167
	Managing exceptions	168
	How are exceptions organized?	170
	Adding or editing an exception	170
	Overriding an exception	172
	If multiple exceptions could apply, which takes precedence?	173
	Editing multiple exceptions at the same time	174
	Exception shortcuts	174
	How do I block or permit a URL for everyone?	175
	How do I block or permit a URL for one person?	175
	How do I block or permit a URL for my entire role?	176
	How do I block or permit a URL for one of my managed clients?	177
Topic 9	Block Page Management	179
	Blocking graphical advertisements	181
	Blocking embedded pages	181
	Working with block pages	182
	Customizing the block message	184
	Changing the size of the message frame	185
	Changing the logo that displays on the block page	186
	Using block page content variables	186
	Reverting to the default block pages	188
	Creating alternate block messages	188
	Using an alternate block page on another machine	189
	Determining why a request was blocked	190
	Request blocked by Filtering Service	190
	Request blocked by the hybrid service	191
Topic 10	Configure the Hybrid Service	193
	Activate your hybrid service account	194
	Define hybrid service filtered locations	195
	Adding or editing hybrid service filtered locations	196
	Managing hybrid service explicit proxies	198
	Adding or editing a hybrid service explicit proxy	198
	Configuring failover to the hybrid service	199
	Specify sites not managed by the hybrid service	200
	Adding or editing hybrid service unfiltered destinations	201
	Configure user access to the hybrid service	202

	Adding domains for hybrid self-registration	204
	Editing domains for hybrid self-registration	204
	Customizing hybrid block pages	205
	Configuring certificate verification bypass	206
	Enabling hybrid HTTPS notification pages	206
	What is the hybrid PAC file?	207
	Send user and group data to the hybrid service	209
	Configure Directory Agent settings for the hybrid service	209
	Configure how data is gathered for the hybrid service	211
	Oracle (Sun Java) Directory Server and the hybrid service	212
	Novell eDirectory and the hybrid service	213
	Adding and editing directory contexts for the hybrid service	213
	Optimizing directory search results for the hybrid service	215
	Schedule communication with the hybrid service	216
	Define custom authentication settings for the hybrid service	218
	Adding custom authentication rules for the hybrid service	219
	Editing custom authentication rules for the hybrid service	220
	Monitor communication with the hybrid service	222
	View hybrid service authentication reports	223
	View the hybrid service User Agent Volume report	224
Topic 11	Manage Off-site Users	227
	Hybrid service management of off-site users	227
	Configuring the hybrid service for off-site users	228
	Off-site user self-registration	228
	Using remote filtering software	229
	Configuring Remote Filtering settings	230
	Configure remote filtering to ignore FTP or HTTPS traffic	231
	Configure the Remote Filtering Client heartbeat interval	232
Topic 12	Combine Web, Data, and Mobile Protection	233
	Integrating web and mobile protection solutions	234
Topic 13	Refine Your Policies	237
	Restricting users to a defined list of URLs	238
	Creating a limited access filter	239
	Editing a limited access filter	240
	Adding sites from the Edit Policy page	241
	Copying filters and policies to roles	242
	Building filter components	243
	Working with categories	244
	Editing categories and their attributes	244
	Reviewing all customized category attributes	245
	Making global category changes	246

Renaming a custom category	246
Creating a custom category	247
Keyword-based policy enforcement	248
Defining keywords	250
Reclassifying specific URLs	251
Prioritizing Security Risk categorization	252
Blocking posts to sites in some categories	253
Protocol-based policy enforcement	255
Editing custom protocols	256
Adding or editing protocol identifiers	257
Renaming a custom protocol	258
Making global protocol changes	258
Creating a custom protocol	259
Adding to a Websense-defined protocol	260
Using Bandwidth Optimizer to manage bandwidth	261
Configuring the default Bandwidth Optimizer limits	262
Managing traffic based on file type	263
Enforcement based on file extension	264
Enforcement based on file analysis	267
Enabling file type blocking in a category filter	269
Working with file type definitions	269
Adding custom file types	270
Adding file extensions to a file type	271
Using regular expressions	271
Using the Toolbox to verify policy enforcement behavior	272
URL Category	272
Check Policy	273
Test Filtering	273
URL Access	274
Investigate User	274
Identifying a user to check policy or test filtering	274
Topic 14 User Identification for Policy Enforcement	277
Identifying on-premises users transparently	278
Transparent identification of remote users	278
Manual authentication	279
Configuring user identification and authentication	280
Setting authentication rules for specific machines	282
Defining exceptions to user identification settings	282
Revising exceptions to user identification settings	283
Secure manual authentication	284
Generating keys and certificates	285
Activating secure manual authentication	286

Accepting the certificate within the client browser	287
DC Agent	288
Configuring DC Agent	289
Reviewing DC Agent polled domains and domain controllers.	291
The dc_config.txt file	292
Logon Agent.	293
Configuring Logon Agent	294
Configuring RADIUS Agent	295
Configuring eDirectory Agent	297
Adding an eDirectory server replica	299
Identification and authentication of hybrid users	300
Authentication priority and overrides	302
TRITON AP-ENDPOINT Web software	303
Integrating the hybrid service with a single sign-on identity provider	305
Directory Agent and hybrid user identification.	306
How Directory Agent works with User Service	307
When hybrid users are not identified	308
Topic 15 Delegated Administration and Reporting	311
The fundamentals of delegated administration.	312
Delegated administration roles	312
Delegated administrators.	313
Delegated administration and reporting permissions	314
Administrators in multiple roles	317
Multiple administrators accessing the TRITON console	318
Preparing for delegated administration	319
Creating a Filter Lock	320
Locking categories.	321
Locking protocols	322
Preparing delegated administrators.	323
Managing delegated administration roles.	324
Adding roles	325
Editing roles	326
Adding Administrators	329
Adding managed clients.	330
Managing role conflicts.	332
Updating delegated administration roles	332
Delete roles	333
Delete managed clients	334
Managing Super Administrator clients.	335
Performing delegated administrator tasks	335
View your user account.	336

	Add clients to the Clients page	337
	Create policies and filters	338
	Reviewing administrator accounts	339
	Enabling network accounts.	339
Topic 16	Server Administration for Web Protection Solutions	341
	Websense web protection components.	342
	Policy enforcement and management components	343
	Reporting components	345
	User identification components.	346
	Interoperability components	347
	Reviewing your web protection deployment	348
	Using the Policy Server map	348
	Using the component list.	349
	Evaluating directory performance.	350
	Review directory server details	351
	Understanding Policy Broker	351
	Reviewing Policy Broker connections	352
	Working with Policy Server	354
	Reviewing Policy Server connections.	354
	Adding or editing Policy Server instances	356
	Working in a multiple Policy Server environment	357
	Changing the Policy Server IP address	357
	Working with Filtering Service	360
	Review Filtering Service details	360
	Review Master Database download status	361
	Resuming Master Database downloads.	362
	Filtering Service support for YouTube in Schools	362
	Policy Server, Filtering Service, and State Server	363
	Integrating with a third-party SIEM solution	366
	Working with Content Gateway.	367
	Managing Content Gateway connections	368
	Viewing and exporting the audit log	368
	Stopping and starting Websense services.	370
	Installation directories for web protection solutions	373
	Alerting.	373
	Flood control	374
	Configuring general alert options	374
	Configuring system alerts	376
	Configuring category usage alerts.	377
	Adding or editing category usage alerts.	378
	Configuring protocol usage alerts	379

	Adding or editing protocol usage alerts	379
	Configuring suspicious activity alerts	380
	Reviewing current system status	381
Topic 17	Reporting Administration	383
	Assigning categories to risk classes	384
	Configuring reporting preferences	385
	Configuring how requests are logged	386
	Configuring Log Server	388
	Testing the Log Database connection	393
	Introducing the Log Database for web protection solutions	394
	Web protection reporting database jobs	395
	Log Database administration settings	396
	Configuring database partition options	397
	Configuring Log Database maintenance options	400
	Configuring how URLs are logged	402
	Configuring Internet browse time options	403
	Configuring trend and application data retention	404
	Log Database sizing guidance	406
	Configuring Dashboard reporting data	407
	Configuring investigative reports	410
	Database connection and report defaults	410
	Display and output options	412
	Self-reporting	414
Topic 18	Configure Network Agent	417
	Configuring Network Agent global settings	418
	Configuring Network Agent local settings	419
	Configuring Network Agent NIC settings	421
	Configuring Network Agent monitoring settings for a NIC	422
	Adding or editing IP addresses during Network Agent configuration	423
Topic 19	Troubleshooting	425
	Web protection installation and subscription issues	425
	Web protection Master Database issues	426
	The initial database is being used	426
	The Master Database is more than 1 week old	427
	The Master Database does not download	427
	Verify Filtering Service Internet access	428
	Verify Filtering Service firewall or proxy server settings	429
	Contacting Technical Support for database download issues	430
	Filtering Service alert messages	430
	Filtering Service is not running	431

High CPU usage on the Filtering Service machine.	431
Insufficient disk space on the Filtering Service machine	432
Insufficient memory on the Filtering Service machine.	433
Filtering Service can't communicate with a transparent identification agent.	434
Network Agent issues.	435
Network Agent is not installed	435
Network Agent is not running.	435
Network Agent is not monitoring any NICs	436
Network Agent can't communicate with Filtering Service.	436
Insufficient memory on the Network Agent machine.	437
High CPU usage on the Network Agent machine.	438
User configuration and identification issues	438
User and group-based policies are not applied	438
Unusually high directory server connection latency.	439
User Service is not available.	440
DC Agent has insufficient permissions.	440
DC Agent unable to access required file.	441
DC Agent Domains and Controllers page is blank.	442
I cannot add users and groups to the Web module of the TRITON Manager.	443
Directory service connectivity and configuration	443
Directory service configuration	444
User identification and Windows Server.	444
Turning on the Computer Browser service	445
Changing DC Agent, Logon Agent, and User Service permissions	446
User Service on a Websense appliance or Linux server	446
Health alerts and Usage Monitor issues	447
Where do I find error messages for Websense components?	448
Websense Health alerts	448
Usage Monitor is not available	450
Usage Monitor is not running	450
Policy Server and Policy Broker issues	451
I forgot my password.	451
The Websense Policy Database service fails to start	452
Policy Server stops unexpectedly	452
A Policy Broker replica cannot synchronize data.	453
Log Server and Log Database issues	453
Log Server is not running	454
Log Server has not received log files from Filtering Service	454
Low disk space on the Log Server machine	456
No Log Server is installed for a Policy Server	457
More than one Log Server is installed for a Policy Server.	458

Log Database was not created	459
Log Database is not available	460
Log Database size causes reporting delays	461
More than 100 files in the Log Server cache directory	461
Last successful ETL job ran more than 4 hours ago	463
Log Server is not recording data in the Log Database	464
Log Server cannot connect to the directory service	464
Wrong reporting page displayed	465
Investigative report and presentation report issues	465
Presentation Reports Scheduler not connected to Log Database	465
Inadequate disk space to generate presentation reports	466
Scheduled jobs in presentation reports failed	466
Trend data is missing from the Log Database	467
Trend reports are not displaying data	467
All reports are empty	468
Verify your Log Server configuration	468
Error generating presentation report, or report does not display	469
Investigative reports search issues	470
General investigative reports issues	470
Other reporting issues for web protection solutions	470
Low memory on the Real-Time Monitor machine	471
Real-Time Monitor is not running	471
Real-Time Monitor is not responding	472
No charts appear on the Status > Dashboard page	472
There is a forensics data configuration problem	473
The forensics repository location could not be reached	473
Forensics data will soon exceed a size or age limit	473
Websense Multiplexer is not running or not available	474
TRITON AP-WEB interoperability issues	475
Content Gateway is not running	475
Content Gateway is not available	476
Content Gateway non-critical alerts	476
Administrator unable to access other TRITON modules	479
Sync Service is not available	479
Sync Service has been unable to download log files	480
Sync Service has been unable to send data to Log Server	481
Hybrid policy enforcement data does not appear in reports	481
Disk space is low on the Sync Service machine	481
The Sync Service configuration file	482
Directory Agent is not running	483
Directory Agent cannot connect to the domain controller	484
Directory Agent communication issues	485
Directory Agent does not support this directory service	485

	The Directory Agent configuration file.	486
	Directory Agent command-line parameters	487
	Alerts were received from the hybrid service	488
	Unable to connect to the hybrid service	489
	Hybrid service unable to authenticate connection.	489
	Missing key hybrid configuration information	490
Topic 20	Copyrights.	491
	Trademarks.	491
	Other Acknowledgements	492

1

Getting Started

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Get the most from Websense® TRITON® AP-WEB by using the instructions, troubleshooting tips, and overviews in this guide. Unless otherwise indicated, the topics in this guide also apply to Web Filter & Security.

To get started, you can browse the guide or select a topic from the table below:

First steps	Troubleshooting
<ul style="list-style-type: none">• Working in the TRITON Manager• Your subscription• The Status Dashboards	<ul style="list-style-type: none">• Web protection installation and subscription issues• Web protection Master Database issues• Filtering Service alert messages• User configuration and identification issues• Health alerts and Usage Monitor issues• Log Server and Log Database issues• Investigative report and presentation report issues• TRITON AP-WEB interoperability issues
Policy enforcement	
<ul style="list-style-type: none">• Managing access to categories and protocols• Adding a client• Working with policies• Assigning a policy to clients• Configuring Content Gateway analysis	
Using reports	
<ul style="list-style-type: none">• Presentation reports• Investigative reports• Real-Time Monitor• Reporting on advanced analysis activity• Using the Toolbox to verify policy enforcement behavior	
Advanced tools	
<ul style="list-style-type: none">• Exceptions to Web Protection Policies• Reclassifying specific URLs• Delegated Administration and Reporting	

Overview

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use your Websense solution to develop and enforce policies to protect your network. Together, a series of Websense components (described in [Websense web protection components, page 342](#)) provide security for web-based transactions, as well as management, user identification, alerting, reporting, and troubleshooting capabilities.

Initially, your software uses its **Default** policy to monitor Internet use without blocking requests. (See [The Default policy, page 74](#), for more information.)

- ◆ The Default policy governs Internet access for all clients in the network until you define your own policies and assign them to clients.
- ◆ You are encouraged to edit the Default policy so that it can be used for enforcement, rather than just monitoring.
- ◆ After you have created custom policies, the Default policy is applied to any request not governed by another policy.

To get started with policy enforcement, see:

1. [Internet Usage Filters, page 31](#)
2. [Web Protection Clients, page 55](#)
3. [Web Protection Policies, page 73](#)

A single, browser-based tool—the TRITON Manager—provides a central, graphical interface to the general configuration, policy management, and reporting functions of all on-premises TRITON solutions. See [Working in the TRITON Manager, page 2](#), for more information.

You can define levels of access to the TRITON Manager to allow specified administrators to manage one or more TRITON modules. Within the Web module, you can further refine access permissions to allow administrators to manage policies, perform reporting tasks, and more. See [Delegated Administration and Reporting, page 311](#), for more information.

Working in the TRITON Manager

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Navigating the Web module of the TRITON Manager, page 4](#)
- ◆ [The Status Dashboards, page 17](#)

The TRITON Manager (also called the TRITON console) is the central configuration interface used to manage TRITON Web, Email, and Data solutions. Use its Web

module to customize policies, generate reports, monitor the system, and manage configuration and settings.

At installation, the TRITON Manager 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.

Until a subscription key has been entered, when a Super Administrator logs on and selects the Web module, an **Initial Setup Checklist** is displayed. Use the checklist to enter your subscription key and perform basic initial configuration tasks.

Once a key has been entered and validated, administrators connecting to the Web module are taken to the **Status > Dashboard** page.

- ◆ A quick tutorial is available for administrators new the Web module. Click **Help**, then **Getting Started**, and select **New Admin Tutorial**.
- ◆ On first logon, when an administrator navigates away from the dashboard, the **Save and Deploy** button activates. This allows initial default dashboard settings to be saved for that administrator account.
- ◆ If you are using an account with permissions to access multiple TRITON modules, use the TRITON toolbar to switch between modules. See [Navigating the Web module of the TRITON Manager](#), page 4.
- ◆ 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 311.

When you log on to the TRITON Manager, the Web module 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 module toolbar.

A TRITON Manager 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 changes, and **Save and Deploy** to record and implement those changes.
- ◆ If the console 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 the console is open in multiple browser windows on the same computer, the instances share the same session **unless** you:
 - Launch multiple Internet Explorer windows independently of one another.
 - Use the File > New Session command to open a new Internet Explorer window.
 - Use Internet Explorer to open one connection to the TRITON Manager, and then use Firefox or Chrome to open another connection.

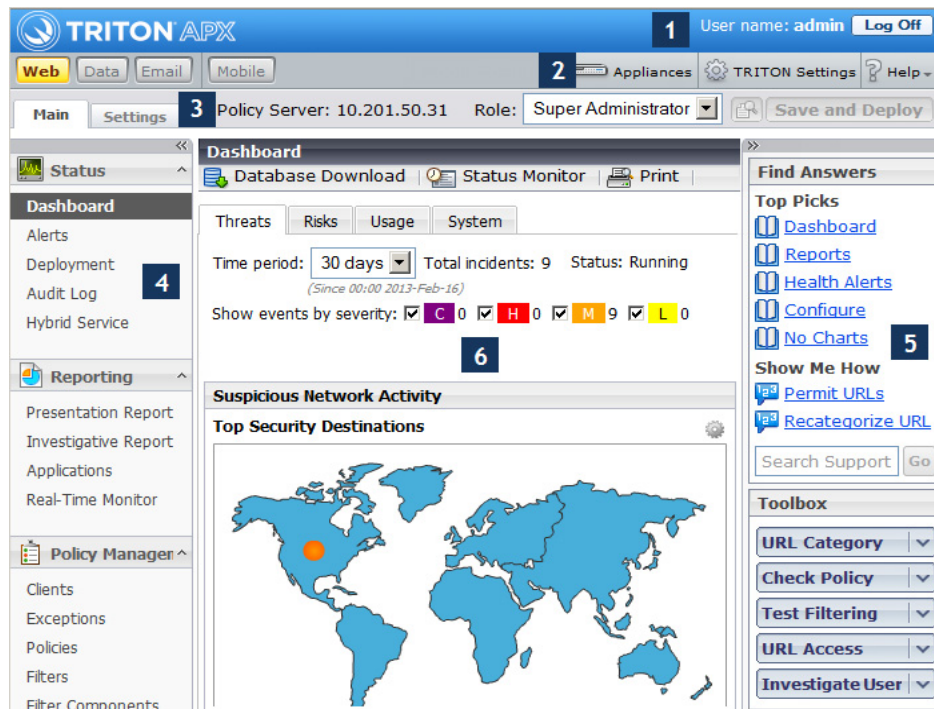
If you close the browser without logging off of the TRITON Manager, or if the remote machine from which you are accessing a TRITON module shuts down unexpectedly, you may be temporarily locked out of the console. The management components

typically detect this issue within about 2 minutes and end the interrupted session, allowing you to log on again.

Navigating the Web module of the TRITON Manager

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Web module of the TRITON Manager can be divided into 6 main areas:



1. Banner
2. TRITON toolbar
3. Web module 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 311, 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 Manager.
- ◆ Connect to the Appliance manager for any V-Series **Appliances** deployed in your network.
- ◆ Configure global **TRITON Settings** that affect all installed modules.
- ◆ Access **Help**, tutorials, product information, and Websense Technical Support resources.

The Web module toolbar



The Web module 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 switch between Policy Server instances, if applicable (see [Working with Policy Server](#), page 354).
- ◆ View your administrative **Role**, switch between roles, or release policy permissions for the current role.



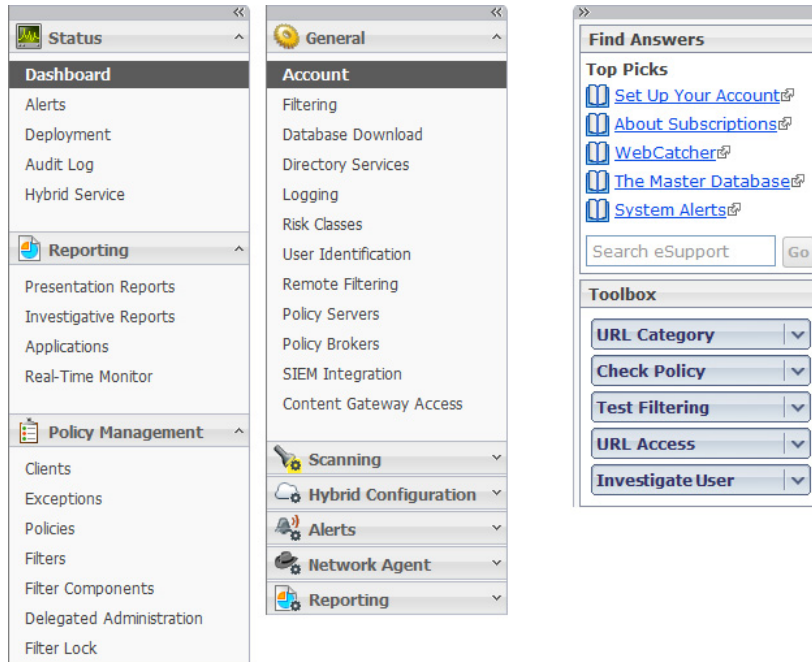
Tip

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 and Deploy** pending changes. If there are no cached changes waiting to be saved, these buttons are disabled.

See [Reviewing, saving, and discarding changes](#), page 6, 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. (Note that the Settings tab displays different options depending on your subscription level.)

The right shortcut pane contains links to useful tools and information.

- ◆ **Find Answers** provides links to articles, webinars, videos, worksheets, and tutorials to help you complete your tasks. Use the Search box to find more information in the Websense eSupport Knowledge Base.
- ◆ The **Toolbox** contains quick lookup tools that you can use to verify your configuration. See [Using the Toolbox to verify policy enforcement behavior](#), page 272, 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When you make a change in Web module of the TRITON Manager, you must typically click **OK** at the bottom of the page to cache the change, then click **Save and**

Deploy to save the change to the Policy Database, which causes the change to take effect.

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

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 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 Changes or Cancel All Changes, you are returned to the last page you selected. There is no undo for either option.

Use the Audit Log to review the details of changes made in the Web module. See [Viewing and exporting the audit log, page 368](#), for more information.

Your subscription

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Websense subscriptions are issued on a per-client (IP address) basis.

To activate your software, enter a valid subscription key (see [Configuring your account information, page 8](#)). This lets you:

- ◆ Download the Master Database (see [The Websense Master Database, page 11](#)), enabling policy enforcement.
- ◆ (*TRITON AP-WEB only*) Download the analytic databases to support real-time analysis of web content.

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

A component called Websense Filtering Service maintains a subscription table of clients generating Internet requests each day. In TRITON AP-WEB deployments, there is no change in policy enforcement when the number of subscribed clients is exceeded, but if the subscription is consistently exceeded, you may be asked to increase your subscription limit. In Web Filter & Security deployments, when the number of subscribed clients is exceeded, requests from users who exceed the subscription count are permitted or blocked based on the setting **Block users when**

subscription expires, found on the **Settings > General > Account** page in the Web module of the TRITON Manager.

In all deployments, when a subscription expires, all requests are permitted or blocked, depending on the same configurable setting. When the expiration date approaches, administrators are notified through a combination of email alerts and TRITON console health alerts.

- ◆ To configure how Internet requests are handled when a subscription expires, see [Configuring your account information, page 8](#).
- ◆ To have an alert message sent when the subscription approaches or exceeds its limit, see [Configuring system alerts, page 376](#).

Managing your account through the MyWebsense Portal

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.

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, hotfixes, and upgrades for your Websense product and version.

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

Configuring your account information

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Your subscription, page 7](#)
- ◆ [Configuring database downloads, page 12](#)
- ◆ [Protocol-based policy enforcement, page 255](#)

Use the **Settings > General > Account** page in to enter or review subscription information, and to determine how your product responds when the subscription expires or, for Web Filter & Security deployments, when the subscription count is exceeded.

Also use the page to direct Websense components to send category and protocol usage data to Websense, Inc., anonymously. This information may be used to help optimize the Websense Master Database (see [The Websense Master Database, page 11](#)) and contribute to the Websense ThreatSeeker® Intelligence Cloud (see websense.com/content/Threatseeker.aspx).

After installation, or any time you receive a new subscription key, you can 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 357).
- ◆ If the key syntax is correct, but the Master Database download fails because the key is invalid or expired, a health alert message is displayed on the **Status > Alerts** page. By default, the message also appears on the System dashboard.

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 analyzing traffic and enforcing policies.
Subscribed users	<i>When you add the Web Hybrid module to TRITON AP-WEB, this shows the sum of users managed by on-premises components and the hybrid service.</i>
Subscribed network users	Number of in-network users whose Internet requests may be managed.
Subscribed remote users	<i>When you add the Remote Filter module to Web Filter & Security, this shows the number of users whose requests may be handled when they are outside the network.</i>
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** to block all Internet access for all users when the subscription expires. In Web Filter & Security deployments, selecting this option will also block all requests from users who exceed the subscription count.

Leave the option unselected to give users unrestricted Internet access when the subscription expires.

2. Mark **Send category and protocol data to Websense, Inc.**, to have Websense components 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 its security capabilities.

3. Under WebCatcher, mark **Send URL information to Websense** to help Websense, Inc., improve URL categorization and security effectiveness. See *What is WebCatcher?*, page 14, for more information about this tool.
 - To submit uncategorized URLs to be evaluated for categorization, mark **Send uncategorized URLs to improve URL categorization**.
 - To send in security-related URLs to help track malicious website activity, mark **Send security URLs to improve security effectiveness**.
 - To keep a local copy of the information sent to Websense, Inc., for your review, mark **Save a copy of the data being sent to Websense**.

When this option is enabled, WebCatcher saves the data as unencrypted XML files in the **Websense\Web Security\bin** directory on the Log Server machine. These files are date and time stamped.
 - Select the **Country of origin** for your organization. This should be the country where the majority of Internet activity is being logged.
 - Specify a **Maximum upload file size**. When the maximum size is reached, collected WebCatcher data is sent automatically and a new file is started.
 - Use the **Daily start time** field to indicate a time each day when WebCatcher should send the data it has collected if the maximum file size has not been reached.
4. (*Web Hybrid module for TRITON AP-WEB*) To activate or update the connection between the on-premises and hybrid portions of your software:
 - Enter the **Contact email address** for your TRITON AP-WEB administrators. This is typically a group email alias that is monitored frequently. Alerts about hybrid service 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.

User requests are not managed by the hybrid service until this information has been provided and validated. For more information, see *Configure the Hybrid Service*, page 193.
5. When you are finished making changes, click **OK**. Changes are not implemented until you click **Save and Deploy**.

The Websense Master Database

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Real-time updates](#), page 12
- ◆ [Configuring database downloads](#), page 12
- ◆ [Review Master Database download status](#), page 361
- ◆ [Resuming Master Database downloads](#), page 362

The Websense Master Database houses the category and protocol definitions that provide the first step in enhancing your Internet security (see [Managing access to categories and protocols](#), page 32).

- ◆ **Categories** are used to group websites (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.

With TRITON AP-WEB, the Master Database provides an initial sorting mechanism to help prioritize real-time analysis, as well as many of the names and descriptions that appear in reports on analyzed content.

A limited version of the URL database is installed with Websense Filtering Service, but it is a good idea to download the full Master Database as soon as possible to enable comprehensive policy management and reporting.

To download the Master Database for the first time, enter your subscription key in the **Initial Setup Checklist** that appears when you open the Web module of the TRITON console. (If Filtering Service must go through a proxy to perform the download, also configure proxy settings in the checklist.)

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, Filtering Service downloads database changes on a schedule that you establish (see [Configuring database downloads](#), page 12). 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, your Websense software stops policy enforcement.

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 the **System** dashboard and click **Database Download** in the toolbar at the top of the content pane.

Real-time updates

In addition to scheduled downloads of the full database, 2 types of smaller, partial updates occur when needed.

- ◆ Real-time database updates
- ◆ Real-Time Security Updates™

Real-time database updates

A real-time database update might be used, for example, to recategorize a site that was temporarily miscategorized. These updates ensure that sites and protocols are managed appropriately.

Websense Filtering Service checks for database updates every hour. The most recent updates are listed on the **Status > Alerts** page (see [Reviewing current system status](#), page 381).

Real-Time Security Updates

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 websites or applications.

With TRITON AP-WEB, these updates allow malicious sites to be blocked based on Master Database categorization, rather than real-time analysis, reducing load on the analytic components and potentially improving performance.

Filtering Service checks for security updates every 5 minutes. Because the updates tend to be small, they do not disrupt normal network activity.

Use the **Settings > General > Database Download** page to enable Real-Time Security Updates (see [Configuring database downloads](#), page 12).

Configuring database downloads

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring your account information](#), page 8
- ◆ [The Websense Master Database](#), page 11
- ◆ [Review Master Database download status](#), page 361

Use the **Settings > General > Database Download** page to establish the schedule for automatic Master Database downloads. If you did not already enter the information in the Initial Setup Checklist, you can use this page to configure any proxy server or firewall settings that Websense Filtering Service must use to download the database.

1. Select **Enable real-time security updates** (default) to have Websense Filtering Service 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 help reduce vulnerability to threats like new phishing (identity fraud) scams, rogue applications, and malicious code infecting a mainstream website or application while lightening the load on real-time analysis.

2. Select the **Download days** for automatic downloads.
 - All download days are selected when Real-Time Security Updates are enabled. Downloads are automatically performed every day to assure that the most up-to-date standard database is available for the security updates.
 - You must download the Master Database at least once every 14 days for policy enforcement to continue without interruption.
 - If you deselect all download days, Filtering Service automatically attempts a download when the database is 7 days old.
3. Next to **Download between**, select a start time and end time between which Filtering Service attempts to download Master Database updates. By default, download occurs between 21:00 (9 p.m.) and 06:00 (6 a.m.), according to the time on the Filtering Service machine.
 - Filtering Service selects a random time during this period to contact the Master Database server. To configure alerts for download failures, see [Configuring system alerts](#), page 376.
 - Any time Filtering Service is restarted, it checks for available Master Database updates. The update may begin immediately, rather than waiting for the defined period.

**Note**

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

4. Select **Use proxy server or firewall** if Filtering Service must access the Internet through a proxy server or a proxying firewall to download the Master Database. Then provide:
 - The **IPv4 address or hostname** of the proxy server or firewall.
 - The **Port** through which the database download must pass (8080, by default).
5. If the proxy server or firewall configured above requires authentication to reach the Internet, select **Use authentication**, and then enter the **User name** and **Password** that Filtering Service 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 63).

What is WebCatcher?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

WebCatcher is an optional feature that collects unrecognized and security-related URLs, and submits them to Websense Security Labs. Uncategorized URLs are reviewed for categorization, and security-related URLs are analyzed for what they can reveal about active Internet threats. (Full URL logging is not required for WebCatcher processing.) The results of the analysis are used to update the Master Database, resulting in improved performance.



Note

If you have multiple Websense Log Server instances, enable WebCatcher only once, on the **Web > Settings > General > Accounts** page in the TRITON console.

The information sent to Websense Security Labs contains only URLs and does not include user information. For example:

```
<URL HREF="http://www.ack.com/uncategorized/" CATEGORY="153"
IP_ADDR="200.102.53.105" NUM_HITS="1" />
```

The IP address in the example reflects the IPv4 or IPv6 address of the machine hosting the URL, not the requester's IP address.



Note

Intranet sites are not sent by WebCatcher. This includes all sites with IP addresses in the 10.xxx.xxx.xxx, 172.16.xxx.xxx, and 192.168.xxx.xxx ranges.

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.

Websense Technical Support

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

- ◆ the searchable Websense Knowledge Base
- ◆ Webinars and show-me videos
- ◆ 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.

The contact page includes information for finding solutions, opening an online support case, and calling Websense Technical Support.

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

For telephone requests, please have ready:

- ◆ Websense subscription key
- ◆ Access to the management console for your solutions (for example, the TRITON Manager, Content Gateway manager, and Appliance manager)
- ◆ 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

The Status Dashboards

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The **Threats** tab of the **Status > Dashboard** page appears first when you log on to the TRITON Manager and connect to the Web module. It shows information about suspicious activity that may be related to malware threats in your network (see [Threats dashboard](#), page 19).

Dashboard elements are visible to Super Administrators and those delegated administrators with permission to view reports on the Status > Dashboard page (see [Editing roles](#), page 326).

- ◆ Delegated administrator access to the Risks, Usage, and System dashboards is configured separately from Threats dashboard access.
- ◆ With TRITON AP-WEB, delegated administrators with Threats dashboard access can also be granted permission to view forensics details associated with advanced malware threats. See [Reviewing threat-related forensic data](#), page 25.

The first time an administrator logs on to the Web module, then navigates away from the **Status > Dashboard** page, the **Save and Deploy** button activates. This happens regardless of whether any changes were made, in order to save default dashboard settings for each administrator account.


After the initial defaults are saved, navigating away from the dashboard activates the **Save and Deploy** button only when charts are added, removed, or edited.

The dashboard includes 3 additional tabs:

- ◆ **Risks** shows information about blocked and permitted requests for URLs that fall into the Security Risk class. See [Risks dashboard](#), page 26.
- ◆ **Usage** shows information about traffic patterns in your network, including bandwidth information and summaries of blocked and permitted requests. See [Usage dashboard](#), page 27.
- ◆ **System** shows alert messages, status information, and graphical charts that show the current state of your deployment, focusing on Internet activity in your network. See [System dashboard](#), page 27.

The Risks, Usage, and System dashboards can each display up to 12 elements (charts, status summaries, or counters) at a time. Most dashboard charts can be customized to change their time period (today, last 7 days, last 30 days, and so on) and their display format (stacked column, stacked area, multi-series line, and so on). You can include

multiple versions of the same chart on a tab (for example, showing different time periods).

- ◆ Dashboard elements are updated every 3 minutes.
All elements on a tab are also updated when any element on the tab is modified. For example, if the time period for one chart is changed, data is refreshed in all of the elements on the page.
- ◆ The available set of dashboard elements depends on your subscription type. Charts related to the hybrid service, for example, require TRITON AP-WEB and the Web Hybrid module.
- ◆ To add an element to the tab, click **Add Charts**, then see [Adding elements to a dashboard](#), page 28, for instructions.
- ◆ To remove an element from the tab, click the Options icon () in the element title bar, then select **Remove**.
- ◆ To access all editing options for an element, click the Options icon in the element title bar, then select **Edit**.
- ◆ Clicking a pie, bar, or line chart typically opens an investigative report with more details. Some security-related charts link instead to the Threats dashboard.

Up to 4 buttons appear in the dashboard toolbar:

- ◆ **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 361).
- ◆ **Status Monitor** releases the current administrator's policy permissions and enters a monitoring mode that allows access to the following pages without timing out:
 - Status > Dashboard
 - Status > Alerts
 - Reporting > Real-Time MonitorSee [Status Monitor mode](#), page 30.
- ◆ **Add Charts** allows administrators to customize their view of the selected dashboard tab by adding elements to the page. See [Adding elements to a dashboard](#), page 28.
- ◆ **Print** opens a secondary window with a printer-friendly version of the charts displayed on the page. Use browser options to print the page.

Threats dashboard

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [The Status Dashboards](#), page 17
- ◆ [Investigate threat event details](#), page 21
- ◆ [How severity is assigned to suspicious activity](#), page 23
- ◆ [Reviewing threat incident details](#), page 23
- ◆ [Reviewing threat-related forensic data](#), page 25

Use the **Threats** tab of the **Status > Dashboard** page to monitor and investigate suspicious activity in your network.

- ◆ TRITON AP-WEB is required to display information about outbound threats and to provide detailed forensic data about the threats.
- ◆ You cannot add elements to, nor remove elements from, the Threats dashboard.

The initial view of the Threats dashboard shows:

- ◆ **Top Security Destinations** shows the top countries to which suspicious traffic is being sent, or in which sites associated with suspicious activity are hosted.
- ◆ **Security Events by Type** shows the number of blocked requests, permitted requests, or both for sites (destinations) in the top security categories associated with malware threats.
- ◆ **Suspicious Event Summary** lists information about threat-related events in your network.

A **Status** control in the top, right corner of the tab indicates whether Threats data is being updated automatically.

- ◆ If the status is **Running**, click **Pause** to prevent data from being updated while you examine current results.
- ◆ If the status is **Paused**, click **Start** to update the dashboard with any new data collected while updates were halted.

Additional controls at the top of the tab let you restrict the information in the charts and summary table to the specified:

- ◆ Time period (Today, 7 days, 30 days, and so on)
 - Date details under the drop-down list shows the start date and time used to calculate the selected period.
 - Configure the maximum time period available on the **Settings > Reporting > Dashboard** page (see [Configuring Dashboard reporting data](#), page 407).
With Microsoft SQL Server Express, the maximum time period is 30 days, and cannot be changed.

- ◆ Severities (Critical, High, Medium, or Low)
Click the **Severity Mapping** link for more information about the categories associated with each severity level.
- ◆ Action (All, Permitted, or Blocked)
- ◆ Direction (All, Inbound, or Outbound)

You can also use the Top Event Destinations map and Security Events by Category chart to further refine the information that appears in the summary table at the bottom of the page.

- ◆ Click a dot on the map to display only traffic associated with that country in the Suspicious Event Summary table.
The size of the dot reflects the number of incidents associated with that country. Hover over a dot to see a tooltip showing the country name. (Hovering over a blue area without a dot displays the name of the continent.)
- ◆ Click a category in the chart to display only traffic associated with that category in the table.
Each category is represented by a different color in the chart; hover over a bar or segment in the chart to see a tooltip showing the category name.

By default:

- ◆ The Top Event Destinations map shows the top 20 countries from which suspicious activity originates, or to which suspicious traffic is being sent.
- ◆ The Security Events By Category chart shows the top 5 categories associated with suspicious activity in the network, displayed in stacked column format.



To modify the information in the map or the chart:

- ◆ Click the **Options** icon, then select **Edit**.
- ◆ Use the **Top** list (both elements) or **Chart type** list (Security Events by Category chart) to update the display.
Changing the “top” value or chart type does not affect the information displayed in the summary table.

The Suspicious Event Summary table offers a variety of options to help you identify specific events to investigate.

- ◆ Use the Search box to find events for a user name, IP address, or hostname (if available; requires Content Gateway).
To stop filtering the table based on the term in the Search box, click **Clear**.
- ◆ Each of the filters (time, severity, action, direction, country, category) currently applied to the summary table is listed. Clear the check box next to a filter to remove it and expand the information shown in the table.
- ◆ Click a user name, IP address or hostname (if available) to see a detailed report. See [Investigate threat event details](#), page 21.

The Suspicious Event Summary can be customized to show or hide any of the following columns. The columns displayed by default are marked with an asterisk (*).

Column	Description
Severity*	Indicated by an “S” icon with a blue background (). Shows the severity (Critical, High, Medium, or Low) assigned to the event.
Forensics*	(<i>TRITON AP-WEB only</i>) Indicated by a magnifying glass icon (). Indicates whether the event included an attempt to send files.
User*	The user name (if any) associated with the activity.
IP address	The IP address of the machine on which the activity occurred.
Device*	(<i>TRITON AP-WEB only</i>) The name of the machine on which the activity occurred.
Category*	The Master Database category assigned to the activity.
Last Attempt*	The timestamp of the most recent event sharing all of the characteristics displayed in the row.
Country*	Indicated by the abbreviation “CC” (for country code). Shows the 2-letter country code for the event destination (target). If more than one destination is associated with an event, “Multiple” is displayed.
Direction	Whether the suspicious activity involved inbound or outbound traffic. Outbound threat detection requires TRITON AP-WEB.
Incidents*	The number of incidents sharing all of the characteristics displayed in the row except for “Last Attempt.”

To add columns to the chart, or to remove columns, click the **Customize** link above the table. Mark or clear the check box next to a column name to add or remove the column from the table.

To export the contents of the table to a CSV file, click **Export to CSV**. Select the time period for which to export event data, then click **Export**.

Investigate threat event details

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Dashboard > Threats > Event Details** page to research suspicious activity incidents. The page can show incidents related to:

- ◆ A specific user name, IP address, or device, selected from the Suspicious Event Summary table on the Threats dashboard. (Device names are provided by Content Gateway, and are not available in Web Filter & Security deployments.)

- ◆ A specific severity level, selected by clicking the link in a suspicious activity alert email notification (see [Configuring suspicious activity alerts](#), page 380).

At the top of the page, a table lists each incident associated with the selected user, IP address, hostname, or severity level. The table shows 10 rows of data per page.

- ◆ Use the Search field to narrow results to a specific incident or group of related incidents. Click **Clear** to remove the search filter.
- ◆ Refer to the information on the top, right portion of the page to see the time period covered in the table, and when the table was last updated.
- ◆ Click **Customize** in the toolbar at the top of the content pane to change the columns shown in the table. The detail table has the same column options as the summary table on the Threats dashboard.
- ◆ Click a row in the table to update the bottom portion of the page with additional details about the selected incident, its associated threats, and the detection methods used (see [Reviewing threat incident details](#), page 23).

The incident details section includes a link to Websense ACEInsight. Use this link to view current information about the URL and threats associated with the incident.

- ◆ If there are more than 10 incidents, use the paging controls at the bottom of the table to navigate through the data.

With TRITON AP-WEB, files associated with attempts to either infect your network or send sensitive data out of your network may be captured. File-related data is referred to collectively as **forensic data**, and it is stored in a special database, called the **forensics repository**.

- ◆ Forensics capture and storage is enabled by default.
- ◆ Configure forensics capture and storage on the **Settings > Reporting > Dashboard** page (see [Configuring Dashboard reporting data](#), page 407).

When forensics capture is enabled and there are files (like spreadsheets, documents, or compressed files) associated with an incident, an icon appears in the Forensics column of the Event Details table. When you select an incident that includes forensics data, information about any files associated with the incident is displayed in the Forensic Data section of the page (see [Reviewing threat-related forensic data](#), page 25).



Warning

Use caution when opening a file associated with a threat incident. If the file is infected with malware, it could infect the machine you use to investigate the incident.

Captured files may also contain sensitive data.

If a user agent header was captured for the incident, the User Agent String field includes a link that you can use to **Search for other instances** of the user agent. Click the link to see results on the Search tab of the **Reporting > Applications** page. See [Application reporting](#), page 156, for more information about application reports and user agents.

To export event information to a CSV file, click **Export** in the toolbar at the top of the content pane. All threat-related events logged in the selected **time period** are exported; not just those for the user, IP address, hostname, or severity level currently displayed on the page.

How severity is assigned to suspicious activity

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Websense Master Database assigns a severity level to threat-related events based on the category assigned to the request.

- ◆ Severity levels are mapped to categories in the Websense Master Database, and may change when the Master Database is updated.
- ◆ Websense Web Filter & Security subscriptions do not include all categories with High and Critical severity levels. These categories may appear on the Threats dashboard, but cannot be managed in category filters.

Click the **Severity Mapping** link near the top of the Threats dashboard for a current list of the categories that have an associated severity ranking. The list indicates any categories that are not available for policy enforcement with your subscription.

Reviewing threat incident details

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When an administrator selects an incident in the table at the top of the **Threats > Event Details** page, the area below the table is populated with all available details about the incident. The available details may vary based on:

- ◆ What type of incident occurred. For example:
 - An outbound request for a URL that is blocked by its Master Database category is unlikely to include a threat name, intent, or type, because the request is blocked before Content Gateway analysis occurs.
 - A request that does not include a file transfer does not include forensic data.
- ◆ Your subscription level. For example:
 - Only Content Gateway (a component of TRITON AP-WEB) passes hostname, threat name, threat intent, threat type, and scanning category information.
 - Not all Web Filter & Security integrations pass protocol, method, or content type information.
- ◆ Whether any file transfer attempts were associated with the incident. (Only Content Gateway provides this type of forensic data.) See [Reviewing threat-related forensic data](#), page 25.

The following incident details may be displayed on the page:

Field	Description
Severity	Critical, High, Medium, or Low. See How severity is assigned to suspicious activity , page 23.
Category	The Master Database or custom category assigned to the destination URL.
Threat Name	The name associated with the malicious software, bot traffic, or other threat activity (if applicable).
Threat Intent	What the threat would attempt to do (log keystrokes, open a back door into the network, and so on).
Platform	The operating system targeted by the threat (Windows, Android, and so on).
Threat Type	The classification of the malicious software (Trojan, worm, advanced persistent threat, and so on).
Action	The action assigned to the request (Permit or Block).
Reason	The reason the permit or block action was applied (for example, the category assigned to the URL).
Incident Time	The date and time the incident occurred.
ACEInsight Link	A link to ACEInsight.com to enable further research on the URL or threat.
User	The user requesting the URL (if a user is identified).
Source IP Address	The IP address from which the request originated.
Device	(<i>TRITON AP-WEB only</i>) The name of the machine from which the request originated. (When a hostname is not available, the source IP address is repeated).
Destination IP Address	The IP address of the requested URL.
Port	The port used to communicate with the requested URL.
Protocol	The protocol used to request the URL.
Direction	Whether the incident involved an inbound or outbound connection.
Method	Whether the request was a GET or a POST.
Content Type	The value reported in the “Content-Type” field of the HTTP header associated with the request (for example, text/html, image/gif, or application/javascript).
Bytes Sent	The number of bytes sent out from the source machine.
Bytes Received	The number of bytes returned by the target (destination) URL. If the request was blocked, this is 0.
Country	The country hosting the destination URL.
Full URL	The full URL (domain, path, CGI string, and file) of the target site.

Field	Description
Active Policy	The policy used to manage the request.
Database Category	The category assigned to the request by the Websense Master Database.
Scanning Category	The category assigned to the request by Content Gateway analysis (may match the Master Database category).
Role	The delegated administration role responsible for the policy used to manage the request.

Reviewing threat-related forensic data

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When an administrator selects an incident on the **Threats > Event Details** page that includes forensic data, the Forensic Data area below the table is populated with details about the attempted file transfer. Forensic details include:

Field	Description
Source	The user or IP address making the request.
Destination	The IP address of the target machine.
Incident ID	<p>The Websense TRITON AP-DATA ID number associated with the incident. Can be used to further investigate the incident in the Data module of the TRITON Manager.</p> <p>This feature requires either the Web DLP module or a TRITON AP-DATA solution.</p>
Files	<p>The name and size of the file or files associated with the incident. The file name is a link that can be used to open the actual file.</p> <p>WARNING: Use caution when opening a captured file. The file might contain malware that could infect the machine used for investigation. The file could also contain sensitive data.</p>
Parameters and Body	<p>Shows CGI parameters and HTML body details for the request used to send or retrieve the file.</p> <p>The number of parameters and the details included in the body of the request may vary widely from incident to incident.</p>

Risks dashboard

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [System dashboard](#), page 27
- ◆ [Threats dashboard](#), page 19
- ◆ [Usage dashboard](#), page 27
- ◆ [Adding elements to a dashboard](#), page 28

Use the **Risks** tab of the **Status > Dashboard** page to monitor permitted and blocked requests for URLs in the Security Risk class. By default, the following charts are displayed:

- ◆ **30-Day Risk Trends** shows blocked request trends for specific security and legal liability categories over a 30-day period that includes today. When you click a spark line:
 - For security-related categories (like Malicious), the Threats dashboard is displayed to allow further investigation.
 - For other categories (like Adult), an investigative report with more detailed information is displayed.
- ◆ **Clients with Security Risks** shows which computers have accessed Security Risk sites. You may want to check these machines to make sure they are not infected with any viruses or spyware.
- ◆ **Top Security Risk Categories** shows which Security Risk categories have received the most requests to help you determine whether your current policies are providing the right protection for your network.
- ◆ **Risk Classes** shows how many requests to each risk class have been permitted and blocked (see [Risk classes](#), page 36) to help you evaluate whether the current policies are effective.
- ◆ **Top Uncategorized** shows which URLs not categorized by the Websense Master Database have been accessed most. Go to **Filter Components > Edit Categories** to assign a URL to a category.
- ◆ (*TRITON AP-WEB only*) **Analytics: Security Risks** shows how many requests were assigned to new categories by Content Gateway analysis because the content had been changed or the site was compromised.

Click any chart on the page to open an investigative report with more detailed information.

Usage dashboard

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [System dashboard, page 27](#)
- ◆ [Risks dashboard, page 26](#)
- ◆ [Threats dashboard, page 19](#)
- ◆ [Adding elements to a dashboard, page 28](#)

Use the **Usage** tab of the **Status > Dashboard** page to monitor general Internet activity trends for your organization. By default, the following charts are displayed:

- ◆ **Top Blocked Users** shows which users have requested the most blocked sites.
- ◆ **Top Requested Categories** shows the categories that are being accessed most to provide a high-level overview of potential security, bandwidth, or productivity concerns. Click the chart to see an investigative report with more detailed information.
- ◆ **Enforcement Summary** provides an overview of recently permitted requests, blocked requests for sites in the Security Risk class, and other blocked requests.
- ◆ *(TRITON AP-WEB only)* **Web 2.0 Categories** shows the top categories assigned to requested Web 2.0 URLs, measured by requests.
- ◆ *(TRITON AP-WEB only)* **Web 2.0 URL Bandwidth** shows the Web 2.0 URLs using the most bandwidth.
- ◆ *(TRITON AP-WEB only)* **Analytics: Top Categories** shows the top categories to which requested URLs were assigned after real-time analysis determined that they no longer fit their original category.

Click any chart or element except the 30-Day Activity Summary to open an investigative report with more detailed information.

System dashboard

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Threats dashboard, page 19](#)
- ◆ [Risks dashboard, page 26](#)
- ◆ [Usage dashboard, page 27](#)
- ◆ [Adding elements to a dashboard, page 28](#)

Use the **System** tab of the **Status > Dashboard** page to monitor the status of your deployment. By default, the following dashboard elements are displayed:

- ◆ **Health Alert Summary** shows component alert and status messages. 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 381).

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

- ◆ **User Activity: Zoom Trend** shows the volume of Internet requests processed into the Log Database in the selected time period.

- Click and drag the cursor to select a section of the chart for closer examination. This can be done multiple times to select increasingly narrower time periods for review.

- At maximum zoom, a data point is shown for each 10 minute period (for example, 12:00:00, 12:10:00, 12:20:00).

In the chart default (macro) view is shown, each data point may be based on sampling of multiple 10-minute interval data points within the selected area of the chart. As a result, the numbers shown in the macro view may not correlate exactly to the numbers shown when the chart is zoomed in.

- Click **Zoom Out** to return to the previous level of focus.
- Click **Reset Chart** to return to the default level of detail.

- ◆ **Protocol Bandwidth Use** shows which protocols are using the most bandwidth in your network.

- ◆ **Filtering Service Status** 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 360.

- ◆ *(Requires the Web Hybrid module)* **Hybrid Bandwidth Summary** shows the bandwidth consumed by Internet requests managed by the hybrid service.
- ◆ *(Requires the Web Hybrid module)* **Hybrid Requests** shows how many requests by users from your organization were permitted and blocked by the hybrid service.

Adding elements to a dashboard

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Status > Dashboard > Add Chart** page to add elements to the Risk, Usage, or System dashboard.

Note that you can neither add elements to nor remove elements from the Threats dashboard.

To start, use the **Add elements to tab** drop-down list to select a tab, then select the element that you want to add from the **Dashboard Elements** list.

- ◆ You can add an element to any configurable tab.
- ◆ Each tab can show a maximum of 12 elements.
- ◆ Elements currently displayed on the selected tab are marked by a blue circle icon.
- ◆ You can add multiple copies of the same element to a tab (for example, each might show a different time period).

When you select an element in the list, a sample is displayed in the **Preview** pane. You can use the preview pane to make changes to the chart **Name** and, if applicable, **Chart type**, **Time period**, and **Top** value (for example, top 1-5 categories, or top 16-20 users).

- ◆ **Chart type:** Many charts can be displayed as a multi-series bar, column, or line chart, or as a stacked area or column chart. Some can be displayed as bar, line, or pie charts. Which types are available depends on the data being displayed.
- ◆ **Time period:** Most charts can display a variable time period: Today (the period since midnight of the current day), the last 7 days, or last 30 days. If the maximum time period for dashboard charts is extended, charts may also be able to show the last 180 or 365 days.
 - With Microsoft SQL Server Express, the maximum time period for dashboard charts is 30 days, and cannot be changed.
 - Using the default maximum time period (30 days) may improve dashboard performance.

See [Configuring Dashboard reporting data, page 407](#), for information about extending the time period for dashboard charts.
- ◆ **Top:** Charts displaying information about the top users, categories, URLs, and so on can display up to 5 values. Select whether to show the top 5 values, 6-10 values, 11-15 values, or 16-20 values.

When you are finished making changes, click **Add**. The dashboard tab is updated immediately.

If you have been editing a chart and would like to start over, click **Restore Defaults** to reset the chart to its default time period, type, and top value (if any).

Two dashboard elements do not appear on any tab by default, but are available to be added:

- ◆ **Activity Today** provides examples of how your software has protected your network today. The information varies based on your subscription. With TRITON AP-WEB, it includes information about sites analyzed by Content Gateway. This element 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.
- ◆ **30-Day Value Estimates** provide a way to estimate time and bandwidth savings afforded by your software over a 30-day period that includes today.

Mouse over the **Time** or **Bandwidth** item (under Saved) for an explanation of how the estimate was calculated. The calculation can be customized on the Add Charts page.

Status Monitor mode

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

- ◆ You must log off of other TRITON modules to enter Status Monitor mode in the Web module.
- ◆ In Status Monitor mode, information on the **Status > Dashboard**, **Status > Alerts**, **Status > Deployment**, and **Reporting > Real-Time Monitor** pages 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 module toolbar.
- ◆ Click the **Status Monitor** button in the toolbar at the top of the **Status > Dashboard** or **Status > Alerts** page.

To exit Status Monitor mode, log off of the console or close the browser.

3

Internet Usage Filters

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing access to categories and protocols, page 32](#)
- ◆ [Working with filters, page 42](#)
- ◆ [Configuring filtering settings, page 50](#)
- ◆ [Web Protection Policies, page 73](#)
- ◆ [Refine Your Policies, page 237](#)

Policies govern user Internet access. A policy is a schedule that determines how and when clients are able to access websites and Internet applications. At their simplest, policies consist of:

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



Note

If you have the Web Hybrid module, note that the hybrid service does not enforce protocol filters.

- ◆ A schedule that determines when each filter is enforced

Policies let 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, the **Default** policy is used to monitor Internet requests as soon as a subscription key is entered (see [The Default policy, page 74](#)). Initially, the Default policy permits all requests.

To apply different levels of access to different clients, start by defining category filters. You might define:

- ◆ One category filter that blocks access to all websites except those in the Business and Economy, Education, and News and Media categories

- ◆ A second category filter that permits all websites except those that represent a security risk and those containing adult material
- ◆ A third category filter that monitors access to websites without blocking them (see [Creating a category filter](#), page 43)

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 46)

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 [Web Protection Policies](#), page 73).

Managing access to categories and protocols

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [When a category or protocol is blocked](#), page 33
- ◆ [New Master Database categories and protocols](#), page 34
- ◆ [Special categories](#), page 34
- ◆ [Risk classes](#), page 36
- ◆ [Security protocol groups](#), page 39

The Websense Master Database organizes similar websites (identified by URLs and IP addresses) into **categories**. Each category has a descriptive name, like Adult Material, Peer-to-Peer File Sharing, or Spyware. You can also create your own, custom categories to group sites of particular interest to your organization (see [Creating a custom category](#), page 247). Together, the Master Database categories and user-defined categories form the basis for policy enforcement.

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 from websense.com.

To suggest that a site be added to the Master Database, or that a site be moved from one category to another, go to **support.websense.com** and click **Site Lookup Tool**. You are prompted to log on to MyWebsense, then given access to the tool, which allows you to verify the current category assigned to a site and request a new category.

When you create a **category filter**, 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 policies.

The up-to-date list of Master Database protocols is available from websense.com.

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

**Note**

With TRITON AP-WEB, it is possible to filter non-HTTP protocols that tunnel over HTTP ports using Content Gateway (see [Configuring tunneled protocol detection](#), page 92). You can also use Network Agent to enable policy enforcement for additional protocols.

In Websense Web Filter & Security deployments, Network Agent must be installed to enable protocol-based policy enforcement.

The hybrid service 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.

When a category or protocol is blocked

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When a user requests a URL in a blocked category, the browser displays a **block page**, rather than displaying the requested site. The block page is a customizable HTML page with a brief explanation of why the requested URL has been blocked.

See [Block Page Management](#), page 179, for a detailed description of the block page, along with information about customizing block pages.

When a user attempts to use an application that relies on a blocked protocol (for example, a chat or torrent program), no blocking message is displayed. The application may display an error message, or it may simply appear to hang.

To minimize error reports from users who are attempting to access blocked protocols, make sure that users understand which applications they are and are not allowed to use on your organization's equipment.

New Master Database categories and protocols

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When new categories and protocols are added to the Master Database, each is assigned a default action, like **Permit** or **Block** (see [Actions](#), page 40).

- ◆ The default action is applied in all active category and protocol filters (see [Working with filters](#), page 42). 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 changes](#), page 246, and [Making global protocol changes](#), page 258.
- ◆ The default action is based on feedback regarding whether or not the sites or protocols in question are generally considered business-appropriate.

You can have a system alert generated whenever new categories or protocols are added to the Master Database. See [Alerting](#), page 373, for more information.

Special categories

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Master Database contains special categories to help you manage specific types of Internet usage.

- ◆ The **Special Events** category is used to classify bandwidth-oriented content related to hot topics to help you manage event-related surges in Internet traffic. For example, the video pages offering live stream of the World Cup might generally appear in the Internet Radio and TV 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 **Security** category focuses on Internet sites containing malicious code, which can bypass virus-detection software programs.
 - Advanced Malware Command and Control (*TRITON AP-WEB only*)

- Advanced Malware Payloads (*TRITON AP-WEB only*)
- Bot Networks
- Compromised Websites
- Custom-Encrypted Uploads (*TRITON AP-WEB only*)
- Files Containing Passwords (*TRITON AP-WEB only*)
- Keyloggers
- Malicious Embedded iFrame
- Malicious Embedded Link
- Malicious Websites
- Mobile Malware
- Phishing and Other Frauds
- Potentially Exploited Documents (*TRITON AP-WEB only*)
- Potentially Unwanted Software
- Spyware
- Suspicious Embedded Link
- Unauthorized Mobile Marketplaces
- ◆ The **Productivity** category focuses on preventing time-wasting behavior.
 - Advertisements
 - Application and Software Download
 - Instant Messaging
 - Message Boards and Forums
 - Online Brokerage and Trading
 - Pay-to-Surf
- ◆ The **Bandwidth** category focuses on saving network bandwidth.
 - Educational Video
 - Entertainment Video
 - Internet Radio and TV
 - Internet Telephony
 - Peer-to-Peer File Sharing
 - Personal Network Storage and Backup
 - Streaming Media
 - Surveillance
 - Viral Video
- ◆ The **Extended Protection** category focuses on potentially malicious websites.
 - **Dynamic DNS** includes sites that mask their identity using Dynamic DNS services, often associated with advanced persistent threats.
 - **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.
- **Newly Registered Websites**
- **Suspicious Content** includes sites likely to contain little or no useful content.

The Extended Protection group filters potentially malicious websites 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Assigning categories to risk classes, page 384](#)
- ◆ [Presentation reports, page 113](#)
- ◆ [Investigative reports, page 134](#)

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 dashboards 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.

There are 5 risk classes. By default, each risk class contains the categories listed below.

- ◆ 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. When you receive notice that a new category has been added to the Master Database, it is a good idea to check its default risk class assignment.

Legal Liability

Adult Material (includes Adult Content, Lingerie and Swimsuit, Nudity, Sex)
Bandwidth > Peer-to-Peer File Sharing
Gambling
Illegal or Questionable
Information Technology > Hacking and Proxy Avoidance
Intolerance
Militancy and Extremist
Tasteless
Violence
Weapons

Network Bandwidth Loss

Bandwidth (includes Educational Video, Entertainment Video, Internet Radio and TV, Internet Telephony, Peer-to-Peer File Sharing, Personal Network Storage and Backup, Streaming Media, Surveillance, Viral Video)
Entertainment > Media File Download Services
Productivity > Advertisements, Application and Software Download
Social Web Controls - Facebook > Facebook Video Upload
Social Web Controls - YouTube > YouTube Video Upload

Business Usage

Bandwidth > Educational Video
Business and Economy (includes Financial Data and Services, Hosted Business Applications)
Collaboration - Office (includes Office - Mail, Office - Drive, Office - Documents, and Office - Apps)
Education > Educational Materials, Reference Materials
Government (includes Military)
Social Web Controls - LinkedIn (includes LinkedIn Connections, LinkedIn Jobs, LinkedIn, Mail, LinkedIn Updates)
Information Technology (includes Computer Security, Search Engines and Portals, Web Collaboration, Web Translation, Web Analytics, and Web and Email Marketing)
Travel
Vehicles

Security Risk

Bandwidth > Peer-to-Peer File Sharing

Extended Protection (includes Dynamic DNS, Elevated Exposure, Emerging Exploits, Newly Registered Websites, Suspicious Content)

Information Technology > Hacking, Proxy Avoidance, Web and Email Spam

Parked Domain

Productivity > Application and Software Download

Security (includes Bot Networks, Compromised Websites, Keyloggers, Malicious Embedded iFrame, Malicious Embedded Link, Malicious Web Sites, Phishing and Other Frauds, Potentially Unwanted Software, Spyware, Suspicious Embedded Link)

(TRITON AP-WEB only) Also includes Advanced Malware Command and Control, Advanced Malware Payloads, Custom-Encrypted Uploads, Files Containing Passwords, and Potentially Exploited Documents.

Productivity Loss

Abortion (includes Pro-Choice, Pro-Life)

Adult Material > Sex Education

Advocacy Groups

Bandwidth > Entertainment Video, Internet Radio and TV, Peer-to-Peer File Sharing, Streaming Media, Surveillance, Viral Video

Collaboration - Office (includes Office - Mail, Office - Drive, Office - Documents, and Office - Apps)

Drugs (includes Abused Drugs, Marijuana, Prescribed Medications, Nutrition)

Education (includes Cultural Institutions, Educational Institutions)

Entertainment (includes Media File Download Services)

Gambling

Games

Government > Political Organizations

Health

Information Technology > Web and Email Spam, Web Hosting, and Web and Email Marketing

Internet Communication (includes General Email, Organizational Email, Text and Media Messaging, Web Chat)

Job Search

News and Media (includes Alternative Journals)

Parked Domain

Productivity (includes Application and Software Download, Instant Messaging, Message Boards and Forums, Online Brokerage and Trading, Pay-to-Surf)

Religion (includes Non-Traditional Religions, Traditional Religions)

Shopping (includes Internet Auctions, Real Estate)

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

Productivity Loss

Social Web Controls - Facebook (includes Facebook Apps, Facebook Chat, Facebook Commenting, Facebook Events, Facebook Friends, Facebook Games, Facebook Groups, Facebook Mail, Facebook Photo Upload, Facebook Posting, Facebook Questions, Facebook Video Upload)

Social Web Controls - LinkedIn (includes LinkedIn Connections, LinkedIn Jobs, LinkedIn Mail, LinkedIn Updates)

Social Web Controls - Twitter (includes Twitter Follow, Twitter Mail, Twitter Posting)

Social Web Controls - Various (includes Blog Commenting, Blog Posting, Classifieds Posting)

Social Web Controls - YouTube (includes YouTube Commenting, YouTube Sharing, YouTube Video Upload)

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

Special Events

Sports (includes 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 384).

Security protocol groups

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

In addition to the Security and Extended Protection categories, there are 2 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 47). The Malicious Traffic (Cannot block) protocols can be logged for reporting, but no other action can be applied.

Actions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Category and protocol filters assign an **action** to each category or protocol. This is the action that Websense components take 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 261.

Additional actions can be applied only to categories.

- ◆ **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 41, for more information.



Important

In multiple Filtering Service deployments, Websense State Server is required for correct application of the Confirm and Quota actions. See [Policy Server, Filtering Service, and State Server](#), page 363, 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 [Keyword-based policy enforcement](#), page 248.
- ◆ **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 263.

Using quota time to limit Internet access

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.

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, Filtering Service uses a priority list to determine how to respond when a user requests a site in a quota category. It 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, quota time is allotted according to the **Use more restrictive blocking** setting on the **Settings > General > Filtering** page (see [Configuring filtering settings](#), page 50).

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, Filtering Service sees the communication, then blocks the request according to applicable quota restrictions.

Search filtering

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, your 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, search engine feature stops thumbnail images associated with blocked sites from being displayed in search results.

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 373).

Search filtering is activated via the **Web > Settings > General > Filtering** page in the TRITON Manager. See [Configuring filtering settings](#), page 50, for more information.

Working with filters

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing access to categories and protocols](#), page 32
- ◆ [Web Protection Policies](#), page 73
- ◆ [Creating a category filter](#), page 43
- ◆ [Creating a protocol filter](#), page 46
- ◆ [Creating a limited access filter](#), page 239

Use the **Policy Management > Filters** page to view, create, and modify category, protocol, and limited access filters.

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.
Although Network Agent must be installed to enable full protocol-based policy enforcement, with TRITON AP-WEB, Content Gateway can manage non-HTTP protocols that tunnel over HTTP ports. See [Configuring tunneled protocol detection](#), page 92, for more information.
The hybrid service does not provide protocol-based policy enforcement.
- ◆ **Limited Access Filters** define a restrictive list of permitted websites (see [Restricting users to a defined list of URLs](#), page 238).

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 44

- [Editing a protocol filter](#), page 47
- [Editing a limited access filter](#), page 240
- ◆ To create a new category, protocol, or limited access filter, click **Add**. For more information, see:
 - [Creating a category filter](#), page 43
 - [Creating a protocol filter](#), page 46
 - [Creating a limited access filter](#), page 239

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 311), 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 242, for more information.

Creating a category filter

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with filters](#), page 42
- ◆ [Editing a category filter](#), page 44
- ◆ [When a category or protocol is blocked](#), page 33

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 49.

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing access to categories and protocols](#), page 32
- ◆ [Actions](#), page 40
- ◆ [Using quota time to limit Internet access](#), page 41
- ◆ [Working with filters](#), page 42
- ◆ [Working with categories](#), page 244

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 action associated with the selected category.
2. Before making changes to the action applied to a category, use the details section (to the right of the Categories list) to review any special attributes associated with the category.
 - To list recategorized URLs assigned to the category, if any, click **See custom URLs in this category**. See [Reclassifying specific URLs](#), page 251.

- To list keywords assigned to the category, click **See keywords in this category**. See [Keyword-based policy enforcement](#), page 248.
 - To list regular expressions used to define custom URLs or keywords for the category, click **See regular expressions in this category**.
3. Use the buttons to the right of the category list to change the action applied to the selected category. For more information about the available actions, see [Actions](#), page 40.

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 actions to the selected category:
- To change the way that keywords are used for assigning requests to the selected category, mark or clear **Block keywords**. See [Keyword-based policy enforcement](#), page 248
 - 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 263.

If you have chosen to block file types, select one or more file types to block.

To apply the selected file type settings to all permitted categories in the filter, click **Apply to All Categories**.



Warning

With TRITON AP-WEB, applying file type blocking to all categories may have a serious performance impact.

All files with an extension that does not match the blocked type are scanned to find their true file type, including text files, like HTML and CSS files.

- 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 261.
- If you have chosen to block based on bandwidth, specify which threshold limits to use.
- To apply the selected bandwidth settings to all permitted categories in the filter, click **Apply to All Categories**.
5. Repeat steps 1 through 3 to make changes to the 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 and Deploy**.

To activate a new category filter, add it to a policy and assign the policy to clients. See [Web Protection Policies](#), page 73.

Creating a protocol filter

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing access to categories and protocols](#), page 32
- ◆ [Actions](#), page 40
- ◆ [Editing a protocol filter](#), page 47
- ◆ [Protocol-based policy enforcement](#), page 255
- ◆ [When a category or protocol is blocked](#), page 33

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 49) 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing access to categories and protocols](#), page 32
- ◆ [Creating a protocol filter](#), page 46
- ◆ [Actions](#), page 40
- ◆ [Protocol-based policy enforcement](#), page 255
- ◆ [Using Bandwidth Optimizer to manage bandwidth](#), page 261

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 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 261](#), 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 and Deploy**.

To activate a new protocol filter, add it to a policy and apply the policy to clients (see [Web Protection Policies, page 73](#)).

**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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Your software includes several sample category and protocol filters. You can use these filters as they are, or modify them. 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
- ◆ Strict Security

The Block All and Permit All category filters are not listed on the Filters page, though they can be added to policies. These filters are handled differently than the others, and cannot be deleted or edited. When Filtering Service receives an Internet request, it first checks to see if the Block All or Permit All filter applies, before performing any additional checks (see [Responding to a URL request](#), page 81).

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 during the policy enforcement process.

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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**.

Your software includes 7 category filter templates:

- ◆ **Monitor Only** and **Permit All** permit 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.
- ◆ **Strict Security** extends the Default template by blocking 2 additional security categories, and adding file-type blocking for executables to a third category.

- ◆ **Basic Security** blocks only the default categories in the Security Risk class (see [Risk classes](#), page 36).

Your software also includes 4 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 and Malicious Traffic.
- ◆ **Default** blocks the Instant Messaging / Chat protocols, as well as the P2P File Sharing, Proxy Avoidance, Instant Messaging File Attachments, and Malicious Traffic.

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 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 policy 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 43, or [Creating a protocol filter](#), page 46.

Configuring filtering settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related Topics:

- ◆ [Managing access to categories and protocols](#), page 32
- ◆ [Block Page Management](#), page 179
- ◆ [Password override](#), page 69
- ◆ [Account override](#), page 70
- ◆ [Using Bandwidth Optimizer to manage bandwidth](#), page 261
- ◆ [Keyword-based policy enforcement](#), page 248

Use the **Settings > General > Filtering** page to establish basic settings for how Internet requests are handled.

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

1. To determine how user requests are handled when multiple group policies apply, mark or clear **Use most restrictive group policy** (see [Enforcement order](#), page 79).
 - When the option is selected, the policy that applies the most restrictive action 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 [Keyword-based policy enforcement](#), page 248).

CGI only	Blocks sites when keywords appear in CGI query strings (after the “?” in a Web address). Example: search.yahoo.com/search?p=test Filtering Service 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, Filtering Service 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, Filtering Service 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 69).
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 [Actions](#), page 40).
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 70).
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 41).
A session begins when the user clicks the Use Quota Time button.
7. Enter the **Default quota time per day** (up to 1440 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.

Under **State Server**, provide **IPv4 address or hostname** and **Port** information if:

- ◆ Your environment includes multiple Websense Filtering Service instances, and
- ◆ You use the Quota or Confirm actions, password override, or account override.

State Server tracks clients' quota, confirm, password override, and account override sessions to ensure that session time is allocated correctly across multiple Filtering Service instances (see [Policy Server, Filtering Service, and State Server](#), page 363).

After entering State Server connection details, click **Check Status** to verify the connection. Configure State Server connection information for each Policy Server instance in your deployment.

Under **Bandwidth Optimizer**, enter the information needed to filter Internet usage based on available bandwidth. For more information about enforcing bandwidth-based Internet access, see [Using Bandwidth Optimizer to manage bandwidth](#), page 261.

**Note**

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 actions are enforced. Note that when the thresholds are set, but no category or protocol filters include bandwidth-base actions, no bandwidth usage restriction occurs.
 - **Network:** When total network traffic reaches this percentage of total available bandwidth, start limiting access based on bandwidth, 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. (*TRITON AP-WEB only*) 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 **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 188), or to configure TRITON AP-WEB to include a link to ACEInsight on block pages.

- ◆ 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.

If you have created custom block pages, and want to use those block pages for all protocols, you can also use the fields in this section blank (see [Customizing the block message, page 184](#)).

- ◆ With the Web Hybrid module for TRITON AP-WEB, custom block messages specified in the fields above are not applied to requests handled by the hybrid service.

Instead, use the **Settings > Hybrid Configuration > User Access** page to customize the hybrid block page (see [Customizing hybrid block pages, page 205](#)).

- ◆ (*TRITON AP-WEB only*) When a user clicks the ACEInsight link, the URL the user attempted to access is sent to ACEInsight and a web page is displayed showing ACEInsight analysis.

The URL sent to ACEInsight is truncated, to omit the CGI string (which could include a user name or password). As a result, ACEInsight does not analyze password-protected content, and may return different results than Content Gateway.

The ACEInsight link does not appear on hybrid block pages.

Under **Search Filtering**, select **Enable search filtering** to 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 41](#)).

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

When you have finished configuring settings on this page, click **OK** to cache the changes. Changes are not implemented until you click **Save and Deploy**.

4

Web Protection Clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

You can customize how your Websense software manages requests from specific users or machines by adding them as **clients** in the Web module of the TRITON Manager. 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

The hybrid service can apply policies to users or groups, and to filtered locations, but not to individual clients or networks. See [Working with hybrid service clients](#), page 72.

Initially, all client requests are managed by the **Default** policy (see [The Default policy](#), page 74). Once you add a client to the Clients page, you can assign that client a specific policy.

When multiple policies could apply, such as when one policy is assigned to the user and another is assigned to the machine, by default, Websense Filtering Service uses the following enforcement order:

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 Filtering Service considers all group policies that apply (see [Enforcement order](#), page 79).
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 Filtering Service processes requests, see [Responding to a URL request](#), page 81.

For information about configuring Filtering Service to prioritize group and domain policies over IP address-based (computer and network) policies, see [Prioritizing group and domain policies](#), page 80.

For information about how the hybrid service applies policies to clients, see [Enforcement order](#), page 79.

Working with clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Clients](#), page 55
- ◆ [Working with computers and networks](#), page 57
- ◆ [Working with users and groups](#), page 58
- ◆ [Adding a client](#), page 66
- ◆ [Changing client settings](#), page 68

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 66, 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 58).
- ◆ **Networks**, IPv4 or IPv6 address ranges within the network that can be governed by a single policy (see [Working with computers and networks](#), page 57).
- ◆ **Computers**, individual machines in the network, identified by IPv4 or IPv6 address (see [Working with computers and networks](#), page 57).

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 [Web Protection Policies](#), page 73).

- ◆ Whether or not the client can use a **password override** (see [Password override, page 69](#)) or **account override** (see [Account override, page 70](#)) 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 41](#)).

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 68](#), 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 66](#), for more information.

If you have created delegated administration roles (see [Delegated Administration and Reporting, page 311](#)), 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 71](#), for more information.

If you have configured Websense User Service 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 65](#)).

To remove a client from the Clients page, select the client and click **Delete**.

Working with computers and networks

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with clients, page 56](#)
- ◆ [Working with users and groups, page 58](#)
- ◆ [Adding a client, page 66](#)
- ◆ [Assigning a policy to clients, page 78](#)

In the Web module of the TRITON Manager, a **computer** is the IP address (for example, 10.201.3.1 or fd3a:918a:71a1:bcaa::0011) associated with a filtered machine. A **network** is the IP address range (for example, 10.201.3.2 - 10.201.3.44 or fd3a:918a:71a1:bcaa::1111 - fd3a:918a:71a1:bcaa::1211) that corresponds to a group of filtered machines.

- ◆ If you have the Web Hybrid module, note that the hybrid service does not apply policies to individual computer and network clients. See [Working with hybrid](#)

[service clients](#), [page 72](#), for information about applying policies to filtered locations.

- ◆ Before applying policies to IPv6 computer and network clients, disable temporary IPv6 addresses on the affected machines. See support.websense.com for details.

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 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 client machine, **unless** you have assigned a policy to the logged-on user. When on-premises components enforce policy, the computer or network policy takes precedence over any group policies that may apply to the user. (The hybrid service applies the group policy before applying a computer or network policy. See [Working with hybrid service clients](#), [page 72](#).)

Working with users and groups

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with clients](#), [page 56](#)
- ◆ [Connecting web protection software to a directory service](#), [page 59](#)
- ◆ [Working with custom LDAP groups](#), [page 65](#)
- ◆ [Working with computers and networks](#), [page 57](#)
- ◆ [Adding a client](#), [page 66](#)
- ◆ [Assigning a policy to clients](#), [page 78](#)

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

User Service 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 Active Directory or Oracle Directory Server, user names with blank passwords are not supported. Assign passwords to all users.

User Service conveys information from the directory service to Filtering Service for use in applying policies. As a best practice, install User Service on a Windows machine (though it can reside on Linux).

To configure directory service communication, see [Connecting web protection software to a directory service](#).

Connecting web protection software to a directory service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

A directory service stores information about a network's users and resources. Before you can add directory clients (users, groups, or OUs) in the Web module of the TRITON Manager, you must configure Websense User Service 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

If you have the Web Hybrid module, information from the Directory Services page is also used to populate the **Hybrid Configuration > Shared User Data** page. This allows the hybrid service to apply user and group-based policies. See [Send user and group data to the hybrid service](#), page 209.

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:

- ◆ [Connecting to Windows Active Directory \(Mixed Mode\)](#), page 60
- ◆ [Connecting to Windows Active Directory \(Native Mode\)](#), page 61
- ◆ [Connecting to Novell eDirectory or Oracle Directory Server](#), page 62



Important

If you have the Web Hybrid module, the hybrid service supports Windows Active Directory (Native Mode), Oracle Directory Server, and Novell eDirectory.

Once configuration is complete, User Service communicates with the directory service so that users, groups, and OUs can be added as clients and assigned policies.

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 policy enforcement 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 the TRITON Manager, 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.

Connecting to Windows Active Directory (Mixed Mode)

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Supply additional information on this screen if Websense User Service resides on a Websense appliance or a Linux server, and either of the following is true:

- ◆ DC Agent is being used for transparent identification (see [DC Agent](#), page 288)
- ◆ Logon Agent is being used for transparent identification with Active Directory in **native** mode.

If this matches your configuration, User Service and DC Agent or Logon Agent must communicate with a Windows Internet Name Server (WINS) to resolve domain names to domain controller IP addresses (see [User Service on a Websense appliance or Linux server](#), page 446).

To enable that communication, use the fields under Windows Active Directory (Mixed Mode) to provide:

1. The account name of an **Administrative user** that can access the directory service.
2. The **Password** for the account.
3. **Domain** information for the account.
4. The IP address or hostname of a **WINS server** in your network.

Note that:

- ◆ If you are using DC Agent, you can also perform these steps on the **Settings > User Identification > DC Agent** page when you are configuring a DC Agent instance. You do not need to perform the configuration in both places.
- ◆ If you are using Logon Agent, even though you are connecting to the directory in native mode, you must perform these steps on the Windows Active Directory (Mixed Mode) page (including caching and saving your changes). Then, once the WINS connection is configured, you can complete your directory service setup on the Windows Active Directory (Native Mode) page.

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

Connecting to Windows Active Directory (Native Mode)

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x



Important

If User Service resides on a Websense appliance or Linux server, and you are using Logon Agent to identify users, set up a WINS server connection on the Active Directory (Mixed Mode) page first (and save your changes). Then return to the Windows Active Directory (Native Mode) page and configure your directory service connection.

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 User Service 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. Provide the **IPv4 address or hostname** of 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 IPv4 address or hostname (if name resolution is enabled in your network) of the server to add.
3. Enter the **Port** that User Service should use to communicate with the global catalog (by default, **3268**).
4. Optionally, enter the **Root context** for User Service to use when associating user and group information with Internet requests. Note that this context is used for policy enforcement, but not for adding clients in the management console.
 - 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 there is no root context, User Service begins searching at the top level of the directory service.
 - If the specified port is 389 or 636, you must provide a root context.



Note

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

5. Specify which administrative account User Service 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 **Test Connection** to verify that User Service can connect to the directory using the account information provided.
 7. Click **OK** to return to the Directory Services page.
 8. Repeat the process above for each global catalog server.
 9. Click **Advanced Directory Settings**, and then go to [Advanced directory settings](#), page 63.

Connecting to Novell eDirectory or Oracle Directory Server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

To retrieve information from the directory, User Service needs the distinguished name, root context, and password for a user account with administrative privileges.

1. Enter the **IPv4 address or hostname** of the directory server.
2. Enter the **Port** number that User Service 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**.
4. Enter the **Root Context** that User Service should use when searching for user information. For example, *o=domain.com*.
 - Providing a root context is mandatory for Oracle Directory Server, but optional for Novell eDirectory.

- Narrowing the context increases speed and efficiency in retrieving user information.
- User Service uses the context when searching for user and group information to aid in policy enforcement. It is not used for adding clients to the management console.

**Note**

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

5. Provide a **Password** for the administrator account entered above.
6. Click **Test Connection** to verify that User Service can connect to the directory server using the information provided.
7. Click **Advanced Directory Settings**, and then go to [Advanced directory settings](#), page 63.

Advanced directory settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Connecting to Windows Active Directory \(Native Mode\)](#), page 61
- ◆ [Connecting to Novell eDirectory or Oracle Directory Server](#), page 62

These settings can be used to define:

- ◆ How Websense User Service searches the directory service to find user, group, and domain information
- ◆ Whether User Service uses an encrypted connection to communicate with the directory service
- ◆ Which character set User Service 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, mark **Use custom filters**. The default filter strings are listed below the check box.
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
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 User Service and your directory service, check **Use SSL**.
4. To determine which character set User Service 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 and Deploy**.

Working with custom LDAP groups

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with users and groups, page 58](#)
- ◆ [Connecting web protection software to a directory service, page 59](#)
- ◆ [Adding or editing a custom LDAP group, page 66](#)

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 User Service to communicate with an LDAP-based directory service.



Important

When you add custom LDAP groups, 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, connect to each Policy Server in turn 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 66](#)).
- ◆ 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 and Deploy**.

Adding or editing a custom LDAP group

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Add Custom LDAP Group** page to define a group 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 and Deploy**.

Adding a client

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with clients, page 56](#)
- ◆ [Working with computers and networks, page 57](#)
- ◆ [Working with users and groups, page 58](#)
- ◆ [Searching the directory service from the TRITON Manager, page 67](#)
- ◆ [Changing client settings, page 68](#)

Use this page to add user, group, computer, and network clients to:

- ◆ Your Clients page, so that you can assign them a policy (**Clients > Add Clients**)

- ◆ A policy exception that blocks or permits specific URLs (**Exceptions > Add Other Clients to Exception**)

If you are logged on to a delegated administration role, you can only add clients that appear in your managed clients list to the Clients page or exception.

In policy management and reporting roles, the process of adding managed clients to the Clients page requires assigning them a policy. (Investigative reporting roles do not have this requirement.)

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 from the TRITON Manager, page 67](#)).
 - To add a computer or network client, enter an **IP address** or **IP address range** in either IPv4 or IPv6 format.
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. If you are adding clients to the Clients page, 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 and Deploy**.

The clients you selected are displayed either on the Clients page or in your exception.

After adding clients to the Clients page, you can select one or more client entries and click **Edit** to change policy assignments and other client configuration settings. See [Changing client settings, page 68](#), for more information.

Searching the directory service from the TRITON Manager

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you have configured Websense User Service to communicate with an LDAP-based directory service, you can use a search function to find the directory clients you want to identify for policy or exception assignment.

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 OU **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. Use the **Search for** list to specify how to perform the search:
 - Select **Entries containing search string** to find all directory entries that contain the search term you entered.
 - Select **Exact search string only** to find only the directory entry that precisely matches the search term.
5. Browse the **Search Context** tree to specify which portion of the directory to search. A more precise context helps to speed the search.
6. Click **Go**.

A list of search results is displayed.
7. Select one or more entries in search results, and then click the right arrow (>) to add each selection as a client or administrator.
 - Click **New Search** to enter another set of search criteria.
 - Click **Browse** to stop using search and instead navigate through the directory tree to identify users.
8. When you are finished making changes, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Changing client settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.
 - (*Super Administrators only*) 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 69.

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 70.

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 and Deploy**.

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

Password override

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Password override lets clients in the Super Administrator role that have valid passwords access sites in blocked categories. Password override can be granted to individual users, groups, computers, or networks, but not OUs.

When a Super 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.

This option is not available to delegated administrators, because it would effectively provide a method for overriding the Filter Lock (see [Creating a Filter Lock](#), page 320).



Important

In multiple Filtering Service deployments, Websense State Server is required for correct allocation of password override time. See [Policy Server, Filtering Service, and State Server](#), page 363, for more information.

Configure how long clients with password override privileges can access blocked sites per password entry on the **Settings > General > Filtering** page (see [Configuring filtering settings](#), page 50).

Grant password override privileges to specific clients via the **Policy Management > Clients** page (see [Adding a client](#), page 66, or [Changing client settings](#), page 68).

Account override

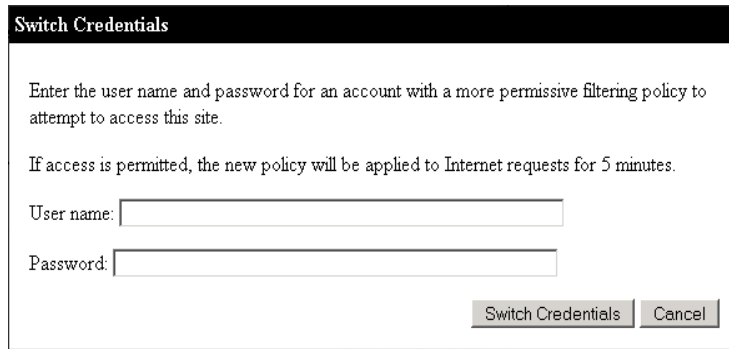
Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 OUs).

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.

A screenshot of a web-based dialog box titled "Switch Credentials". The dialog has a black header bar with the title in white. Below the header, the text reads: "Enter the user name and password for an account with a more permissive filtering policy to attempt to access this site." followed by "If access is permitted, the new policy will be applied to Internet requests for 5 minutes." There are two input fields: "User name:" and "Password:". At the bottom right, there are two buttons: "Switch Credentials" and "Cancel".

Switch Credentials

Enter the user name and password for an account with a more permissive filtering policy to attempt to access this site.

If access is permitted, the new policy will be applied to Internet requests for 5 minutes.

User name:

Password:

Once the user clicks **Switch Credentials**, Websense Filtering Service 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 filtering settings](#), page 50.



Important

In multiple Filtering Service deployments, Websense State Server is required for correct allocation of account override time. See [Policy Server, Filtering Service, and State Server](#), page 363, for more information.

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 or OU is added to a role as a managed client, delegated administrators in that role can assign policies to individual users in the group or OU.
- ◆ 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 popup message indicates that the selected clients are being moved. The move process may take a while.

3. Changes are not implemented until you click **Save and Deploy**.

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

Working with hybrid service clients

Administrator Help | TRITON AP-WEB | Version 8.1.x

If you have the Web Hybrid module, the hybrid service can manage 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.

The hybrid service can apply policies (created in the Web module of the TRITON Manager) 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 and authentication of hybrid users](#), page 300).

- ◆ 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.

The hybrid service does **not** apply policies to individual client machines in your network

Directory clients (users, groups, and OUs) managed by the hybrid service are identified on the **Policy Management > Clients** page, just like those whose requests are managed 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 hybrid service filtered locations](#), page 195).
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 57).
3. Apply a policy to the IP address or range.

Any time no user, group, or location policy applies, the Default policy is used.

5

Web Protection Policies

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Internet Usage Filters, page 31](#)
- ◆ [Web Protection Clients, page 55](#)
- ◆ [The Default policy, page 74](#)
- ◆ [Working with policies, page 74](#)
- ◆ [Enforcement order, page 79](#)

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

- ◆ Category filters, used to apply actions (permit, block) to URL categories (see [Managing access to categories and protocols, page 32](#))
- ◆ Limited access filters, used to permit access to only a restricted list of URLs (see [Restricting users to a defined list of URLs, page 238](#))
- ◆ Protocol filters, used to apply actions to Internet protocols (see [Managing access to categories and protocols, page 32](#))
- ◆ A schedule that determines when each category or limited access filter and protocol filter is enforced

Your software includes 3 predefined policies:

- ◆ **Default** filters Internet access for all clients not governed by another policy. This policy becomes active as soon as you enter a subscription key (see [The Default policy, page 74](#)).
- ◆ **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 Internet access at different times. This policy is used in the New Admin 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Policies, page 73](#)
- ◆ [Working with policies, page 74](#)
- ◆ [Enforcement order, page 79](#)

After installation, when you enter a valid subscription key, the **Default** policy begins monitoring Internet activity. Initially, the Default policy permits all requests.

As you create and apply additional policies, the Default policy continues to govern Internet access for any clients not assigned another policy.

The Default policy must provide coverage (enforce a combination of category or limited access filters and protocol filters) 24 hours a day, 7 days a week.

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

Working with policies

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Policies, page 73](#)
- ◆ [Creating a policy, page 75](#)
- ◆ [Editing a policy, page 76](#)
- ◆ [Internet Usage Filters, page 31](#)
- ◆ [Refine Your Policies, page 237,](#)

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 75](#), for more information.

- ◆ To edit a policy, click the policy name in the list, and then see [Editing a policy, page 76](#), 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 popup window.

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 policy information.

If you have created delegated administration roles (see [Delegated Administration and Reporting, page 311](#)), 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 242](#), for more information.

Creating a policy

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Policies, page 73](#)
- ◆ [Working with policies, page 74](#)
- ◆ [Editing a policy, page 76](#)
- ◆ [Working with filters, page 42](#)
- ◆ [Restricting users to a defined list of URLs, page 238](#)

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 76.

Editing a policy

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Policies](#), page 73
- ◆ [Working with policies](#), page 74
- ◆ [Creating a policy](#), page 75
- ◆ [Working with filters](#), page 42
- ◆ [Restricting users to a defined list of URLs](#), page 238

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 78.

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 filters 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 demonstrates how to define a time 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 43](#), or [Creating a limited access filter, page 239](#), for instructions.

5. 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 46](#), for instructions.

6. 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 Actions, page 40). 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 URLs, page 238). 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 Actions, page 40). 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 and Deploy**.

Assigning a policy to clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web Protection Policies](#), page 73
- ◆ [Creating a policy](#), page 75
- ◆ [Editing a policy](#), page 76
- ◆ [Web Protection Clients](#), page 55
- ◆ [Adding a client](#), page 66

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.

1. Mark the check box next to each client that you want to assign to the selected policy.
2. Click **OK** to return to the Edit Policy page.
3. Click **OK** again to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Enforcement order

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

For each request, Websense components:

1. Verify subscription compliance, making sure that the subscription is current.
2. Determine which exception or policy applies, searching in this order:
 - On-premises software (Websense Filtering Service):
 - a. Policy or exceptions assigned to the **user**
 - b. Policy or exceptions assigned to the **IP address** (computer or network) of the machine being used
 - c. Policies or exceptions assigned to **groups** the user belongs to
 - d. Policies or exceptions assigned to the user's **domain (OU)**
 - e. The **Default** policy



Note

You can configure Filtering Service to prioritize group and domain-based policies over IP address-based policies, if needed. See [Prioritizing group and domain policies](#), page 80.

- For users whose requests are managed by the hybrid service:
 - a. Policy or exceptions assigned to the **user**
 - b. Policy or exceptions assigned to **groups** the user belongs to
 - c. Policy or exceptions assigned to the user's **domain (OU)**
 - d. Policy or exceptions assigned to the external **IP address** (filtered location) from which the request originates
 - e. The **Default** policy

The first applicable exception or policy found is used.

3. Filter the request according to the exception or 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, Websense components check the policies assigned to each of the user's groups.

- ◆ If all the groups have the same policy, Websense software enforces that policy.
- ◆ If one of the groups has a different policy, Websense software uses the **Use more restrictive blocking** selection on the **Settings > General > Filtering** page to determine which policy to enforce.
 - If **Use more restrictive blocking** is checked, and any of the applicable policies blocks access to the requested category, the site is blocked.
 - If the option is not checked, and any of the applicable policies permits access to the requested category, the site is permitted.

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 enforcement order](#), page 238.

- ◆ 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.

Prioritizing group and domain policies

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

In some cases, organizations may prefer that policies applied to users, groups, and OUs take precedence over policies applied to IP addresses (computers and networks).

This might occur, for example, if both of the following are true:

1. Group-based policies are used widely in the organization.
2. The Account Override option (see [Account override](#), page 70) is applied to IP addresses in the network.

When the default enforcement order is used, the IP address-based policy overrides any group-based policies, which could cause account override to fail frequently. When group and domain policies take precedence, the problem is avoided.

You can configure Websense Filtering Service to prioritize directory policies (in other words, use the search order User > Group > Domain > Computer > Network to identify the policy to apply to a request).

When Filtering Service is installed on a Windows or Linux server:

1. Navigate to the Websense **bin** directory on the Filtering Service machine (C:\Program Files\Wesense\Web Security\bin or /opt/Wesense/bin/, by default).
2. Open the **eimserver.ini** file in a text editor.
3. Locate the **[FilteringManager]** section of the file, and add the following parameter:

```
UserGroupIpPrecedence=true
```
4. Save and close the file.
5. Restart Filtering Service.

- *Windows*: Use the Windows Services tool to restart **Websense Filtering Service**.
- *Linux*: Use the `/opt/Websense/WebsenseDaemonControl` command to restart **Filtering Service**.

When Filtering Service is on a Websense appliance:

1. Log on to the Appliance manager.
2. Navigate to the **Administration > Toolbox** page.
3. In the Appliance Command Line section, under Command Line Utility, click **Launch Utility**.
4. Select the Web module (TRITON AP-WEB or Web Filter & Security) from the Module drop-down list.
5. Enter **user-group-ip-precedence** in the Command field.
6. Enter **enable** in the Action field, then click **Run**.

Filtering Service is stopped and restarted automatically to implement the change.

Responding to a URL request

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When a user requests a site, Websense Filtering Service is responsible for determining whether to block or permit the request.

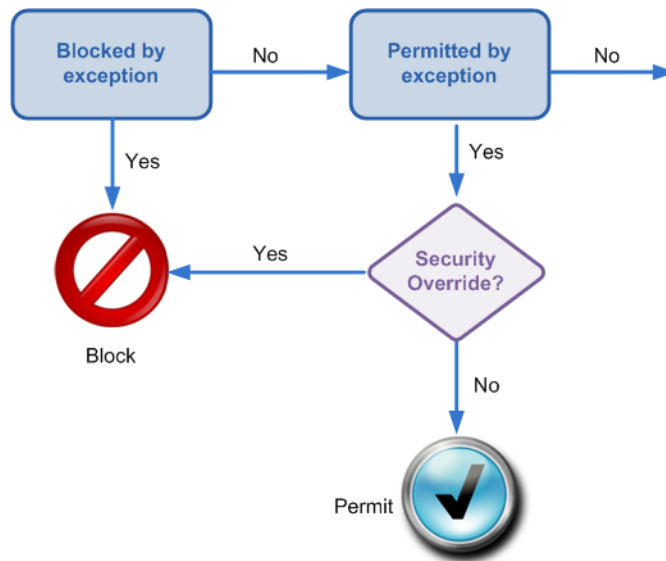


Important

With the Web Hybrid module, requests from users outside the network are managed by the hybrid service, which has its own rules for determining whether to block or permit a request.

If a request triggers Content Gateway analysis, Filtering Service uses the category returned as a result of the analysis to determine whether to block or permit the request.

Filtering Service determines which action to take as follows:

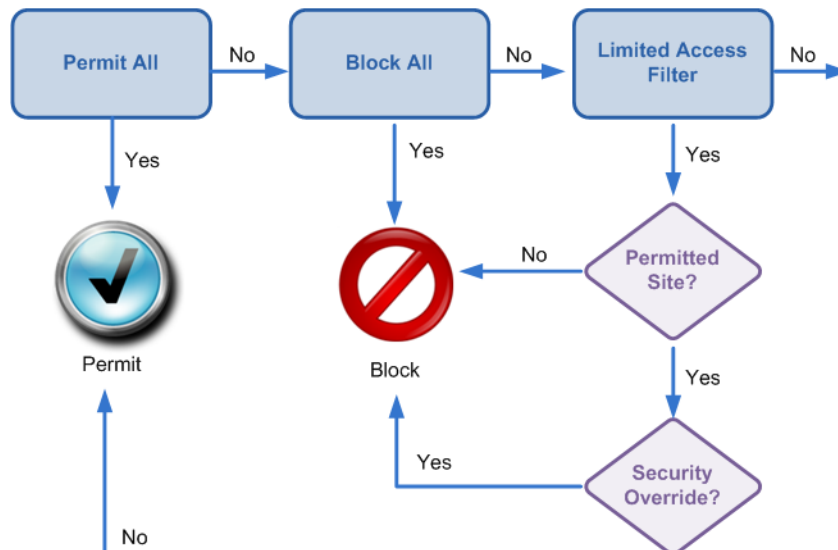


1. Check to see whether the site is listed in an **exception**.

- If there is a **block exception**, block the site.
- If there is a **permit exception**, permit the site.

Note that if the URL is permitted by exception, but classified as a security risk, the default action is to block it via security override. Administrators can disable the security override, but this is not recommended.

- If there is no exception for the site, continue to Step 2.



2. 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.

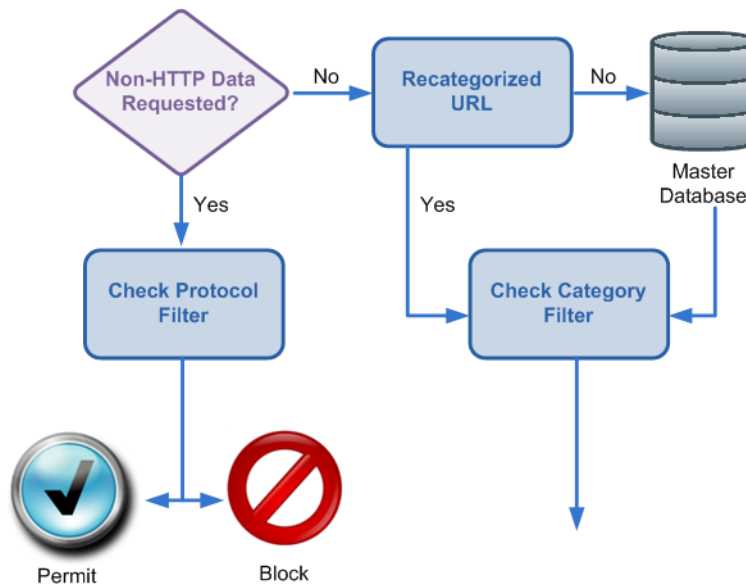
Note that if the URL is permitted by the limited access filter, but classified as a security risk, the default action is to block it via security override. Administrators can disable the security override, but this is not recommended.

- If any other category filter applies, continue to Step 3.



Note

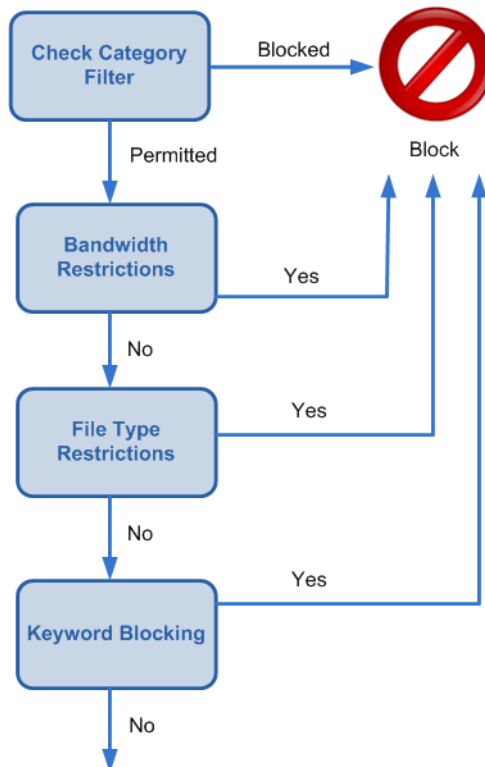
Filtering Service handles URLs accessed from search engine's cache like any other URL. They are blocked or permitted according to the applicable policies. Log records for cached URLs show the entire cached URL, including any search engine parameters.



3. Checks the active **protocol filter** and determines whether any non-HTTP protocols are associated with the request.
 - If so, apply the appropriate action, as defined in the protocol filter.
 - 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.

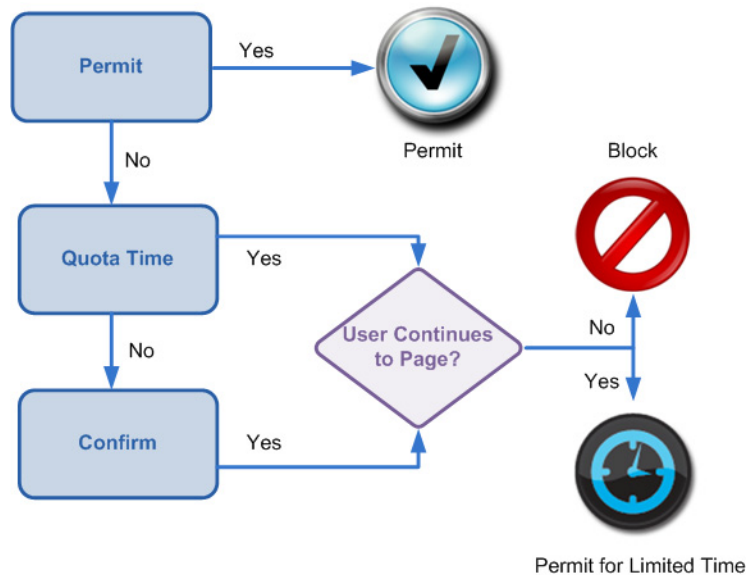
- If a match is not made, categorize the site as Miscellaneous/Uncategorized and continue to Step 6.

With TRITON AP-WEB, sites not categorized by the Master Database are analyzed by Content Gateway. If this returns a new category for the site, Filtering Service applies an action based on the new category (rather than continuing to classify the site as Uncategorized).



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 261](#)).
 - 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 263](#)).
 - 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 [Keyword-based policy enforcement](#), page 248).
- 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.

Filtering Service proceeds until the requested site is either blocked or explicitly permitted. At that point, no further investigation is attempted. For example, if a requested site belongs to a blocked category and contains a blocked keyword, Filtering Service blocks the site at the category level without checking the keyword. 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 websites regardless of why the site was blocked.

6

Content Gateway Analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring tunneled protocol detection, page 92](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content Gateway outbound security analysis, page 99](#)
- ◆ [Content Gateway advanced analysis options, page 100](#)
- ◆ [Configuring exceptions to Content Gateway analysis, page 103](#)
- ◆ [Data files used with Content Gateway analysis, page 105](#)
- ◆ [Reporting on advanced analysis activity, page 105](#)
- ◆ [SSL decryption bypass, page 108](#)

Advanced analysis and SSL decryption bypass features are available with TRITON AP-WEB. These features are not available with Web Filter & Security.

Websense Content Gateway performs advanced analysis of web traffic as it flows through the on-premises proxy. Only sites that are not already blocked, based on the active policy, are analyzed.

- ◆ [Configuring content categorization, page 90](#), categorizes content from URLs that are not in the Websense Master Database and from sites with dynamic content, as identified by Websense Security Labs. Analysis returns a category for use in policy enforcement.
- ◆ [Configuring tunneled protocol detection, page 92](#), analyzes traffic to discover protocols tunneled over **HTTP** and **HTTPS**. Such traffic is reported to Filtering Service for protocol policy enforcement. Analysis is performed on both inbound and outbound traffic.
- ◆ [Configuring content security, page 93](#), analyzes inbound content to find security threats such as malware, viruses, phishing, URL redirection, web exploits, proxy avoidance, and others.

- ◆ [Configuring file analysis, page 94](#), can apply as many as 3 methods of inspection to detect security threats.
 - **Websense Advanced Detection** to discover malicious content, such as viruses, Trojan horses, and worms, returning a threat category for policy enforcement.
 - **Traditional antivirus (AV) definition files** to find virus-infected files.
 - **File Sandboxing** uploads suspicious files to a cloud-hosted sandbox for analysis and can be configured to send alerts via email, SNMP, or both when a file is found to contain malicious content.

When either **Advanced Detection** or **Antivirus Scanning** 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. This setting does not apply to File Sandboxing.

- ◆ [Content Gateway outbound security analysis, page 99](#), provides 2 types of outbound analysis. The first performs outbound content analysis that mirrors your inbound Security Threats content analysis and file analysis configuration. The second performs data theft analysis, looking for and blocking outbound custom encrypted files, password files, and other sensitive data.
- ◆ The **Content Categorization and Scanning Sensitivity** control allows you to tune the Content Categorization and Content Analysis sensitivity thresholds ([Content Gateway advanced analysis options, page 100](#)).
- ◆ For large, streaming, or slow transactions, the **Content Delay Handling** option provides some control over how long to wait before releasing a portion of buffered content to the client ([Content Gateway advanced analysis options, page 100](#)).
- ◆ The **Scanning Timeout, File Size Limit** and **Content Stripping** Advanced Options apply to all traffic transiting the proxy ([Content Gateway advanced analysis options, page 100](#)).

Several presentation reports can provide details about how advanced analysis features protect your network from attempts to access sites containing threats. See [Reporting on advanced analysis activity, page 105](#).

SSL decryption bypass options support the specification of clients, websites, and website categories that are **not** subject to decryption and analysis as they flow through the proxy. These options apply only if SSL support is enabled in Content Gateway. See [SSL decryption bypass, page 108](#).

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 [Configuring exceptions to Content Gateway analysis, page 103](#).

Enabling analysis and SSL decryption bypass features

Administrator Help | TRITON AP-WEB | Version 8.1.x

To enable the advanced analysis and SSL decryption bypass features that are available with TRITON AP-WEB, an appropriate subscription key must be entered in the TRITON Manager. You can enter the key:

- ◆ In the Initial Setup Checklist
- ◆ On the **Settings > General > Account** page
- ◆ On the **Settings > General > Policy Servers** page, after selecting a Policy Server instance to edit.

Review current key information on the Account or Policy Servers page.

The key is automatically passed to all Content Gateway instances associated with the current Policy Server. See [Reviewing Policy Server connections, page 354](#), and [Managing Content Gateway connections, page 368](#), for more information.

For information about configuring advanced analysis options, see [Configuring Content Gateway analysis, page 89](#). For information about SSL decryption bypass options, see [SSL decryption bypass, page 108](#).

Configuring Content Gateway analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring tunneled protocol detection, page 92](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content Gateway outbound security analysis, page 99](#)
- ◆ [Content Gateway advanced analysis options, page 100](#)
- ◆ [Configuring exceptions to Content Gateway analysis, page 103](#)
- ◆ [Reporting on advanced analysis activity, page 105](#)

The analysis options available with TRITON AP-WEB 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 Content Gateway, see [Content Gateway Analysis, page 87](#).

Use the **Settings > Scanning > Scanning Options** page to configure the following:

- ◆ [Configuring content categorization, page 90](#)

- ◆ [Configuring tunneled protocol detection](#), page 92
- ◆ [Configuring content security](#), page 93
- ◆ [Configuring file analysis](#), page 94
- ◆ [Content Gateway outbound security analysis](#), page 99
- ◆ Analytic sensitivity, analysis timeout, size limits, content delay handling, and content stripping ([Content Gateway advanced analysis options](#), page 100)

Basic settings are:

- ◆ **Off** – No analysis.
- ◆ **On** (default) – Analyze content or files with elevated risk profiles, as determined by Websense Security Labs.
- ◆ **Aggressive analysis** – Analyze content and files with elevated risk profiles and content and files with lower risk profiles. Aggressive analysis consumes more resources. For best results, monitor system performance and scale system resources to meet demand.

In addition to the On/Off/Aggressive analysis 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 [Configuring exceptions to Content Gateway analysis](#), page 103.



Warning

Sites on the Never Scan list are not analyzed under any circumstances. If a site on the Never Scan list is compromised, the malicious code is not analyzed and detected.

When you have completed configuration on the current page, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring content categorization

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis](#), page 89
- ◆ [Configuring tunneled protocol detection](#), page 92
- ◆ [Configuring content security](#), page 93
- ◆ [Configuring file analysis](#), page 94
- ◆ [Content Gateway outbound security analysis](#), page 99
- ◆ [Content Gateway advanced analysis options](#), page 100
- ◆ [Configuring exceptions to Content Gateway analysis](#), page 103
- ◆ [Reporting on advanced analysis activity](#), page 105

When a web page 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
- ◆ The URL has an elevated risk profile, as identified by Websense Security Labs

The category that is determined by content categorization is forwarded to Filtering Service for policy enforcement.

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 113, for more information.



Important

If you plan to generate reports of advanced analysis activity, enable full URL logging (see [Configuring how URLs are logged](#), page 402). 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 [What is WebCatcher?](#), page 14), URLs categorized through content categorization are forwarded for inclusion in the Master Database.

To configure content categorization:

1. Go to the **Settings > Scanning > Scanning Options** page.
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 analysis. 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 and Deploy**.

The algorithms used to perform content categorization are tuned by Websense Security Labs to provide the best results for most organizations. However, if the Optimized setting does not produce the results you expect, you can adjust the sensitivity level to influence more restrictive or more permissive results. See the

[Content Gateway advanced analysis options](#) section of this screen.

Configuring tunneled protocol detection

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content Gateway outbound security analysis, page 99](#)
- ◆ [Content Gateway advanced analysis options, page 100](#)
- ◆ [Configuring exceptions to Content Gateway analysis, page 103](#)
- ◆ [Reporting on advanced analysis activity, page 105](#)

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 Filtering Service for protocol-based policy enforcement. When tunneled protocol detection is enabled, analysis is performed on both inbound and outbound traffic, regardless of other 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 Filtering Service 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 [Managing access to categories and protocols, page 32](#), for information about protocol-based policy enforcement.



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 Filtering Service for policy enforcement.
3. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Use the **Settings > Scanning > Scanning Exceptions** page to specify trusted sites that are never analyzed ([Configuring exceptions to Content Gateway analysis](#), page 103).

Configuring content security

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis](#), page 89
- ◆ [Configuring content categorization](#), page 90
- ◆ [Configuring tunneled protocol detection](#), page 92
- ◆ [Configuring file analysis](#), page 94
- ◆ [Content Gateway outbound security analysis](#), page 99
- ◆ [Content Gateway advanced analysis options](#), page 100
- ◆ [Configuring exceptions to Content Gateway analysis](#), page 103
- ◆ [Reporting on advanced analysis activity](#), page 105

Content Security performs web page content analysis to discover security threats and malicious code in HTTP and HTTPS content (HTTPS when Content Gateway SSL support is enabled).

Use the **Settings > Scanning > Scanning Options** page to enable and configure content security.

1. Select **Off** to disable content analysis.
2. Select **On** (default) to enable content analysis for uncategorized sites and sites with elevated risk profiles, as identified by Websense Security Labs.
3. Select **Aggressive analysis** to analyze content from sites with elevated risk profiles and also sites with lower risk profiles. This option consumes additional system resources.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Use the **Settings > Scanning > Scanning Exceptions** page to specify untrusted or trusted sites that are always analyzed or never analyzed ([Configuring exceptions to Content Gateway analysis](#), page 103).

Content analysis sensitivity is tuned by Websense Security Labs to provide the best results for most organizations. However, if the Optimized setting does not produce the results you expect, you can adjust the sensitivity in the [Content Gateway advanced](#)

[analysis options](#) section.

Configuring file analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis](#), page 89
- ◆ [Configuring content categorization](#), page 90
- ◆ [Configuring content security](#), page 93
- ◆ [Content Gateway outbound security analysis](#), page 99
- ◆ [Content Gateway advanced analysis options](#), page 100
- ◆ [Configuring exceptions to Content Gateway analysis](#), page 103
- ◆ [Reporting on advanced analysis activity](#), page 105

File analysis inspects files that users attempt to download or open remotely for viruses and other malicious content. File analysis returns a category to Filtering Service for policy enforcement.

There are 5 types of file analysis. They can be used together.

Four types of analysis are done by Content Gateway.

- ◆ [Advanced Detection](#) applies techniques developed by Websense to discover known and emerging threats, including viruses, Trojan horses, worms, and other malicious content.
- ◆ [Antivirus Scanning](#) uses antivirus definition files to identify virus-infected files.
- ◆ [Rich Internet application scanning](#) examines Flash files for malicious content.
- ◆ [FTP file scanning](#) examines inbound FTP files for malicious content.

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, the best remedy is to create an exception for that site. See [Configuring exceptions to Content Gateway analysis](#).

Use the **Settings > Scanning > Scanning Exceptions** page to specify untrusted or trusted sites that are always analyzed or never analyzed ([Configuring exceptions to Content Gateway analysis](#), page 103).

Use the **Settings > Scanning > Scanning Options** page to enable and configure file analysis.

The fifth type of file analysis is *File Sandboxing*, which sends files that fit a profile defined by Websense Security Labs to a cloud-hosted sandbox for activation and observation. If a file is found to be malicious by the sandbox, an email alert is sent to the configured administrator that contains a description of the threat, a link to a detailed File Sandboxing report, and a link to an investigative report built from your Log Database.

To enable File Sandboxing, TRITON AP-WEB subscribers can purchase the Web Sandbox module. A full description is included in the step-by-step configuration section, below.

Advanced Detection

1. Select **Off** to disable file analysis.
2. Select **On** (default) to enable file analysis on files from uncategorized sites and files from sites with elevated risk profiles, as identified by Websense Security Labs.
3. Select **Aggressive analysis** to analyze inbound files from sites with elevated risk profiles and from sites with lower risk profiles. This option consumes additional system resources.

Antivirus Scanning

1. Select **Off** to disable antivirus analysis.
2. Select **On** (default) to enable antivirus analysis of files from uncategorized sites and files from sites with elevated risk profiles, as identified by Websense Security Labs.
3. Select **Aggressive analysis** to apply antivirus analysis to inbound files from sites with elevated risk profiles and from sites with lower risk profiles. This option consumes additional system resources.

Rich Internet application scanning

Select **Scan rich Internet applications** to analyze Flash files for malicious content.

FTP file scanning

Select **Scan FTP files** to analyze files that are downloaded with the FTP protocol. (FTP over HTTP file downloads and uploads are subject to the HTTP/HTTPS file

scanning settings.) To be meaningful, this option requires that Content Gateway be configured to proxy FTP traffic. See the Content Gateway Online 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 analysis feature is turned off, these options are disabled and the check boxes are cleared.

File Type Options

1. To specify the types of files Content Gateway is 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 and Deploy**.

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 113, for more information.

See [Managing traffic based on file type](#), page 263, for information about blocking files based on type and URL category.

File Sandboxing

**Note**

At least one of the four Content Gateway analysis options must be enabled for files to be sent to the File Sandbox.

With the Web Sandbox module, you have the option to:

1. Select **Off** (default) to disable File Sandboxing analysis.
2. Select **On** to send qualified executable files to the cloud-hosted sandbox for analysis.
3. Control the types of files sent to the file sandbox. Check the box next to the general file types listed to keep those file types from being sent to the file sandbox. By default, none of the boxes are checked; all suspicious files are sent.
Note that analysis is performed to determine a file's true type.

When a file type is selected for “Do not submit”, both the true file type and the file extension are used to determine that the file will not be sent to the file sandbox.

Caution: Electing not to send file types to the sandbox may expose the network to unknown risk. Select the file types based on proper risk assessment. Balance the privacy risks involved in sending files to the sandbox against the security risks involved in not sending them.

4. To not send files having a specific extension, check **Files with the following extensions**, enter file extensions in the input box provided, and click **Add**. Multiple file extensions can be added in a comma separated list.
To remove an entry from the list, highlight a file extension and click **Delete**.
5. When you are done configuring File Sandboxing options, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

A file that qualifies for File Sandboxing:

- Is **not** classified as “malicious” in the Websense Master Database.
- Passes all selected **Security Threats: File Analysis** analytics.
- Fits the Websense Security Labs profile for suspicious files.
- Is a supported file type. Executable files are always supported. See the knowledge base article titled: [File Sandboxing supported file types](#).



Note

Because the file was **not** detected as malicious, it was **not blocked** and has been delivered to the requester.



Important

To receive File Sandboxing alerts, which is the only mechanism used to report files found to be malicious by sandbox analysis, **you must enable and configure email or SNMP alerts.**

Go to **Settings > Alerts > Enable Alerts**. Select **Enable email alerts** and specify an **Administrator email address**. Also confirm that your SMTP settings are correct. Select **Enable SNMP alerts** and provide information about your SNMP Trap system.

Then, enable File Sandboxing Alerts on the **Settings > Alerts > Suspicious Activity Alerts** page.

**Important**

The Content Gateway web proxy manages on-premises File Sandboxing traffic.

Traffic is sent to:

- ◆ *.websense.net
- ◆ *.blackspider.com

The User-Agent is **ssbc**.

File Sandboxing traffic must not be subject to man-in-the-middle decryption, and cannot be challenged for authentication by any device in the network.

Filter.config rules are configured, by default, in Content Gateway. If Content Gateway is in a proxy chain or behind a firewall, those devices may have to be configured to meet the requirements described above.

You can test your configuration to ensure that File Sandboxing is properly configured in your deployment via the **Real-time Analysis Test Pages** section of <http://testdatabasewebsense.com/>

What does a File Sandboxing transaction look like?

1. An end user browses to a website and explicitly or implicitly downloads a file.
2. The URL is **not** categorized as “malicious” and **Security Threats: File Analysis** does **not** find the file to be malicious.
3. The file is delivered to the requester.
4. However, the file fits the Websense Security Labs profile for suspicious files and is sent to File Sandboxing in the cloud for analysis.
5. File Sandboxing analyzes the file, which may take as long as 5 to 10 minutes, but is typically much quicker.
6. If the file is found to be malicious, Content Gateway sends a File Sandboxing malicious file detection message to the configured alert recipient. The alert email includes links to the File Sandboxing report and an investigative report created from your log records (examples below).
7. Upon receipt of the message, administrators should:
 - a. Access and evaluate the File Sandboxing report for the file
 - b. Examine the investigative report for the incident
 - c. Assess the impact of the intrusion in their network
 - d. Plan and begin remediation
8. Separately, File Sandboxing updates the ThreatSeeker[®] Intelligence Cloud with information about the file, the source URL, and the command and control targets.
9. ThreatSeeker updates the Websense Master Database, ACE analytic databases, and other security components, which are then pulled by Websense deployments.

10. The next time someone tries to browse the site, they and the organization are protected by their TRITON AP-WEB deployment.

File Sandboxing alert messages and reports

When Content Gateway learns that File Sandboxing has detected a malicious file, it sends a plain-text alert email to the configured administrator.



Important

To receive File Sandboxing email messages, which is the only mechanism used to report files found to be malicious by sandbox analysis, **you must enable and configure email alerts.**

In the body, the **User** field includes the user name only if Content Gateway user authentication was used to identify the client. Otherwise, the client IP address appears in the field.

Two links are included. The first links to a detailed File Sandboxing report on the file and its malicious contents. The second launches an investigative report, using your log records, for the time period in which the file download occurred. Depending on your browser, you may have to enable popups to allow the report to be displayed. Also note that you may receive the File Sandboxing alert message before TRITON AP-WEB has written all of the transaction records in the Log Database. Periodically refresh the report to include pending records.

Content Gateway outbound security analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Configuring exceptions to Content Gateway analysis, page 103](#)
- ◆ [Reporting on advanced analysis activity, page 105](#)
- ◆ [Content Gateway advanced analysis options, page 100](#)

Outbound security:

- ◆ Provides outbound analysis that mirrors your inbound Security Threats configuration. This option also supports social web controls.

- ◆ Performs specialized data theft protection, analyzing for and blocking outbound custom encrypted files, password files, and other forms of sensitive data (see number 2, below).
- 1. Enable **Analyze for and block outbound security threats** (default) to analyze outbound content for threats like bot and spyware phone home traffic. This option performs outbound analysis that mirrors your inbound Security Threats configuration.



Important

This option must be enabled to support social web controls.

- 2. Enable **Data theft protection** (default) to analyze and block:
 - a. Outbound custom encrypted files that are posted to **uncategorized** sites and suspicious destinations, as defined by Websense Security Labs.
 - b. Password files and files containing sensitive or suspicious data, regardless of the destination.

The results of analysis are reported to the Threats dashboard, and are included in transaction logs and reports.

Content Gateway advanced analysis options

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content Gateway outbound security analysis, page 99](#)
- ◆ [Configuring exceptions to Content Gateway analysis, page 103](#)
- ◆ [Reporting on advanced analysis activity, page 105](#)

Use these options to:

- ◆ Set the sensitivity level of Content Categorization and Content Security analysis
- ◆ Set the analysis time limit*
- ◆ Set the analysis size limit*
- ◆ Enable stripping of specific types of code from HTML content*

*These settings apply to all incoming traffic.

Content categorization and scanning sensitivity level

The algorithms used to perform content categorization and content analysis 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 influence more restrictive or more permissive results.

There are 5 sensitivity levels.

- ◆ **Optimized** is the sensitivity level tuned by Websense Security Labs.
- ◆ **More Stringent** and **Most Stringent** raise analytic sensitivity.
- ◆ **Less Stringent** and **Least Stringent** reduce analytic sensitivity.

When you are finished, click **OK** to cache your changes.

Changes are not implemented until you click **Save and Deploy**.

Scanning timeout

Each content or file analysis 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).

Scan size limit

The scan size limit is the threshold to which analysis is performed. Analysis stops when the threshold is reached. The default is 10 MB. To change the value, select **Custom** and enter a size in megabytes.

Content delay handling

Depending on the Content Gateway configuration and load conditions, very large files, streamed transactions, and slow origin servers can leave clients waiting for content.

The options in this section provide a tool for delivering a portion of buffered content to the client **before analysis is performed**. Analysis begins when all data is received or the scan size limit is exceeded.

Use **Begin returning data to the client after** to specify a time period after which a percentage of buffered data is released to the client. The default is 30 seconds. Select **Custom** to enter another value.

Use **Specify how much data to return to the client** to specify the percentage of buffered data to release to the client. The default is 80 percent. Select **Custom** to enter a different value, up to 90 percent.

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 [Configuring Content Gateway analysis](#), page 89).

Content is removed only after the advanced analysis options have categorized the site and Filtering Service 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 and Deploy**.



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 the 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.
-

Configuring exceptions to Content Gateway analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring content security, page 93](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content Gateway outbound security analysis, page 99](#)
- ◆ [Content Gateway advanced analysis options, page 100](#)

Scanning exceptions are lists of trusted or untrusted sites (hostnames and URLs) that are **never analyzed** or **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 options, see [Content Gateway Analysis, page 87](#).

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 Security, or File Analysis options are **On**, sites on the **Always Scan** list are always analyzed, and sites on the **Never Scan** list are never analyzed (see [Configuring Content Gateway analysis, page 89](#)).
- ◆ When the Tunneled Protocol Detection option is **On** or **Aggressive analysis** 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, TRITON AP-WEB 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 labels are supported. For example: **www.bbc.co.uk**
- You can use the wild card “*” to match leading subdomains only.

For example: *.yahoo.com.

- You can enter a complete or partial hostname or URL. The leading scheme “HTTPS://” is not required. An exact match is performed on the specified string.

For example: www.example.com/media/

Or: www.youtube.com/watch?v=

2. After entering a single or group of hostnames/URLs, select the scanning 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.

Click **OK** to add the entry.

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 and Deploy**.

To change the scanning 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 and Deploy**.

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 and Deploy**.

Data files used with Content Gateway analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Analysis uses a set of data files to support its work. These files are updated regularly by Websense Security Labs and made available on the Websense download server. Websense Content Gateway checks for updated analytic data files at regular intervals. The name and version of each file is displayed in the Content Gateway manager on the **Monitor > MyProxy > Summary** page.

Data file updates occur independent 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:

```
/opt/WCG/WCGAdmin restart
```

Reporting on advanced analysis activity

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring Content Gateway analysis, page 89](#)
- ◆ [Configuring content categorization, page 90](#)
- ◆ [Configuring file analysis, page 94](#)
- ◆ [Content stripping, page 102](#)

After you install Websense Content Gateway and enter a key that enables the advanced analysis features, you can see and analyze the effects of these features on the dashboard, and with presentation and investigative reports.

On the Usage dashboard, by default, 2 charts tally requests to Web 2.0 sites over the past 30 days:

- ◆ Web 2.0 Categories
- ◆ Web 2.0 URL Bandwidth

See [The Status Dashboards, page 17](#), for information about customizing the charts or moving them to a different dashboard tab.

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 402](#)) to ensure that reports of analysis activity 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 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 analysis activities, as well as general policy enforcement. 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 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 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 134](#), for details on using all the investigative reports features.

How analysis activity is logged

Administrator Help | TRITON AP-WEB | Version 8.1.x

There are important differences in the way that general Internet activity and advanced analysis activity are logged.

For general Internet activity, you have several options to reduce the size of the Log Database.

- ◆ Enable **visits** to log only one record for each website requested. See [Configuring Log Server](#), page 388.
- ◆ Enable **consolidation** to combine into a single log record multiple requests with certain common elements. See [Configuring Log Server](#), page 388.
- ◆ 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 402.



Note

If your organization needs reports that include the full URL of each site visited, you should leave full UR 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 requests are logged](#), page 386.



Note

Enabling **visits**, **consolidation**, or **selective category logging**, will impact the accuracy of Internet Browse Time.

Advanced analysis features, however, are bound only partially by these settings. When a site is analyzed, 2 separate log records are created.

- ◆ **Standard log records** take advantage of any size reduction settings that have been implemented, and are available for all reporting tools.
- ◆ **Advanced analysis records** ignore most size reduction settings. Every separate hit is logged, requests to all categories are logged, and no records are consolidated. These records are 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. Advanced analysis records are used to populate the Threats dashboard and presentation reports that focus on the results of Content Gateway analysis (like those described in [Reporting on advanced](#)

[analysis activity, page 105](#)).

If you have enabled any Log Database size reduction options, the numbers that appear on the Threats dashboard and in presentation reports on Content Gateway analysis may **not** match those that appear in standard investigative and presentation 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 reports, but may show as multiple hits in advanced analysis reports.

To see comparable data for standard activity and advanced analysis **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 383](#), for more information on configuring size reduction settings. See [Presentation reports, page 113](#), and [Investigative reports, page 134](#), for information on generating reports.

SSL decryption bypass

Administrator Help | TRITON AP-WEB | Version 8.1.x

When SSL support is enabled in Content Gateway to manage encrypted traffic:

- ◆ Category settings can be used to specify categories of websites for which decryption and inspection are bypassed.
- ◆ A list of client IP addresses and IP address ranges can be created to specify trusted clients for which decryption and inspection are bypassed.
- ◆ A list of destination hostnames, IP addresses, and IP address ranges can be created to specify trusted destination servers for which decryption and inspection are bypassed.



Note

There is a known limitation with Internet Explorer version 8 (IE8) that prevents some sites from being bypassed as expected. IE8 does not send a Server Name Indicator (SNI) and when the hostname in the origin server certificate includes a wildcard (*), the common name and the hostname don't match. As a result, the category lookup is performed on the destination IP address.

Category settings

For Category settings, a predefined **Privacy Category** group includes categories that may be subject to regulatory requirements.

Default privacy categories include:

- ◆ Education
- ◆ Financial Data and Services
- ◆ Government
- ◆ Health
- ◆ Online Brokerage and Trading
- ◆ Prescribed Medications

Traffic that involves websites 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 website they are viewing is not decrypted by verifying that the certificate is the original for that site.

Use the **Settings > Scanning > SSL Decryption Bypass** page to select the default privacy categories for SSL decryption bypass:

1. Click the **Select Privacy Categories** button. Check boxes for the website 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 website categories for which decryption is not performed:

1. Click a check box to select a category or subcategory for bypass.
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, select the category and click the **Remove** button.

Client list

To identify a client IP address or IP address range for SSL decryption bypass:

1. Click **Add** and enter the client IP address or IP address range in the **Add Client Entry** box, one entry per line.
When specifying an IP address range, use a “-” (hyphen) to separate the first address from the last.
2. To facilitate maintenance of the list, add a description that identifies the entry.
3. Click **OK** to add the entries to the list.

To modify an entry, click on the IP address and modify the entry in the **Edit Client Entry** box. Click **OK** to save your changes or **Cancel** to close the dialog box without saving your changes.

To remove an entry from the list, select the check box adjacent to the entry and click **Delete**. Confirm the action.

When you are finished, click **OK** to cache your changes.

Changes are not implemented until you click **Save and Deploy**.

Destination list

To specify a destination hostname, IP address, or IP address range for SSL decryption bypass:

1. Click **Add** and enter the hostname, IP address, or IP address range in the **Add Destination Entry** box, one entry per line. For example: **thissite.com**.
 - Be sure to enter both the hostname and the TLD (top level domain). For example, **thissite.com** and **thissite.net** are distinct hosts.
 - Hosts with subdomains are supported. For example: **media.example.com**.
 - Include the wild card "*" to match leading subdomains. For example: ***.example.com**.
 - The protocol (HTTPS://) is not needed.
 - Use a "-" (hyphen) to separate the first and last address in an IP address range.
2. To facilitate maintenance of the list, add a description that clearly identifies the entry.
3. Click **OK** to add the entries to the list.

To modify an entry, click on the hostname or IP address and modify the entry in the **Edit Destination Entry** dialog box. Click **OK** to save your changes or **Cancel** to close the dialog box without saving your changes.

To remove an entry, select the check box adjacent to the entry and click **Delete**. Confirm the action.

When you are finished, click **OK** to cache your changes.

Changes are not implemented until you click **Save and Deploy**.

7

Use Reports to Evaluate Internet Activity

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports](#), page 113
- ◆ [Investigative reports](#), page 134
- ◆ [Accessing self-reporting](#), page 156
- ◆ [Application reporting](#), page 156
- ◆ [Real-Time Monitor](#), page 163

The Web module of the TRITON Manager provides several reporting tools that can help you evaluate the effectiveness of your web 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 311.

- ◆ **Dashboard charts** provide threat, risk, usage, and system information to help you review Internet activity in your network at a glance. For most charts, the time period, chart style, and set of results shown can be customized. See [The Status Dashboards](#), page 17.
- ◆ **Presentation reports** offer a list of predefined reports, custom reports, and report templates. Reports are available in bar chart, trend chart, and tabular formats. Copy any predefined report to apply your own filters to create a custom report, or use a report template to create your report from scratch. See [Presentation reports](#), page 113, 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 134, for details on the many ways you can view Internet use data.

- ◆ **Application reports** provide information about the browsers and platforms from which Internet requests are originating, with a search item that lets you investigate activity associated with a specific user agent string.

See [Application reporting](#), page 156, for more information.

- ◆ **Real-Time Monitor** shows current Internet activity in your network, including the URLs being requested and the action applied to each request. With TRITON AP-WEB, the monitor also shows which sites were analyzed by Content Gateway. If a site is dynamically recategorized based on real-time analysis, both the original category and current category are shown.

See [Real-Time Monitor](#), page 163, for more information.

What is Internet browse time?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web protection reporting database jobs](#), page 395
- ◆ [Configuring Internet browse time options](#), page 403

You can generate both presentation and investigative reports showing **Internet browse time**, the estimated amount of time a user spent accessing websites. No software program can tell the exact amount of time that someone spends viewing a site once it is open. One person 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 a site in detail before moving to the next one.

A Log Database job (see [Web protection reporting database jobs](#), page 395) calculates browse time based on configurable parameters. This job runs once a day, so browse time information can lag the actual log data.

For browse time calculations:

- ◆ An Internet session begins when a user opens a browser and continues as long as that user requests additional websites at least every 3 minutes (by default).
If you want to change the read time threshold, see [Configuring Internet browse time options](#), page 403.
- ◆ The Internet session ends when more than 3 minutes pass before the user requests another site.
- ◆ 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.

The database job calculates the total time of each session, starting with the time of the first request and ending 3 minutes after the last request.


Presentation reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a new presentation report, page 115](#)
- ◆ [Working with presentation report Favorites, page 123](#)
- ◆ [Running a presentation report, page 124](#)
- ◆ [Scheduling presentation reports, page 125](#)
- ◆ [Viewing the presentation reports scheduled jobs list, page 131](#)

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



Available reports and templates are found in the Report Catalog, which organizes them into related **report categories**. Your subscription determines which report categories and predefined reports appear in the catalog. For example, report categories like Real Time Security Threats and Scanning Activity require TRITON AP-WEB.

- ◆ Expand a category to see the reports or templates 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 124](#).
 - 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 template, predefined report, or custom report in the Report Catalog as the basis for a new report:

1. Select a report or template name in the catalog.

If you select a report template:

 - A **New Trend Report** shows Internet activity trends over time.
 - A **New Top N Report** shows top levels of Internet activity with the characteristics you specify.
2. Click **Save As**.
3. Provide a name, title, and report category for the new file.

If you are using a report template, also define the report dimensions (i.e., what is measured and the unit of measurement).

For instructions, see [Creating a new presentation report, page 115](#).
4. To refine the report, edit the report filter. The report filter controls elements such as which users, categories, protocols, and actions are to be included in your report.

For instructions, see [Defining the presentation report filter, page 117](#).

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 or report templates.

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 presentation reports scheduled jobs list, page 131](#), 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 presentation report Favorites, page 123](#)). 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 125](#).

- ◆ Click **Job Queue** to see and manage a list of existing scheduled jobs, along with the status of each job. See [Viewing the presentation reports scheduled jobs list, page 131](#).
- ◆ Click **Review Reports** to see and manage a list of reports that were successfully scheduled and run. See [Reviewing scheduled presentation reports, page 133](#).

Creating a new presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports, page 113](#)
- ◆ [Defining the presentation report filter, page 117](#)
- ◆ [Running a presentation report, page 124](#)

Use the **Save As New Report** page to create:

- ◆ An editable version of any predefined report.
- ◆ A copy of an existing custom report, in order to apply different report filters.
- ◆ A new report, based on a report template.

The options available on the page depend on which option you have selected.

If you are making a copy of a predefined or custom report:

1. Replace the **Report 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. Enter a **Report title**. This is the title that will appear on at the top of the page when the report is generated.
3. Select a **Report category**. This determines how the report is grouped in the Report Catalog. The default is User-Defined Reports.
4. Do one of the following:
 - Click **Save** to save the new version of the report and return to the Report Catalog.
 - Click **Save and Edit** to edit the report filter for the new report (see [Defining the presentation report filter, page 117](#)).
 - Click **Cancel** to abandon your changes and return to the Report Catalog.

If you are using a report template to create a new report:

1. Enter a unique **Report name**. This is the name that will appear in the Report Catalog.

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

2. Enter a **Report title**. This is the title that will appear on at the top of the page when the report is generated.
3. Select a **Report category**. This determines how the report is grouped in the Report Catalog. The default is User-Defined Reports.
4. If you are creating a top N report, continue with step 5.

If you are creating a trend report, indicate the **Time unit** for the trend report's X-axis. You can create a report showing trends by day (default), week, month, or year.

- To ensure that the data you want appears in a trend report, make sure the **first day** of the first week, month, or year that you want to include is set as the first date in the range. (By default, the first day of the week is Sunday, but this may vary based on your Microsoft SQL Server configuration and locale.)
- When user information is updated in the directory service, user group information may also change. This can affect weekly, monthly, and yearly group trend reports, because to be included in a group report, the user must be in the group at least one day before the start of the selected period.

For example, for a user's activity to be included in a monthly group trend report for August 2012, the user must be in the group as of July 31, 2012. A user joining the group on August 23, 2012 (a Wednesday) would be included in daily trend reports starting on the following day, in weekly trend reports starting the following Saturday, Sunday, or Monday (depending on your Microsoft SQL Server configuration), and in monthly trend reports starting September 01, 2012.

5. Use the **Internet activity per** drop-down list to select the focal area of the report. You can show Internet activity per category (default), protocol, risk class, action (like permit or block), user, or group.
6. Use the **Measure by** drop-down list to select how the focal area is measured. You can measure by requests (default), bandwidth, or browse time.
7. Do one of the following:
 - Click **Save** to save the report and return to the Report Catalog. The new report is now listed in the report category that you selected in step 5.
 - Click **Save and Edit** to edit the report filter for the new report. The process of editing the report filter is the same as for any custom report (see [Defining the presentation report filter](#), page 117).
 - Click **Cancel** to abandon your changes and return to the Report Catalog.

Defining the presentation report filter

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a new presentation report, page 115](#)
- ◆ [Running a presentation report, page 124](#)

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 actions (like permit or block). You also can give the report a new name and description, change the report title, select a custom logo, and set other general options through the report filter.



Note

To use a custom logo, you must create the image in a supported format and place the file in the appropriate location before updating the report filter. See [Customizing the presentation report logo, page 122](#).

The options available in the filter vary:

- ◆ If you are editing a predefined report or a custom report based on a predefined report, the 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.
- ◆ If you are editing a report created using a the New Top N Report or New Trend Report template, **all** options are shown in the filter, even if they are not applicable in the custom report.

Be careful to select only options relevant to your report.

The filter for predefined reports 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 presentation report, page 118](#)
- ◆ [Selecting categories for a presentation report, page 119](#)
- ◆ [Selecting protocols for a presentation report, page 120](#)
- ◆ [Selecting actions for a report, page 120](#)
- ◆ [Setting presentation report options, page 121](#)

On the **Confirm** tab, choose whether to run or schedule the report, and save the report filter. See [Confirming presentation report filter definitions](#), page 123.

Selecting clients for a presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Selecting categories for a presentation report](#), page 119
- ◆ [Selecting protocols for a presentation report](#), page 120
- ◆ [Selecting actions for a report](#), page 120
- ◆ [Setting presentation report options](#), page 121
- ◆ [Confirming presentation report filter definitions](#), page 123

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 a combination of users and IP addresses for the same report.



Note

If you select the IP address client type, you can identify clients by IPv4 or IPv6 address.

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 or OUs 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 (OUs) 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 presentation report](#), page 119.

Selecting categories for a presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Selecting clients for a presentation report](#), page 118
- ◆ [Selecting protocols for a presentation report](#), page 120
- ◆ [Selecting actions for a report](#), page 120
- ◆ [Setting presentation report options](#), page 121
- ◆ [Confirming presentation report filter definitions](#), page 123

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 36.

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 presentation report](#), page 120.

Selecting protocols for a presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Selecting clients for a presentation report, page 118](#)
- ◆ [Selecting categories for a presentation report, page 119](#)
- ◆ [Selecting actions for a report, page 120](#)
- ◆ [Setting presentation report options, page 121](#)
- ◆ [Confirming presentation report filter definitions, page 123](#)

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 120](#).

Selecting actions for a report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Selecting clients for a presentation report, page 118](#)
- ◆ [Selecting categories for a presentation report, page 119](#)
- ◆ [Selecting protocols for a presentation report, page 120](#)
- ◆ [Setting presentation report options, page 121](#)
- ◆ [Confirming presentation report filter definitions, page 123](#)

The **Actions** tab of the Presentation Reports > Edit Report Filter page lets you control which precise 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 presentation report options](#), page 121.

Setting presentation report options

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Customizing the presentation report logo](#), page 122
- ◆ [Selecting clients for a presentation report](#), page 118
- ◆ [Selecting categories for a presentation report](#), page 119
- ◆ [Selecting protocols for a presentation report](#), page 120
- ◆ [Selecting actions for a report](#), page 120
- ◆ [Setting presentation report options](#), page 121
- ◆ [Confirming presentation report filter definitions](#), page 123

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 presentation report logo](#), page 122.
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 presentation report filter definitions](#), page 123.

Customizing the presentation report logo

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 **ReportTemplates\images** directory. The default path is:

```
C:\Program Files (x86)\Websense\Web Security\Manager\  
ReportTemplates\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 presentation report options](#), page 121.)



Note

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

Confirming presentation report filter definitions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Selecting clients for a presentation report, page 118](#)
- ◆ [Selecting categories for a presentation report, page 119](#)
- ◆ [Selecting protocols for a presentation report, page 120](#)
- ◆ [Selecting actions for a report, page 120](#)
- ◆ [Setting presentation report options, page 121](#)

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 presentation report options, page 121](#).)

2. Indicate how you want to proceed:

Option	Description
Save	Saves the report filter and returns to the Report Catalog. See Presentation reports, page 113 .
Save and Run	Saves the report filter and opens the Run Report page. See Running a presentation report, page 124 .
Save and Schedule	Saves the report filter and opens the Schedule Report page. See Scheduling presentation reports, page 125 .

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

Working with presentation report Favorites

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports, page 113](#)
- ◆ [Running a presentation report, page 124](#)
- ◆ [Scheduling presentation reports, page 125](#)

You can mark presentation reports as Favorites to identify the reports you generate most frequently and want to be able to locate quickly.

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, simply select the report again and click **Favorite** to remove the star symbol.

Running a presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports](#), page 113
- ◆ [Scheduling presentation reports](#), page 125

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 125).

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. Viewing requires Adobe Reader 7.0 or later.
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. Viewing requires Microsoft Excel 2003 or later.

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.
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, the report will be displayed. 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. If you selected HTML, click **Presentation Reports** to return to the Report Catalog. If you selected PDF or XLS, you can use the Run Reports window again to generate the same report.

If you selected HTML, click **Presentation Reports** to return to the Report Catalog. If you selected PDF or XLS, you can use the Run Reports window again to generate the same report.

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 and use the print options in Adobe Reader.

Scheduling presentation reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports, page 113](#)
- ◆ [Running a presentation report, page 124](#)
- ◆ [Viewing the presentation reports scheduled jobs list, page 131](#)

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 133](#)).

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 presentation report filter, page 117](#)).

- ◆ 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 presentation reports schedule, page 127](#)
- ◆ [Selecting presentation reports to schedule, page 128](#)
- ◆ [Setting the date range for a scheduled presentation report, page 129](#)
- ◆ [Selecting output options for scheduled presentation reports, page 130](#)

After creating jobs, use the Job Queue to review job status and find other helpful information (see [Viewing the presentation reports scheduled jobs list, page 131](#)).

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 management server machine (C:\Program Files (x86)\Websense\Web Security\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 management server machine. If the ReportingOutput directory grows too large before the files are automatically deleted, you can delete the files manually.

The report is generated in the format you choose: PDF (Adobe Reader 7.0 or later), XLS (Microsoft Excel 2003 or later), or HTML. If you choose HTML format, the report may display in the TRITON Manager 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.

Setting the presentation reports schedule

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Scheduling presentation reports](#), page 125
- ◆ [Selecting presentation reports to schedule](#), page 128
- ◆ [Selecting output options for scheduled presentation reports](#), page 130
- ◆ [Setting the date range for a scheduled presentation report](#), page 129

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).

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

The job begins according to the time on the TRITON management server.

**Note**

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

4. 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 presentation reports scheduled jobs list , page 131.
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 presentation reports scheduled jobs list , page 131.
End by	Set the date when the job stops running. It does not run on or after this date.

5. Click **Next** to open the Reports tab. See [Selecting presentation reports to schedule](#), page 128.

Selecting presentation reports to schedule

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Scheduling presentation reports](#), page 125
- ◆ [Setting the presentation reports schedule](#), page 127
- ◆ [Selecting output options for scheduled presentation reports](#), page 130
- ◆ [Setting the date range for a scheduled presentation report](#), page 129

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

1. Highlight a report for this job in the Report Catalog tree.
2. Click the right arrow (>) button to move that report to the **Selected** list.
3. Repeat steps 1 and 2 until all reports for this job appear in the **Selected** list.

4. Click **Next** to open the Date Range tab. See [Setting the date range for a scheduled presentation report](#), page 129.

Setting the date range for a scheduled presentation report

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Scheduling presentation reports](#), page 125
- ◆ [Setting the presentation reports schedule](#), page 127
- ◆ [Selecting presentation reports to schedule](#), page 128
- ◆ [Selecting output options for scheduled presentation reports](#), page 130

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	Reports include all dates available in the Log Database. No additional entries are required. When this option is used for repeating jobs, there may be duplicate information on reports in separate runs.
Specific Dates	Choose the exact start (From) and end (To) dates for the reports in this job. This option is ideal for jobs that run only one time. Choosing this option for a repeating schedule results in duplicate reports.
Relative Dates	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. 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. 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.

After setting the date range for the job, click **Next** to display the Output tab. See [Selecting output options for scheduled presentation reports](#), page 130.

Selecting output options for scheduled presentation reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Scheduling presentation reports, page 125](#)
- ◆ [Setting the presentation reports schedule, page 127](#)
- ◆ [Selecting presentation reports to schedule, page 128](#)
- ◆ [Setting the date range for a scheduled presentation report, page 129](#)

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 presentation reports scheduled jobs list, page 131](#).

Viewing the presentation reports scheduled jobs list

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports](#), page 113
- ◆ [Scheduling presentation reports](#), page 125
- ◆ [Selecting output options for scheduled presentation reports](#), page 130
- ◆ [Scheduling investigative reports](#), page 150

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 153.)

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 Viewing the presentation reports scheduled job history , page 132.
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 Scheduling presentation reports, page 125 .
Add Job	Opens the Scheduler page where you can define a new job. See Scheduling presentation reports, page 125 .
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.
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 the presentation reports scheduled job history

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Scheduling presentation reports, page 125](#)
- ◆ [Viewing the presentation reports scheduled jobs list, page 131](#)

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Presentation reports](#), page 113
- ◆ [Running a presentation report](#), page 124
- ◆ [Scheduling presentation reports](#), page 125

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 385).

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 presentation reports scheduled jobs list](#), page 131, for more information about using the job queue.

To schedule a new report job, click **Scheduler** (see [Scheduling presentation reports](#), page 125).

Investigative reports

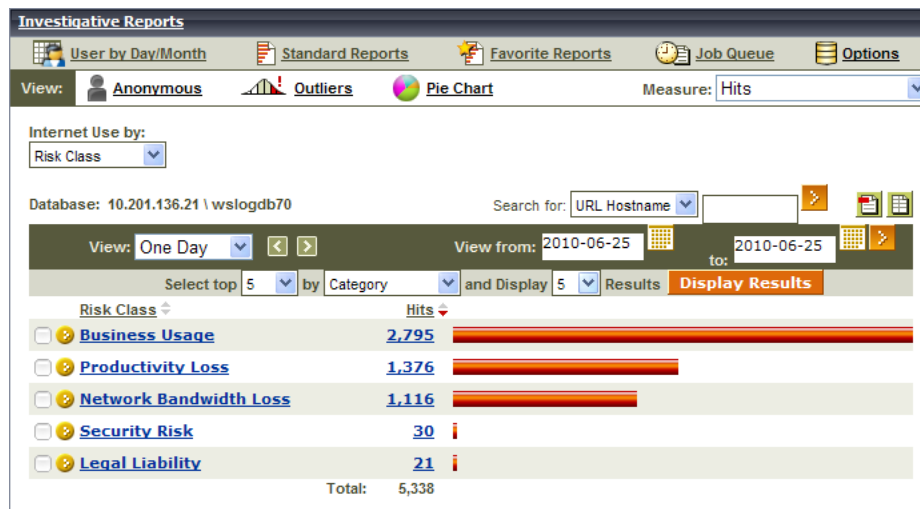
Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Summary investigative reports](#), page 136
- ◆ [Multi-level summary investigative reports](#), page 141
- ◆ [Flexible detail investigative reports](#), page 142
- ◆ [User Activity Detail investigative reports](#), page 146
- ◆ [Standard investigative reports](#), page 148
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150
- ◆ [Outliers investigative reports](#), page 153
- ◆ [Output options for investigative reports](#), page 155
- ◆ [Database connection and report defaults](#), page 410

Use the **Reporting > Investigative Reports** page to analyze Internet activity in an interactive way.



Initially, the main Investigative Reports page shows a summary report of activity by risk class (see [Risk classes](#), page 36).



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 investigative reports](#), page 136).

Multi-level summary reports (see [Multi-level summary investigative reports](#), page 141) and flexible detail reports (see [Flexible detail investigative reports](#), page 142) 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 investigative reports , page 146.
Standard Reports	Displays a list of predefined reports so you can quickly see a specific combination of data. See Standard investigative reports , page 148.
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 149.
Job Queue	Displays the list of scheduled investigative reports jobs. See Scheduling investigative reports , page 150.
Outliers	Displays reports showing Internet usage that is significantly different from average. See Outliers investigative reports , page 153.
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 410.
	Click this button, at the right of the Search fields, to export the current report to a spreadsheet file compatible with Microsoft Excel 2003 or later. You are prompted to either open or save the file. See Output options for investigative reports , page 155.
	Click this button, at the right of the Search fields, to export the current report to a PDF file compatible with Adobe Reader v7.0 or later. You are prompted to either open or save the file. See Output options for investigative reports , page 155.

Investigative reports support IPv6 for source and destination IP addresses. Also, anywhere an IP address can be entered in an investigative reports feature, both IPv4 and IPv6 formats are accepted.

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 requests are logged](#), page 386), that information cannot be included.

- ◆ Similarly, if you disable logging for certain protocols (see [Editing a protocol filter](#), page 47), 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 402).
- ◆ If your directory service does not include the first and last name of the user, reports cannot display user name information.

Investigative reports are limited by the processor and available memory of the TRITON management server, as well as some network resources. Some large reports may take a very long time to generate. The progress message includes an option to save the report as a Favorite so you can schedule it to run at another time. See [Scheduling investigative reports](#), page 150.

Summary investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Multi-level summary investigative reports](#), page 141
- ◆ [Flexible detail investigative reports](#), page 142
- ◆ [User Activity Detail investigative reports](#), page 146
- ◆ [Standard investigative reports](#), page 148
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150
- ◆ [Outliers investigative reports](#), page 153
- ◆ [Output options for investigative reports](#), page 155

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 410.

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.

Option	Description
Bandwidth [KB]	The amount of data, in kilobytes, contained in both the initial request from the user and the response from the website. This is the combined total of the Sent and Received values.
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 website.)
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 Content Gateway, as a result of analysis, the kilobytes received represents the size of the page analyzed. See Content Gateway Analysis, page 87.</p> <p>If a third-party 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	An estimate of the amount of time spent viewing the site. See What is Internet browse time? , page 112.

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 412.)

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.

Option	Action
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 requests are logged, page 386.</p> <p>For more information about hiding user-identifying information, see Anonymizing investigative reports, page 140.</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. 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> <p>See Using search to generate a summary investigative report, page 139, for more details.</p>

- 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 investigative reports](#), page 141.
- 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 investigative reports](#), page 142.

Using search to generate a summary investigative report

Use the **Search for** box on the main Investigative Reports page to quickly find information about Internet traffic or client activity of interest.

First, select a report element from the list, then enter all or part of the string that you want to report on.

The elements available for search are:

- ◆ The **URL Hostname** of the requested website
- ◆ A **Group** defined in your directory service
- ◆ A **User** defined in your directory service

If you select User, but enter an IP address, you will get results only for requests from the selected IP address for which no user was identified.
- ◆ The **Source IP** address of the computer from which a request originated
- ◆ The **Destination IP** address of the requested website
- ◆ The **Port** used for the request
- ◆ A **Source IP Range** from which requests originated
- ◆ **Multiple Source IP Ranges** from which requests originated, in a comma-separated list

When you enter multiple IP address ranges, you can also specify individual IP addresses or sub-ranges to exclude from the search, by preceding the IP address or range with an exclamation point (bang) character. For example:

10.21.1.1-10.21.1.10,10.22.55.1-10.22.55.50,!10.22.55.5

Anonymizing investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, source IP addresses, and hostnames. 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 requests are logged, page 386](#), 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 311](#), 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, as described below.

The Anonymous option

By default, clicking **Anonymous** hides only user names, continuing to show source IP addresses in reports. You can configure investigative reports to instead hide both user names and source IP addresses when Anonymous is selected:

1. On the TRITON management server, open the **wse.ini** file in a text editor. (By default, this file is located in C:\Program Files (x86)\Websense\Web Security\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 investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Summary investigative reports](#), page 136
- ◆ [Flexible detail investigative reports](#), page 142
- ◆ [User Activity Detail investigative reports](#), page 146
- ◆ [Standard investigative reports](#), page 148
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150
- ◆ [Outliers investigative reports](#), page 153
- ◆ [Output options for investigative reports](#), page 155

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.

The screenshot shows a configuration bar for a report. It contains the following elements from left to right: a label 'Select top', a text input field containing the number '5', a dropdown arrow, a label 'by', a text input field containing the word 'User', a dropdown arrow, a label 'and Display', a text input field containing the number '10', a dropdown arrow, a label 'Results', and an orange button labeled 'Display Results'.

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 investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Summary investigative reports](#), page 136
- ◆ [Multi-level summary investigative reports](#), page 141
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150
- ◆ [Outliers investigative reports](#), page 153
- ◆ [Output options for investigative reports](#), page 155
- ◆ [Database connection and report defaults](#), page 410
- ◆ [Columns for flexible detail investigative reports](#), page 144

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 TRITON management server, 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 investigative reports](#), page 136, or [Multi-level summary investigative reports](#), page 141.)
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.

3. 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 popup 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 149.

4. 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 410.)

5. 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.

6. 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 investigative reports](#), page 144, for a list of the columns available, and a description of each.

7. 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.

8. 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 412.
 - Click the URL to open the requested website 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 [Working with presentation report Favorites](#), page 123).

Columns for flexible detail investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Flexible detail investigative reports](#), page 142
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150

The table below describes the columns available for detail reports (see [Flexible detail investigative reports](#), page 142).

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 384.
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 taken 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 the Web Hybrid module, 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 the Web Hybrid module, you can use this option to review requests coming from a specific hybrid filtered location. See Define hybrid service filtered locations , page 195.
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 402.
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 website. This is the combined total of the Sent and Received values.
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 website, 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 Content Gateway, as a result of analysis, the bytes received represents the size of the page analyzed. See Content Gateway Analysis, page 87.</p> <p>If a third-party 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 investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134

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, as described below. From that report, you can generate a report of the same user's activity for a full month.

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.)

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 412 . 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.

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.

7. 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.

8. Click **Prev** or **Next** to display this user's Internet activity for the previous or the next month.

Standard investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Favorite investigative reports](#), page 149
- ◆ [Scheduling investigative reports](#), page 150

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 149.

Favorite investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Scheduling investigative reports](#), page 150

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.

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 default name.

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 Favorite Reports list.

From the Favorite Reports list, 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.
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 150, for more information.
 - Click **Delete** to remove the report from the Favorites list.

From the Favorite Reports page, you can also create a new Favorite report that is similar to an existing one:

1. Click **Favorite Reports** to display a list of reports saved as favorites.

2. Select and run the existing Favorite report that most closely resembles the new report you want to create.
3. Modify the displayed report as desired.
4. Click **Favorite Reports** to save the revised display as a Favorite report with a new name.

Scheduling investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Favorite investigative reports, page 149](#)
- ◆ [Favorite investigative reports, page 149](#)
- ◆ [Managing scheduled investigative reports jobs, page 153](#)

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 compressed into zip files and 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\Other\

The default installation path is C:\Program Files (x86)\ Websense\Web Security.



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.

For the best performance when running scheduled jobs, a minimum of 4G of RAM is recommended. Companies with large amounts of data should consider increasing to 8G or more.

1. Save one or more reports as Favorites. (See [Favorite investigative reports, page 149](#)).

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 address 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.) Highlight an email address and click Delete to remove it from the list.
Additional email address	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

Field	Description
Email text	<p>Enter text to be added to the email message for distributing scheduled reports.</p> <p>The email reads as shown below, with your text in place of <CUSTOM TEXT>.</p> <p>Report scheduler generated the attached file or files on <date time>.</p> <p><CUSTOM TEXT></p> <p>To view the generated report(s), click on the following link(s).</p> <p>Note: The link will not work if the recipient does not have access to the web server from which the job was sent.</p>
Scheduled job name	<p>Assign a unique name for the scheduled job. The name identifies this job in the Job Queue. See Managing scheduled investigative reports jobs, page 153.</p>
Output format	<p>Choose the file format for the scheduled reports:</p> <p>PDF: Portable Document Format files are viewed in Adobe Reader.</p> <p>Excel: Excel spreadsheet files are viewed in Microsoft Excel.</p>
Date range	<p>Set the date range to be covered by reports in this job.</p> <p>All Dates: all available dates in the Log Database.</p> <p>Relative: Choose a time period (Days, Weeks, or Months) and the specific period to include (This, Last, Last 2, and so on).</p> <p>Specific: set specific dates or a date range for the reports in this job.</p>

- Click **Next** to display the Schedule Confirmation page.
- Click **Save** to save your selections and go to the Job Queue page (see [Managing scheduled investigative reports jobs](#), page 153).

Managing scheduled investigative reports jobs

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Scheduling presentation reports](#), page 125

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 investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Summary investigative reports](#), page 136

An Outliers report shows which users have the most unusual Internet activity in the database. A report query 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 options for investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

After 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 the TRITON console, 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 the TRITON console, 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>

You can also print investigative reports, as follows:

- ◆ Use the browser print function while the report is displayed.
- ◆ Create a PDF or XLS file, as described above, and then use the print function in Adobe Reader or Microsoft Excel.

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 investigative reports](#), page 142), 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 412.)

Reports that exceed 1,000,000 records should be saved as a Favorite and added to a scheduled job.

Accessing self-reporting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Investigative reports](#), page 134
- ◆ [Configuring reporting preferences](#), page 385
- ◆ [Self-reporting](#), page 414

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 administrator 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 administrator for assistance.
3. Enter the **User name** and **Password** you use to log on to the network.
4. Click **Log On**.

The TRITON Manager displays 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.

Application reporting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [How is user agent data collected?](#), page 159
- ◆ [Browser use details](#), page 160
- ◆ [Platform use details](#), page 161

With TRITON AP-WEB or standalone Web Filter & Security, you can use the **Reporting > Applications** page to review the browsers and operating systems used to make web requests in your network. You can also use the Search tab to investigate activity based on user agent strings. (The user agent string is an HTTP header that identifies the client software from which a request originates.)

- ◆ Find instances of older browsers that may present a security vulnerability.
- ◆ Identify which machines in your network may be vulnerable when a zero-day exploit is discovered.
- ◆ Track adoption of new browsers and operating systems.
- ◆ Use the user agent string associated with an application to identify the machines in your network on which that application is running.
- ◆ Search for user agent strings associated with malware or suspicious activity to identify machines that may be at risk.

When Web Filter & Security is integrated with a third-party proxy, cache, firewall, or other device, the integration product does not send application data to Websense Filtering Service. As a result, no data is available on the Applications page.



Important

After a new installation or upgrade, the Applications page will not show data on the Browser or Source Platform tab until a nightly database job runs. As users connect to the Internet, new user agent strings appear on the Search tab, but the strings are not parsed into recognized browsers and platforms until the job runs.

See [How is user agent data collected?](#), page 159, for more information about how user agent data is logged, processed, and made available for use in reports.

The Applications page is made up of 3 tabs:

- ◆ Reports on the **Browser** tab show which supported browser families (including desktop and mobile versions of Microsoft Internet Explorer, Mozilla Firefox, Google Chrome, Safari, and Opera) and versions are being used to access the Internet from your network.
The Browser tab is selected by default when you navigate to the Applications page.
- ◆ Reports on the **Source Platform** tab show which supported operating systems (including Windows, Linux, UNIX, OS X, iOS, Android, BlackBerry, Symbian, and Java ME) the browsers accessing the Internet are running on.
- ◆ The **Search** tab lets you search for specific strings in user agent headers detected in your network. The search results show the top qualifying user agents by requests or bandwidth.

On any tab, you can select an alternate **Time period** from the drop-down list at the top of the tab. By default, 30 days worth of information (if available) is shown in the charts and tables.

Different time periods are available depending on whether you use a standard or enterprise version of Microsoft SQL Server, or Microsoft SQL Server Express.

On the Browser and Source Platform tabs, you can also select a **Device type** (desktop or mobile) to limit the information that appears on the page. By default, information for both desktop and mobile browsers or platforms is shown.

Both the Browser and Source Platform tabs include a table that lists browsers or platforms and versions.

- ◆ An icon in the **Type** column indicates whether the browser or platform is for desktop machines or mobile devices.
- ◆ The next column gives the name of the browser family or operating system platform.
- ◆ The **Lowest Version** and **Highest Version** values give the range of versions being used in your network during the selected period.
- ◆ **Number** of browsers of the specified type actively being used to make requests, or of source platforms from which requests are originating. The count is made based on the number of unique client IP addresses associated with the browser or operating system.

Click a link in the table to open a detail report with more information about the selected browser family, platform, or version number. See [Browser use details, page 160](#), or [Platform use details, page 161](#), for more information.

Use the **Browser Family** or **Platform Comparison** chart to view the top browser families or operating systems used in your network, and the **Browser** or **Platform Use Trend** chart to track use of the different browsers or operating systems over time.

For any chart, select a different **Chart type** to change the way the information is displayed.

The **Search** tab initially displays the top 10 user agents, based on number of requests. To search the database for specific user agents, enter a string in the **User agent** field and click **Search**. The string can be all or part of a user agent header, up to a maximum of 128 characters.

The top (up to 200) results that match your search string are displayed in the **User Agent Search Results** table, which includes:

- ◆ The actual **User Agent** that matches the search criteria. If the string is truncated, mouse over it to see the full string.
- ◆ The last column header shows whether results are sorted by **Requests** or **Bandwidth**. (Use the **Sort by** list to choose a measurement.)

After performing a search, click **Clear** to return to showing the default Top 10 User Agents table.

Click a user agent in the Top 10 or Results table to display a **User Agent Detail** table at the bottom of the page. The details table shows:

- ◆ The **User** whose browsing included the user agent.
- ◆ The **Client IP Address** from which the request originated.
- ◆ The **Source Server IP Address** for the component (Content Gateway or Network Agent) processing the requests.

- ◆ The number of **Requests** that included the selected user agent.
- ◆ The volume of **Bandwidth** for all requests that include the user agent from the specified user and client machine.

Click **Export to CSV** to export the report detail information to a CSV file that can be manipulated using spreadsheet software like Microsoft Excel.

**Note**

If there are more records than your system can handle, the output file will not contain actual CSV-format data. If this occurs, select a shorter timeframe to reduce the data set and export the data again.

How is user agent data collected?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Application reporting, page 156](#)
- ◆ [Browser use details, page 160](#)
- ◆ [Platform use details, page 161](#)

The user agent is an HTTP header that web browsers and other web applications use to identify themselves and their capabilities. Your Websense software captures and logs user agent data when users browse the Internet. If the user agent data includes browser and platform information, that information is parsed and displayed in application reports.

- ◆ If a browser or platform is installed in your network, but is **not** used for Internet access, it does not appear in application reports.
- ◆ Because there are no widely-adopted standards for user agent headers, your software is not able to identify all Internet-accessing applications.

Some applications, in fact, deliberately disguise their identity in the user agent header in an attempt to avoid detection.

The application browsing data that Websense Log Server receives includes the user agent header, user name, and source IP address. All requests that share the identical user agent, user, and source IP address during a 60-second period are combined into a single record that provides the total number of requests and the volume of bandwidth associated with those requests. That record is then forwarded to the Log Database. How soon browser and platform reports are updated with data about current Internet activity depends on whether a user agent has previously been seen and analyzed:

- ◆ If the user agent corresponds to a browser, browser version, or platform that has not previously been parsed and identified, information about requests from that browser and platform do not appear in application reports until after the nightly

Trend job (see [Web protection reporting database jobs](#), page 395).

In other words, there is a delay of up to 24 hours before information about new browsers, browser versions, and platforms appears in browser and platform reports.

As a result, the Browser and Source Platform tabs do not initially show any reports after a new installation or upgrade.

- ◆ If the user agent corresponds to a browser, browser version, and platform that have previously been parsed and identified, information about requests from that browser and platform appear in browser and platform reports as soon as they are recorded in the Log Database.

Data on the Search tab is not subject to the same delays as the browser and platform reports. User agent strings are available for search as soon as they are recorded in the Log Database. This includes both the strings associated with browsers and platforms, and strings used by other types of web apps.

Browser use details

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Application reporting](#), page 156
- ◆ [How is user agent data collected?](#), page 159
- ◆ [Platform use details](#), page 161

When you click a browser family or version on the Browser tab of the Applications page, a **Browser Detail Report** is displayed.

The **Browser Inventory** gives a visual overview of:

- ◆ For a selected browser family, the top versions in use, and usage trends for those versions.
- ◆ For a selected browser version, the top users of the version, and usage trends for the version.

You can position your mouse over different chart elements to see additional details, and use the **Chart type** options under each chart to change the way the data is displayed.

Below the charts, a **Users Sending Requests** table lists up to the top 200 active users of the selected browser family or version. The table includes:

- ◆ The name of the **User** making the Internet requests.
- ◆ The **Client Hostname**, if available, and **Client IP Address** of the machine used to browse the Internet.
- ◆ The **Source Server IP Address** corresponding to the component (Content Gateway or Network Agent) that processes HTTP requests.

- ◆ The **Browser** name and version.
- ◆ The **Type** (mobile or desktop) of browser used.
- ◆ The volume of requests made through the browser by **Count** and **Bandwidth**.
- ◆ The **User Agent** associated with this browser and version. Click the icon to see the full user agent.

Click **Export to CSV** to export the available data to a CSV file for manipulation in a spreadsheet program like Microsoft Excel.



Note

If there are more records than your system can handle, the output file will not contain actual CSV-format data. If this occurs, select a shorter timeframe to reduce the data set and export the data again.

Use the paging options at the bottom of the table to navigate through the data. Each page can display up to 20 lines of information.

Click **Close** to return to the summary data on the Browser tab of the Applications page.

Platform use details

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Application reporting, page 156](#)
- ◆ [How is user agent data collected?, page 159](#)
- ◆ [Browser use details, page 160](#)

When you click a platform family or version on the Source Platform tab of the Applications page, a **Platform Detail Report** is displayed.

The **Platform Inventory** gives a visual overview of:

- ◆ For a selected operating system, the top versions in use, and usage trends for those versions.
- ◆ For a selected operating system version, the top users of the version, and usage trends for the version.

You can position your mouse over different chart elements to see additional details, and use the **Chart type** options under each chart to change the way the data is displayed.

Below the charts, a **Users Sending Requests** table lists up to the top 200 active users of the selected operating system or version. The table includes:

- ◆ The name of the **User** making the Internet requests.
- ◆ The **Client Hostname**, if available, and **Client IP Address** of the machine used to browse the Internet.
- ◆ The **Source Server IP Address** corresponding to the component (Content Gateway or Network Agent) that processes HTTP requests.
- ◆ The operating system **Platform** name and version.
- ◆ The **Type** (mobile or desktop) of operating system used.
- ◆ The volume of requests made through the browser by **Count** and **Bandwidth**.
- ◆ The **User Agent** associated with this operating system and version. Click the icon to see the full user agent.

Click **Export to CSV** to export the available data to a CSV file for manipulation in a spreadsheet program like Microsoft Excel.



Note

If there are more records than your system can handle, the output file will not contain actual CSV-format data. If this occurs, select a shorter timeframe to reduce the data set and export the data again.

Use the paging options at the bottom of the table to navigate through the data. Each page can display up to 20 lines of information.

Click **Close** to return to the summary data on the Source Platform tab of the Applications page.

Real-Time Monitor

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Use Reports to Evaluate Internet Activity, page 111](#)
- ◆ [Real-Time Monitor in Multiple Policy Server Deployments, page 165](#)

Use the **Reporting > Real-Time Monitor** page to review current Internet activity in your network.



Important

If Real-Time Monitor does not display correctly in Internet Explorer, 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 policies are 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.

By default, 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.

Click **Customize** in the toolbar at the top of the page, then select **Show the full URL** to change this behavior.
- ◆ 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.

Hover the mouse over an entry to see the policy or policies used to determine the action. Multiple policies may be listed if, for example:

- Multiple group policies could be applied to the same user.
- A policy is assigned to both the IP address and the user or group.

When multiple policies are listed, you can use the Test Filtering tool to see which policy takes precedence for a request from the user or IP address shown in Real-Time Monitor.

- ◆ 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.

- ◆ By default, data is refreshed every 15 seconds. To change the update rate, click **Customize** in the toolbar at the top of the page, then select a new **Data refresh rate** value.

Depending on your current settings, Real-Time Monitor holds a set number of records (250, 500, or 1000), and always displays 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 are not available for display in the monitor. (The requests are, however, stored in the Log Database, and appear in investigative and presentation reports.)

To change how many records are displayed, click **Customize** in the toolbar at the top of the page, then select a new **Number of records shown** value.

Using filters to show specific Real-Time Monitor 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.

Understanding timeout behavior

By default, TRITON Manager 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 165, for considerations and instructions.

Real-Time Monitor in Multiple Policy Server Deployments

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When you go to the Reporting > Real-Time Monitor page, 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 clients.

If you want Real-Time Monitor to continue monitoring traffic for a specific Policy Server, regardless of which Policy Server instance the Web module of the TRITON Manager 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 Internet activity.
- ◆ You can have multiple Real-Time Monitor instances running in full-screen mode, each showing data for a different Policy Server:
 1. Log on to the TRITON Manager and select the Web module. It connects 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 the TRITON Manager and use the Policy Server Connection button in the 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.

8

Exceptions to Web Protection Policies

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing exceptions, page 168](#)
- ◆ [Adding or editing an exception, page 170](#)
- ◆ [Exception shortcuts, page 174](#)

Exceptions give administrators a way to quickly permit URLs and IP addresses in blocked categories, or block URLs and IP addresses in permitted categories.

Creating an exception does not require changing the category of a URL, nor does it change the policy assigned to affected clients. It simply allows a flexible and rapid response to user requests, changes in company policies, spikes in Internet activity, or other changes in circumstance.

For example:

- ◆ Permit access to an approved vendor's website for all employees, even though the Default policy blocks access to the Shopping category.
- ◆ Block all clients in the Students role from accessing an uncategorized URL that is experiencing a suspicious spike in traffic while the website is investigated.
- ◆ Permit access to a design blog for 3 members of the Web Marketing team, while continuing to block general access to the Blogs and Personal Sites category.
- ◆ Block a specific user from accessing a list of URLs at the request of the Human Resources department.

For streamlined instructions for common tasks, see [Exception shortcuts, page 174](#).

Managing exceptions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Adding or editing an exception, page 170](#)
- ◆ [How are exceptions organized?, page 170](#)
- ◆ [If multiple exceptions could apply, which takes precedence?, page 173](#)
- ◆ [Editing multiple exceptions at the same time, page 174](#)

Use the **Policy Management > Exceptions** page to review, edit, or delete existing exceptions, or to add new exceptions.

Super Administrators see all exceptions, regardless of the role in which they were created. Delegated administrators see all exceptions that affect their current role. For more information about how exceptions are ordered in the list, see [How are exceptions organized?, page 170](#).

- ◆ If a single URL or regular expression is blocked or permitted by the exception, the URL or expression is listed. Otherwise, click the link in the URLs column to see a complete list of affected URLs.
- ◆ If the exception affects:
 - A single client, the client's IP address, address range, or display name is listed.
 - A single role, the role name is displayed in the format "Role [Role_Name]"
 - All clients in all roles, the word "Global" is shown.

Global exceptions that can be overridden by delegated administrators are marked with an icon in the Clients column (see [Overriding an exception, page 172](#)).

 - Multiple, specific clients, the number of clients is shown. Click the link to see a complete list of affected clients.

The exceptions list also shows:

Column	Description
Type	Displays an icon to indicate whether URLs in the exception are: <ul style="list-style-type: none"> • Blocked (🚫) • Permitted (🟢) • Permitted with security override disabled (🟢🚫)
Last Modified	Shows the date that the exception was last edited.

Column	Description
Expires	Indicates whether or not the exception has an expiration date, and if so, displays the date.
Active	Shows whether the exception is currently being enforced (Active) or not (Inactive).

Use the **Filter** drop-down list to display only exceptions with specified characteristics. The following filters are available:

Filter	Description
Permitted	Exceptions that permit URLs.
Blocked	Exceptions that block URLs.
Active	Exceptions currently being enforced.
Inactive	Exceptions not currently used.
Will Expire	Exceptions for which an expiration date is specified.
Expired	Exceptions that are inactive because their expiration date has passed.
Never Expires	Exceptions set to remain active indefinitely.
Global	Exceptions that apply to all clients in all roles.
All Clients in a Role	Exceptions that apply to all clients in a specific delegated administration role (including the Super Administrator role).
Specific Clients	Exceptions that apply to one or more specific clients.

You can also use the **Search** fields to limit which exceptions are displayed:

1. Use the drop-down list to indicate which table columns you want to search.
2. Enter all or part of the string you want to identify.
3. Click **Search**.
4. To return to your previous view, click **Clear Search Results**.

To create a new exception, click **Add**. See [Adding or editing an exception, page 170](#), for instructions.

To edit an existing exception, click the exception name, or mark the check box next to one or more exceptions, and then click **Edit**. See [Adding or editing an exception, page 170](#), or [Editing multiple exceptions at the same time, page 174](#), for instructions.

To remove an exception, mark the check box next to the exception name, and then click **Delete**.

How are exceptions organized?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The order in which exceptions are displayed on the Policy Management > Exceptions page depends on the administrator's role.

For Super Administrators, exceptions are grouped as follows:

1. Global exceptions (affecting all clients in all roles)
2. Exceptions that affect specific clients from the Clients page in the Super Administrator role
3. Exceptions that include one or more clients that are not explicitly assigned to a role (do not appear on any Clients page or in any Managed Clients list)
4. Exceptions applied to the entire Super Administrator role
5. Exceptions applied to specific clients in another delegated administration role
6. Exceptions applied to an entire delegated administration role

For delegated administrators in other roles, exceptions are grouped as follows:

1. Exceptions that affect specific clients in the role
2. Exceptions that affect the entire role (including global exceptions)

Within each grouping, exceptions are shown in alphanumeric order.

Adding or editing an exception

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Policy Management > Exceptions > Add Exception** or **Edit Exception** page to create or update an exception that overrides standard policy enforcement to block or permit specific websites for specific clients.

1. Enter or update the unique, descriptive **Name** for the exception. The name must be between 1 and 50 characters long, and cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
2. In the **URLs** field, list the URLs or IP addresses to be permitted or blocked by the exception.
 - If you enter a URL in the format **domain.com**, both the domain and its subdomains (*www.domain.com*, *subdomain.domain.com*) are matched.
 - If you enter a URL in the format **www.domain.com**:
 - *http://www.domain.com* **is** matched
 - *http://domain.com* **is not** matched
 - *http://subdomain.domain.com* **is not** matched

Enter one URL or IP address per line.

3. Specify which **Clients** are affected by this exception.

Super Administrators can create:

- **Global** exceptions that apply to all clients in all roles.
If you select this option, also specify whether or not to **Allow delegated administrators to create exceptions that override this exception** (see [Overriding an exception](#), page 172).
- Exceptions that apply to **All clients in a role**.
After selecting this option, select a role from the drop-down list.
- Exceptions that apply to **Specific clients in any role**.
After selecting this option, you are offered 2 lists. One (on the left) shows all clients that have been **Defined**: added as managed clients in a delegated administration role, added to the Clients page in any role, or added to an exception. The other (on the right) shows clients **Selected** for this exception.

Search boxes appear above each list to help you quickly find clients to add or remove.

To add a client to the exception that does not appear in the list on the left, click **Add Other Clients**, then add user, group, computer (IPv4 or v6 address), or network (IPv4 or v6 address range) clients.



Important

If you select specific clients that belong to multiple roles, when the exception is created, it is automatically split so that a new exception is created for each affected role.

For example, if you define an exception called “Permit Craigslist” that applies to clients in the Super Administrator, HR, and Facilities roles, when you click OK, 3 exceptions are created.

- ◆ The exceptions for the HR role and Facilities role are marked with an icon. Move the mouse over the icon to see which role is affected by the exception.
 - ◆ The exception for the Super Administrator role is not annotated.
-

- Delegated administrators can create exceptions that apply to **All managed clients in this role** or **Specific clients in this role**.

If you select the latter option, you are offered 2 lists. One (on the left) shows all clients **Defined** in your Managed Clients list and Clients page. The other (on the right) shows the clients **Selected** for this exception.

- Search boxes appear above each list to help you quickly find the clients that you want to add.
- If a client does not appear in the **Defined clients** list, that individual is likely a member of a group, OU, or network (IP address range) defined as a managed client in your role. To add such a client, click **Add Other Clients**, then specify the user, group, or IPv4 or v6 address that you want to add.

4. Specify the exception **Type**. This determines whether to **Block** or **Permit** the listed URLs for the specified clients.
5. Indicate when the exception **Expires**.
 - If you select **Never**, the exception is used until you delete it, or edit it to add an expiration date.
 - If you select **After**, enter an expiration date in the format mm/dd/yyyy, or click the calendar icon to select a date. The exception expires at midnight (based on the time set on the Filtering Service machine), when the selected day ends.
6. Determine the exception **State**. By default, the exception is **Active**, and is immediately enforced after you cache and save your changes. If you do not want the exception to be used at this time, clear the check box.
7. By default, if a URL is associated with a Security Risk category (like Malicious Web Sites or Spyware), any permitted exception is ignored, and the URL is filtered based on the active policy (see [Prioritizing Security Risk categorization, page 252](#)):
 - If a category filter blocks the category, the request is blocked.
 - If a category filter permits the category, the request is permitted.
 - If a limited access filter is being used, the request is blocked.

To override this security feature, click **Advanced**, then clear the **Block URLs that become a security risk, even if they are permitted by exception** check box.

Making this change is not recommended.
8. To use regular expressions to define URLs that are permitted or blocked by exception, click **Advanced**, then enter one expression per line in the **Regular expressions** box.

To validate the expressions that you create, click **Test Regular Expression**.

Note that using large numbers of regular expressions, or using poorly-formed or overly-broad expressions, can lead to a significant decrease in performance.
9. When you are finished making changes, click **OK** to cache your changes and return to the Exceptions page. Changes are not implemented until you click **Save and Deploy**.

Overriding an exception

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Adding or editing an exception, page 170](#)
- ◆ [If multiple exceptions could apply, which takes precedence?, page 173](#)

By default, when a Super Administrator creates an exception, the exception takes precedence over any exceptions that a delegated administrator might create.


For example:

- ◆ A Super Administrator global exception blocks **mysite.com** and a delegated administrator exception for some managed clients permits **mysite.com**.
The URL is blocked by default.
- ◆ A Super Administrator global exception permits **anothersite.com** and a delegated administrator exception blocks the same site.
The URL is permitted by default.

When creating an exception, however, Super Administrators have the option to **Allow delegated administrators to create exceptions that override this exception**. If this option is selected, delegated administrator exceptions take precedence over the Super Administrator exception.

For example:

- ◆ A Super Administrator global exception permits **samplesite.com**, and a delegated administrator exception blocks **samplesite.com** for the delegated administration role.
The URL is blocked for clients in the delegated administration role.
- ◆ A Super Administrator global exception blocks **example.com**, and a delegated administrator exception permits **example.com** for a managed client.
The URL is permitted for the specified managed client.

Super Administrator exceptions that can be overridden are marked by an icon () in the Clients column on the Policy Management > Exceptions page.

If multiple exceptions could apply, which takes precedence?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

By default, Super Administrator exceptions take precedence over exceptions created by delegated administrators. So if a Super Administrator exception blocks a URL, and a delegated administrator exception permits the same URL, the request is **blocked**.

If, however, the Super Administrator configures an exception to allow delegated administrator overrides (see [Overriding an exception, page 172](#)), then the delegated administrator exception takes precedence. So if a Super Administrator exception blocks a URL, and a delegated administrator exception permits the same URL, the request is **permitted**.

If multiple equivalent exceptions could apply to a request (for example, if multiple Super Administrator exceptions include the same URL):

- ◆ Filtering Service checks for blocked exceptions first, so if there is a blocked exception and a permitted exception, the request is **blocked**.
- ◆ If there are multiple blocked exceptions, the first one found is applied.
- ◆ If there are no blocked exceptions and multiple permitted exceptions, the first permitted exception is applied.

After creating an exception, use the Test Filtering tool (see [Test Filtering, page 273](#)) to verify that client requests are filtered as expected.

Editing multiple exceptions at the same time

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Policy Management > Exceptions > Edit Exceptions** page to edit multiple exceptions at the same time.

When you edit multiple exceptions, you can edit only the exception type (permitted or blocked), expiration setting (never expires or expiration date), state (active or inactive), or security override setting (whether URLs in a permitted exception are permitted or blocked if Websense software finds a security risk).

Click the **View details of each selected exception** link near the top of the page for more information about the exceptions you are editing.

1. Verify the exception **Type** (Block or Permit). To make a change, click **Change**, then make a new selection.
2. To update the **Expires** setting for the exception, click **Change**, then
 - If you select **Never**, the exception is used until you delete it, or edit it to add an expiration date.
 - If you select **After**, enter an expiration date in the format mm/dd/yyyy, or click the calendar icon to select a date.
3. To update the exception **State**, click **Change**, then mark or clear the **Active** check box. Inactive exceptions are not used.
4. By default, if Websense components determine that a URL is a security risk (hosts malicious software or spyware, for example), the URL is blocked, even if it has been permitted by exception.

To update the current security settings for a permitted exception, click **Advanced**, then click **Change**. Mark or clear the **Block URLs that become a security risk, even if they are permitted by exception** check box.

Disabling the default security override protection is not recommended.

5. When you are finished making changes, click **OK** to cache your changes and return to the Exceptions page. Changes are not implemented until you click **Save and Deploy**.

Exception shortcuts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use these shortcuts to find the fastest way to perform common tasks.

For Super Administrators:

- ◆ [How do I block or permit a URL for everyone?, page 175](#)

- ◆ [How do I block or permit a URL for one person?, page 175](#)

For delegated administrators:

- ◆ [How do I block or permit a URL for my entire role?, page 176](#)
- ◆ [How do I block or permit a URL for one of my managed clients?, page 177](#)

How do I block or permit a URL for everyone?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Super Administrators can use the following steps to block or permit a URL for everyone in the network:

1. Go to the **Policy Management > Exceptions** page and click **Add**.
2. Enter a unique **Name** for the exception.
3. Enter the **URL** that you want to permit or block.
4. By default, the exception is set to apply to all clients (**Global** is selected).
5. By default, the exception is set to **Block** the URL. To change this, set the **Type** to **Permit**.
6. Set an expiration date, if applicable.
7. Click **OK** to cache the change, then click **Save and Deploy** to implement it.

How do I block or permit a URL for one person?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Super Administrators can use the following steps to block or permit a URL for a single client in the network, regardless of the client's role.

1. Go to the **Policy Management > Exceptions** page and click **Add**.
2. Enter a unique **Name** for the exception.
3. Enter the **URL** that you want to permit or block.
4. To specify the client affected by this exception, select **Specific clients in any role**.
5. Enter all or part of the user name or IP address in the search box above the **Defined clients** list, then press **Enter**.
 - If the client appears in the search results, select the client and click the right arrow (>) button to place the client in the **Selected** list.
 - If the client does not appear in the search results, click **Add Other Clients**, then:
 - Select a user or group name from the list, or click **Search** to find a user or group in your user directory.
 - Enter an IP address or range in either IPv4 or IPv6 format.

When you have identified the client that you want to add, use the appropriate right arrow (>) button to move the client to the Selected list, then click **OK**.

6. By default, the exception is set to **Block** the URL. To change this, set the **Type** to **Permit**.
7. Set an expiration date, if applicable.
8. Click **OK** to cache the change, then click **Save and Deploy** to implement it.

How do I block or permit a URL for my entire role?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Delegated administrators can use the following steps to block or permit a URL for all managed clients in the role they manage:



Important

Exceptions created by a Super Administrator may take precedence over exceptions created by a delegated administrator.

If you create an exception that does not seem to be applied to your managed clients, use the **Test Filtering** tool to see if another exception is overriding the one that you created (see [Test Filtering](#), page 273).

1. Go to the **Policy Management > Exceptions** page and click **Add**.
2. Enter a unique **Name** for the exception.
3. Enter the **URL** that you want to permit or block.
4. By default, the exception is set to apply to **All managed clients in this role**.
5. By default, the exception is set to **Block** the URL. To change this, set the **Type** to **Permit**.
6. Set an expiration date, if applicable.
7. Click **OK** to cache the change, then click **Save and Deploy** to implement it.

How do I block or permit a URL for one of my managed clients?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Delegated administrators can use the following steps to block or permit a URL for one of their managed clients.



Important

Exceptions created by a Super Administrator may take precedence over exceptions created by a delegated administrator.

If you create an exception that does not seem to be applied to your managed clients, use the **Test Filtering** tool to see if another exception is overriding the one that you created (see [Test Filtering](#), page 273).

1. Go to the **Policy Management > Exceptions** page and click **Add**.
2. Enter a unique **Name** for the exception.
3. Enter the **URL** that you want to permit or block.
4. To specify the client affected by this exception, select **Specific clients in this role**.
5. Enter all or part of the user name or IP address in the search box above the **Defined clients** list, then press **Enter**.
 - If the client appears in the search results, select the client and click the right arrow (>) button to place the client in the **Selected** list.
 - If the client is a member of a group, OU, or network (IP address range) defined as a managed client in your role, but does not explicitly appear in your Managed Clients list or on your Clients page, that client will not appear in your search results.

In this case, cancel creation of the exception, add the client to your Clients page, then create the exception. This time, the client will appear in your search results on the Add Exceptions page.
6. By default, the exception is set to **Block** the URL. To change this, set the **Type** to **Permit**.
7. Set an expiration date, if applicable.
8. Click **OK** to cache the change, then click **Save and Deploy** to implement it.

9

Block Page Management

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Blocking graphical advertisements, page 181](#)
- ◆ [Blocking embedded pages, page 181](#)
- ◆ [Working with block pages, page 182](#)
- ◆ [Creating alternate block messages, page 188](#)
- ◆ [Using an alternate block page on another machine, page 189](#)
- ◆ [Determining why a request was blocked, page 190](#)

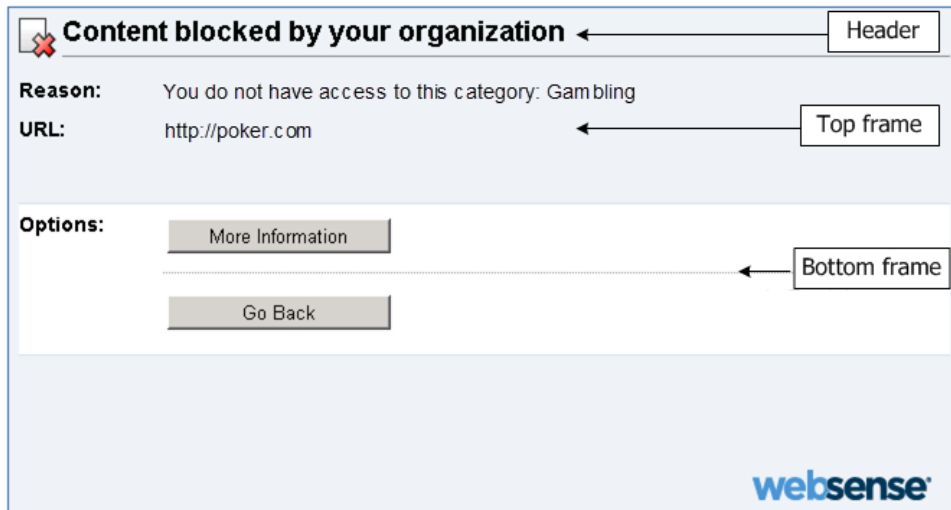
When Websense components block a website, a block page appears in the client's browser.



Note


The block page does not display correctly for IPv6-only clients. The request is blocked, but the user receives a browser error rather than a block page. Dual-stack (IPv6 and IPv4) clients receive the block page as expected.

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 36](#)), the block page has a special security banner.



With TRITON AP-WEB, Super Administrators can enable an enhanced version of the block page that includes a link to Websense ACEInsight.

- ◆ Enable the link on the **Settings > General > Filtering** page.
- ◆ Users can click the link find more information about URLs blocked for security reasons.

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



Note

With the Web Hybrid module, changes to the on-premises block pages do not affect hybrid block pages. See [Customizing hybrid block pages, page 205](#).

- ◆ Customize the default files to change the block message (see [Working with block pages, page 182](#)).

- ◆ 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 189).

Blocking graphical advertisements

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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\Wesense\Web Security\BlockPages\Images or /opt/Wesense\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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 (popup) 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.

Working with block pages

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Customizing the block message, page 184](#)
- ◆ [Creating alternate block messages, page 188](#)
- ◆ [Using an alternate block page on another machine, page 189](#)

The files used to create Websense block pages are stored in the following directory:

- ◆ Windows:

```
C:\Program Files\Websense\Web Security\BlockPages\  
<language_code>\Default
```

or

```
C:\Program Files(x86)\Websense\Web Security\BlockPages\  
<language_code>\Default
```

- ◆ Linux:

/opt/Websense/BlockPages/<language_code>/Default



Note

With the Web Hybrid module, these block pages are applied only to requests managed by on-premises components. To customize the pages provided by the hybrid service, see [Customizing hybrid block pages, page 205](#).

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 popups (like the account override popup)
popup.html	When an embedded page is blocked (see Blocking embedded pages, page 181), this file is used to display the full-sized block page popup.
block.inl	Provides tools used in constructing the block frame of the block page
blockframe.inl	Provides additional information for standard block pages

File Name	Description
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

If you have the Web DLP module, an additional file, **policyViolationDefaultPage.html**, provides block page content when Web DLP components block content from being posted or downloaded.

Customizing the block message

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Changing the size of the message frame, page 185](#)
- ◆ [Changing the logo that displays on the block page, page 186](#)
- ◆ [Using block page content variables, page 186](#)
- ◆ [Reverting to the default block pages, page 188](#)

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 organizations 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:
Websense/Web Security/BlockPages/en/Default

2. Copy the block page files to the custom block page directory. For English:
`Websense/Web Security/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 Websense Filtering Service (see [Stopping and starting Websense services](#), page 370, for instructions).

Changing the size of the message frame

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 370).

Changing the logo that displays on the block page

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 370).

Using block page content variables

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Variable Name	Content Displayed
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 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_QUOTACGIDATA	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
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

Variable Name	Purpose
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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 370).

Creating alternate block messages

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with block pages](#), page 182
- ◆ [Customizing the block message](#), page 184

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 189).

After creating alternate block message files, you must configure Websense software to display the new messages (see [Configuring filtering settings](#), page 50). During this process, you can specify which message is used for each of the configurable protocols.

Using an alternate block page on another machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with block pages](#), page 182
- ◆ [Customizing the block message](#), page 184
- ◆ [Creating alternate block messages](#), page 188

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 370).
2. On the Filtering Service machine, navigate to the Websense **bin** directory (C:\Program Files\Wesense\Web Security\bin or /opt/Websense/bin/, by default).
3. Create a backup copy of the **eimserver.ini** file and store it in another directory.
4. Open **eimserver.ini** file in a text editor, and locate the **[WebsenseServer]** section (at the top of the file).

5. Enter either the hostname or the IP address of the server hosting the block page in the following format:

UserDefinedBlockPage=http://<hostname 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 requests handled by the appliance or on-premises software), click **More information**. Next, right-click anywhere in the message text and select **View Source**. See [Request blocked by Filtering Service](#), page 190.



Note

With Internet Explorer 10, the View Source option is not always available. If the View Source option does not appear, click **Page Tools** and select **View on the desktop**.

- ◆ If the block page was sent by the hybrid service, right-click anywhere in the block message and select **View Source**. See [Request blocked by the hybrid service](#), page 191.

Request blocked by Filtering Service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The HTML source for the more information block page shows information about who requested the site, and what criteria were used to manage 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 filtering](#)

[settings](#), page 50.

- ◆ 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, Content Gateway analysis, and so on).

For example:

```
User Name: LDAP://subdomain.company.com
OU=Users,OU=HQ,DC=company,DC=copy/Lastname\, First
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 block action is based on the Default policy because no policy is assigned to the user or IP address. The policy assignment was performed in the Super Administrator role, and the requested site was categorized by the Master Database.

Request blocked by the hybrid service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The HTML source for the block page sent by the hybrid service 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 312.
- ◆ 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
- ◆ 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 was assigned 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.

10

Configure the Hybrid Service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Adding the Web Hybrid module to TRITON AP-WEB gives you a flexible, comprehensive security solution that combines on-premises and cloud-based policy enforcement, and lets you configure everything in your on-premises TRITON Manager.

An organization might use the robust on-premises software to protect a main office or campus, while smaller regional offices or satellite locations send their Internet requests through the hybrid service in the cloud. The hybrid service is also useful for users who are off-network, such as telecommuters, those who travel for business, and so on (see [Hybrid service management of off-site users](#), page 227).

After installation, complete the following steps to enable hybrid policy enforcement:

1. [Activate your hybrid service account](#), page 194
2. [Define hybrid service filtered locations](#), page 195
3. [Specify sites not managed by the hybrid service](#), page 200 (if any)
4. [Configure user access to the hybrid service](#), page 202
5. [Identification and authentication of hybrid users](#), page 300
6. [Send user and group data to the hybrid service](#), page 209

In order to ensure that the hybrid service has current policy, user, and group information, and that the on-premises reporting software has reporting data from users managed by the hybrid service, see [Schedule communication with the hybrid service](#), page 216.

Activate your hybrid service account

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Define hybrid service filtered locations, page 195](#)
- ◆ [Specify sites not managed by the hybrid service, page 200](#)
- ◆ [Configure user access to the hybrid service, page 202](#)
- ◆ [Send user and group data to the hybrid service, page 209](#)
- ◆ [Schedule communication with the hybrid service, page 216](#)

Before you can configure the hybrid service to start managing 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 components of TRITON AP-WEB and the hybrid service in the cloud.

Use the Hybrid Service section of the **Settings > General > Account** page to provide the contact email address and country for your TRITON AP-WEB administrators.

The email address is typically an alias monitored by the group responsible for managing web protection for your organization. It is very important that 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 the hybrid service.
- ◆ 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 the hybrid service for your account, you can specify which locations (identified by IP address, IP address range, or subnet) are managed by the hybrid service, how information is exchanged between on-premises and cloud components, how users managed by the hybrid service are authenticated, and more.

Define hybrid service filtered locations

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Specify sites not managed by the hybrid service, page 200](#)
- ◆ [Configure user access to the hybrid service, page 202](#)
- ◆ [Schedule communication with the hybrid service, page 216](#)

Use the **Settings > Hybrid Configuration > Filtered Locations** page to review, add, or edit information about the locations whose Internet traffic is managed by the hybrid service.

A **filtered location** is the external IP address, IP address range, or subnet from which Internet requests appear to originate when seen by the hybrid service. Hybrid policy enforcement can be applied to off-site users, regardless of how requests from those users are managed when they are in-network.

- ◆ For users managed by the hybrid service both in and outside the network, enter their in-network location details and specify that the location managed by the hybrid service. When off-site users make an Internet request, they are prompted to log on to the hybrid service so that the appropriate user or group-based policy can be applied.

Because the hybrid service is hosted outside your network, any locations managed by the hybrid service must be **external** addresses, visible from the Internet.

Locations managed by the hybrid service:

- Are public-facing IP addresses
- 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
- ◆ For users whose requests are managed 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 forwarding an Internet request.

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. Specify that these users are managed by local Websense software and define whether their on-premise policy enforcement 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

(hostname 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 [Adding or editing hybrid service filtered locations](#), page 196.
- ◆ To define a new location, click **Add**, and then see [Adding or editing hybrid service filtered locations](#), page 196.
- ◆ 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 hybrid service explicit proxies](#), page 198.

If you have added or edited a location entry, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Adding or editing hybrid service filtered locations

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Filtered Locations > Add Filtered Location** or **Edit Filtered Locations** page to:

- ◆ Define a location either managed the hybrid service (like a branch office, remote site, or satellite campus), or that contains users managed by the hybrid service when off site.
- ◆ Change the way a location managed by the hybrid service is defined.

To define a filtered location, or update an existing entry:

1. Enter, review, or update the location **Name**. The name must be unique, and have between 1 and 50 characters. It cannot include any of the following characters:
* < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
Names can include spaces, dashes, and apostrophes.
2. Enter, review, or update the 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. Select or verify the **Time zone** of the filtered location. Time zone information is used in applying policies, to ensure that the correct filters are applied at the appropriate time.

Each location whose requests go through the hybrid service can have a different time zone setting. Locations with 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 or verify the method used to define this location: **IP address**, IP address **Range**, 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.

5. Enter, verify, or update the external IP address, range, or subnet of the firewall or firewalls through which filtered clients at this location access the Internet.
 - 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 managed 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.

- Do not include the IP address of any Content Gateway machine.
 - 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. Specify, verify, or update how the requests from the location are managed: using the hybrid service, or using local Websense software.
 7. If the site is managed by local Websense software, select, verify, or update the proxy mode for this location: using 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 hybrid service explicit proxies](#), page 198, 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 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.

8. 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 and Deploy**.

Managing hybrid service explicit proxies

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 hostname, 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 [Adding or editing a hybrid service explicit proxy, page 198](#).
- ◆ To define a new explicit proxy, click **Add**, and then see [Adding or editing a hybrid service explicit proxy, page 198](#).
- ◆ 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 the filtered locations that reference it to remove it from their Explicit Proxy Configuration table.

Adding or editing a hybrid service explicit proxy

Administrator Help | TRITON AP-WEB | Version 8.1.x

When managing explicit proxies, use the **Add Explicit Proxy** or **Edit Explicit Proxy** page to define or update information about an on-premises explicit proxy to be used for your filtered locations.

1. Enter, verify, or update the proxy **Name**. The name must be unique, and have between 1 and 50 characters. It cannot include any of the following characters:

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

Names can include spaces, dashes, and apostrophes.

2. Enter, 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 hostname (for example my.example.com)

The IP address or name can include a port number, for example 123.45.67.89:443.

3. Enter or update the proxy port or ports. There must be 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.

Configuring failover to the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

It is possible to configure failover to the hybrid service for filtered locations that use explicit proxies. This ensures that users are able to access the Internet and policy enforcement always occurs, even if on-premises proxies are unavailable.

Failover to the hybrid service for a filtered location must be approved, to ensure that Websense services can provision the correct number of users at the data center nearest to your location. Once failover for a filtered location has been approved, it does not need to be re-approved if you change the failover details or later disable and then re-enable failover.

To configure failover to the hybrid service:

1. On the **Hybrid Configuration > Filtered Locations** page, select a filtered location name to edit it. This must be a location managed by local Websense software with the proxy mode set to Explicit.
2. Click **Advanced**.
3. Mark **Enable failover to hybrid service**.
4. Enter the **Number of users filtered by this filtered location**.
5. Select the **Nearest data center** to the filtered location.
6. 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 and Deploy**.

When failover for a filtered location is approved, an alert appears on the System dashboard and on the Status > Alerts page. You can view the approval status of all failover requests on the Status > Hybrid Service page.



Note

If automatic proxy caching is disabled in Internet Explorer, end users may notice a delay on every page they visit as the browser checks the list of proxies. When automatic proxy caching is enabled, the browser checks the proxy list only on startup. For more information, see the Microsoft article at <http://support.microsoft.com/kb/271361>.

Specify sites not managed by the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Define hybrid service filtered locations, page 195](#)
- ◆ [Send user and group data to the hybrid service, page 209](#)
- ◆ [Schedule communication with the hybrid service, page 216](#)

Use the **Settings > Hybrid Configuration > Unfiltered Destinations** page to review, add, or edit information about target sites to which you want to grant clients unrestricted 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 service password can retrieve it via email.
-

Destinations listed here are added to the Proxy Auto-Configuration (PAC) file that defines how users' browsers connect to the hybrid service (see [Configure user access to the hybrid service, page 202](#)). By default, the PAC file excludes all non-routable and multicast IP address ranges from policy enforcement. 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 [Adding or editing hybrid service unfiltered destinations, page 201](#).
- ◆ To define a new location, click **Add**, and then see [Adding or editing hybrid service unfiltered destinations, page 201](#).
- ◆ 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 and Deploy**.

Adding or editing hybrid service unfiltered destinations

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Unfiltered Destinations > Add Unfiltered Destination** or **Edit Unfiltered Destination** page to define or change the URL or URLs that users can access directly, without sending a request to the hybrid service or an on-premises explicit proxy.

1. Enter, verify, or update the destination **Name**. The name must be unique, and have between 1 and 50 characters. It cannot include any of the following characters:

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

Names can include spaces, dashes, and apostrophes.

2. Enter, verify, or update the short **Description** of the destination. This appears next to the unfiltered destination name on the Unfiltered Destinations page, and should clearly identify the 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, verify, or update 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, verify, or update 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 or verify the **Proxy** type that this unfiltered destination applies to.
 - Select **Hybrid** to enable all hybrid 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 managed 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 and Deploy**.

Configure user access to the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

To use the hybrid service for policy enforcement, you must configure how users connect to and are managed 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 users' browsers retrieve the PAC file (see [What is the hybrid PAC file?](#), page 207).

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 managed by the hybrid service](#), page 200). The PAC file also contains information about filtered locations, and the proxy configuration for any locations that manage Internet access for 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.

The default PAC file is retrieved over port 8082. If users request this PAC file from a location where port 8082 is locked down, they cannot access it. In this case, use the second PAC file address in this section, which enables the user to access the PAC file and hybrid service over port 80. Remote users should also use the PAC file address for port 80 if requesting access from a network that has port 8081 locked down. Even if they can access the PAC file on port 8082, port 8081 is the standard port required to be able to use the hybrid service.



Important

If you are using PingFederate as the identity provider for single sign-on, the PAC file defined for port 8082 is the only PAC file that can be used.

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 hybrid service filtered locations](#), page 195)
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 205).

Use the **Certificate Verification Bypass for HTTPS Sites** section to choose whether or not to use certificate verification and, when enabled, whether and how end users can bypass certificate verification failures (see [Configuring certificate verification bypass](#), page 206).

Use the **HTTPS Notification Pages** section to enable users making HTTPS requests to view the appropriate Websense notification pages (see [Enabling hybrid HTTPS notification pages](#), page 206).

If the hybrid service 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 (see [Send user and group data to the hybrid service](#), page 209). 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 receive the hybrid service Default policy.

- ◆ Click **Add** to add a domain (see [Adding domains for hybrid self-registration](#), page 204).
- ◆ Click a domain entry to edit the domain or its attributes (see [Editing domains for hybrid self-registration](#), page 204).

You can also apply hybrid policy enforcement 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 the hybrid service for off-site users**.

If you clear this check box, any user connecting from an unknown IP address will not be filtered.

See [Hybrid service management of off-site users](#), page 227 for more information.

By default, end user web traffic is routed to the nearest cloud data center based on the egress IP address of your Domain Name Server (DNS). This may mean that traffic for users in a geographic location different from the DNS is not optimally routed, causing

some latency issues. Select **Route traffic based on end users' egress IP** on the **Settings > Hybrid Configuration > User Access** to re-route your web traffic to data centers based on the location of the end user, rather than your DNS.

Adding domains for hybrid self-registration

Administrator Help | TRITON AP-WEB | Version 8.1.x

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) to the hybrid service. This is typically enabled only in organizations that do not use Directory Agent to send user information to the hybrid service.

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 service 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 and Deploy**.

Editing domains for hybrid self-registration

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **User Access > Edit Domain** page to make changes to the domain entries that allow users to self-register for the hybrid service.

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 and Deploy**.

Customizing hybrid block pages

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 block page, create a directory named **logo** in the Websense **ssdata** directory (by default, C:\Program Files\Wellsense\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 0 KB and smaller than 50 KB 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 222).

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 and Deploy**.

Configuring certificate verification bypass

Administrator Help | TRITON AP-WEB | Version 8.1.x

The hybrid service verifies certificates for HTTPS sites that it has decrypted and analyzed. Certificate verification checks apply to all certificates in the trust chain. Use of certificate verification is recommended in order to avoid security risks from malicious sites with certificates that misrepresent their identity.

If certificate verification fails, a notification page displays indicating that a certificate error has been detected. End users can be given the option to bypass certificate errors for specified sites. They can proceed to the site or go back.

You can create a list of sites that do not return a notification page for certificate errors. Instead the user is given access to the site. This option is useful, for example, for sites that you trust even if the certificate is expired, is not yet valid, or is self-signed.

1. Click **On** under **Perform certificate verification** to enable the feature. Click **Off** to disable it.
2. Select **Provide end users an option to bypass all certificate errors** to provide all users with the notification page that includes an option to bypass a certificate error and proceed to the site.
3. If you have selected **Perform certificate verification**, you can maintain a list of domains or IP addresses for which certificate verification errors are automatically bypassed. The end user receives no notification page and is given access to the site.

Enter the domains and IP addresses in the entry field provided.

A comma-separated list can be used, but IP address ranges are not supported.

Click **Add** to populate the list.

Select an item on the list and click **Delete** to remove it.

4. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Enabling hybrid HTTPS notification pages

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 the hybrid service, the hybrid proxy can establish SSL channels with most browsers 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 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**

With single-sign on, end users require this root certificate to ensure seamless authentication to HTTPS sites. If the certificate is not installed for single sign-on users, they must authenticate using NTLM identification or manual authentication, depending on the settings on the Hybrid User Identification page. See [Integrating the hybrid service with a single sign-on identity provider](#).

To install the hybrid root certificate on all clients using the hybrid service:

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 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 and Deploy**.

What is the hybrid PAC file?

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 policy enforcement contains a number of global settings and allows you to configure sites (for example, intranet sites or organizational webmail) that users can access directly, without sending the request to the hybrid service (see [Specify sites not managed by the hybrid service](#), page 200).

If you want to use the hybrid service 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 (see [Configure user access to the hybrid service](#), page 202).

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, 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, 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 the Websense hybrid service, and comprises default settings 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_`.



Important

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 222).

If you are unfamiliar with PAC files, Wikipedia has a good introductory article, and a good website for more information and several example PAC files is <http://www.findproxyforurl.com/>.

Send user and group data to the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 the hybrid service.
- ◆ **Websense Sync Service** transports policy, reporting, custom PAC file information, and user/group data between the on-premises and hybrid systems.

When the hybrid service is configured properly, the information from Directory Agent can be used to apply user- and group-based policies.

If your organization uses Windows Active Directory in mixed mode, user and group data cannot be collected and sent to the hybrid service.

If the hybrid service 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.

For this reason, be sure that your organization's webmail address has been added as an unfiltered destination. See [Specify sites not managed by the hybrid service](#), page 200.

Configure Directory Agent settings for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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, also do the following manual step to ensure that the server is fully removed from Directory Agent settings:

- ◆ Software deployments: Delete all files in the **Websense/Web Security/bin/snapshots** directory. Then go to **Settings > Hybrid Configuration > Scheduling**, and click **Send** under Send Update Now.
 - ◆ Appliance deployments: Contact Websense Technical Support for assistance.
-

To refine the way that Directory Agent searches the directory and packages results for the hybrid service, click an IP address or hostname in the table. See [Configure how data is gathered for the hybrid service, page 211](#).

To view the global catalog directory contexts defined for identifying hybrid users, click **View Context** under Contexts in the table. See [Adding and editing directory contexts for the hybrid service, page 213](#).

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 hostname 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 and Deploy**.

Configure how data is gathered for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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.

Under Root Context for Hybrid Service 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 directory contexts for the hybrid service](#), page 213.



Warning

There is a limit to how many groups the hybrid service can support. The limit is affected by a number of factors, but if it is exceeded, the service fails open, permitting all requests.

If your organization has a large directory forest with thousands of groups, be sure to configure Directory Agent to upload only the information required to manage the users whose requests are sent to the hybrid service. You might select only specific groups to upload, or set a specific and narrowed root context.

It is best to provide contexts that include only users managed 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 directory search results for the hybrid service](#), page 215, for more information.

Oracle (Sun Java) Directory Server and the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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
```

Remove the comment symbol (#) from the GroupMembershipAttribute and MemberOfAttribute parameters, then save the file and restart Directory Agent.

1. Under Root Context for Hybrid Service 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 directory contexts for the hybrid service](#), page 213.

Provide a context that includes only users managed 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 IPv4 address or hostname 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 directory search results for the hybrid service](#), page 215, for more information.

Novell eDirectory and the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 Service 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 directory contexts for the hybrid service](#), page 213.

Provide a context that includes only users managed 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).

When Sync Service and Directory Agent connect to the same Policy Server, these fields are populated automatically. In other cases, update the fields manually.

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 IPv4 address or hostname 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 directory search results for the hybrid service](#), page 215, for more information.

Adding and editing directory contexts for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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.



Warning

There is a limit to how many groups the hybrid service can support. The limit is affected by a number of factors, but if it is exceeded, user requests are not handled properly.

If your organization has a large directory forest with thousands of groups, be sure to configure Directory Agent to upload only the users whose requests are sent to the hybrid service.

You can select multiple contexts within the directory. It is best to include contexts that include only users managed 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 select 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 service context, for example to include only users that are members of a group managed by the hybrid service.

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, such as an administrator group with multiple email addresses in a record. You can only set a context as an exclude context if it is within an included directory context.

1. Expand the Directory Entries tree to locate the context you want to use when gathering user and group data from the directory. Narrow the context to increase speed and efficiency.
Use the search field to locate the context name if required. You can search on OUs, groups, users, or all directory entries. If multiple contexts appear in the search results, select a context and click **Show in Tree** to see the context's location in the Directory Entries tree.
2. Mark the context, then click **Set as Include Context**.
3. In the popup window that appears, indicate how far below the root context Directory Agent looks for users and groups.
 - Select **Context Only** 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 groups or OUs to **Set as Include Context**, and then selected One Level or All Levels for group searches, the **Include all users in selected groups, regardless of context** option is enabled. Check the box 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.
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. Click **OK** to save the directory context.
8. When you specify that a context is included, by default any contexts below that context in the tree are also included. To exclude a context within an included context, mark the context that should not be sent to the hybrid service, and click **Set/Edit/Remove Exclude Context**. You can select multiple contexts if required.
9. In the popup window that appears, note that **Set as exclude context** is selected. The **Remove exclude context** option is available only when you select an existing excluded context and click **Set/Edit/Remove Exclude Context** to edit it.

10. Indicate how far below the excluded context Directory Agent looks for users and groups.

- Select **Context Only** to limit searches to the specified context only.
- Select **One Level** to limit searches to the specified context and one level below.
- Select **All Levels** to expand searches to the specified context and all levels below.

Note that the user and group levels for an excluded context cannot be greater than the defined levels for its root context. For example, if the root context's Directory Search level for either users or groups is set to Context Only, the corresponding users or groups search level for the excluded context are also set to Context Only and cannot be changed.

If you select All Levels for both users and groups, everything below the selected context is excluded and you cannot browse further levels of the Directory Entries tree.

11. If only groups are specified as **exclude** contexts, and **One** or **All** levels have been selected for exclusion, use the **Exclude all users in selected groups, regardless of context** option to determine whether:

- (Check box marked) Users in exclude contexts are always excluded, regardless of whether they are also defined in other (included) contexts.
- (Check box cleared) Users in exclude contexts are not excluded when they are also defined in other (included) contexts.

12. 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 Service Users table. You must also click **OK** on the Shared User Data page to cache the change.

Click a link on the Root Context for Hybrid Service Users table to access the **Edit Context** page for the selected context.

Optimizing directory search results for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 the hybrid service when they are off site (see [Configuring the hybrid service for off-site users](#), page 228).

Any search filters that you create in the Web module of the TRITON Manager 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 popup 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 the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 and Deploy** in the Web module of the TRITON Manager, 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. A success message indicates 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 requests are logged, page 386](#), for details.

Threat data is not provided by the hybrid service.

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 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 the hybrid service makes Internet requests and when those requests appear in 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 hostname 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 and Deploy**.

Define custom authentication settings for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Settings > Hybrid Configuration > Custom Authentication** page to add and edit custom rules to change the default authentication behavior for specific applications or sites. These rules only apply to users browsing from a filtered location.

Occasionally, some Internet applications and websites cannot authenticate with the hybrid service. This might occur with, for example, instant messaging programs, antivirus updates, or software update services.

To allow particular applications that do not properly handle authentication challenges to bypass authentication, you can specify user agents, domains, or URLs, or a combination of these options.

A user agent is a string sent from your browser or Internet application to the server hosting the site that you are visiting. This string indicates which browser or application you are using, its version number, and details about your system, such as the operating system and version. The destination server then uses this information to provide content suitable for your specific browser or application.

For example, this is a user agent for Firefox:

```
Mozilla/5.0 (Windows NT 6.1; WOW64; rv:27.0) Gecko/20100101  
Firefox/27.0)
```

In this example, Windows NT 6.1 indicates that the operating system is Windows 7.

To get the user agent string for your browser, enter the following in the browser's address bar:

```
javascript:alert(navigator.userAgent)
```

You can view the user agents that have made authentication requests via the hybrid service in the User Agents by Volume report, available from the Custom Authentication page and also on the **Main > Status > Hybrid Service** page. If a user agent in this report has a high number of authentication requests, it may be experiencing authentication problems. You can select a user agent in the report and click **Create Rule** to add a new custom authentication rule for that agent. See [View the hybrid service User Agent Volume report](#), page 224.

To define a custom authentication rule, click **Add**, and then see [Adding custom authentication rules for the hybrid service](#), page 219.

To edit an existing rule, click the rule **Name**, and then see [Editing custom authentication rules for the hybrid service](#), page 220.

To remove a custom authentication rule, mark the check box next to the rule name, and then click **Delete**.

If you have added or edited a custom authentication rule, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Adding custom authentication rules for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Custom Authentication > Add Custom Authentication Rule** page to define one or more user agents, domains, or URLs that are failing to authenticate with the hybrid service.

1. Enter a **Name** for the rule. 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. Define the **User agents**, if any, for the rule:
 - To match against all user agent strings, select **All user agents**. You might do this to set up a custom rule that applies to all browsers on all operating systems in your organization.
 - If the application does not send a user agent string to the Internet, select **No user agent header sent**.
This option matches against all applications that do not send a user agent. In this case, refine the rule by entering one or more URLs or domains in the **Destinations** field.
 - To apply the custom authentication to one or more user agents, select **Custom user agents**. Enter each user agent on a separate line. Use the asterisk wildcard to match one line to multiple user agent strings, for example Mozilla/5.0*.



Note

If you are creating a new rule directly from the User Agents by Volume report, the user agents you selected in the report are already entered in this field.

3. Define the URLs or domains (if any) for the rule in the **Destinations** field:
 - To match against all URLs and domains, select **All destinations**. You might want to do this if you are setting up a custom rule that applies to a specific user agent that accesses multiple sites.
 - To apply the custom authentication to one or more specific domains or URLs, select **Custom destinations**. Enter each URL or domain on a separate line.
URLs must include the protocol portion (http://) at the beginning and a forward slash (/) at the end (for example, http://www.google.com/). If these elements are not present, the string is treated as a domain. Domains cannot include a forward slash at the end (for example, mydomain.com).

Use the asterisk wildcard to match one line to multiple destinations: for example, entering *.mydomain.com would match against all domains ending in “mydomain.com.”

4. Select the **Authentication method** for the custom rule.

**Note**

The authentication method you select must be enabled on the Hybrid User Identification page.

- **Default:** Uses your default authentication method.
 - **NTLM:** Uses NTLM identification for the specified user agents and destinations. If an application is not NTLM-capable, basic authentication is used instead.
 - **Secure form authentication:** Uses secure form authentication to display a secure logon form to the end user. For more information, see [Identification and authentication of hybrid users](#), page 300.
 - **Basic authentication:** Uses the basic authentication mechanism supported by many Web browsers. No welcome page is displayed. For more information on basic authentication, see [Identification and authentication of hybrid users](#), page 300.
 - **Welcome page:** Displays a welcome page to users before they use basic authentication to proceed.
 - **None:** Bypasses all authentication and identification methods in the hybrid service. Select this option for Internet applications that are incapable of authentication.
5. Optionally, select **Bypass content scanning** to bypass all filtering for the specified user agents and destinations.

**Important**

Select this option **only** for applications and sites that for some reason do not work well with the hybrid service, and that you trust implicitly. Selecting this option could allow viruses and other malware into your network.

6. Click **OK** to return to the Custom Authentication page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Editing custom authentication rules for the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Custom Authentication > Edit Custom Authentication Rule** page to edit user agents, domains, or URLs that are failing to authenticate with the hybrid service.

1. If you make changes to the rule **Name**, ensure it is between 1 and 50 characters long, and does not include any of the following characters:
 * < > { } ~ ! \$ % & @ # . " | \ & + = ? / ; : ,
 Names can include spaces, dashes, and apostrophes.
2. Define or update the **User agents**, if any, for the rule:
 - To match against all user agent strings, select **All user agents**. You might want to do this if you are setting up a custom rule that applies to all browsers on all operating systems in your organization.
 - If the application does not send a user agent string to the Internet, select **No user agent header sent**.
 This option will match against all applications that do not send a user agent. In this case, we recommend you refine the rule by entering one or more URLs or domains in the **Destinations** field.
 - To apply the custom authentication to one or more user agents, select **Custom user agents**. Enter each user agent on a separate line. Use the asterisk wildcard to match one line to multiple user agent strings, for example Mozilla/5.0*.
3. Define or update the URLs or domains (if any) for the rule in the **Destinations** field:
 - To match against all URLs and domains, select **All destinations**. You might want to do this if you are setting up a custom rule that applies to a specific user agent that accesses multiple sites.
 - To apply the custom authentication to one or more specific domains or URLs, select **Custom destinations**. Enter each URL or domain on a separate line.
 URLs must include the protocol portion (http://) at the beginning and a forward slash (/) at the end (for example, http://www.google.com/). If these elements are not present, the string is treated as a domain. Domains cannot include a forward slash at the end (for example, mydomain.com).
 Use the asterisk wildcard to match one line to multiple destinations: for example, entering *.mydomain.com would match against all domains ending in 'mydomain.com.'
4. Verify or update the **Authentication Method** for the custom rule.
 - **Default:** Uses your default authentication method.
 - **NTLM:** Uses NTLM identification for the specified user agents and destinations. If an application is not NTLM-capable, basic authentication is used instead.

**Note**

You must have NTLM identification enabled for your account to use this option.

- **Form Authentication:** Uses secure form authentication to display a secure logon form to the end user. For more information, see [Identification and authentication of hybrid users](#), page 300.

- **Basic Authentication:** Uses the basic authentication mechanism supported by many Web browsers. No welcome page is displayed. For more information on basic authentication, see [Identification and authentication of hybrid users, page 300](#).
 - **Welcome Page:** Displays a welcome page to users before they use basic authentication to proceed.
 - **None:** Bypasses all authentication and identification methods in the hybrid service. Select this option for Internet applications that are incapable of authentication.
5. Optionally, select **Bypass content scanning** to bypass all filtering for the specified user agents and destinations.

**Important**

Select this option **only** for applications and sites that for some reason do not work well with the hybrid service, and that you trust implicitly. Selecting this option could allow viruses and other malware into your network.

6. Click **OK** to return to the Custom Authentication page, and then click **OK** again to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Monitor communication with the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

You can view the status of the hybrid service on the **Web > Status > Dashboard > Hybrid Service** page in the TRITON Manager. 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 on-premises and hybrid components, 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 and user agent reports from the hybrid service (see [View hybrid service authentication reports, page 223](#), and [View the hybrid service User Agent Volume report, page 224](#)), 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 the hybrid PAC file?, page 207](#))
- ◆ 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

Select **View Report** under Authentication 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. **TRITON AP-ENDPOINT**, **Single sign-on**, **NTLM identification**, **Form authentication**, and **Manual authentication** are all set up for clients on the Settings > Hybrid Configuration > Hybrid User Identification page (see [Identification and authentication of hybrid users, page 300](#)).

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 the last authenticated browsing 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 the TRITON console.

To display authentication reports in XLS format, Microsoft Excel 2003 or later must be installed on the machine from which you are accessing the console.

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**.

View the hybrid service User Agent Volume report

Administrator Help | TRITON AP-WEB | Version 8.1.x

Select **View Report** under User Agent Volume Report on the **Main > Status > Hybrid Service** page to view user agents that have made authentication requests via the hybrid service.

The report output consists of a table, showing the number of authentication requests and total requests made by each user agent. If a user agent already has a custom authentication rule associated with it, you can hover over the **Rule** column to see details of the custom rule.

You can filter the report results as follows:

- ◆ Enter a search term and click **Search**
- ◆ Select a **Time range** from the drop-down list. If you select Custom date range, select a time period between 1 and 14 days.

- ◆ Mark **View only user agents with rules** to see only the user agents that have custom authentication rules associated with them.

If there is more than one page of results, click the arrow buttons at the bottom of the report to view previous or subsequent pages.

If a user agent in this report has a high number of authentication requests, it may be experiencing authentication problems. To add a new custom authentication rule for one or more user agents in the report, mark the check box for each agent and click **Create Rule**. The user agents you select are automatically entered in the **Custom user agents** field on the Add Custom Authentication Rule page. See [Adding custom authentication rules for the hybrid service](#), page 219.

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 the TRITON console.

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

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**.

11

Manage Off-site Users

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Using remote filtering software, page 229](#)
- ◆ [Hybrid service management of off-site users, page 227](#)

Both TRITON AP-WEB and Web Filter & Security offer a means of enforcing policy when users access the Internet from outside your network.

- ◆ TRITON AP-WEB customers can use the Web Hybrid module to manage Internet activity for users outside the network, regardless of how their requests are handled when they are in the network. See [Hybrid service management of off-site users, page 227](#).
- ◆ Web Filter & Security customers can use the Remote Filter module to manage Internet activity for users outside the network. See [Using remote filtering software, page 229](#).

These methods can be used, for example, to provide policy enforcement for users who work from home, users who travel using company laptops, or students who use institutional laptops on and off campus.

Hybrid service management of off-site users

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Configuring the hybrid service for off-site users, page 228](#)
- ◆ [Off-site user self-registration, page 228](#)

With the Web Hybrid module for TRITON AP-WEB, the hybrid service in the cloud can be configured to manage off-site users, regardless of how those users requests' are handled when they are in-network.

- ◆ For users whose requests are handled 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 forwarding an Internet request.

If you are using the PAC file generated by the hybrid service, this configuration occurs automatically based on the settings that you provide in the TRITON Manager.

- ◆ For users managed by the hybrid service 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 the hybrid service so that the appropriate user or group-based policy can be applied.

Configuring the hybrid service for off-site users

Administrator Help | TRITON AP-WEB | Version 8.1.x

To configure the hybrid service to manage users outside a filtered location:

- ◆ If the hybrid service 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 209](#)), 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 228](#)).
- ◆ 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 the hybrid service, page 202](#).
- ◆ Once you have established an identification policy for off-site users, mark **Enable the hybrid service off-site users** on the **Settings > Hybrid Configuration > User Access** page in the Web module of the TRITON Manager. See [Configure user access to the hybrid service, page 202](#).
- ◆ Select **Route traffic based on end users' egress IP** on the **Settings > Hybrid Configuration > User Access** to re-route your web traffic to data centers based on the location of the end user, rather than your DNS. See [Configure user access to the hybrid service, page 202](#).

Off-site user self-registration

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 for their requests to be handled correctly 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 the Web module of the TRITON Manager (see [Configure user access to the hybrid service, page 202](#)).

Users connecting to the hybrid service 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. The hybrid service 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. The hybrid service then applies your organization's Default policy to their Internet requests.

Using remote filtering software

Administrator Help | Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Remote Filtering settings, page 230](#)

When you add the Remote Filter module to Websense Web Filter & Security, your remote filtering software components are configured to monitor HTTP, SSL, and FTP traffic by default, and to apply a user-based policy or the Default policy. Remote filtering software does not apply policies to IP addresses (computers or network ranges).

- ◆ Bandwidth restrictions are 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 231](#).

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 supported client machines that are used outside the network.

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 handled 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.

Complete information about planning for, deploying, and configuring remote filtering software is available in the [Remote Filtering Software](#) technical paper, available from the support.websense.com.

Configuring Remote Filtering settings

Administrator Help | Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configure remote filtering to ignore FTP or HTTPS traffic, page 231](#)
- ◆ [Configure the Remote Filtering Client heartbeat interval, page 232](#)

Use the **Settings > General > Remote Filtering** page to configure what happens if any Remote Filtering Client instance cannot communicate with 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.
- ◆ 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).

When Remote Filtering Client is configured to 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 Internet access at startup, during the timeout period, 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 all requests are permitted (fail open).

This timeout period allows users who pay for Internet access when traveling 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, the user must restart the computer to begin the timeout interval again.

For detailed information about how remote filtering works, which components are involved, and how to deploy components, see the [Remote Filtering Software](#) technical paper.

Configure remote filtering to ignore FTP or HTTPS traffic

Administrator Help | Web Filter & Security | Version 8.1.x

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.

1. Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) on the Remote Filtering Server machine.
2. Open the **securewisproxy.ini** file in a text editor.
3. To have this Remote Filtering Server instance ignore FTP traffic, add the following line to the file:

```
FilterFTP=0
```

If you want to later turn FTP management back on, change the parameter value from “0” to “1”.

4. To have this Remote Filtering Server instance ignore HTTPS traffic, add the following line to the file:

```
FilterHTTPS=0
```

If you want to later turn HTTPS management 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

Administrator Help | Web Filter & Security | Version 8.1.x

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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) on the Remote Filtering Server machine.
2. Open the **securewispproxy.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.

12

Combine Web, Data, and Mobile Protection

Administrator Help | TRITON AP-WEB | Version 8.1.x

TRITON AP-WEB 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? To protect against such data loss, deploy the Web DLP module. You can also help protect your end users' mobile devices from potential data loss, the possible theft of intellectual property, and from mobile malware and other threats by combining TRITON AP-WEB with TRITON AP-MOBILE integrated with AirWatch® Mobile Device Management.

Protecting against data loss

With the Web DLP module, not only can you protect sensitive data from leakage over the Web, but you can also monitor removable media devices, printers, instant messages, copy/paste operations, or email for the such data. To protect against data loss over other channels, in addition to the Web, you can purchase TRITON AP-DATA to integrate with TRITON AP-WEB.

Websense web and data protection interoperate in fundamental ways, giving the data protection software access to user information (collected by User Service) and URL categorization information (from the Master Database).

By combining web and data protection, 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](#).

For instructions on creating data security policies, see the [Help](#) for the Data module of the TRITON Manager.

Protecting end users' devices

Websense TRITON AP-MOBILE protects your end users' devices from potential data loss and the possible theft of intellectual property, plus from mobile malware, web threats, phishing attacks, spoofing, and more—all of which helps them safely access corporate resources.

When integrated with AirWatch Mobile Device Management (MDM), you can provision iOS and Android mobile devices to send traffic to the Websense cloud service for analysis and policy enforcement. You can also enroll devices in your enterprise environment quickly, configure and update device settings over the air, create different policies for corporate versus personal devices, and secure mobile devices through actions such as locking and wiping them.



Important

AirWatch Mobile Device Management (MDM) integration is available for TRITON AP-MOBILE at this time on an Early Access as-is basis only.

Integrating web and mobile protection solutions

Administrator Help | TRITON AP-WEB | Version 8.1.x

If you have TRITON AP-WEB and TRITON AP-MOBILE, set up this account to integrate with AirWatch Mobile Device Management (MDM) in the Web module of the TRITON Manager, as follows:

1. Navigate to **Settings > Hybrid Configuration > Mobile Integration > Mobile Device Management Account Setup**.

2. Select the **Integrate with MDM provider** check box.

To disable integration with AirWatch MDM, deselect the box and click **Save Now**.

3. Enter the API URL and API key. You need to obtain these from the AirWatch Console. See Step 4, Log on to the AirWatch Console in the [Getting Started Guide](#).

For the API URL, remove the “/API” from the end of the URL. For example, change:

`https://orgname.airwlab.com/API`

to

`https://orgname.airwlab.com`

4. Enter the user name and password that you use to log on to your AirWatch administrator account.



Important

If the password for the **AirWatch administrator** account changes or expires, you must enter the new password on the **Mobile Device Management Account Setup** page to maintain the integration of AirWatch MDM with the cloud service.

An alternative to using the administrator account is to create a service account in Active Directory with the password set to never expire. Use the logon name and password for this account instead of the AirWatch administrator account logon credentials.

5. Click **Save Now**.
6. After clicking Save Now and the settings are confirmed and saved successfully, this page then displays a user name and password that have been automatically generated for your hybrid account, along with a connection URL.
Copy and paste these three items into the VPN connection information section of the AirWatch Console.

Should you need to change the credentials for your hybrid account, for example, if they've been compromised, you can generate a new user name and password by clicking **Advanced Options** and then **Generate New User Name and Password**.



Important

After clicking Generate New User Name and Password but before clicking Save Now, you **must re-enter the password that you use to log on to the AirWatch Console**.

You **must also enter the new user name and password** generated for your *hybrid account* into the **VPN connection information section of the AirWatch Console** to maintain the integration of AirWatch MDM with the hybrid solution.

1. For an overview of the mobile integration process, see the [Getting Started Guide](#).

13

Refine Your Policies

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

At its simplest, Websense software requires a single policy that applies one category filter and one protocol filter 24 hours a day, 7 days a week. Your software includes tools, however, for going far beyond this basic safety net, 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 URLs](#), page 238).
- ◆ Create **custom categories** to redefine how selected sites are treated (see [Working with categories](#), page 244).
- ◆ **Recategorize URLs** to move specific sites from their default, Master Database category to another Websense-defined or custom category (see [Reclassifying specific URLs](#), page 251).
- ◆ 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 261).
With the Web Hybrid module, bandwidth-based restrictions are not enforced for requests managed by the hybrid service.
- ◆ Define **keywords** used to block sites in otherwise permitted categories when keyword blocking is enabled and activated (see [Keyword-based policy enforcement](#), page 248).
- ◆ 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 263).

Restricting users to a defined list of URLs

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Limited access filters and enforcement order, page 238](#)
- ◆ [Creating a limited access filter, page 239](#)
- ◆ [Editing a limited access filter, page 240](#)

Limited access filters provide a very precise method of granting Internet access. Each limited access filter is a list of individual URLs, IP addresses, or regular expressions. 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 websites 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 enforcement order

When multiple group policies apply to a user, the **Use more restrictive blocking** setting (see [Enforcement order, page 79](#)) determines which one is used to manage the user's requests. By default, this setting is off.

Filtering Service determines which setting is less restrictive at the filter level. In cases where a user might be assigned to 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 policy enforcement 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)

Creating a limited access filter

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with filters, page 42](#)
- ◆ [Restricting users to a defined list of URLs, page 238](#)
- ◆ [Editing a limited access filter, page 240](#)

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Restricting users to a defined list of URLs, page 238](#)
- ◆ [Limited access filters and enforcement order, page 238](#)
- ◆ [Creating a limited access filter, page 239](#)
- ◆ [Editing a policy, page 76](#)

A limited access filter is a list of URLs, IP addresses, and regular expressions, used to identify specific websites 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 URL 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 252](#).

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 URLs, IP addresses, and regular expressions 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. IP addresses may use IPv4 or IPv6 format.

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 managed 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 website'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 271](#), for detailed information about using regular expressions for policy enforcement.
 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 and Deploy**.

Adding sites from the Edit Policy page

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Restricting users to a defined list of URLs, page 238](#)
- ◆ [Limited access filters and enforcement order, page 238](#)
- ◆ [Creating a limited access filter, page 239](#)
- ◆ [Editing a policy, page 76](#)

Use the **Policies > Edit Policy > Add Sites** page to add URLs and IP addresses 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 and Deploy**.

Changes made to a limited access filter affect all policies that enforce the filter.

Copying filters and policies to roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a category filter, page 43](#)
- ◆ [Creating a protocol filter, page 46](#)
- ◆ [Creating a limited access filter, page 239](#)
- ◆ [Creating a policy, page 75](#)

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 apply the filters or policies to 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 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 popup 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 and Deploy**.

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 the TRITON

Manager. 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Policy Management > Filter Components** page to access tools used to refine and customize your policies. The 3 buttons on the screen are associated with the following tasks:

Edit Categories	<ul style="list-style-type: none"> • Recategorize a URL (see Reclassifying specific URLs, page 251). For example, if the Shopping category is blocked by your 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 247). 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 Keyword-based policy enforcement, page 248). 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 271), patterns or templates that can be used to match multiple URLs and assign them to a category.
Edit Protocols	<p>Create or edit custom protocol definitions (see Creating a custom protocol, page 259, and Editing custom protocols, page 256). 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.</p>
File Types	<p>Create or edit file type definitions, used to block files with specific extensions within otherwise permitted categories (see Managing traffic based on file type, page 263).</p>

Working with categories

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Editing categories and their attributes, page 244](#)
- ◆ [Creating a custom category, page 247](#)
- ◆ [Keyword-based policy enforcement, page 248](#)
- ◆ [Reclassifying specific URLs, page 251](#)

Websense software provides multiple methods for managing sites that are not in the Master Database, and for changing the way that individual URLs in the Master Database are handled.

- ◆ Create **custom categories** for more precise policy enforcement 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 requests are logged, page 386](#). If a category is not logged, client requests for that category do not appear in reports.

Editing categories and their attributes

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a custom category, page 247](#)
- ◆ [Reviewing all customized category attributes, page 245](#)
- ◆ [Making global category changes, page 246](#)
- ◆ [Keyword-based policy enforcement, page 248](#)
- ◆ [Reclassifying specific URLs, page 251](#)

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 245](#), for more

information.

- ◆ To create a new category, click **Add**, and then go to [Creating a custom category, page 247](#), 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 246](#)).
- ◆ To change the action associated with a category in all category filters, click **Override Action** (see [Making global category changes, page 246](#)).
- ◆ 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 [Reclassifying specific URLs, page 251](#), 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 [Keyword-based policy enforcement, page 248](#), 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 271](#)).
 - 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 and Deploy**.

Reviewing all customized category attributes

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 and Deploy**.

Making global category changes

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 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 [Actions](#), page 40, 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 [Keyword-based policy enforcement](#), page 248, for more information.
3. Specify whether or not to **Block File Types**, and customize blocking settings. See [Managing traffic based on file type](#), page 263, for more information.
4. Specify whether or not to **Block with Bandwidth Optimizer** to manage access to HTTP sites, and customize blocking settings. See [Using Bandwidth Optimizer to manage bandwidth](#), page 261, 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 244). The changes are not cached until you click **OK** on the Edit Categories page.

Renaming a custom category

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Editing categories and their attributes, page 244](#)
- ◆ [Keyword-based policy enforcement, page 248](#)
- ◆ [Reclassifying specific URLs, page 251](#)

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 policy enforcement 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 [Reclassifying specific URLs, page 251](#), 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 [Keyword-based policy enforcement](#), page 248, for more information.

You can also edit this list after creating the category.

6. Define a default **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 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** 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 and Deploy**.

The new category is added to the Categories list and custom URL and keyword information for the category is displayed.

Keyword-based policy enforcement

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Reclassifying specific URLs](#), page 251
- ◆ [Configuring filtering settings](#), page 50
- ◆ [Creating a category filter](#), page 43
- ◆ [Editing a category filter](#), page 44
- ◆ [Working with categories](#), page 244

Keywords are associated with categories, and then used to offer protection against URLs 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 filtering settings](#), page 50).
2. Define keywords associated with a category (see [Defining keywords](#), page 250).
3. Enable keyword blocking for the category in an active category filter (see [Editing a category filter](#), page 44).

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, 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Editing a category filter, page 44](#)
- ◆ [Working with categories, page 244](#)
- ◆ [Keyword-based policy enforcement, page 248](#)
- ◆ [Using regular expressions, page 271](#)

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 and Deploy**.

In order for keyword blocking to be enforced, you must also:

1. Enable keyword blocking via the **Settings > General > Filtering** page (see [Configuring filtering settings, page 50](#)).
2. Enable keyword blocking in one or more active category filters (see [Editing a category filter, page 44](#)).

Reclassifying specific URLs

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Exceptions to Web Protection Policies, page 167](#)
- ◆ [Creating a custom category, page 247](#)
- ◆ [Keyword-based policy enforcement, page 248](#)

Your software offers the option to manually change the category assigned to a URL. URLs that have been added to a new category are called custom URLs or recategorized URLs.

- ◆ Use the **Policy Management > Filter Components > Edit Categories > Recategorize URLs** page to add sites to a new category.
- ◆ Make changes to existing recategorized sites on the **Edit URLs** page.

To change the category of a URL, you can add it to:

- ◆ A different Websense-defined category
- ◆ Any custom category (see [Creating a custom category, page 247](#))

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.



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 252](#).

When you recategorize sites:

- ◆ Enter each URL or IP address on a separate line.
 - 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.
 - With recategorized URLs, the URL is not automatically matched to its equivalent IP address. To ensure that a request for a site is handled properly, specify both its URL and IP address.
- ◆ 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 return to the Edit Categories page. You must also click **OK** on the Edit Categories page to cache your changes. Changes are not implemented until you click **Save and Deploy**.

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 recategorized URL.

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 policy enforcement behavior](#), page 272.

Prioritizing Security Risk categorization

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

By default, when a site belongs to a Security Risk category, Filtering Service applies an action based on the site's Security Risk classification, even when the site:

- ◆ Is added as a recategorized URL in a permitted category
- ◆ 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 analysis):

- ◆ 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 the Web module of the TRITON Manager.

If you want requests managed based on custom categorization, regardless of whether the URL is classified as a Security Risk:

1. Navigate to the Websense **bin** directory on the Filtering Service machine (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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 tool to restart Websense Filtering Service.
 - *Linux:* Use the /opt/Websense/WebsenseDaemonControl command to stop and then start Filtering Service.

If you have the Web Hybrid module, you can also disable this feature for the hybrid service:

1. Navigate to the Websense **bin** directory on the Sync Service machine (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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 tool 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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 tool to restart Websense Filtering Service.
- *Linux*: Use the /opt/Websense/WebsenseDaemonControl command to stop and then start Filtering Service.

If you have the Web Hybrid module, you can also enable this feature for the hybrid service:

1. Navigate to the Websense **bin** directory on the Sync Service machine (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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 tool to restart Websense Sync Service.
- *Linux*: Use the /opt/Websense/WebsenseDaemonControl command to stop and then start Sync Service.

Protocol-based policy enforcement

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Editing custom protocols, page 256](#)
- ◆ [Creating a custom protocol, page 259](#)
- ◆ [Adding or editing protocol identifiers, page 257](#)
- ◆ [Adding to a Websense-defined protocol, page 260](#)

The Websense Master database includes protocol definitions used to manage 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 manage 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.

- ◆ With TRITON AP-WEB, Content Gateway can be configured to detect non-HTTP protocols that tunnel over HTTP ports. See [Configuring tunneled protocol detection, page 92](#), for more information. Network Agent can also be used to manage non-HTTP protocols.
- ◆ In Websense Web Filter & Security deployments, Network Agent must be installed to enable protocol-based policy enforcement.

How protocol requests are managed

When a protocol request is made, the following steps are used 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 any of this information cannot be determined, all content associated with the protocol is permitted.

If the protocol is FTP, HTTPS, or gopher, a check is first performed to see if the protocol is blocked. If the protocol is permitted, Filtering Service performs a URL lookup to see if the requested site is permitted or blocked.

Defining custom protocols

In addition to using Websense-defined protocol definitions, you can define custom protocols. 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**. When ports are used to define a protocol, blocking the 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.

To work with custom protocol definitions, go to **Policy Management > Filter Components**, and then click **Protocols**. See [Editing custom protocols](#), page 256, and [Creating a custom protocol](#), page 259, for details.

Editing custom protocols

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Protocol-based policy enforcement](#), page 255
- ◆ [Creating a custom protocol](#), page 259
- ◆ [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 259.

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 action applied to this protocol in all protocol filters (see [Making global protocol changes](#), page 258).
3. Click **Add Identifier** to define additional protocol identifiers for this protocol (see [Adding or editing protocol identifiers](#), page 257).
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 and Deploy**.

To delete a protocol definition, select an item in the Protocols list, and then click **Delete**.

Adding or editing protocol identifiers

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 a TRITON AP-WEB 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 and Deploy**.

Renaming a custom protocol

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 changes

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 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 [Actions, page 40](#), 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 261](#), 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 256](#)). You must also click **OK** on the Edit Protocols page to cache the changes.

Creating a custom protocol

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Protocol-based policy enforcement](#), page 255
- ◆ [Editing custom protocols](#), page 256
- ◆ [Adding to a Websense-defined protocol](#), page 260

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 260, 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

In Web Filter & Security deployments, use caution when defining a protocol on port 80 or 8080. Network Agent listens for Internet requests over these ports.

Since custom protocols take precedence over Websense protocols, if you define a custom protocol using port 80, all protocols that use port 80 (potentially including HTTP) are managed according to the custom protocol definition.

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	

- Under Default Action, specify the 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 261).
- When you are finished, click **OK** to return to the Edit Protocols page. The new protocol definition appears in the Protocols list.
- Click **OK** again to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Adding to a Websense-defined protocol

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with categories, page 244](#)
- ◆ [Protocol-based policy enforcement, page 255](#)
- ◆ [Configuring the default Bandwidth Optimizer limits, page 262](#)

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

The hybrid service 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.

When installed, 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 all TRITON AP-WEB deployments, Content Gateway collects bandwidth data for FTP, HTTP, and, when enabled, the individual protocols that tunnel over HTTP (see [Configuring tunneled protocol detection](#), page 92). Measurement and reporting parallel that used by Network Agent. You can specify that this data be used to determine bandwidth-based policy enforcement for protocols in the Bandwidth Optimizer settings.

1. In the TRITON Manager, go to the **Web > Settings > General > Filtering** page.
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 and Deploy**.

When bandwidth options are active, enforcement starts 10 minutes after initial configuration, and 10 minutes after each Websense Policy Server restart. This delay ensures accurate measurement of bandwidth data.

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 Page Management](#), page 179.

Configuring the default Bandwidth Optimizer limits

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Editing a category filter](#), page 44
- ◆ [Editing a protocol filter](#), page 47
- ◆ [Using Bandwidth Optimizer to manage bandwidth](#), page 261

Before specifying bandwidth settings in policies, verify the default bandwidth thresholds that trigger bandwidth-based enforcement:

- ◆ 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 the TRITON Manager, go to the **Web > Settings > General > Filtering** page.
2. Enter the bandwidth usage thresholds that will trigger bandwidth-based enforcement, when enabled.
 - When a category or protocol is blocked based on traffic for the entire network, **Default bandwidth for network** defines the default threshold.
 - When a category or protocol is blocked based on traffic for the protocol, the **Default bandwidth per protocol** defines the default 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 and Deploy**.

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Enforcement based on file extension](#), page 264
- ◆ [Enforcement based on file analysis](#), page 267
- ◆ [Working with file type definitions](#), page 269
- ◆ [Adding custom file types](#), page 270
- ◆ [Adding file extensions to a file type](#), page 271

When you create or edit a category filter, you can configure file type blocking for permitted categories. This allows your organization to restrict access to particular file types from websites in some or all permitted categories. For example, you could permit the category Sports, but block multimedia (audio and video) files from sites in the Sports category.

How file type blocking is implemented depends on your software.

- ◆ TRITON AP-WEB enable 2-part file type blocking, based on a combination of file extension (see [Enforcement based on file extension](#), page 264) and analysis of requested files (see [Enforcement based on file analysis](#), page 267).

For example:

1. The General Email category is permitted in the active category filter, but file type blocking is enabled for Compressed Files in the category.
2. An end user attempts to download a file with a file with a .zip extension (like "myfile.zip").

3. The user receives a block page indicating that the download was blocked by file type, because the “.zip” file extension is associated with the Compressed Files file type.
 4. The user attempts to download another file from email. This file does not have a known file extension (for example, “myfile.111”).
 5. The file is scanned to find its file type.
 - If analysis determines that the file is in a compressed format, the user receives a block page indicating that the download is blocked by file type.
 - If analysis determines that the file is not compressed, the download request is permitted.
- ◆ Websense Web Filter & Security allows blocking based solely on file extension (see [Enforcement based on file extension](#), page 264).

For example:

1. The General Email category is permitted in the active category filter, but file type blocking is enabled for Compressed Files in the category.
2. An end user attempts to download a file with a file with a .zip extension (like “myfile.zip”).
3. The user receives a block page indicating that the download was blocked by file type, because the “.zip” file extension is associated with the Compressed Files file type.

Combine protocol-based policy enforcement with file type enforcement to better manage Internet audio and video media. Protocol filters handle streaming media, while file type enforcement handles files that can be downloaded and then played.

Enforcement based on file extension

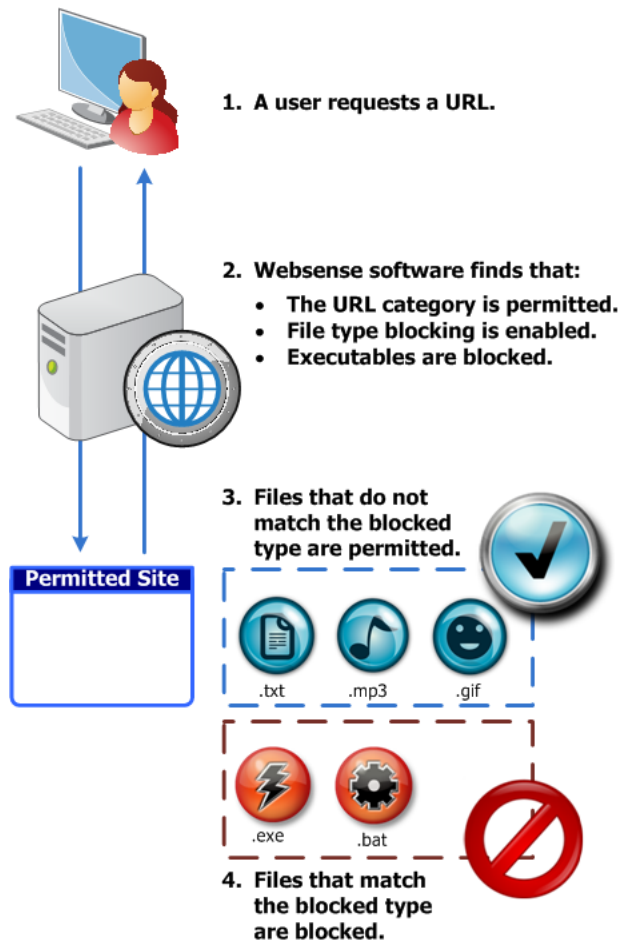
Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Enabling file type blocking in a category filter](#), page 269
- ◆ [Working with file type definitions](#), page 269
- ◆ [Adding custom file types](#), page 270
- ◆ [Adding file extensions to a file type](#), page 271

When a user requests a URL in a permitted category for which file type blocking is enabled, Filtering Service checks the files associated with the URL to see if any of them has a file extension that is assigned to a blocked file type. If so, the request is

blocked, and the user receives a block page that indicates that the request was blocked by file type.



If the file extension is not associated with a blocked file type, what happens next depends on your product:

- ◆ *TRITON AP-WEB*: The file is analyzed to determine its true file type, and permitted or blocked based on that analysis (see [Enforcement based on file analysis](#), page 267).

If you have the Web Hybrid module, the file is analyzed either by Content Gateway or the hybrid service, depending on which proxy handles the user's request.

- ◆ *Websense Web Filter & Security*: The file is permitted.

Several predefined file types (groups of file extensions) are included with the product. These file type definitions are maintained in the Master Database, and may be changed as part of the Master Database update process.

You can use the predefined file types, modify the existing file type definitions, or create new file types. You cannot, however, delete Websense-defined file types, or delete the file extensions associated with them.

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.

File type definitions may contain as many or as few file extensions as are useful for enforcement purposes. Pre-defined file types, for example, include the following file extensions:

File Type	Associated Extensions
Compressed files	.ace, .arc, .arj, .b64, .bhz, .cab, .gz, .gzip, .hqx, .iso, .jar, .lzh, .mim, .rar, tar, taz, .tgz, .tz, .uu, .uue, .xxe, .z, .zip
Documents	.ade, .adp, .asd, .cwk, .doc, .docx, .dot, .dotm, .dotx, .grv, .iaf, .lit, .lwp, .maf, .mam, .maq, .mar, .mat, .mda, .mdb, .mde, .mdt, .mdw, .mpd, .mpp, .mpt, .msg, .oab, .obi, .oft, .olm, .one, .ops, .ost, .pa, .pdf, .pip, .pot, .potm, .potx, .ppa, .ppam, .pps, .ppsm, .ppsx, .ppt, .pptm, .pptx, .prf, .pst, .pub, .puz, .sldm, .sldx, .snp, .svd, .thmx, .vdx, .vsd, .vss, .vst, .vsx, .vtx, .wbk, .wks, .wll, .wri, .xar, .xl, .xla, .xlb, .xlc, .xll, .xlm, .xls, .xlsb, .xlsm, .xlsx, .xlt, .xltm, .xltx, .xlw, .xsf, .xsn
Executables	.bat, .exe
Images	.bmp, .cmu, .djvu, .emf, .fbm, .fits, .gif, .icb, .ico, .jpeg, .jpg, .mgr, .miff, .pbk, .pbm, .pcx, .pdd, .pds, .pix, .png, .psb, .psd, .psp, .rle, .sgi, .sir, .targa, .tga, .tif, .tiff, .tpic, .vda, .vst, .zif
Multimedia	.aif, .aifc, .aiff, .asf, .asx, .avi, .ivf, .m1v, .m3u, .mid, .midi, .mov, .mp2, .mp2v, .mp3, .mpa, .mpe, .mpg, .mpv2, .ogg, .qt, .ra, .ram, .rmi, .snd, .wav, .wax, .wm, .wma, .wmp, .wmv, .wmx, .wxv
Rich Internet Applications	.swf
Text	.htm, .html, .txt, .xht, .xhtml, .xml
Threats	.vbs, .wmf

When a user requests a site, Websense software:

1. Determines the URL category.
2. Checks the file extension.
3. (*TRITON AP-WEB*) If not blocked by extension, the file is analyzed to find its true file type.

**Note**

When multiple group policies could apply to a user request, file type blocking is not performed.

When a user tries to access a blocked file type, the **Reason** field on the Websense block page indicates that the file type was blocked (see [Block Page Management](#), page 179).

The standard block page is not displayed if a blocked 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.

To view existing file type definitions, edit file types, or create custom file types for enforcement by extension, go to **Policy Management > Filter Components**, and then click **File Types**. See [Working with file type definitions](#), page 269, for more information.

To enable file type blocking, see [Enabling file type blocking in a category filter](#), page 269.

Enforcement based on file analysis

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Enabling file type blocking in a category filter](#), page 269
- ◆ [Configuring file analysis](#), page 94

If user traffic passes through Websense Content Gateway or the hybrid service, requested files are analyzed to define their type when all of the following are true:

1. A user requests a URL in a permitted category.
2. File type blocking is enabled for the category in the active category filter.
3. There is no file extension match in a blocked file type (see [Enforcement based on file extension](#), page 264).

In this case, the file type returned for policy enforcement describes the purpose or behavior of similar files, independent of extension. So attempts to disguise an executable by giving it a “.txt” or other innocuous file extension are prevented by file type analysis.

File type definitions are maintained in the analytics databases, and may be changed as part of the Content Gateway database or hybrid service update process.

The file types identified by file analysis are:

File Type	Description
Compressed files	Files that have been packaged to take up less space, like ZIP, RAR, or JAR archives.
Documents	Binary document formats, like DOCX or PDF.
Executables	Programs that can be run on your machine, like EXE or BAT files.
Images	Picture formats, like JPG, BMP, and GIF.
Multimedia	Audiovisual formats, like MP3, WMV, and MOV.
Rich Internet Applications	Web applications that run in a browser, like Flash.

File Type**Description****Text**

Unformatted textual material, like HTML and TXT files.

Threats

Malicious applications that could harm your machine or network, like spyware, worms, or viruses.

When a user requests a website, on-premises or hybrid components first determine the site category, and then check for blocked file types (first by extension, then by analysis).

**Note**

When multiple group policies could apply to a user request, file type blocking is not performed.

If compressed files are permitted, when a compressed file is selected for download, its contents are analyzed. Policy enforcement is then based on the file type assigned to the content of the compressed archive. For example, if compressed files are permitted, but executable files are blocked, when a user attempts to download a compressed file, the contained files are analyzed. If the compressed file contains an executable file, the download is blocked based on the executable file type. Or if the compressed file contains a file that is determined to be malicious, the download is blocked.

**Note**

The .xz file format is not supported for compressed file analysis.

When a user tries to access a blocked file type, the **Reason** field on the Websense block page indicates that the file type was blocked (see [Block Page Management](#), page 179).

The standard block page is not displayed if a blocked 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.

To view existing file extensions in a file type, edit file types, or create custom file types for enforcement by extension, go to **Policy Management > Filter Components**, and then click **File Types**. See [Working with file type definitions](#), page 269, for more information.

To enable file type blocking, see [Enabling file type blocking in a category filter](#), page 269.

Enabling file type blocking in a category filter

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

To prevent users from accessing some file types in otherwise permitted categories:

1. Go to the **Policy Management > Filters** page and click on a category filter name.
Note that you can also edit category filters from within a policy.
2. Select a category in the **Categories** list.
3. Mark the **Block file types** check box under Advanced Filtering on the right-hand side of the page.
A list of file types is displayed.
4. Use the check boxes to select one or more file types to block.
5. If you want to block the select file types in all categories permitted by this category filter, click **Apply to All Categories**.
6. Click **OK**, then **Save and Deploy** to implement your changes.

Working with file type definitions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing traffic based on file type, page 263](#)
- ◆ [Editing a category filter, page 44](#)
- ◆ [Responding to a URL request, page 81](#)

Use the **Policy Management > Filter Components > Edit File Types** page to create and manage up to 32 file types (groups of file extensions) that can be explicitly blocked in category filters (see [Managing traffic based on file type, page 263](#)).



Important

Custom file types and custom additions to predefined types are used in extension-based enforcement, but not true file type analysis. See [Enforcement based on file extension, page 264](#), and [Enforcement based on file analysis, page 267](#), for more information.

- ◆ 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 271](#), for further instructions.
- ◆ To create a new file type, click **Add File Type**, and then see [Adding custom file types, page 270](#), 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 and Deploy**.

Adding custom file types

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Filter Components > Edit File Types > Add File Type** page to define custom file types.



Important

Custom file types and custom additions to predefined types are used in extension-based enforcement, but not true file type analysis. See [Enforcement based on file extension, page 264](#), and [Enforcement based on file analysis, page 267](#), for more information.

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 and Deploy**.

Adding file extensions to a file type

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Filter Components > Edit File Types > Add File Extensions** page to add file extensions to the selected file type.



Important

Custom file types and custom additions to predefined types are used in extension-based enforcement, but not true file type analysis. See [Enforcement based on file extension, page 264](#), and [Enforcement based on file analysis, page 267](#), for more information.

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 and Deploy**.

Using regular expressions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 care. They provide a powerful tool, but they need to be constructed well. Poorly constructed regular expressions can result in excessive overhead, over-blocking, or under-blocking. Using regular expressions as policy enforcement criteria may increase CPU usage.

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 policy enforcement behavior

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The right shortcut pane in the Web module of the TRITON Manager includes a **Toolbox** that allows you to perform quick checks of your policy 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 272](#)
- ◆ [Check Policy, page 273](#)
- ◆ [Test Filtering, page 273](#)
- ◆ [URL Access, page 274](#)
- ◆ [Investigate User, page 274](#)

You can also click **Support Portal** to access the Websense Technical Support website in a new browser tab or window. From the Support Portal, you can search the knowledge base to find articles, tips, tutorials, videos, and product documentation.

URL Category

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 274](#)).
 - 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 [Enforcement order, page 79](#)).

Test Filtering

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 274](#)).
 - 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 (if user identification has been configured) or IP address (for machines on which users are not identified).

The IP address search shows only results for which no user name has been logged.

3. Click **Go**.

An investigative report shows the client's usage history.

Identifying a user to check policy or test filtering

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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**.

14

User Identification for Policy Enforcement

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 and authentication methods are available for the on-premises software:

- ◆ A Websense transparent identification agent works in the background to communicate with a directory service and identify users (see [Identifying on-premises users transparently](#)).
- ◆ Websense software prompts users for their network credentials, requiring them to log on when they open a web browser (see [Manual authentication](#), page 279).
- ◆ (*TRITON AP-WEB only*) Content Gateway uses one or more several supported methods (including Integrated Windows Authentication, Legacy NTLM, LDAP, and RADIUS) to authenticate user requests (see the Content Gateway manager [Online Help](#)).

This option may be used in conjunction with a transparent identification agent to provide a fallback method for applying user-based policies when user authentication is unavailable.

- ◆ (*Web Filter & Security only*) A third-party integration product identifies or authenticates users, and then passes user information to Websense software.

A list of supported integration products is available in the [Deployment and Installation Center](#).

With the Web Hybrid module, the hybrid service must likewise be able to identify or authentication users to apply user and group based policies.

- ◆ A component called Websense Directory Agent collects the information used to identify users (see [Identification and authentication of hybrid users](#), page 300).
- ◆ The TRITON AP-WEB client software is installed on client machines to provide transparent authentication, enforce use of the hybrid service, and pass authentication details to the hybrid service.
- ◆ Single sign-on provides authentication using an identity provider that communicates with your directory service.

Ping Federate is currently the only supported identity provider.

Identifying on-premises users transparently

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)

In general, **transparent identification** describes any method used to identify users in your directory service without prompting them for login information. This includes any of the optional transparent identification agents available when user requests are managed by on-premises Websense software.

- ◆ Websense [DC Agent, page 288](#), is used with a Windows-based directory service. The agent periodically queries domain controllers for user login sessions and polls client machines to verify login status. It runs on a Windows server and can be installed in any domain in the network.
- ◆ Websense [Logon Agent, page 293](#), identifies users as they log on to Windows domains. The agent runs on a Linux or Windows server, and its associated logon application runs on Windows or Mac clients.
- ◆ Websense [Configuring RADIUS Agent, page 295](#), 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 [Configuring eDirectory Agent, page 297](#), is used with Novell eDirectory. The agent uses Novell eDirectory authentication to map users to IP addresses.

Agent can be used alone, or in certain combinations.

Both general user identification settings and specific transparent identification agents are configured on the **Web > Settings > General > User Identification** page of the TRITON Manager.

See [Configuring user identification and authentication, page 280](#), for detailed configuration instructions.

In some instances, transparent identification agents may not be able to provide correct user information to other components. 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 279](#)).

Transparent identification of remote users

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

In certain configurations, Websense software can transparently identify users logging on to your network from remote locations:

- ◆ 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 288](#)).
- ◆ 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 policies based on users or groups (see [Configuring RADIUS Agent, page 295](#)).
- ◆ (*Web Filter & Security*) With the Remote Filter module, Websense software can identify any off-site user logging on to a cached domain using a domain account. For more information, see [Manage Off-site Users, page 227](#).

Manual authentication

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Setting authentication rules for specific machines, page 282](#)
- ◆ [Secure manual authentication, page 284](#)
- ◆ [Configuring user identification and authentication, page 280](#)

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 apply 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 and authentication, page 280](#) and [Configure user access to the hybrid service, page 202](#)).

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 282](#)).

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 and authentication

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Working with users and groups, page 58](#)

Use the **Settings > General > User Identification** page to manage when and how on-premises 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 Content Gateway, a transparent identification agent, or an integration product.
- ◆ 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 [Identifying on-premises users transparently, page 278](#), 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 289](#)
- ◆ [Configuring Logon Agent, page 294](#)
- ◆ [Configuring RADIUS Agent, page 295](#)
- ◆ [Configuring eDirectory Agent, page 297](#)

To remove an agent instance from the list, mark the checkbox next to the agent information in the list, and then click **Delete**.

If you have one or more DC Agent instances, under DC Agent Domains and Controllers, click **View Domain List** for information about which domain controllers the agents are currently polling. See [Reviewing DC Agent polled domains and domain controllers, page 291](#), for more information.

Under **User Identification Exceptions**, list the IP addresses of machines that should use different user identification settings than the rest of your network.

For example, if you use Content Gateway, a transparent identification agent, or a third-party 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 **Add**, and then see [Setting authentication rules for specific machines, page 282](#). To remove an exception, mark the check box next to an IP address or range, then click **Delete**.

Under **Additional Authentication Options**, specify the default response of Websense software when users are not identified transparently:

- ◆ 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 279](#)).

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, this default domain context is used, even if the global rule is to apply a computer or network-based policy.

When you are finished making changes on this page, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Setting authentication rules for specific machines

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring user identification and authentication, page 280](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Secure manual authentication, page 284](#)

Selective authentication lets you determine whether users requesting Internet access from a specific client machine (identified by IPv4 or IPv6 address) are prompted to provide their login 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 > General > 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 282](#), 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 and Deploy**.

Defining exceptions to user identification settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)

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 network **Range** in IPv4 or IPv6 format to identify clients 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 **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 receive the appropriate user-based policy.
 - Select **Prompt user for logon information** to require users to provide logon credentials. The **Default domain context** entered for **Additional Authentication Options** is displayed and will be used.
If “Try to identify user transparently” is also selected, users receive a browser prompt only if they are not identified transparently.
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 and Deploy**.

Revising exceptions to user identification settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)

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 **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.
 - Select **Prompt user for logon information** to require users to provide logon credentials. The **Default domain context** entered for **Additional Authentication Options** is displayed and will be used.

If “Try to identify user transparently” is also selected, users receive a browser prompt only if they are not identified transparently.
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 and Deploy**.

Secure manual authentication

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring user identification and authentication, page 280](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Setting authentication rules for specific machines, page 282](#)
- ◆ [Activating secure manual authentication, page 286](#)

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

(*Web Filter & Security only*) 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 285).
2. Enable secure manual authentication (see [Activating secure manual authentication](#), page 286) and secure communication with the directory service.
3. Import certificates into the browser (see [Accepting the certificate within the client browser](#), page 287).

Generating keys and certificates

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Manual authentication](#), page 279
- ◆ [Setting authentication rules for specific machines](#), page 282
- ◆ [Secure manual authentication](#), page 284
- ◆ [Activating secure manual authentication](#), page 286
- ◆ [Accepting the certificate within the client browser](#), page 287

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://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 Service 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Manual authentication, page 279](#)
- ◆ [Setting authentication rules for specific machines, page 282](#)
- ◆ [Secure manual authentication, page 284](#)
- ◆ [Generating keys and certificates, page 285](#)
- ◆ [Accepting the certificate within the client browser, page 287](#)

1. Stop Websense Filtering Service (see [Stopping and starting Websense services, page 370](#)).
2. Navigate to the Websense installation directory on the Filtering Service machine (by default, **C:\Program Files\Websense\Web Security\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 > General > Directory Services** page. See [Advanced directory settings, page 63](#), for details.

Accepting the certificate within the client browser

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Manual authentication, page 279](#)
- ◆ [Setting authentication rules for specific machines, page 282](#)
- ◆ [Secure manual authentication, page 284](#)
- ◆ [Generating keys and certificates, page 285](#)
- ◆ [Activating secure manual authentication, page 286](#)

The first time you try to browse to a website, 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 website.
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

1. Open the browser and go to a website.
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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Configuring DC Agent, page 289](#)

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](#).

In order to perform domain discovery (automatic domain and domain controller detection) and computer polling (to verify the logged-on user), DC Agent must run with **domain admin** or **enterprise admin** permissions. If you do not plan to use either of these features, DC Agent can run as any network user with read privileges on the domain controller. Note that when domain discovery is disabled, you must maintain the domain and domain controller list for each DC Agent instance manually (see [The *dc_config.txt* file, page 292](#)).

2. Configure DC Agent to communicate with other Websense components and with domain controllers in your network (see [Configuring DC Agent](#)).
3. Use the TRITON Manager to assign policies to users, groups, and OUs (see [Adding a client, page 66](#)).

Your Websense software can prompt users for identification if DC Agent is unable to identify users transparently. For more information, see [Manual authentication, page 279](#).

Configuring DC Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently](#)
- ◆ [Manual authentication](#)
- ◆ [Configuring user identification and authentication](#)
- ◆ [DC Agent](#)

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 **IPv4 address or hostname** of the machine on which the agent is installed.



Note

Hostnames must start with an alphabetical character (a-z), not a numeric or special character.

Hostnames 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, select **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.

Some of these settings can, however, be overridden in a configuration file (see the [Using DC Agent for Transparent User Identification](#) technical paper).

1. Under Domain Discovery, mark or clear **Enable automatic domain discovery** determine whether DC Agent automatically finds domains and domain controllers in your network.
2. If domain discovery is enabled, also specify:

- How often to **Discover domains**. Domain discovery occurs at 24-hour intervals, by default.
- Whether **DC Agent** or **User Service** is responsible for performing domain discovery.

In many environments, it is preferable to use User Service for domain discovery.

If DC Agent is used for domain discovery, the service must run with **domain** or **enterprise admin** privileges.

3. When User Service is installed on a Websense appliance or Linux server, the page includes a **Linux WINS Server Information** section. A WINS server is required to resolve domain names to domain controller IP addresses.

If you have not already provided WINS information on the Settings > Directory Services page, enter:

- a. The account name of an **Administrative user** that can access the directory service.
 - b. The **Password** for the account.
 - c. **Domain** information for the account.
 - d. The IP address or hostname of a WINS server in your network.
4. In the Domain Controller Polling section of the DC Agent Communication box, mark **Enable domain controller polling** to enable DC Agent to query domain controllers for user logon sessions.

To perform domain controller polling, the DC Agent service needs only read privileges on the domain controller. Automatic domain discovery (steps 1 and 2) and computer polling (step 7) require that the service run with elevated permissions.

You can specify which domain controllers each instance of DC Agent polls in a configuration file (see [The *dc_config.txt* file](#), page 292).
 5. 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.
 6. 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.
 7. 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**.

If DC Agent performs computer polling, the service must run with **domain** or **enterprise admin** privileges.

8. 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.

9. 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.

10. Click **OK** to return to the User Identification page, then click **OK** again to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Reviewing DC Agent polled domains and domain controllers

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **User Identification > DC Agent Domains and Controllers** page to review which domain controllers each DC Agent instance in your network is currently polling.

**Important**

If the DC Agent Domains and Controllers page displays text explaining that DC Agent is not “polling any domain controllers at this time,” see [DC Agent Domains and Controllers page is blank](#), page 442.

Typically, the page shows the **Domains** and **Domain Controllers** detected by each of the **DC Agent Instances** in your network.

By default, DC Agent performs its **domain discovery** process (identifying domains and domain controllers) at startup, and at 24 hour intervals thereafter. Domain and controller information is stored in a file called **dc_config.txt** (see [The dc_config.txt file](#), page 292).

Information displayed on the DC Agent Domains and Controllers page is compiled from each **dc_config.txt** file in your deployment.

- ◆ The list includes only domains and controllers that are actively being queried.
 - If you have disabled queries to a domain controller in the `dc_config.txt` file, that domain controller is not shown.
 - Likewise, if you have disabled queries to all domain controllers within a domain, neither the domain nor its controllers are listed.
- ◆ Information is shown for all of the DC Agent instances in your network.
 - If the same domain controller is polled by multiple DC Agent instances, each is listed.
 - To configure different DC Agent instances to poll different domains, update the `dc_config.txt` file for each instance. See [The `dc_config.txt` file](#), page 292.
- ◆ The TRITON Manager checks for the latest domain and controller information each time you navigate to the DC Agent Domains and Controllers page. This means that if domain discovery is underway while you are viewing the page, you must navigate away, then return to the page to see updates.

The `dc_config.txt` file

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.

- ◆ By default, DC Agent performs domain discovery (identifying domains and domain controllers) at startup, and every 24 hours thereafter.
- ◆ Either DC Agent or User Service can be used to perform domain discovery.

For information about enabling domain discovery and setting the discovery interval, see [Configuring DC Agent](#), page 289.

DC Agent stores domain and domain controller information in a file called **`dc_config.txt`**. Edit the file to change which domain controllers DC Agent polls:

1. Go to the Websense **bin** directory (by default, `C:\Program Files\Websense\Web Security\bin`) on the DC Agent machine.
2. Make a backup copy of the **`dc_config.txt`** file in another location.
3. Open the original **`dc_config.txt`** file in a text editor (like Notepad).
4. Confirm that all of your domains and domain controllers are listed. For example:

```
[WEST_DOMAIN]
dcWEST1=on
dcWEST2=on
[EAST_DOMAIN]
dcEAST1=on
dcEAST2=on
```

5. If there are domain controllers in the list that DC Agent should not poll, change the entry value from **on** to **off**. For example:

```
dcEAST2=off
```

- If you configure DC Agent to avoid polling an active domain controller, the agent cannot transparently identify users logging on to that domain controller.
 - If DC Agent's automatic domain discovery has detected a domain controller that should not be used to identify users, set the entry to **off**, rather than removing it. Otherwise, the next discovery process will re-add the controller.
6. If there are domain or domain controller entries missing from the list, you can add them manually. Before adding entries, run the **net view /domain** command on the DC Agent machine to make sure that the agent can see the new domain.
 7. Save your changes and close the file.
 8. Restart the **Websense DC Agent** service.

Logon Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Configuring Logon Agent, page 294](#)

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) on Windows and Mac 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 [Using Logon Agent for Transparent User Identification](#) technical paper.

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 a Websense appliance or Linux server, page 446](#), for additional configuration steps.

Configuring Logon Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)
- ◆ [Logon Agent, page 293](#)

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 **IPv4 address or hostname** of the Logon Agent machine.



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 (30602, by default).
3. To establish an authenticated connection between Filtering Service and Logon Agent, mark **Enable authentication**, and then enter a **Password** for the connection.

Next, customize global Logon Agent communications settings. By default, changes that you make here affect all Logon Agent instances.

1. Under Logon Application Communication, specify the **Connection port** that the logon application uses to communicate with Logon Agent (15880, by default).
2. Enter the **Maximum number of connections** that each Logon Agent instance allows (200, by default).

If your network is large, you may need to increase this number. Increasing the number does increase network traffic.

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 operate in **persistent mode** or **nonpersistent mode** (default). (More information is

available in the [Using Logon Agent for Transparent User Identification](#) technical paper.)

- ◆ 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 return to the Settings > User Identification page, then click **OK** again to cache your changes. Changes are not saved until you click **Save and Deploy**.

Configuring RADIUS Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)

Websense RADIUS Agent lets you apply user and group-based policies using authentication provided by a RADIUS server. This allows transparent identification of users who access your network using a dial-up, Virtual Private Network (VPN), Digital Subscriber Line (DSL), or other remote connection.

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 **IPv4 address or hostname** of the Logon Agent machine.



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 (30800, by default).
3. To establish an authenticated connection between Filtering Service and RADIUS Agent, mark **Enable authentication**, and then enter a **Password** for the connection.

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 the [Using RADIUS Agent for Transparent User Identification](#) technical paper).

1. Under RADIUS Server, enter the **RADIUS server address or name**. If you provide the IP address, use IPv4 address format.
RADIUS Agent forwards authentication requests to the RADIUS server, and must know the identity of this machine.
2. If your network includes a RADIUS client, enter the **RADIUS client address or name**. If you provide the IP address, use IPv4 address format.
Websense software queries this machine for user logon sessions.
3. 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.
4. 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 (authentication default 1645; accounting default 1646)
 - RADIUS Agent and the RADIUS client (authentication default 12345; accounting default 12346)
5. When you are finished making configuration changes, click **OK** to return to the Settings > User Identification page, then click **OK** again to cache your changes. Changes are not saved until you click **Save and Deploy**.

For information about configuring your RADIUS client and RADIUS server to communicate with Websense RADIUS Agent, see the [Using RADIUS Agent for Transparent User Identification](#) technical paper.

Configuring eDirectory Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Identifying on-premises users transparently, page 278](#)
- ◆ [Manual authentication, page 279](#)
- ◆ [Configuring user identification and authentication, page 280](#)

Websense eDirectory Agent gathers user logon session information from Novell eDirectory, which authenticates users logging on to the network. The agent then:

1. Associates each user with an IP address.
2. Records user name-to-IP-address pairings to a local user map.
3. Communicates the map to Filtering Service.

Filtering Service uses the information to apply policies to users, groups, or OUs.



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.

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.

For detailed information eDirectory Agent deployment, including configuration options not available via the TRITON Manager, see the [Using eDirectory Agent for Transparent User Identification](#) technical paper.

To add a new instance of eDirectory Agent:

1. Under Basic Agent Configuration, enter the **IPv4 address or hostname** of the eDirectory Agent machine.



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 (30700, by default).
3. To establish an authenticated connection between Filtering Service and eDirectory Agent, select **Enable authentication**, and then enter a **Password** for the connection.

Next, customize global eDirectory Agent communication settings:

1. 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.
2. 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 the [Using eDirectory Agent for Transparent User Identification](#) technical paper for details.

3. 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 299.

When you are finished making configuration changes, click **OK** to return to the User Identification page, then click **OK** again to cache your changes. Changes are not saved until you click **Save and Deploy**.

Adding an eDirectory server replica

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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. Enter the eDirectory master or replica **Server IP address**.
2. Enter the **Port** that eDirectory Agent uses to communicate with the eDirectory machine. The valid values are **389** (default) 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 return to the Settings > User Identification page, then click **OK** again to cache your changes.
6. Click **Save and Deploy** to implement the changes.
7. Stop and start eDirectory Agent so that the agent can begin communicating with the new replica. See [Stopping and starting Websense services](#), page 370, for instructions.

Identification and authentication of hybrid users

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Directory Agent and hybrid user identification, page 306](#)
- ◆ [When hybrid users are not identified, page 308](#)
- ◆ [Authentication priority and overrides, page 302](#)
- ◆ [Working with hybrid service clients, page 72](#)

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 users, whether from a filtered location or when off-site, you have the following options for identifying or authenticating the users transparently:

- ◆ **TRITON AP-ENDPOINT Web** is installed on client machines to provide transparent authentication, enforce use of the hybrid service, and pass authentication details to the hybrid service. See [TRITON AP-ENDPOINT Web software, page 303](#).
- ◆ **Single sign-on** provides clientless transparent authentication via a gateway hosted on your network. See [Integrating the hybrid service with a single sign-on identity provider, page 305](#).
- ◆ Users at filtered locations (see [Define hybrid service filtered locations, page 195](#)) can be identified transparently via **NTLM**. This option is not available for off site users.
- ◆ 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 the hybrid service, page 202](#)).

If you do not enable any form of transparent identification or authentication:

- ◆ Off-site users without TRITON AP-ENDPOINT Web or single sign-on are prompted for an email address and password when they open a browser and connect to the Internet.
- ◆ Other hybrid users are assigned policies based on their IP address if TRITON AP-ENDPOINT Web, single sign-on, or NTLM identification are not available.

Indicate how the hybrid service should identify users requesting Internet access. These options are also used as a fallback if either the endpoint client software or single sign-on fails.

- ◆ Mark **Always authenticate users on first access** to enable transparent NTLM identification, secure form authentication, or manual authentication when users first connect to the hybrid service.

If you do not select this option and you have not enabled any other authentication methods for users in filtered locations, those users receive an IP address-based policy, and their identity does not appear in reports

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.



Important

When NTLM is used to identify users, **do not** use self-registration (configured on the User Access page under Registered Domains).

- ◆ Mark **Use secured form authentication to identify users** to display a secure logon form to the end user. When the user enters their email address and hybrid service password, the credentials are sent over a secure connection for authentication.



Note

If Ping Federate is used as the identity provider, single sign-on cannot fall back to secured form authentication.

If you select this option, define how often users' credentials are revalidated for security reasons under Session Timeout. The default options are 1, 7, 14, or 30 days. The same session timeout applies to single sign-on, if enabled.



Note

It is possible to extend the Session Timeout options to 3 months, 6 months, and 12 months. To enable this extended feature, contact Support.

If the users have not previously registered to use the service, they can do so by clicking **Register** on the logon form. To use this option, enable self-registration (configured on the User Access page under Registered Domains). Advise end users **not** to use the same password for hybrid service access that they use to log on to the network.

If you do not select either the NTLM or the secured form authentication option, but **Always authenticate users on first access** is selected, users who could not be identified via another means see a logon prompt every time they access the Internet. Basic authentication is used to identify users who receive a logon prompt.

- ◆ Specify whether or not a Welcome page is displayed when users who have not been identified via NTLM or who are not using secured form authentication 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 and Deploy**.

Once you have set up the hybrid service 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 service 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.

Authentication priority and overrides

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Identification and authentication of hybrid users, page 300](#)
- ◆ [When hybrid users are not identified, page 308](#)

You can select multiple authentication options for your end users on the **Settings > Hybrid Configuration > Hybrid User Identification** page. The options are prioritized as follows:

- ◆ TRITON AP-ENDPOINT Web is always used, if installed.
- ◆ If the endpoint client software is not installed or fails, single sign-on is used if both of the following are true:
 - a. It has been deployed in your network.
 - b. It has been selected on the Hybrid User Identification page for an end user whose requests are managed by the hybrid service.
- ◆ If neither the endpoint client software nor single sign-on is available, the end user is authenticated via secure form-based authentication, if both of the following are true:
 - a. It has been selected on the Hybrid User Identification page.

- b. The user agent or application requesting authentication supports form-based authentication via an HTML page.
- ◆ Applications that do not support form-based authentication use either NTLM identification or basic authentication. Basic authentication is always used if **Always authenticate users on first access** is selected and none of the other options are either selected or available.

You can also enforce a specific authentication option for certain end users, for example all users in a branch office, by deploying a PAC file URL in the following format:

```
http://hybrid-web.global.blackspider.com:8082/proxy.pac?a=X
```

The a= parameter controls the authentication option, and X can be one of the following:

Parameter	Description
a=n	NTLM identification or basic authentication is used, depending on the policy settings and the browser or application capability.
a=f	Authentication is performed using secure form-based authentication.

TRITON AP-ENDPOINT Web software

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **Settings > Hybrid Configuration > Hybrid User Identification** page to configure user identification and authentication methods for users whose requests are managed by the hybrid service.

The Web Hybrid module includes TRITON AP-ENDPOINT Web software, which can be installed on client machines to enforce the use of the hybrid service. The client software passes authentication information to the hybrid service, enabling secure transparent authentication.

- ◆ Detailed deployment and installation instructions are available in the [Websense endpoint solutions technical paper](#).
- ◆ If you have purchased additional Websense endpoint solution modules, be sure to refer to your Websense endpoint solutions documentation for information about combining modules.

To enable TRITON AP-ENDPOINT Web software deployment:

1. Under TRITON AP-ENDPOINT Web, mark **Enable TRITON AP-ENDPOINT Web installation and update on client machines**.

Selecting this option allows you to configure deployment and automatic update settings. If you later deselect this option, any installed client software instances continue to work until uninstalled, though they no longer receive automatic updates.

2. Select a deployment method:

- Click **Deploy Manually** if you want to install the endpoint client software by hand on individual machines or via your preferred distribution method. (This is the only option available for Mac.)

Note the WSCONTEXT value displayed on screen. If you plan to use GPO to distribute the client software, you will use this value in your deployment script to ensure that users are correctly associated with your organization.

Click **View Files** to find the appropriate client software. Select a client operating system, then click on a version of the client software 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 client software directly to Windows machines from the hybrid service, mark the **Deploy the client software from the hybrid service** check box.

Choose whether the client software is deployed to **all users** whose requests go 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 client software 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 **Customize Installation Page**, 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. Enter and confirm your anti-tampering password. The password must be between 4 and 25 characters. Anti-tampering provisions:

- Protect endpoint client software files and folders from being deleted or renamed.
- Restart the client software if it is stopped or killed.
- Require a password is required to uninstall or stop the client software.
- Prevent hybrid client software registry settings from being modified or deleted.
- Block the Service Control command to delete the client software.

Until you define an anti-tampering password, you cannot download the endpoint client installation file or enable deployment from the hybrid service.



Important

For security reasons, Websense does not retain a copy of your anti-tampering password. If you forget your password, enter and confirm a new password. All endpoint client software installations will be updated to use the new password next time they connect to the Internet.

4. Under **Version Update**, select the operating systems for which you want to **Apply automatic updates**.

With these selections, you ensure that client machines always have the latest version of the endpoint client software when it is available.

If you later remove the check from one or both boxes, endpoint updates will no longer be applied to client machines using that operating system. Existing endpoint installations will, however, continue to work.

5. Create a list of applications that should bypass endpoint policy enforcement.

Some applications do not work properly with endpoint enforcement. Application Bypass allows you to add a list of applications that may be causing problems.

- Click **Add** to open the **Add Applications** window.
- Specify the operating system for the Applications you wish to add and enter the applications in the field provided.
 - Enter a single application or a comma-separated list of applications.
 - Include the file extension for each application. If no extension is entered, the application name is treated like a regular expression.
 - An asterisk (*) wildcard can be used in application names. For example, appl.*.

- Click **Add** to return to the **Hybrid User Identification** page and add your entry to the list.

If there are any errors found in your entry, correct them and click **Add** again.

- Remove an entry in the Application Bypass list by checking the box next to the application name and clicking **Delete**.

Note that this feature does not work for applications that use system browser settings to determine a proxy. Also, you may need to update your endpoint deployments. End users must have at least endpoint build 1138 (Windows) or 1566 (Mac) to use application bypass.

6. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Integrating the hybrid service with a single sign-on identity provider

Administrator Help | TRITON AP-WEB | Version 8.1.x

Single sign-on uses an identity provider to authenticate user identity, attributes, and roles with enterprise directories. All communications between components are secured.

When single sign-on is installed on your network, clients connecting to the hybrid proxy are redirected to an identity provider. The identity provider must be configured if off-site users are to be authenticated. Once single sign-on has authenticated a user against your directory service, they are directed back to the hybrid proxy and the appropriate policy is applied. Clients who have authenticated once do not then have to authenticate again for subsequent browsing sessions.

Currently, only PingFederate is supported as a single sign-on identity provider. For information on how to deploy PingFederate, please visit [their website](#).

To integrate a single sign-on identity provider:

1. 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 single sign-on users, they receive a certificate error when they browse to an HTTPS site. If they then select the “Continue to this website (not recommended)” link, they must authenticate using NTLM identification or manual authentication, depending on the settings on the Hybrid User Identification page. See [Enabling hybrid HTTPS notification pages](#), page 206.
2. Mark **Use PingFederate as the identity provider for single sign-on** to activate single sign-on for all client machines.
3. Once single sign-on is configured and the SSL certificate is installed on clients, copy the metadata URL from the identity provider’s metadata and enter it in the Metadata URL field on the Hybrid User Identification page.
4. Under Session Timeout, define how often users’ credentials are revalidated for security reasons. The default options are 1, 7, 14, or 30 days.



Note

It is possible to extend the Session Timeout options to 3 months, 6 months, and 12 months. To enable this extended feature, contact Websense Technical Support.

5. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Directory Agent and hybrid user identification

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Identification and authentication of hybrid users](#), page 300
- ◆ [Working with users and groups](#), page 58
- ◆ [Connecting web protection software to a directory service](#), page 59
- ◆ [Send user and group data to the hybrid service](#), page 209

With the Web Hybrid module, an interoperability component called **Websense Directory Agent** is required if you want to enable user, group, and domain (OU) based policy enforcement 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 the TRITON Manager to configure Directory Agent to collect data from your directory service (see [Send user and group data to the hybrid service, page 209](#)). 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 the hybrid service, page 216](#)), Sync Service sends the user and group information collected by Directory Agent to the hybrid service. Sync Service compresses large files before sending them.

How Directory Agent works with User Service

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Identification and authentication of hybrid users, page 300](#)
- ◆ [Working with users and groups, page 58](#)
- ◆ [Connecting web protection software to a directory service, page 59](#)
- ◆ [Send user and group data to the hybrid service, page 209](#)

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 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 the TRITON console, 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 hostname 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 and Deploy** 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 58).
- ◆ 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 209).

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 hybrid users are not identified

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Identification and authentication of hybrid users](#), page 300
- ◆ [Working with hybrid service clients](#), page 72

With the Web Hybrid module, when users are not identified or authenticated transparently, only 3 types of policies can be applied to requests:

- ◆ 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 service 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 service account.

User and group policies cannot be applied to self-registered users. Self-registered users always receive the Default policy (see [Off-site user self-registration](#), page 228).

15

Delegated Administration and Reporting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Delegated administration provides an effective way to distribute responsibility for TRITON AP-WEB and Web Filter & Security configuration, policy management, 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 the Web module of the TRITON Manager 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 312*](#)
- ◆ [*Preparing for delegated administration, page 319*](#)
- ◆ [*Managing delegated administration roles, page 324*](#)
- ◆ [*Updating delegated administration roles, page 332*](#)
- ◆ [*Performing delegated administrator tasks, page 335*](#)
- ◆ [*Enabling network accounts, page 339*](#)

The fundamentals of delegated administration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Delegated administration roles, page 312](#)
- ◆ [Delegated administrators, page 313](#)
- ◆ [Delegated administration and reporting permissions, page 314](#)
- ◆ [Administrators in multiple roles, page 317](#)
- ◆ [Multiple administrators accessing the TRITON console, page 318](#)

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 312](#).
- ◆ **Administrators** are individuals or groups given responsibility for configuring settings, managing policies for clients, running Internet activity reports, or auditing the system in the Web module of the TRITON Manager. An administrator's set of responsibilities is determined by the **role** and **permissions** that the administrator is assigned. See [Delegated administrators, page 313](#).
- ◆ **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 314](#).

Delegated administration roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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**.

The Web module of the TRITON Manager 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Administrators are the individuals who can access the TRITON Manager. Depending on their permissions, in the Web module they may be able to:

- ◆ Log on and view some elements of the Status > Dashboard page, but take no other actions.
- ◆ Access all configuration and management features of the Web module, 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 the Web module.

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 312](#).

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 the TRITON Manager. 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 module 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 314).

- ◆ **Access**, which allows the administrator to log on and view limited portions of the **Status > Dashboard** 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 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 module access via TRITON Settings can be added to roles.

Delegated administration and reporting permissions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The permissions available to an administrator depend on whether the administrator is assigned to the Super Administrator role, a policy management and reporting role, or an investigative reporting role.

Super Administrator permissions

The Super Administrator role can contain 2 types of administrators: unconditional Super Administrators and conditional Super Administrators.

When you create a Global Security Administrator account on the TRITON Settings > Administrators page, or select the **Grant access and the ability to modify access permissions for other accounts** option for the Web module, the account is automatically added to the Super Administrator role in the Web module with unconditional permissions.

Unconditional Super Administrators can:

- ◆ Access all system configuration settings in the Web module (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 320.
- ◆ 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 component status and stop or start components from the Status > Deployment page.
- ◆ Review the audit log, which records administrator access to and actions within the Web module.

- ◆ (*TRITON AP-WEB only*) Open the Content Gateway manager via a button on the Settings > General > Content Gateway Access page and be logged on automatically, without having to provide credentials.

When an unconditional Super Administrator adds additional administrators to the Super Administrator role (via the Policy Management > Delegated Administration page), the new administrators are granted conditional permissions.

Unlike unconditional Super Administrators, whose permissions cannot be changed, conditional Super Administrators can be granted a combination of policy management, reporting, and access permissions.

- ◆ **Full policy** permissions allow conditional Super Administrators to:
 - Create and edit delegated administration roles, filter components, filters, policies, and exceptions, 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.
- ◆ **Exceptions only** permissions allow conditional Super Administrators to create and edit exceptions. (Exceptions permit or block URLs for specified users, regardless of which policy normally governs their Internet access.)
Policies, filters, filter components, the Filter Lock, and all Settings pages are hidden for Super Administrators with exceptions only permissions.
- ◆ **Reporting** permissions allow conditional Super Administrators to:
 - Access Status > Dashboard page charts.
 - Run investigative and presentation reports on all users.
 If an administrator is granted reporting permissions only, the Check Policy tool does not appear in the Toolbox.
- ◆ **Real-Time Monitor** permissions allow Super Administrators to monitor all Internet activity for each Policy Server associated with the TRITON Manager.
- ◆ **Content Gateway direct access** permissions allow Super Administrators to be logged on to the Content Gateway manager automatically via a button on the Settings > General > Content Gateway Access page in the TRITON Manager.

Only one administrator at a time can log on to a role with **full policy** or **exceptions only** 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 module 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:

- ◆ **Full policy** permissions allow delegated administrators to create and manage filter components (including custom categories and recategorized URLs), filters (category, protocol, and limited access), policies, and exceptions (black and white lists) 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. (As part of enforcing the Filter Lock, delegated administrators cannot give their managed clients password override permissions.)

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.

- ◆ **Exceptions only** permissions allow delegated administrators to create and manage exceptions for managed clients in their role. (Exceptions permit or block URLs for specified users, regardless of which policy normally governs their Internet access.)

Policies, filters, and filter components are hidden for delegated administrators with exceptions only permissions.

- ◆ **Deployment status** permissions allow delegated administrators to review component status on the Status > Deployment page. Delegated administrators with deployment status permissions can also be granted permission to start components, stop components, or both.
- ◆ 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 Status > Dashboard page, investigative reports, and the Settings pages used to manage Log Server and 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 activity for each Policy Server associated with the TRITON Manager.

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 > Dashboard page.

Auditors

Any conditional Super Administrator or delegated administrator account can be granted **Auditor** permissions. An auditor can see most Web module features and functions, 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 and Deploy button is disabled.

Administrators in multiple roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Delegated administration roles, page 312](#)
- ◆ [Delegated administrators, page 313](#)
- ◆ [Delegated administration and reporting permissions, page 314](#)

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 management:**
 - **Full policy:** You can add and edit filters and policies for the role selected during logon, and apply policies to that role's managed clients.
 - **Exceptions only:** You can create and manage exceptions for the role selected during logon, and apply exceptions 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 charts on the Status > Dashboard page, 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 the TRITON console

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Administrators in **different** roles can access the Web module of the TRITON Manager 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 full policy or exceptions-only permissions in the shared role. If a second administrator tries to log on with full policy or exceptions-only 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 > Dashboard and Alerts pages, as well as Real-Time Monitor (if applicable). Their session does not time out.

If administrators in Status Monitor mode try to go to a page other than Dashboard, Alerts, or Real-Time Monitor, they are prompted to log on again.

Preparing for delegated administration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [The fundamentals of delegated administration, page 312](#)
- ◆ [Creating a Filter Lock, page 320](#)
- ◆ [Preparing delegated administrators, page 323](#)
- ◆ [Managing delegated administration roles, page 324](#)

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 320](#).)
 - 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 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 the TRITON Manager 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 module 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 324.
3. The Super Administrator notifies the delegated administrators that they have been granted administrative access to the Web module of the TRITON Manager, and explains their level of permissions. See [Preparing delegated administrators](#), page 323.

Creating a Filter Lock

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Locking categories](#), page 321
- ◆ [Locking protocols](#), page 322

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 321.
- ◆ Click the **Protocols** button to block and lock protocols, or to specify protocols that are always logged. See [Locking protocols](#), page 322.

Locking categories

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a Filter Lock, page 320](#)
- ◆ [Locking protocols, page 322](#)

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.

1. 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.

2. 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.
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 and Deploy**.

Locking protocols

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Creating a Filter Lock](#), page 320
- ◆ [Locking categories](#), page 321

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 379.

1. Select a protocol in the tree.
Delegated administration roles do have access to custom protocols created by the Super Administrator. Therefore, custom protocols do appear in this tree.
2. 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 websites 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 and Deploy**.

Preparing delegated administrators

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [The fundamentals of delegated administration, page 312](#)
- ◆ [Preparing for delegated administration, page 319](#)
- ◆ [Performing delegated administrator tasks, page 335](#)

After assigning individuals as administrators in any administrative role, make sure to give them the following information:

- ◆ The URL for logging on to the TRITON Manager. By default:
`https://<TRITON_location>:9443/triton/`
Substitute the IP address or hostname of the TRITON management server.
- ◆ 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 module 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 the TRITON console. 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 the Web module of the TRITON Manager and manage policy for that role.

- ◆ How to find the list of clients managed by their role. Administrators can go to the **Policy Management > Delegated Administration** page, 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 335](#).

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [The fundamentals of delegated administration, page 312](#)
- ◆ [Preparing for delegated administration, page 319](#)
- ◆ [Managing role conflicts, page 332](#)

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 325 .
Role	Click a role name to view or configure the role. See Editing roles, page 326 .
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 333 , for information about how a role's clients are managed after the role is deleted.
Advanced	Click to access the Manage Role Priority function.
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 332 .
View Administrator Accounts	Click to see the local and network administrator accounts with Web module access, and review their permission level and role assignments. See Reviewing administrator accounts, page 339 .

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 326 .

Adding roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Preparing for delegated administration, page 319](#)
- ◆ [Managing delegated administration roles, page 324](#)
- ◆ [Editing roles, page 326](#)

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 326](#).

- If you created a policy management and reporting role, it is added to the Role drop-down list in the Web module 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 317](#)).

Editing roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing delegated administration roles](#), page 324
- ◆ [Adding roles](#), page 325
- ◆ [Managing role conflicts](#), page 332

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).
Reporting	Give the administrator permission to use reporting tools.
Real-Time Monitor	Give the administrator permission to monitor all Internet 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 314.
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 329.
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 335.
Add	Opens the Add Managed Clients page. See Adding managed clients , page 330.
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 334, for more information.

4. Use the **Deployment Status Permissions** area to indicate whether administrators in this role can **Access the Status > Deployment page** to see information about the status of the components in your deployment.

If you grant delegated administrators access to the page, also select whether they can **Start components** or **Stop components**.

5. 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 Dashboard page.

- 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 113.
Access the Status > Dashboard page	Enables display of charts showing Internet activity on the Risks, Usage, and System dashboards. See The Status Dashboards , page 17. If this option is deselected, administrators can view only the Health Alert and Value Estimates (if displayed) sections of the System dashboard.
Access the Threats dashboard	Allows administrators to access charts, summary tables, and event details related to advanced malware threat activity in your network. See Threats dashboard , page 19.
Access forensics data in the Threats dashboard	(TRITON AP-WEB only) Allows administrators to view files associated with threat activity, and review information about attempts to send the files. See Configuring forensics data storage , page 409.
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 134.
View user names in investigative reports	Allows administrators in this role to view user names, if they are logged. See Configuring how requests are logged , page 386. 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	Allows administrators in this role to create favorite investigative reports. See Favorite investigative reports , page 149. This option is available only when administrators are granted access to investigative reports.
Schedule investigative reports	Allows administrators in this role to schedule investigative reports to run at a future time or on a repeating cycle. See Scheduling investigative reports , page 150. This option is available only when administrators are granted permissions to save investigative reports as favorites.

Option	Description
Manage the Log Database	Allows administrators to access the Settings > Reporting > Log Database page. See Log Database administration settings , page 396.
Access application reports	Allows administrators to see browser, platform, and user agent data on the Reporting > Applications page. See Application reporting , page 156.

- 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 and Deploy**.

Adding Administrators

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Delegated administrators](#), page 313
- ◆ [Editing roles](#), page 326

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 368.

- 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 [Connecting web protection software to a directory service](#), page 59) 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.
- 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.

- 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.

- 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 traffic in real time. See Real-Time Monitor , page 163.
Auditor	Give the administrator access to view all features available to other administrators in the role, without the ability to save changes.

- When you are finished making changes, click **OK** to return to the Edit Role page.
- Click **OK** on the Edit Role page to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Adding managed clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing delegated administration roles](#), page 324
- ◆ [Editing roles](#), page 326

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 IPv4 or v6 addresses), and networks (IPv4 or v6 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 414).

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 from the TRITON Manager](#), page 67.
 - Under **Computer**, enter the IP address to be added to this role in IPv4 or IPv6 format.
 - Under **Network**, enter the first and last IP addresses in a range in IPv4 or IPv6 format.
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 and Deploy**.

Managing role conflicts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Managing delegated administration roles, page 324](#)
- ◆ [Adding managed clients, page 330](#)

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 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 and Deploy**.

Updating delegated administration roles

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Delete roles, page 333](#)
- ◆ [Delete managed clients, page 334](#)

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 326).

Super Administrators can also move clients (see [Moving clients to roles](#), page 71) and policies ([Copying filters and policies to roles](#), page 242) 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 [Enforcement order](#), page 79). 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 and Deploy**.

The deleted role is cleared from Role drop-down list in the banner the next time you log on to the TRITON console.

Delete managed clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, 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 module 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 module 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 and Deploy**.

Managing Super Administrator clients

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 66. 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Any delegated administrator who uses a Websense account (not their network credentials) to log onto the TRITON console can review account their account information and change their password. See [View your user account](#), page 336.

Delegated administrators who have **policy** permissions can perform the following tasks.

- ◆ View their role definition.
Navigate to the **Policy Management > Delegated Administration** page and click the role name. This brings up the Edit Role page, which lists the role's managed clients and shows the reporting features available to administrators who have reporting permissions in the role.
- ◆ [Add clients to the Clients page](#), page 337.
- ◆ [Create policies and filters](#), page 338.
- ◆ Apply policies to clients on the Clients page (see [Assigning a policy to clients](#), page 78).

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.

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. For information about using any of those features, see:

- ◆ [The Status Dashboards](#), page 17
- ◆ [Presentation reports](#), page 113
- ◆ [Investigative reports](#), page 134
- ◆ [Application reporting](#), page 156
- ◆ [Real-Time Monitor](#), page 163

View your user account

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Performing delegated administrator tasks](#), page 335
- ◆ [Add clients to the Clients page](#), page 337
- ◆ [Create policies and filters](#), page 338

If you log on to the TRITON console 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 Web module **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.

Add clients to the Clients page

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Performing delegated administrator tasks, page 335](#)
- ◆ [View your user account, page 336](#)
- ◆ [Create policies and filters, page 338](#)

After Super Administrators assign managed clients to a role, delegated administrators can add them to the Clients page and assign them policies. See [Adding a client, page 66](#), for instructions.

When clients are added to a managed clients list, their Internet requests are immediately subject to 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 receive 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 that appears in the Managed Clients list 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 OU
- ◆ Individual computers that are members of the network

Because a user may be part of multiple groups or OUs, adding individuals from a larger client grouping has the potential to create conflicts when different roles manage groups or OUs with common members. If administrators in different roles access the Web module at the same time, they might add the same client (individual member of a group, for instance) to their Clients page. In that situation, policy enforcement for that client is governed by the priority established for each role. See [Managing role conflicts, page 332](#).

Create policies and filters

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Performing delegated administrator tasks, page 335](#)
- ◆ [View your user account, page 336](#)
- ◆ [Add clients to the Clients page, page 337](#)

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 74](#). For instructions on creating filters, see [Working with filters, page 42](#).

You can edit filter components for your role, with some limitations.

- ◆ **Categories:** Add or edit custom categories; assign custom URLs and keywords to custom or Master Database categories; change the action applied by default in category filters. (Changes to a category's default action are implemented only if the category is not locked by the Filter Lock.)
- ◆ **Protocols:** Change the action applied by default in protocol filters in your role. (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.

For more information, see [Building filter components](#), page 243.

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.

Reviewing administrator accounts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Delegated Administration > View Administrator Accounts** page to:

- ◆ See a list of local and network accounts that have been given Web module 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 the TRITON console with their network credentials.

This task is done **in addition to** the configuration done by Web module Super Administrators to define the directory service used to identify user and group clients.



Note

Client directory service information is configured on the **Settings > General > Directory Services** page (see [*Connecting web protection software to a directory service*](#), page 59).

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 Manager, 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.

16

Server Administration for Web Protection Solutions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Websense web protection components, page 342](#)
- ◆ [Understanding Policy Broker, page 351](#)
- ◆ [Working with Policy Server, page 354](#)
- ◆ [Working with Filtering Service, page 360](#)
- ◆ [Integrating with a third-party SIEM solution, page 366](#)
- ◆ [Working with Content Gateway, page 367](#)
- ◆ [Viewing and exporting the audit log, page 368](#)
- ◆ [Stopping and starting Websense services, page 370](#)
- ◆ [Alerting, page 373](#)

Internet policy enforcement requires interaction between several Websense components:

- ◆ User requests for Internet access are received by Content Gateway (TRITON AP-WEB), or by Network Agent or an integrated third-party product or device (Web Filter & Security).
- ◆ The requests are sent to Filtering Service for processing.
- ◆ Filtering Service communicates with Policy Server and Policy Broker to respond appropriately to requests.

The central Policy Broker gives other components access to client, filter, policy, and general configuration information. (It is possible to deploy additional, replica Policy Broker instances with read-only copies of this information, but only the central, or primary, instance is used to make updates to policy or configuration data.)

The TRITON Manager is associated with the central Policy Broker, and can be used to configure any Policy Server in the deployment.

Websense web protection components

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Understanding Policy Broker, page 351](#)
- ◆ [Working with Policy Server, page 354](#)
- ◆ [Working with Filtering Service, page 360](#)
- ◆ [Policy Server, Filtering Service, and State Server, page 363](#)
- ◆ [Stopping and starting Websense services, page 370](#)
- ◆ [Reviewing current system status, page 381](#)

Several Websense components work together to provide Internet security, user identification, and reporting capabilities. This section provides an overview of each component to help you understand and manage your environment.

For a list of components with a description of each, see:

- ◆ [Policy enforcement and management components, page 343](#)
- ◆ [Reporting components, page 345](#)
- ◆ [User identification components, page 346](#)
- ◆ [Interoperability components, page 347](#)

When Websense Web Filter & Security is integrated with Citrix, Microsoft Forefront TMG, or a proxy or proxy-cache that uses ICAP, an additional integration component (integration service, plugin, or ICAP server) is also installed.

Policy enforcement and management components

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Reporting components](#), page 345
- ◆ [User identification components](#), page 346
- ◆ [Interoperability components](#), page 347

Component	Description
Policy Database	Stores configuration 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. <p>Configure Policy Server settings in the TRITON Manager (see Working with Policy Server, page 354).</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 357).</p>
Filtering Service	<p>Provides Internet policy enforcement in conjunction with Content Gateway (TRITON AP-WEB), or with Network Agent or a third-party integration product (Web Filter & Security). 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 managed and logged. ◆ Each Filtering Service instance downloads its own copy of the Websense Master Database. <p>Configure Filtering Service behavior in the TRITON Manager (see Internet Usage Filters, page 31, and Configuring filtering settings, page 50).</p>
Network Agent	<ul style="list-style-type: none"> ◆ Enhances policy enforcement and logging functions ◆ Enables protocol management ◆ Enables policy enforcement in a standalone environment <p>For more information, see Configure Network Agent, page 417.</p>

Component	Description
Master Database	<ul style="list-style-type: none"> • Includes millions of websites, sorted into more than 90 categories and subcategories • Contains more than 100 protocol definitions for use in managing non-HTTP protocols <p>Download the Websense Master Database to activate policy enforcement, and make sure that the database is kept up to date. If the Master Database is more than 2 weeks old, no policy enforcement can occur. See The Websense Master Database, page 11, for more information.</p>
TRITON Infrastructure	<p>The platform that supports and unites the Web, Data, and Email modules of the TRITON Manager.</p> <p>Maintains an internal database of global settings that apply to all TRITON modules.</p>
Web module (<i>part of the TRITON Manager</i>)	<p>Serves as the configuration, management, and reporting interface for your web protection software.</p> <p>Use the Web module to define and customize Internet access policies, configure components, report on Internet activity, and more.</p> <p>The Web module 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 Working in the TRITON Manager, page 2, 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 Alerting, page 373, and Real-Time Monitor, page 163, for more information.</p>
Content Gateway	<ul style="list-style-type: none"> • Provides a robust proxy and cache platform. • Can analyze the content of websites and files in real time to categorize previously uncategorized sites. • Enables protocol management. • 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. <p>See Content Gateway Analysis, page 87.</p>

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 managed, and communicates with Remote Filtering Server. <p>See Manage Off-site Users, page 227, for more information.</p>
Remote Filtering Server	<ul style="list-style-type: none"> Allows policy enforcement for clients outside a network firewall. Communicates with Filtering Service to provide policy enforcement for remote machines. <p>See Manage Off-site Users, page 227, for more information.</p>
State Server	<p>In multiple Filtering Service environments, tracks client quota, confirm, password override, and account override sessions to ensure that access time is allocated correctly. To enable this functionality, deploy one State Server per Policy Server.</p>

Reporting components

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Policy enforcement and management components](#), page 343
- ◆ [User identification components](#), page 346
- ◆ [Interoperability components](#), page 347

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>Log Server is a Windows-only component that must be installed to enable most Web module reporting features. After installing Log Server, configure Filtering Service to pass logging data to the correct location (see Configuring how requests are logged, page 386).</p>
Log Database	<p>Stores Internet request data collected by Log Server for use by Websense reporting tools.</p>

Component	Description
Real-Time Monitor	<p>Displays current Internet 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 analysis) • 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 163.</p>
Multiplexer	<p>When enabled, passes logging data from Filtering Service to:</p> <ul style="list-style-type: none"> • A specified SIEM solution • Log Server <p>Used only when Websense software is integrated with a supported SIEM product. To enable SIEM integration, install one Multiplexer instance per Policy Server.</p>

User identification components

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Policy enforcement and management components](#), page 343
- ◆ [Reporting components](#), page 345
- ◆ [Interoperability components](#), page 347

Component	Description
User Service	<ul style="list-style-type: none"> • Communicates with your directory service. • Conveys user-to-group and user-to-domain relationships, to Filtering Service, for use in applying policies. • Enables display of directory client information in the Web module of the TRITON Manager. <p>For information about configuring directory service access, see Connecting web protection software to a directory service, page 59.</p>
DC Agent	<ul style="list-style-type: none"> • Offers transparent identification of users defined in a Windows-based directory service. • Communicates with User Service to provide up-to-date user logon session information for use in policy enforcement. <p>For more information, see DC Agent, page 288.</p>

Component	Description
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.</p> <p>For more information, see Logon Agent, page 293.</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 Configuring eDirectory Agent, page 297.</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 Configuring RADIUS Agent, page 295.</p>

Interoperability components

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Policy enforcement and management components](#), page 343
- ◆ [Reporting components](#), page 345
- ◆ [User identification components](#), page 346

Component	Description
Directory Agent	With the Web Hybrid module for TRITON AP-WEB, collects user and group information from a supported directory service for use by the hybrid service.
Filtering Plug-In	When Web Filter & Security is integrated with certain firewall, proxy, cache, or similar products, a plug-in may be installed to enable communication between Filtering Service and the integration.
Linking Service	With the Web DLP module for TRITON AP-WEB, or when TRITON AP-WEB integrates with TRITON AP-DATA, gives data protection components access to Master Database categorization information and user and group information collected by User Service.

Component	Description
Sync Service	With the Web Hybrid module for TRITON AP-WEB: <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.

Reviewing your web protection deployment

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Status > Deployment** page to review status information for each Policy Server in your deployment, and for the components that connect to each Policy Server. Also investigate User Service directory connection and lookup speeds.

The Deployment page includes up to 3 tabs:

- ◆ **Policy Server Map** gives a quick graphical and tabular overview of the Policy Server instances in your network. Click a Policy Server icon or IP address to see the status of components associated with the selected Policy Server. See [Using the Policy Server map, page 348](#).

If your deployment only has one Policy Server, this tab is not displayed.

- ◆ **Component List** provides a table listing the web protection components in your network, and allows administrators with appropriate permissions to stop or start components. See [Using the component list, page 349](#).
- ◆ **Directory Performance** provides information about connection and lookup speeds for each LDAP-based directory server that User Service queries for user and group information. See [Evaluating directory performance, page 350](#).

If User Service is not installed, or if your organization uses Windows Active Directory in mixed mode, this tab is not displayed.

Using the Policy Server map

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Websense web protection components, page 342](#)
- ◆ [Understanding Policy Broker, page 351](#)
- ◆ [Working with Policy Server, page 354](#)
- ◆ [Websense Health alerts, page 448](#)

In multiple Policy Server deployments, the **Policy Server Map** tab of the Status > Deployment page gives a graphical representation of all of your Policy Server instances.

- ◆ All additional Policy Server instances are shown connected to the central or base Policy Server for your deployment.
- ◆ Each Policy Server is represented by a server tower or appliance icon with markers that describe its Policy Broker connection.

A legend underneath the map explains the icons.

- ◆ Position the mouse over a Policy Server instance to see its full IP address and description, the IP address of the Policy Broker that it is currently connected to, and the Policy Broker mode (standalone, primary, or replica).

Configuration changes can be written to a standalone or primary Policy Broker, but replica Policy Broker instances are read-only.

Under the map, a table lists the IP address, description, Policy Broker IP address, key type, version, and current status of each Policy Server instance.

Click a Policy Server icon in the map or IP address in the table to see a list of the components (like Filtering Service, Log Server, and User Service) associated with the selected Policy Server instance. Note that in some cases, a single component name (like Real-Time Monitor) is used to represent multiple, interdependent services (like RTM Client, RTM Server, and RTM Database).

For each component, the list displays its name, IP address or hostname, version, and status.

The status column displays one of the following icons:

- ◆ A green icon with a check mark indicates that the Policy Server and its associated components are all running.
- ◆ A red icon with an “x” indicates that the Policy Server or at least one of its associated components is stopped.
- ◆ A yellow icon with an exclamation mark indicates that the Websense Control Service instance on the Policy Server machine is not available, so status information is not available for that Policy Server and its associated components.

For administrators with permissions to start and stop component services or demons, the table also includes a start or stop link.

In some cases, a single entry in the list may represent multiple services. In these cases, all of the services that make up the component are started or stopped when the link is clicked.

An additional link offers the option to show all health alerts associated with the selected Policy Server within the Components pop-up window.

Using the component list

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The **Component List** tab of the **Status > Deployment** page displays a table showing the web protection components deployed in your network. For each component, the table shows its:

- ◆ Name
- ◆ IP address or hostname
- ◆ Policy Broker IP address or hostname
- ◆ Version
- ◆ Status:
 - A green icon with a check mark indicates that the components is running.
 - A red icon with an “x” indicates that the component is stopped.
 - A yellow icon with an exclamation mark indicates that the Websense Control Service is not running, so status information is not available.

For administrators with permissions to start and stop component services or demons, the table also includes a start or stop link.

To export the component data for manipulation in a third-party spreadsheet or reporting tool, the **Export to CSV** link above the table.

Evaluating directory performance

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [User configuration and identification issues, page 438](#)
- ◆ [Working with users and groups, page 58](#)
- ◆ [User Identification for Policy Enforcement, page 277](#)

When User Service is installed and configured to connect to an LDAP-based directory service, the **Directory Performance** tab of the Status > Deployment page displays a table showing directory server performance statistics during the selected period (the last hour, by default).

Select a different **Time period** to see longer-term or more recent data. (The available time periods are last 24 hours, last hour, or last 5 minutes.)

The table contains a separate row for each directory server that User Service has attempted to connect to during the selected period. Each row shows:

- ◆ The IP address of the **Directory Host** machine
- ◆ The **Operation** type (bind or lookup)
- ◆ The **Average**, **Most Recent**, and **Maximum** times for each type of operation during the selected period. The time is shown in milliseconds.
- ◆ The number of attempts User Service made to perform each operation for the specified directory
- ◆ The number of times the operation failed

Click a Directory Host entry for more information about the performance of that directory since midnight, over the last hour, and during the most recent 5-minute

period (see [Review directory server details](#), page 351).

If users in your organization are experiencing browsing delays or sometimes receiving the incorrect policy (especially applied to the first web request of the day, or after a long period without browsing), use the directory performance statistics to identify underperforming directories. If there are persistent problems with specific directory hosts, you may need to take steps to improve:

- ◆ Network connections between User Service and the directory
- ◆ Memory, disk, or CPU speed on the directory server machine

Problems affecting multiple directories may indicate network, DNS, or other configuration issues.

Review directory server details

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Evaluating directory performance](#), page 350
- ◆ [User configuration and identification issues](#), page 438

Use the **Status > Deployment > Directory Server Details** page to review performance data for the specified directory since midnight, over the last hour, and during the most recent 5-minute period.

For each time period, a table displays the following information for bind (connection) and lookup operations:

- ◆ The **Average**, **Most Recent**, and **Maximum** times in milliseconds
- ◆ The number of attempts to perform the operation
- ◆ How many failures occurred

To return to the Directory Performance tab, click **Close**.

Understanding Policy Broker

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with Policy Server](#), page 354
- ◆ [Working with Filtering Service](#), page 360
- ◆ [Policy Server, Filtering Service, and State Server](#), page 363

Websense Policy Broker is responsible for managing access to both policy data (including clients, filters, filter components, and delegated administration settings) and to certain global settings that apply to the entire deployment. Settings specific to a single Policy Server instance (like its Filtering Service and Network Agent connections) are stored separately.

Even in multiple Policy Server environments, the same set of policy and general configuration data is shared throughout the deployment, thanks to Policy Broker.

1. At startup, each Websense component requests applicable configuration information from Policy Broker.
2. Running components frequently check for changes to configuration information.
3. The primary or standalone Policy Broker updates its database each time administrators make changes in the Web module of the TRITON Manager and click Save and Deploy.
4. After a configuration change, each component requests and receives the changes that affect its functioning via Policy Broker.

It is possible to install one or more Policy broker **replicas** in addition to the primary Policy Broker. In a replicated environment, changes made in the Web module of the TRITON Manager are saved to the primary Policy Broker. After the change, each replica synchronizes its copy of the data to receive the latest updates.

- ◆ The Policy Broker mode (standalone, primary, or replica) is set during installation, but can be changed later (for example, to change from a standalone environment to a replicated environment) using a command-line utility. See the [Websense Policy Broker](#) white paper for more information.
- ◆ In a replicated environment, you can configure a Policy Broker connection order for each Policy Server instance in your deployment. This determines where components attached to a Policy Server (like Filtering Service) look first for updates to configuration information. See [Reviewing Policy Broker connections](#), page 352.

Whether you have a single (standalone) Policy Broker or a primary Policy Broker with replicas, be sure to back up your policy and configuration data on a regular basis. See the [Backup and Restore FAQ](#) for more information.

Reviewing Policy Broker connections

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Understanding Policy Broker](#), page 351
- ◆ [Working with Policy Server](#), page 354

If you have a multiple Policy Broker environment (with a primary Policy Broker and one or more replicas), use the **Settings > General > Policy Brokers** page to find a list

of the Policy Broker instances in your deployment. You can also configure which instance each Policy Server in your network attempts to connect to first.

The **Installed Policy Broker Instances** table includes the following information:

- ◆ The **Host** column shows the IP address or hostname of the Policy Broker machine.
- ◆ The **Type** column indicates whether the instance is the primary or a replica. The primary instance always appears first in the list.
- ◆ A **Description** of the instance. Click the pencil icon next to the existing description to update it.
- ◆ When the **Last Policy Sync** occurred for each Policy Broker replica. This is the most recent time the replica received updated policy and configuration information from the primary Policy Broker.

Use the **Policy Server Connections** table to customize how the Policy Server instances in your deployment connect to Policy Broker. The table shows:

- ◆ The IP address or hostname of each Policy Server **Host**
- ◆ A **Description** of the Policy Server instance
- ◆ The **Connection Order** the Policy Server instance uses when it connects to Policy Broker (a list of IP addresses)

To change the connection order, click the Policy Server IP address or hostname. This opens the Policy Broker Connection Order window, with the current connection order listed. To move an instance up or down in the list:

1. Click on a row in the table to select the Policy Broker entry.
2. Click the **Up** or **Down** button to move the entry in the list.
3. Repeat for each entry that you want to move.
4. When you are finished making changes, click **OK** to return to the Policy Brokers page.
5. Click **OK** again on the Policy Brokers page to cache your changes. The changes are not implemented until you click **Save and Deploy**.

Working with Policy Server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Reviewing Policy Server connections, page 354](#)
- ◆ [Adding or editing Policy Server instances, page 356](#)
- ◆ [Working in a multiple Policy Server environment, page 357](#)
- ◆ [Changing the Policy Server IP address, page 357](#)
- ◆ [Working with Filtering Service, page 360](#)
- ◆ [Policy Server, Filtering Service, and State Server, page 363](#)

Policy Server is responsible for identifying other Websense software components and tracking their status.

When you log on to the Web module of the TRITON Manager, you are logging onto a graphical interface to Policy Server.

- ◆ You cannot log on to the Web module until it is configured to communicate with Policy Server.
- ◆ If your Websense software installation includes multiple Policy Servers, you can switch between Policy Server instances after logging on to the Web module.
- ◆ You can add and remove Policy Server instances within the Web module.

Communication between the Web module of the TRITON Manager and one Policy Server instance is established during installation.

Many 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 the Web module (see [Reviewing Policy Server connections, page 354](#)).

Reviewing Policy Server connections

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Settings > General > Policy Servers** page to review Policy Server information for all Policy Server instances associated with the Web module of the TRITON Manager.

If you have multiple Policy Server instances that share a subscription key, you can create one instance as the primary Policy Server. When you add the others as secondary instances, they receive their key information from the primary. This may help to speed up your configuration process and simplify key maintenance (in case you receive a new subscription key in the future).

- ◆ The Web module of the TRITON Manager is associated with a primary Policy Server instance at installation time. This becomes the base Policy Server for the Web module, and its IP address and description cannot be changed.
- ◆ To see the secondary Policy Server instances associated with a primary Policy Server in the list, click the “+” symbol next to the Policy Server name or IP address.
- ◆ To update the information that appears on the page (for example, to see the latest subscription key information or Policy Broker connections, and to see any Policy Server instances that might have recently been automatically added to the console) click the **Refresh** button in the toolbar at the top of the content pane.
- ◆ Policy Server instances that connect to a different Policy Broker than the base Policy Server are flagged with an icon (🚫) indicating that they are not currently configurable.

Each Policy Server entry includes a short description. Primary Policy Server entries also include:

- ◆ Subscription information, including the key associated with the instance and its secondaries and the subscription level (for example, TRITON AP-WEB or Web Filter & Security)
- ◆ The IP address of the Policy Broker that Policy Server is using
In multiple Policy Broker deployments, configure how Policy Server connects to Policy Broker on the **Settings > General > Policy Brokers** page.

Click **Add** to associate an additional Policy Server with the Web module, or click a Policy Server IP address or name to edit configuration information for the selected instance (see [Adding or editing Policy Server instances](#), page 356).

Note that in some cases, Policy Server instances are added to the Web module automatically. For example, when a Policy Server instance is installed on the same machine as a Policy Broker replica, that Policy Server instance appears on the Policy Servers page automatically. You can still edit these instances as needed (for example, to change their description).

Mark one or more Policy Server entries and click **Delete** to remove the connection between the TRITON Manager and the selected Policy Server.

- ◆ This removes the Policy Server instance from the Web module of the TRITON Manager, but does not uninstall or stop the Websense Policy Server service. You cannot delete the base Policy Server instance.
- ◆ Any time you remove a Policy Server instance from your deployment, be sure to also remove the instance from the Policy Servers page in the Web module of the TRITON Manager.

Even if you take down one Policy Server machine, then bring up a new machine and assign it the old IP address, a Policy Server instance installed on the new machine does not automatically inherit the subscription key information from the old instance. You must still delete the old instance from the Web module, then add the new 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 and Deploy**.

Adding or editing Policy Server instances

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Add Policy Server** or **Edit Policy Server** page to associate a new Policy Server instance with the Web module of the TRITON Manager, 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 the management console.
 - A secondary Policy Server uses the same subscription key as another Policy Server that has already been associated with the management console.
4. If this is a **secondary** Policy Server:
 - a. Select the IP address of the primary Policy Server from which the secondary should get its key.
 - b. Indicate whether this secondary should inherit its Directory Services settings from the primary Policy Server.
These are the settings (configured on the Settings > General > Directory Services page) that User Service uses to connect to a directory and retrieve user and group information.
 - c. Click **OK** to return to the Policy Servers page, then click **OK** again on the Policy Servers page to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Note that after adding a secondary Policy Server, you may have to log off of the TRITON console and log on again before you can use the Policy Server **Switch** button to connect to the new Policy Server instance.

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 the management console 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 the management console 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 and Deploy**.

Working in a multiple Policy Server environment

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

In distributed environments, or deployments with a large number of users, it may be appropriate to install multiple Policy Server instances. This entails some special considerations.

- ◆ Because policy information is managed by Policy Broker, policy changes are made available to all Policy Server instances when you click **Save and Deploy**.
- ◆ Many global configuration settings (like risk class definitions and alerting options) are also shared between Policy Server instances.
- ◆ 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.
- ◆ In order to apply time-based actions (Confirm, Quota, Password Override, or Account Override) correctly, one or more instances of Websense State Server is required. State Server allows the timing information associated with these features to be shared, so that clients are granted exactly the Internet access than you intend (see [Policy Server, Filtering Service, and State Server](#), page 363).

To switch between Policy Server instances in the management console:

1. In the Web module toolbar, next to the IP address of the current Policy Server, click **Switch**.
If there are unsaved changes to the current Policy Server instance, a warning prompt appears. To remain connected to the current Policy Server so that you can save your changes, click **Cancel**.
2. Select a Policy Server IP address or hostname from the **Connect to** list.
3. Click **OK**.

You are logged onto the selected Policy Server automatically, and the management console is updated.

Changing the Policy Server IP address

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Before changing the IP address of the Policy Server machine, **stop all Websense services** on the machine.

After changing the IP address, you must manually update Websense configuration files used by the Web module of the TRITON Manager, Policy Server, and other Websense services before policy enforcement resumes.

Step 1: Update management console configuration

Update the management console to use the new IP address to connect to Policy Server.

1. On the TRITON management server, stop the **Websense Web Reporting Tools** and **Websense TRITON - Web Security** services (if necessary).

If the TRITON console 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 any 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 370](#)).
2. Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default).
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. Open the Data Sources tool:
 - Windows Server 2012: Go to **Server Manager > Tools > ODBC Data Sources 64-bit**.
 - Windows Server 2008: Go to **Start > 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 management console 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 the TRITON Manager is installed on the same machine as Policy Server, administrators must use the new IP address to log on.

Working with Filtering Service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Review Filtering Service details, page 360](#)
- ◆ [Resuming Master Database downloads, page 362](#)
- ◆ [Policy Server, Filtering Service, and State Server, page 363](#)
- ◆ [Working with Content Gateway, page 367](#)

Filtering Service is the Websense software component that works with Content Gateway (TRITON AP-WEB), or with Network Agent or a third-party integration product, to provide policy enforcement. When a user requests a site, Filtering Service receives the request, determines which policy applies, and uses the applicable policy to determine whether the site is permitted or blocked.

Each Filtering Service instance downloads its own copy of the Websense Master Database to use in determining how to handle Internet requests.

If you have multiple Filtering Service instances, an additional component, Websense State Server, is required to enable correct application of time-based actions (Confirm, Quota, Password Override, or Account Override). State Server allows the timing information associated with these features to be shared, so that clients are granted exactly the Internet access than you intend (see [Policy Server, Filtering Service, and State Server, page 363](#)).

Filtering Service also sends information about Internet activity to Log Server, so that it can be recorded and used for reporting.

In the Web module of the TRITON Manager, a **Filtering Service Summary** on the System dashboard 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Status > Dashboard > 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 your Websense software is using Content Gateway, Network Agent (Web Filter & Security standalone mode), or a third-party product to detect Internet requests.
- ◆ The IP address and status of any Content Gateway instances with which the selected Filtering Service communicates
- ◆ The IP address and status of any Network Agent instances with which the selected Filtering Service communicates

Click **Close** to return to the Status > Dashboard page.

Review Master Database download status

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Each Filtering Service instance in your network downloads its own copy of the Master Database. When you are working in the management console, the **Status > Alerts** page displays a status message when a Master Database download is in progress, or an alert if a download attempt fails.

For detailed information about recent or ongoing database downloads, click **Database Download** on the toolbar at the top of the Status > Dashboard page. 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 the management console.

Resuming Master Database downloads

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 the Web module of the TRITON Manager, go to the **Status > Dashboard** page 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.

Filtering Service support for YouTube in Schools

Administrator Help | Web Filter & Security | Version 8.1.x8.1.x

Educational institutions with a software deployment of Websense Web Filter & Security can use a Filtering Service configuration parameter to enable YouTube for Schools. This YouTube service provides access to educational videos from inside the school network, even when other YouTube content is blocked.



Note

With TRITON AP-WEB, you can enable YouTube for Schools via Content Gateway, rather than via Filtering Service.

Once you have enrolled in the program and received a school account code or ID, first:

- ◆ In the Web module of the TRITON Manager, navigate to the **Settings > General > Filtering** page, and verify that **Enable search filtering** is selected at the bottom of the page.

You must enable search filtering to use the YouTube in Schools feature.

If search filtering was not previously enabled, click **OK** and **Save and Deploy** to cache and implement the change.

- ◆ Make sure that the YouTube is permitted for the clients that will be granted YouTube in Schools access.

When that configuration is complete, perform the following steps for each Filtering Service instance in your deployment:

1. Navigate to the Websense **bin** directory on the Filtering Service machine (C:\Program Files\Websense\Web Security\bin\ or /opt/Websense/bin/, by default).
2. Make a backup copy of the **eimserver.ini** file in another location.
3. Open the original **eimserver.ini** file and add the following lines:

```
[SafeSearchCustomValues]
YouTubeEDUFilter=<school_account_code>
```

Replace <school_account_code> with the actual code or ID received from YouTube.
4. Save and close the file.
5. Restart Filtering Service:
 - Windows: Use the Windows Services tool to restart **Websense Filtering Service**.
 - Linux: Use the /opt/Websense/WebsenseDaemonControl command to restart **Filtering Service**.

Policy Server, Filtering Service, and State Server


Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Working with Policy Server, page 354](#)
- ◆ [Working with Filtering Service, page 360](#)
- ◆ [Actions, page 40](#)
- ◆ [Password override, page 69](#)
- ◆ [Account override, page 70](#)

If your deployment includes multiple instances of Filtering Service that might handle a request from the same user, an optional component, **Websense State Server**, can be installed to enable proper application of time-based actions (Quota, Confirm) or overrides (Password Override, Account Override).

When State Server is installed, it allows its associated Filtering Service instances to share timing information, so users receive the correct allotment of quota, confirm, or override session time.




State Server is typically installed on a Policy Server machine, and only one State Server instance is required per **logical deployment**. A logical deployment is any group of Policy Server and Filtering Service instances that might handle requests from the same set of users.

- ◆ All Filtering Service instances that communicate with the same State Server instance must share the same time zone, and the time on all machines must be in synch.
- ◆ Each Filtering Service instance can communicate with only one State Server.
- ◆ All Filtering Service instances associated with the same Policy Server must communicate with the same State Server.
- ◆ Multiple Policy Server instances can share a single State Server.

Configure which State Server instance a Policy Server communicates with on the **Web > Settings > General > Filtering** page in the TRITON Manager (see [Configuring filtering settings](#), page 50).

In a geographically dispersed organization, where each location has its own Policy Server and Filtering Service instances, deploy one State Server instance (on the Policy Server machine or V-Series appliance) at each location. For example:



In an organization where all requests are managed through a central location, only one State Server instance is needed.

Integrating with a third-party SIEM solution

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Settings > General > SIEM Integration** page to configure Websense software to send log data from Filtering Service to a supported Security Information and Event Management (SIEM) solution.

Before using this page to enable SIEM integration, make sure an instance of Websense Multiplexer is installed for each Policy Server in your deployment.

Perform these steps for each Policy Server instance in your deployment.

1. Select **Enable SIEM integration for this Policy Server** to turn on the SIEM integration feature.
2. Provide the **IP address or hostname** of the machine hosting the SIEM product, as well as the communication **Port** to use for sending SIEM data.
3. Specify the **Transport protocol** (UDP or TCP) to use when sending data to the SIEM product.
4. Select the **SIEM format** to use. This determines the syntax of the string used to pass log data to the integration.
 - The available formats are syslog/CEF (ArcSight), syslog/key-value pairs (Splunk and others), syslog/LEEF (QRadar), and Custom.
 - If you select Custom, a text box is displayed. Enter or paste the string that you want to use. Click **View SIEM format strings** for a set of sample strings to use as a reference or template.
 - If you select a non-custom option, a sample **Format string** showing fields and value keys is displayed.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

When you save your changes, Websense Multiplexer connects to Filtering Service and takes over the job of distributing log data to both Log Server and the selected SIEM integration.

Note that although the same data is passed from Filtering Service to both Log Server and the SIEM product, Log Server may be configured to perform data reduction processing tasks (like recording visits instead of hits, or consolidating log records). Because the SIEM product does not perform these data reduction tasks, there may be more SIEM entries than records in the Log Database.

For more detailed information about the data passed to the SIEM integration, see [Integrating web protection solutions with third-party SIEM products](#). Subsections of the linked document provide mapping information for category numbers, disposition codes, reason strings, and other information included in the SIEM output.

Working with Content Gateway

Administrator Help | TRITON AP-WEB | Version 8.1.x

Related topics:

- ◆ [Content Gateway Analysis](#), page 87
- ◆ [Working with Filtering Service](#), page 360
- ◆ [Working with Policy Server](#), page 354
- ◆ [Managing Content Gateway connections](#), page 368

Content Gateway is a Linux-only Websense software component that provides high-performance web proxy services for TRITON AP-WEB deployments. Content Gateway is also used as a proxy by TRITON AP-DATA and TRITON AP-EMAIL solutions.

With TRITON AP-WEB, Content Gateway provides:

- ◆ Real-time content analysis and website 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 analysis 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-based policy enforcement

Content Gateway works with Filtering Service to manage 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 the Web module of the TRITON Manager with information about Filtering Service connections to Content Gateway
- ◆ Is used to populate the **Settings > General > Content Gateway Access** page in the Web module of the TRITON Manager, and makes it possible to launch the Content Gateway manager from within the TRITON console

Managing Content Gateway connections

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 the 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, hostname, 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 the 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 and Deploy**.

Viewing and exporting the audit log

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Websense software provides an audit trail showing which administrators have accessed the Web module of the TRITON Manager, 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 314).

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 popup window.

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.
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 and the Previous and Current values for a password change are left blank.

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, the exported file opens. Use options in Excel to save or print the file.

If Microsoft Excel is not installed, follow the on-screen instructions to either locate the software or save the file.

Stopping and starting Websense services

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 Websense TRITON - Web Security and Websense Web Reporting Tools services.
- ◆ 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).

From the management console

Unconditional Super Administrators (including admin) and delegated administrators with appropriate permissions can stop and start services from the **Status > Deployment** page in the Web module of the TRITON Manager.

Services can be started and stopped from either the Policy Server Map or the Component List tabs.

- ◆ On the Policy Server Map tab, click a Policy Server icon or IP address, then click the **Start** or **Stop** link for an associated component.
- ◆ On the Component List tab, find the appropriate component in the list, then click its **Start** or **Stop** link.

Windows

On Windows machines, use the following steps to stop or start **individual** services.

1. Open the Windows Services tool:
 - Windows Server 2012: **Server Manager > Tools > Services**
 - Windows Server 2008: **Start > Administrative Tools > Services**
2. Right-click the Websense service name, and then select **Stop** or **Start**.

To start, stop, or restart all services on the machine:

1. Navigate to the Websense **Web Security** folder (C:\Program Files *or* Program Files (x86)\Websense\Web Security\).
2. Stop, start, or restart the services with one of the following commands:

```
WebsenseAdmin start
WebsenseAdmin stop
WebsenseAdmin restart
```

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.

Websense appliance

On Websense appliances, use the 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 module and click **Restart Module**.
3. When the Network Agent module has restarted, go to the Web module (which displays your product name) and click **Restart Module**.

To stop the services (perhaps while performing a maintenance task):

1. Scroll to the Network Agent module on the **Status > General** page and click **Stop Services**.
2. For the Web module, also click **Stop Services**.
3. When you are ready start the services again:
 - a. For the Web module, click **Start Services**.
 - b. For the Network Agent module, click **Start Services**.

Installation directories for web protection solutions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The installation directory for your web protection solution depends on the machine operating system.

On **Windows** machines **other than** the TRITON management server, the default installation directory is:

```
C:\Program Files\WebSense\Web Security\
```

On the TRITON management server, the default installation directory is:

```
C:\Program Files (x86)\WebSense\Web Security
```

On **Linux** machines, the default installation directory is:

```
/opt/WebSense/
```

Alerting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Flood control, page 374](#)
- ◆ [Configuring general alert options, page 374](#)
- ◆ [Configuring system alerts, page 376](#)
- ◆ [Configuring category usage alerts, page 377](#)
- ◆ [Configuring protocol usage alerts, page 379](#)

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** notify administrators of Filtering Service 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** notify administrators when Internet activity for selected categories or protocols reaches configured thresholds.
Usage alerts can be generated for both Websense-defined and custom categories or protocols.
- ◆ **Suspicious activity alerts** notify administrators when threat-related events of a selected severity level reach a configured threshold or, for TRITON AP-WEB

subscribers who have purchased the Web Sandbox module, when a file that was sent to the sandbox is found to be malicious.

All alerts can be sent to selected recipients via email or SNMP.

Flood control

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Configuring general alert options, page 374](#)

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 374](#), for more information.

You can also set threshold limits for each category and protocol usage alert, and for each suspicious activity 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 377](#), and [Configuring protocol usage alerts, page 379](#), 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Configuring system alerts, page 376](#)
- ◆ [Configuring category usage alerts, page 377](#)
- ◆ [Configuring protocol usage alerts, page 379](#)
- ◆ [Configuring suspicious activity alerts, page 380](#)

Websense software can notify administrators of various kinds of system events, as well as Internet usage or suspicious activity that exceeds defined thresholds

Use the **Settings > Alerts > Enable Alerts** page to specify flood control settings, and to enable and configure one or more alerting notification methods. After enabling

alerting on this page, use the other pages in the Settings > Alerts section to specify which alerts you want to receive.

1. Under Alert Limits per 24 Hours, enter a number to specify the **Maximum daily alerts per type** to be generated for each category usage, protocol usage, and suspicious activity alert.

For example, you might configure a category usage alert 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 the maximum daily alerts per type is 10, administrators would receive alerts about the first 50 requests for Sports sites on a specific day (5 requests per alert multiplied by 10 alerts), but no alerts for subsequent requests for the category on the same day.

2. Mark **Enable email alerts** to deliver alerts and notifications by email. Then, configure these email settings.

SMTP server IPv4 address or name	IPv4 address or hostname 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 **Enable SNMP alerts** 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.
IPv4 address or hostname	The IPv4 address or hostname of the SNMP Trap server.
Port	Port number SNMP messages use. The default is 162.

4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring system alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Configuring general alert options, page 374](#)
- ◆ [Reviewing current system status, page 381](#)

The Web module of the TRITON Manager displays detailed system health and status information via the **Status > Alerts** page, described in [Reviewing current system status, page 381](#).

To assure that administrators are notified of significant system events, configure Websense system alerts to be distributed by email or via SNMP trap.

TRITON AP-WEB administrators have the option to enable system alerts for both Filtering Service 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 boxes to the right of the message summary to indicate how to notify administrators. Depending on what methods are enabled on the Enable Alerts page, you may be able to choose **Email**, **SNMP**, or a combination.

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 Filtering Service 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.

With TRITON AP-WEB, you have the option to enable the following additional system alerts:

- ◆ A domain controller is down.
- ◆ Decryption and inspection of secure content has been disabled.
- ◆ Log space is critically low.

- ◆ Subscription information could not be retrieved.
- ◆ Non-critical alerts have been received. (See [Content Gateway non-critical alerts](#), page 476, 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 and Deploy**.

Configuring category usage alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting](#), page 373
- ◆ [Flood control](#), page 374
- ◆ [Configuring general alert options](#), page 374
- ◆ [Adding or editing category usage alerts](#), page 378

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.

Use the **Settings > 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 or editing category usage alerts](#), page 378) 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 and Deploy**.

Adding or editing category usage alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Configuring general alert options, page 374](#)
- ◆ [Configuring category usage alerts, page 377](#)

Use the **Category Usage > Add Category Usage Alerts** or **Edit Category Usage Alerts** page to:

- ◆ (Add page only) Select new categories for usage alerts
- ◆ Establish or change the threshold for usage alerts
- ◆ Select or update alerting methods (email, SNMP)

If you are creating one or more new alerts, start by marking the check box next to each category that you want to add with the same threshold and alert methods.



Note

You cannot add usage alerts for any category that is excluded from logging. See [Configuring how requests are logged, page 386](#).

The remaining steps are available both for adding and for editing usage alerts:

1. Set or update the **Threshold** by selecting the number of requests that cause an alert to be generated.
2. Mark the check box for each desired alert method (**Email, SNMP**) for these categories.

Only the alert methods that have been enabled on the Alerts page (see [Configuring general alert options, page 374](#)) are available for selection.

3. Click **OK** to cache your changes and return to the Category Usage page (see [Configuring category usage alerts, page 377](#)). Changes are not implemented until you click **Save and Deploy**.

Configuring protocol usage alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Flood control, page 374](#)
- ◆ [Configuring general alert options, page 374](#)
- ◆ [Adding or editing protocol usage alerts, page 379](#)

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 or editing protocol usage alerts, page 379](#)) 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 and Deploy**.

Adding or editing protocol usage alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting, page 373](#)
- ◆ [Configuring general alert options, page 374](#)
- ◆ [Configuring protocol usage alerts, page 379](#)

Use the **Protocol Usage > Add Protocol Usage Alerts** or **Edit Protocol Usage Alerts** page to:

- ◆ (Add page) Select new protocols for usage alerts
- ◆ Establish or update the threshold for usage alerts
- ◆ Select or update alert methods (email, SNMP) for the alerts

If you are creating new protocol usage alerts, start by marking the check box next to each protocol to be added with the same threshold and alert methods.

**Note**

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.

The remaining steps are available both for adding and for editing usage alerts:

1. Set or change the **Threshold** by selecting the number of requests that cause an alert to be generated.
2. Select each desired alert method (**Email**, **SNMP**) for these protocols.
Only the alert methods that have been enabled on the Alerts page (see [Configuring general alert options](#), page 374) are available for selection.
3. Click **OK** to cache changes and return to the Protocol Usage page (see [Configuring protocol usage alerts](#), page 379). Changes are not implemented until you click **Save and Deploy**.

Configuring suspicious activity alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Alerting](#), page 373
- ◆ [Flood control](#), page 374
- ◆ [Configuring general alert options](#), page 374

Your Websense software can notify you when suspicious activity of a specified severity level reaches a defined threshold. You can define alerts for permitted requests and blocked requests of each severity level.

Because Content Gateway is required to detect critical and high severity alerts, it is not possible to configure alerting for those severity levels in Web Filter & Security deployments.

TRITON AP-WEB subscribers who have purchased the Web Sandbox module can enable email or SNMP alerts to be sent when a file submitted to the file sandbox is determined to be malicious.

Use the **Settings > Alerts > Suspicious Activity** page to enable, disable, or change alerting configuration for alerts associated with suspicious events in your network. Detailed information about these events is displayed on the Threats dashboard.

The page displays 2 tables: **Permitted Suspicious Activity Alerts** and **Blocked Suspicious Activity Alerts**. Each table shows:

- ◆ The **Severity** level to be configured. The 4 severity levels are critical, high, medium, and low. Severity level is determined by the threat category associated with the alert. See [How severity is assigned to suspicious activity, page 23](#), for more information.
- ◆ The alerting **Threshold**. By default, the threshold for critical and high severity alerts, both permitted and blocked, is **1**.
- ◆ One or more notification methods. Suspicious activity alerts can be sent via **Email**, **SNMP**, or both.
- ◆ Whether or not the alert is **Enabled**. A green check mark indicates that alerts are being generated for suspicious activity of the selected severity. A red “X” indicates that alerting is disabled for the selected severity.

To update suspicious activity alert settings, you can:

1. Mark the check box to the left of a severity level, then click **Enable** or **Disable** to activate or stop alerts of the selected type.
2. For enabled alerts, enter a number in the **Threshold** field to specify the number of suspicious events that cause an alert to be generated.
3. Select each notification method (**Email**, **SNMP**) to use to deliver suspicious activity alerts.

Only alert methods that have been enabled on the Enable Alerts page (see [Configuring general alert options, page 374](#)) are available for selection.

4. If the Web Sandbox module is included in your configuration, mark the check box or boxes in the File Sandboxing Alerts section to cause an email or SNMP alert to be sent when a file sent for sandbox analysis is found to be malicious.

Each check box is enabled only if the corresponding alert type (email or SNMP) is enabled on the Enable Alerts page.

Note that threats related to file sandboxing are not included on the Threats dashboard.

5. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Reviewing current system status

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 **Hide Persistent Alerts**. If your organization does not use Log Server, Network Agent, or User Service, or if you do not plan to enable WebCatcher, mark the appropriate check box in the **Hide Alert** column of the table. Alerts associated with the selected service are no longer displayed.

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.

The **Real-Time Security Updates** 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 website or application.

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.

17

Reporting Administration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring how requests are logged](#), page 386
- ◆ [Assigning categories to risk classes](#), page 384
- ◆ [Configuring reporting preferences](#), page 385
- ◆ [Configuring Log Server](#), page 388
- ◆ [Log Database administration settings](#), page 396
- ◆ [Configuring investigative reports](#), page 410
- ◆ [Self-reporting](#), page 414

In organizations that use only the default administrator account (admin), everyone who uses the TRITON Manager has access to all reporting settings and tools. In organizations that use delegated administration, access to reporting settings and tools is controlled by members of the Super Administrator role (see [Editing roles](#), page 326).

Administrators with access to reporting settings have many options for customizing reporting in their 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 384.
- ◆ 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 management server. Also configure whether Real-Time Monitor collects data all the time, or only when Real-Time Monitor is open. See [Configuring reporting preferences](#), page 385.

Logging is the process of storing information about Internet activity 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 requests are logged](#), page 386, for more information.

- ◆ Use the **Settings > Reporting > Log Server** page to manage the way the log records are processed and connections to the Log Database. See [Configuring Log Server, page 388](#).
- ◆ Use the **Settings > Reporting > Log Database** page to administer the Log Database, including database partition, URL logging, browse time, and trend data options. See [Log Database administration settings, page 396](#).

An end user who uses the Filtering Service has no direct or indirect influence over the database. Thus, although the log entry is stored in the MSSQL database, the user did not direct its storage and cannot retrieve it.

The only interface to the database itself is from the Log Server, the Reporting services, and the Manager. Filtering Service and Content Gateway do not access the database, but instead send information via the Log Server.

Assigning categories to risk classes

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Risk classes, page 36](#)
- ◆ [Block Page Management, page 179](#)
- ◆ [Use Reports to Evaluate Internet Activity, page 111](#)

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 **Status > Dashboard** page includes charts that track Internet activity by risk class, and you can generate presentation or investigative reports organized by risk class.

Use the **Settings > General > Risk Classes** page to review or change which categories comprise each risk class.

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 and Deploy**.

Configuring reporting preferences

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Self-reporting, page 414](#)
- ◆ [Scheduling presentation reports, page 125](#)
- ◆ [Scheduling investigative reports, page 150](#)

Use the **Settings > Reporting > Preferences** page to provide information used to send completed scheduled reports to selected recipients via email, activate self-reporting, determine how long scheduled presentation reports are stored, and configure when Real-Time Monitor collects data.

1. Under Email Reports, enter the **Email address** to display in the “From” field when scheduled reports are distributed via email.
2. Enter the **SMTP server IPv4 address or name** for the email server to use for distributing scheduled reports.
3. Mark the **Allow self-reporting** check box to let end users in your organization access the TRITON Manager to run investigative reports on their personal Internet activity.

When this option is selected, the URL used to access self-reporting features is displayed. See [Self-reporting, page 414](#).

4. Under Scheduled Presentation Reports, use the **Store reports for** drop-down list to indicate how long reports are stored on the TRITON management server machine (5 days, by default).

As you increase the length of time that reports are stored, you affect the amount of disk space required on the TRITON management server. The management server is not an appropriate location for a long-term reporting archive.

5. Use the **Warn administrators...** drop-down list to indicate how long a warning is displayed on the Review Reports page before a scheduled presentation report is deleted (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 management server.

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. With this option selected, data collection begins when you launch Real-Time Monitor. There may be 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 requests are logged

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Introducing the Log Database for web protection solutions, page 394](#)
- ◆ [Configuring Log Server, page 388](#)

Use the **Settings > General > Logging** page to:

- ◆ Provide the IP address and port that Filtering Service uses to send log records to Log Server.
- ◆ (*Web Hybrid module for TRITON AP-WEB*) 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 377](#)).

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:

- ◆ Each Policy Server can communicate with a single Log Server instance.

- ◆ For reporting data to display, there must be a Log Server associated with the base Policy Server (the Policy Server instance specified during installation, noted on the **Settings > General > Policy Servers** page).

This is typically the Policy Server installed with Policy Broker (for example, the Policy Server on the full policy source appliance).

- ◆ If the Log Server IP address and port are blank for any Policy Server, the Filtering Service instances 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 are 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 **Log Server IPv4 address or hostname**.
2. Enter the **Port** that Filtering Service uses to send log records to Log Server (55805, by default).
3. (*Web Hybrid module for TRITON AP-WEB*) 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 the management console 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 hostname 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**.

- If you are using TRITON AP-WEB, and want Threats dashboard tables to include source device name information, when available, click **Log hostnames**.

Name information is available in threat-related logs only. It is not available for Internet activity to which no severity is assigned.

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 377), no usage alerts can be sent.

Reports cannot include information about categories that are not logged.

- Use the **Find category** search box to quickly jump to a specific category.
 - 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 and Deploy**.

Configuring Log Server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

During installation, you configure certain aspects of Log Server operation, including how Log Server interacts with policy enforcement components. Use the **Settings > Reporting > Log Server** page to update these settings, or to configure other details about Log Server operation.

When you finish your configuration updates, click **OK** to cache your changes. Changes are not saved until you click **Save and Deploy**.

If you make changes to the database connection, after saving and deploying the changes, also restart the **Websense TRITON - Web Security** service on the management server machine to update the database connection for all reporting tools.

In multiple Log Server environments, the settings configured on this page apply to the Log Server instance assigned to the Policy Server whose IP address appears on the Web module toolbar.

Verify basic Log Server details

Under **Location**, verify the Log Server IP address. If necessary, use the **Port** field to update the port over which Log Server communicates with Filtering Service (55805, by default).

This port must match the logging port displayed on the **Settings > General > Logging** page.

Configure the Log Database connection

Under **Log Database Connection**, configure the ODBC connection that Log Server uses to connect to the Log Database.

1. Specify the ODBC **Data source name (DSN)** and enter a unique **Description** for the database connection.
2. Provide the **SQL Server location** (IP address or hostname and instance name, if applicable) for the Microsoft SQL Server installation that hosts the Log Database, as well as the **Connection port** for sending data to the Log Database (1433, by default).
3. If your environment uses SQL Server clustering, enter the virtual IP address for the cluster.
4. Enter the name of the **Default database** (wslogdb70, by default).
5. Indicate whether or not to **Use SSL to connect to the Log Database**. When SSL encryption is enabled:
 - BCP cannot be used to add records to the Log Database.
 - Log Database connections are slower, affecting reporting performance.



Important

When Microsoft SQL Server components are configured so that “Trust Server Certificate” is set to **No** (the default), self-signed SSL certificates are not accepted for encryption of database connections.

In this case, SSL certificates signed by a Certificate Authority must be properly deployed to the SQL Server, TRITON management server, and Log Server machines before you enable the “Use SSL” option in the management console.

See your SQL Server documentation for information about database encryption.

6. Specify a Log Server connection method:
 - By default, **SQL Server authentication** is selected. To use SQL Server authentication, provide the SQL Server **Account** and **Password** to use.
 - Alternatively, you can use a **Windows trusted connect (network logon account)**. The Websense Log Server service must be configured to run as this account.
7. Click **Test Connection** to verify that it is possible to connect to the Log Database using the credentials provided.

For information about the tests performed when you click the button, see [Testing the Log Database connection](#), page 393.

If you make changes to the database connection, after saving and deploying the changes, also restart the **Websense TRITON - Web Security** service on the management server machine to update the database connection for all reporting tools.

Specify how log records are processed into the database

Click **Log Record Creation** to specify how Log Server adds records to the Log Database.

- ◆ **ODBC (Open Database Connectivity)** inserts records into the database individually, using a database driver to manage data between Log Server and Log Database.

If you select this option, also set the **Maximum number of connections** to specify how many internal connections can be made between Log Server and the database engine.

Select a value between 4 and 50, as appropriate for your SQL Server license.



Note

Increasing the number of connections can increase processing speed for log records, but could impact other processes in the network that use the same SQL Server. In most cases, you should set the number of connections to fewer than 20. Contact your Database Administrator for assistance.

- ◆ **BCP (Bulk Copy Program)** (*recommended*) inserts records into the Log Database in batches. This option offers better efficiency than ODBC insertion, and is selected by default if the **bcp.exe** file is found on the machine.

The BCP option is available only if you install the SQL Server Native Client and Command Line Utilities on the Log Server machine.

BCP cannot be used when SQL Server SSL encryption is used.

If you select the BCP option, also specify:

Option	Description
BCP file location	Directory path for storing BCP files. Log Server must have read and write access to the location. (The default folder is C:\Program Files\WebSense\Web Security\bin\Cache\BCP.) After entering the path, click Test Location to verify that the location is accessible.
File creation rate	Maximum number of minutes Log Server spends placing records into a batch file before closing that batch file and creating a new one. This setting works in combination with the batch size setting: Log Server creates a new batch file as soon as either limit is reached.
Maximum batch size	Maximum number of log records before a new batch file is created. This setting works in combination with the creation rate setting: Log Server creates a new batch file as soon as either limit is reached.

After selecting a log record insertion method, click **Log Cache Files** to specify where and how log cache files are created. These provide temporary storage for log records that have not yet been processed into the Log Database or moved to BCP files.

1. For **Cache location**, indicate where on the Log Server machine logging cache files are stored (C:\Program Files\WebSense\Web Security\bin\Cache\, by default).
2. Click **Test Location** to verify that the path is accessible.
3. For **Cache file creation rate**, indicate the maximum number of minutes (1, by default) Log Server should spend sending Internet access information to a log cache file before closing it and creating a new file.
4. For **Maximum cache file size**, specify how large a log cache file should be before Log Server closes it and creates a new one.

The file creation rate and maximum file size settings work in combination: Log Server creates a new log cache file as soon as either limit is reached.

Adjust database sizing settings

Configure **Database Size Management** settings to meet your organization's needs. The higher the level of detail recorded, the larger the Log Database.

1. To minimize the size of the Log Database, mark **Enable log record consolidation**. This combines multiple, similar Internet requests into a single log record, reducing the granularity of reporting data.

If you have enable SIEM integration, note that Log Server applies consolidation to the log records that it processes into the Log Database. Consolidation does not occur for records passed to the SIEM product.

When consolidation is enabled, requests that share all of the following elements are combined into a single log record:

- Domain name (for example: www.websense.com)
- Category
- Keyword
- Action (for example: Category Blocked)
- User/IP address

The log record includes the number of requests combined into the consolidated record, as well as the total bandwidth for all of the consolidated requests.

Reports run faster when the Log Database is smaller. However, consolidation may decrease the accuracy of some detail reports, as separate records for the same domain name may be lost.



Important

To assure consistent reports, create a new database partition whenever you enable or disable consolidation. Also, be sure to generate reports from partitions with the same consolidation setting.

With TRITON AP-WEB, when consolidation is enabled, numbers shown in reports that include traffic blocked by scanning are **lower** than the numbers shown on reports about Content Gateway analysis. This is a side-effect of the way that analytic activity is recorded.

2. If you enable consolidation, also specify the **Consolidation time interval**. This represents the greatest allowable time difference between the earliest and latest records combined to make one consolidation record.

Decrease the interval to increase granularity for reporting. Increase the interval to maximize consolidation. Be aware that a larger interval can also increase usage of system resources, such as memory, CPU, and disk space.

If you enable full URL logging on the **Settings > Reporting > Log Database** page, consolidated log records contain the full path (up to 255 characters) of the first matching site Log Server encounters.

For example, suppose a user visited the following sites and all were categorized in the shopping category.

- www.domain.com/shoeshopping
- www.domain.com/purseshopping
- www.domain.com/jewelryshopping

With full URL logging enabled, consolidation creates a single log entry showing 3 requests for the URL www.domain.com/shoeshopping.

3. Under Hits and Visits, use the **Enable visits** check box to indicate the level of granularity recorded for each user Internet request.

**Note**

It is best to create a new database partition prior to changing the method of logging between visits and hits. See the **Settings > Reporting > Log Database** page to create a new database partition.

When this option is **not** selected, a separate log record is created for each HTTP request generated to display different page elements, including graphics, advertisements, embedded videos, and so on. Also known as logging hits, this creates a much larger Log Database that grows rapidly.

When this option **is** selected, Log Server combines the individual elements that create the web page (such as graphics and advertisements) into a single log record that includes bandwidth information for all elements of the visit.

With TRITON AP-WEB, when visits are enabled, numbers shown in reports that include traffic blocked by real-time analysis are **lower** than the numbers shown on Content Gateway analysis-specific reports. This is a side-effect of the way that analytic activity is recorded.

Configure User Service communication

Click the **User Service Connection** button, then use the **User and group update interval** field to indicate how often Log Server connects to User Service to retrieve full user name and group assignment information (every 12 hours, by default).

Activity for a user whose user name or group information has changed continues to be reported with the original user name or group assignment until the next update occurs. Organizations that update their directory service frequently or have a large number of users should consider updating the user/group information more frequently.

Testing the Log Database connection

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The database connection information used by Log Server and other reporting tools can be updated on the **Settings > Reporting > Log Server** page in the Web module of the TRITON Manager.

The Log Database Connection section of the page includes a **Test Connection** button. When you click the button, Log Server performs the following tests:

1. Log Server retrieves the updated database connection information from the management console.

If Log Server is stopped, or the network is down between the TRITON management server and the Log Server machine, this test fails. If the connection to Log Server fails, an IO exception error is likely to display.

2. Log Server uses ODBC to create a data source name (DSN) for testing purposes.
3. Log Server uses the DSN to establish a connection to the Log Database. Log Server checks to see that:
 - A Websense database exists.
 - The database version is correct.
4. Log Server verifies its database permissions.
See [Log Database Permissions](#) for information about the required database roles and permissions.
5. Log Server deletes the DSN it created for testing.
6. Log Server notifies the management console that its tests succeeded.

If this return notification fails, an IO exception error is likely to display.

In addition, the Web module of the TRITON Manager verifies that it can create a JDBC connection to the database. This test may pass even when a Log Server test fails.

The new database connection information is not used until you cache and save your changes. At that point:

- ◆ The new database connection information is saved to the Policy Server configuration file.
- ◆ Log Server creates a permanent DSN (reproducing the temporary DSN created during the connection test).

Restart the **Websense TRITON - Web Security** service to update reporting tools (like presentation reports) to use the new database connection.

Introducing the Log Database for web protection solutions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Web protection reporting database jobs, page 395](#)
- ◆ [Log Database administration settings, page 396](#)

The Log Database stores the records of Internet activity handled by web protection components. Installation creates the Log Database with a catalog database and one database partition.

The **catalog database** (wslogdb70, by default) provides a single connection point for the various Websense components that need to access the Log Database: dashboards, Log Server, presentation reports, and investigative reports. It contains supporting

information for the database partitions, including the list of category names, risk class definitions, trend data, 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. There are 2 partition types:

- ◆ The standard logging partition (wslogdb70_1, wslogdb70_2, etc.) stores information about all logged Internet requests. Information from the standard logging partition is used to populate investigative and presentation reports, as well as dashboard charts.
- ◆ The threats partition (wslogdb70_amt_1) stores information about requests that have been assigned a severity level (see [How severity is assigned to suspicious activity](#), page 23). Information from the threats partition is used to populate the Threats dashboard.

New standard logging partitions are created based on size or date interval. See [Configuring database partition options](#), page 397, 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 standard logging 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.

Web protection reporting database jobs

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The following database jobs are installed along with the Log Database.



Important

If you are using a full version of Microsoft SQL Server (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 standard logging partition database. When trend data retention is enabled, the ETL job is also responsible for inserting trend data into the catalog 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.
- ◆ When trend data retention is enabled, the **trend job** uses daily trend data created by the ETL job to update weekly, monthly, and yearly trend records for use in presentation reports.

Even when trend data retention is disabled, the trend job processes data from the threats (AMT) partition to provide trend data on the Threats dashboard.

The trend job also parses user agent strings to populate the Browser and Source Platform tabs of the Reporting > Applications page (see [Application reporting, page 156](#)).

This job runs nightly.

- ◆ The **Advanced Malware Threat (AMT) ETL job** receives, processes, and inserts data into the threats partition database. Only log records that include a severity ranking (see [How severity is assigned to suspicious activity, page 23](#)) are recorded in the threats partition. Data from this partition is used to populate the Threats dashboard (see [Threats dashboard, page 19](#)).

Certain aspects of these database jobs can be configured on the Settings > Reporting > Log Database page. See [Log Database administration settings, page 396](#), 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. When trend data retention is enabled, the trend job is run, by default, at 4:30 a.m. Try to avoid starting other jobs at times that might overlap with the trend job.

Log Database administration settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Settings > Reporting > Log Database** page to manage:

- ◆ When, where, and how the Log Database creates new standard logging partitions, and which partitions are used in creating reports ([Configuring database partition options, page 397](#))
- ◆ When and how maintenance jobs are run (see [Configuring Log Database maintenance options, page 400](#))
- ◆ 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 402](#))
- ◆ How Internet browse time is calculated (see [Configuring Internet browse time options, page 403](#))

- ◆ Whether and how long trend and application data should be stored (see [Configuring trend and application data retention](#), page 404).

The name of the active Log Database instance is displayed at the top of the page.

Configuring database partition options

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Log Database administration settings](#), page 396
- ◆ [Log Database sizing guidance](#), page 406
- ◆ [Configuring Internet browse time options](#), page 403
- ◆ [Configuring how URLs are logged](#), page 402
- ◆ [Configuring Log Database maintenance options](#), page 400
- ◆ [Configuring trend and application data retention](#), page 404

Use the **Database Rollover Configuration** 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.

Refer to the **Growth Rates and Sizing** chart at the bottom of the Database Rollover Configuration section for average daily database partition size over time. This may be helpful in planning for future growth, determining how frequently to create new partitions, and in setting partition size and growth options.

- ◆ Use the drop-down list under the chart to configure the **Time period** displayed. (The time period is based on the partition creation date; not the dates that the partition spans.) You can display partitions created in the last 1 week, 1 month, 3 months, 6 months, or show all available partitions.
Note that when a longer time period is selected, each partition may appear as a small dot on the chart.
- ◆ Indicate whether or not to **Show chart legend**. When displayed, the legend indicates which partitions (by name) are mapped in the chart.
The legend is only available when the chart includes 20 or fewer partitions for the selected time period.
- ◆ Select a section of the chart to view it more closely. Click **Zoom Out** or **Reset Chart** to reduce the level of detail shown.

For more help with database sizing, see [Log Database sizing guidance](#), page 406.

To manage database rollover and growth:

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:
 - *SQL Server Standard or Enterprise*: 100-1,000,000 MB, default 5000
 - *Microsoft SQL Server Express*: 100-5000 MB, default 3000
- 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, some organizations set the automatic rollover to a long time period or large maximum size. Then, they perform manual rollovers to prevent the automatic rollover from occurring. See [Configuring Log Database maintenance options](#), page 400, 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, the partition is automatically enabled for use in reporting.

2. Under Partition Management, 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** and **Log** files that make up new database partitions.
 - *SQL Server Standard or Enterprise*: Data file initial size 50-500,000 MB, default 2000; Log file initial size 50-250,000 MB, default 100

- *SQL Server Express*: Data file initial size 50-5000 MB, default 100; Log file initial size 50-4000 MB, default 100



Note

As a best practice, calculate the average partition size over a period of time, then update the initial size to approximate that value. You might, for example, set the initial size to 80% of the average size. This minimizes the number of times the partition must be expanded, and frees resources to process data into the partitions.

Use the information in the Growth Rates and Sizing list (below the list of available partitions) for help in making this calculation.

- c. Under **Growth**, set the increment by which to increase the size of a partition's **Data** and **Log** files when additional space is required.
 - *SQL Server Standard or Enterprise*: Data file growth 100-500,000 MB, default 500; Log file size 1-250,000 MB, default 100
 - *SQL Server Express*: Data file growth 1-1000 MB, default 100; Log file size 1-1000 MB, default 100
3. If you want to create a partition the next time the ETL job runs (see [Web protection reporting database jobs](#), page 395), 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 and Deploy** before clicking **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. The number of Available Partitions that can be included in the list is limited to 500. A SQL Server error may occur if you have more than 500 standard logging partitions in your Log Database. 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.
A maximum of 70 partitions can be enabled on one time.
 - 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.
 - 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.

Only enabled partitions can be deleted. To delete a disabled partition, first enable it, then delete it.



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 400, 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 and Deploy**.

Configuring Log Database maintenance options

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Log Database administration settings](#), page 396
- ◆ [Configuring database partition options](#), page 397
- ◆ [Configuring Internet browse time options](#), page 403
- ◆ [Configuring how URLs are logged](#), page 402
- ◆ [Configuring trend and application data retention](#), page 404

Use the **Database Maintenance** 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 (01:00 hours, by default).

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 403).

2. To permanently delete partitions based on age, select **Automatically delete partitions when data is older than**, and then specify the number of days (from 1 to 1825) after which to delete the partitions.



Warning

After a partition has been deleted, the data cannot be recovered. See [Configuring database partition options, page 397](#), for an alternative way to delete partitions.

3. Select **Enable automatic reindexing of partitions**, and then select a day of the week to have this processing performed automatically each week (Saturday, by default).

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. Select **Process failed batches during the database maintenance job** to have the nightly database maintenance job reprocess any failed batches.

Failed batches occur when there is insufficient disk space or inadequate database permissions to insert log records into the database. Typically, these batches are successfully reprocessed and inserted into the database during the nightly database maintenance job. Reprocessing, however, cannot be successful if the disk space or permission problem has not been resolved.

If this option is unchecked, failed batches are never reprocessed. Instead, they are deleted after the time specified (below), if any.

5. Select **Delete failed batches after** and then enter a number of days (from 0 to 90; 20, by default) after which to delete any failed batches.

If this option is not selected, failed batches are retained indefinitely for future processing.

6. Select **Delete the error log after**, and then enter a number of days (0 to 90; 60, by default) after which to delete database error records from the catalog database.

If this option is not checked, error logs are retained indefinitely.

7. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring how URLs are logged

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Log Database administration settings, page 396](#)
- ◆ [Configuring database partition options, page 397](#)
- ◆ [Configuring Internet browse time options, page 403](#)
- ◆ [Configuring Log Database maintenance options, page 400](#)
- ◆ [Configuring trend and application data retention, page 404](#)

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. Select **Record domain and 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 advanced analysis activity, page 105](#)). 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.

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 [Configuring Log Server, page 388](#), for more information.

2. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring Internet browse time options

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Log Database administration settings, page 396](#)
- ◆ [Configuring database partition options, page 397](#)
- ◆ [Configuring how URLs are logged, page 402](#)
- ◆ [Configuring Log Database maintenance options, page 400](#)
- ◆ [Configuring trend and application data retention, page 404](#)

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** section of the Settings > Reporting > Log Database page.

1. Select an **IBT 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 400](#)), and to select a slow time on the network to minimize any impact on generating reports.

The IBT database job can be 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 **Average browse time per site**, set an average number of minutes for reading the contents of a web page.

This number defines browsing 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 average browse time per site setting 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 average browse time per site value.

Be aware that some websites 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 **Browse time for last site read** value to account for time spent reading the last website before the end of a browse session.

When the time gap of HTTP traffic is longer than the average “per site” browse time threshold, the session is ended and the “last site read” browse time value is added to the session time.

4. To enable detail reports that include browse time using investigative reports, mark **Calculate detailed browse time for use in investigative detail reports**.

If you enable detailed browse time calculations, be sure to create a new database partition. (Create a new partition any time you enable or disable detailed browse time calculations.)



Important

Enabling detailed browse time calculations increases Log Database size, and may also affect database performance. If you use this option, monitor Log Database growth and overall reporting performance carefully.

When detailed browse time is disabled, the IBT job still runs to perform the calculations used to include browse time in summary reports.

5. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring trend and application data retention

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Log Database administration settings, page 396](#)
- ◆ [Configuring database partition options, page 397](#)
- ◆ [Configuring how URLs are logged, page 402](#)
- ◆ [Configuring Log Database maintenance options, page 400](#)
- ◆ [Configuring Internet browse time options, page 403](#)

Optionally, the Log Database can store trend data to enable presentation reporting on Internet activity trends. When trend reporting is enabled, the ETL database job (see [Web protection reporting database jobs, page 395](#)) adds daily trend data to the catalog database, and the trend job runs nightly to store weekly, monthly, and yearly trend information.

The Log Database also stores statistical data (like bandwidth and count) for browsers, operating system platforms, and user agent strings to enable application reporting.

Configuring trend data

Use the **Trend Data Retention** section of the Settings > Reporting > Log Database page to specify how long trend data should be retained in the Log Database.

1. Mark **Store trend data** to prompt the ETL job to store trend data, and to activate the nightly trend job.

Trend data is calculated only for data collected while this option is enabled.

Data stored in the database before trend data retention is enabled, or data collected after the option is disabled, cannot be included in trend reports.

When this option is disabled, the Trend database job runs only to process threat-related data in the AMT partition.

2. Indicate how long to store weekly, monthly, and yearly trend data. Note that increasing the length of time trend data is stored increases the size of the Log Database (see [Log Database sizing guidance](#), page 406).



Note

Because trend data is stored in the catalog database, rather than the partition database, trend data storage periods are not dependent on how long database partitions are retained.

The default storage periods for trend data are:

	SQL Server	SQL Server Express
Daily	90 days	60 days
Weekly	26 weeks	13 weeks
Monthly	18 months	6 months
Yearly	5 years	3 years

The nightly trend job purges data when it is older than the specified retention period.

3. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring application data

Use the **Application Data** section of the Settings > Reporting > Log Database page to determine how long to keep the statistical data used to populate applications reports (on the Browser, Source Platform, and Search tabs of the Applications page).

The time period you select does not affect how long the actual user agent strings are stored. Those are preserved indefinitely. It only affects statistical information, like bandwidth, number of requests, and number of machines.

By default, statistical data for applications reports is stored for 30 days. To select another value:

1. Select a new time period from the **Keep data for** drop-down list. Depending on your database engine, data may be kept for up to 90 days.
2. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Log Database sizing guidance

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring database partition options, page 397](#)
- ◆ [Configuring how requests are logged, page 386](#)
- ◆ [Configuring how URLs are logged, page 402](#)

It is difficult to make precise sizing predictions for the Log Database, because database size is affected by a number of variables, including the number of users and average requests per second. In addition, size is affected by whether the database is configured to:

- ◆ Record hits or visits for each web request (see [Configuring Log Server, page 388](#)).
Recording hits provides a high level of detail, but recording visits can reduce the size of the database by roughly 40%.
- ◆ Consolidate log records (see [Configuring Log Server, page 388](#)).
By default, all requests are logged as separate hits or visits. When you turn on consolidation, similar requests (by the same user, for sites in the same domain, that have the same action applied) in a defined time period are recorded as a single log record. This can reduce the size of the Log Database by roughly 60%.
- ◆ Store the full URL for each logged request (see [Configuring how URLs are logged, page 402](#)).
Recording full URLs provides precise information about which sites a user has visited, but more than doubles Log Database size.
- ◆ Log requests for all categories (see [Configuring how requests are logged, page 386](#)).
By default, requests for sites in all categories are logged. To reduce the size of the Log Database, you can stop logging requests for sites in categories that, for example, present no security risk or legal liability to your organization.
The impact of this change depends on the number of categories that are not logged, and how often users requests sites in those categories.

- ◆ Perform detailed browse time calculations (see [Configuring Internet browse time options](#), page 403).

In order to create investigative detail reports that include browse time, the IBT job must calculate detailed browse time. Storing detailed browse time data, however, increases the size of the database, and may also affect database performance.

- ◆ Store trend data (see [Configuring trend and application data retention](#), page 404).
Storing trend data makes it possible to report on trends in users' Internet activity throughout the course of a day, week, or longer period, but storing trend data increases the size of the Log Database. The longer the data is stored, the greater its effect on database size.

Use the **Growth Rates and Sizing** chart on the Settings > Reporting > Log Database page to monitor the average daily size of your active and inactive standard logging partitions. This information may help you identify trends in traffic volume over time, and make it easier to plan for future growth.

As you collect average sizing information, adjust your rollover **Initial Size** and **Growth** settings (in the Partition Management section of the Settings > Reporting > Log Database page).

As a best practice, set the Initial Size value to approximately 80% of the **average partition size** over the rollover period (week, month, etc.). The idea is to:

- ◆ Minimize the number of times the partition must be expanded.
- ◆ Free resources to process data into the partitions.
- ◆ Prevent unneeded disk space from being allocated when the partition is created.
Unused portions of the initial space allocated to a partition cannot be recovered until the partition is deleted.

Configuring Dashboard reporting data

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Settings > Reporting > Dashboard** page to configure the maximum time period that can be shown in elements of the Threats, Risks, Usage, and System dashboards.

If you have TRITON AP-WEB, also configure whether to create a forensics repository for storing data about files associated with suspicious threat activity in your network.

Configuring the maximum time period for dashboard charts

By default, charts, counters, and tables on all tabs of the Status > Dashboard page show a maximum of **30 days** of data. This limit was chosen to minimize the amount of time it takes to load the Dashboard, optimize TRITON Manager overall performance, and reduce load on the Log Database.

With Standard and Enterprise versions of Microsoft SQL Server, you can configure dashboard charts to show a longer time period. Extending the maximum time period, however, may have serious performance impacts for both the TRITON Manager and the Log Database.

- ◆ To change the maximum time period that can be displayed in Risks, Usage, and System dashboard charts, under General Dashboard Data, select a value from the **Show a maximum of** drop-down list.
 - Increasing the time period does not affect the size of the Log Database, but does increase the time needed to query the database, retrieve information, and update dashboard charts.
 - If you are using Microsoft SQL Server Express, the maximum time period is 30 days, and cannot be changed.
- ◆ To change the maximum time period that can be displayed on the Threats dashboard and Event Details page, under Threats Data, select a value from the **Keep Threats data for** drop down list.
 - Because detailed Threats data is stored in a separate partition from standard logging data, increasing the time period also increases the size of the Log Database.
 - If threat-related forensic data storage is enabled (see below), the forensics repository attempts to store data for the time period selected here. If, however, the maximum repository size is reached, older records are automatically deleted to make room for newer records.
 - If you are using Microsoft SQL Server Express, the maximum time period is 30 days, and cannot be changed.

Note that data may not always be available for the full period selected. If your Websense web protection solution has only been installed for 7 days, for example, 30-day reports show data for only the 7 days that policy enforcement has occurred.

Threats dashboard sample data

If you would like to see examples of the types of data that can appear on the Threats dashboard without generating potentially dangerous network traffic, you can import sample data.

Because the sample data is loaded into the Log Database, where it is mixed with any real data generated in your network, it is best to load the sample data only in a test or evaluation environment.

To clearly flag the sample data, each of the users in the sample database is assigned the middle name **Demo** (for example, Sam Demo Smith and Lisa Demo Brady). In addition, the timestamp on the user activity predates the creation of the Log Database partition holding the data.

To load sample data into the database, click **Sample Data**, then click **Import Sample Data**. When you click **OK** and **Save and Deploy**, the data is loaded into the Log Database. After a few seconds, the Threats dashboard is updated to show the new data.

Configuring forensics data storage

In TRITON AP-WEB deployments, threat-related forensic data can include:

- ◆ Information about the source (IP address, device name, and user) attempting to send the data.
- ◆ Information about the target (IP address, URL, and geographic location) to which the data is being sent.
- ◆ Header information associated with the attempt to send the data.
- ◆ A copy of the actual data being sent (such as a text file, spreadsheet, ZIP file).

If you enable storage of forensics data, also specify where the **forensics repository** (a specialized database) is stored, the maximum size to which the database can grow, and how long to store forensics data.

1. Under Incident Data for Forensic Investigation, mark **Store forensic data about Threats incidents for further investigation** to create the forensics repository.
If your deployment includes TRITON AP-DATA, this new forensics repository is similar to that product's forensics repository. The smaller repository used by web protection components stores information about only those incidents displayed on the Threats dashboard.
2. Indicate whether to store forensics details for **Blocked requests only**, or for **All requests** (both blocked and permitted).
3. Specify the **Path** to the location that will host the forensics repository.
 - The specified directory must already exist.
 - The path can be either local (on the TRITON management server) or remote.
 - Make sure that there is enough free space in the selected location for the repository to grow to the maximum size that you specify (below).
4. Provide credentials for an account with read, write, and delete permissions for the forensics repository directory.
 - Select **Use Local System account** if neither network access nor special permissions are required to access the directory.
 - Select **Use this account** to use a domain account, then enter **User name**, **Password**, and **Domain** for the account.

Click **Test Connection** to verify that the selected account can access the forensics repository location.
5. To specify how large the forensics repository can grow, enter a **Maximum size** in GB (default 20) for the forensics repository.
 - If you are using SQL Server Express, this value cannot be changed.
 - When the maximum size is reached, **or** records reach the age limit specified for Threats data, records are automatically purged from the repository.
6. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Configuring investigative reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Database connection and report defaults](#), page 410
- ◆ [Display and output options](#), page 412

Investigative reports let you interactively delve into the information about your organization's Internet usage. See [Investigative reports](#), page 134.

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 410.

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 412.

Database connection and report defaults

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring investigative reports](#), page 410
- ◆ [Display and output options](#), page 412
- ◆ [Summary investigative reports](#), page 136
- ◆ [Multi-level summary investigative reports](#), page 141

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.
 - Mark **View the catalog database** to connect to the Log Database to which Log Server is currently logging. Proceed to step 2.

- To access a different Log Database, deselect **View the catalog database**, then enter the following information:

Field	Description
Server	Enter the machine name or IP address where the Log Database is located. If your environment uses SQL Server clustering, enter the virtual IP address for the cluster
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.
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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring investigative reports, page 410](#)
- ◆ [Database connection and report defaults, page 410](#)
- ◆ [Output options for investigative reports, page 155](#)

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 (located by default in the C:\Program Files (x86)\Websense\Web Security\webroot\Explorer\ directory).

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	<p>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.</p> <p>Note: There must be 2 or more groups for Group to appear in the Internet Use by list.</p> <p>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.</p>
maxUsersDrilldown	<p>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.</p> <p>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.</p> <p>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.</p>

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	<p>This option is disabled (0) by default. Set it to 1 (enabled) to divide very large, scheduled detail reports into multiple files. The number of records in each file is determined by the value set by recordBatchCount.</p> <p>When this option is disabled, each scheduled detail report is divided into files of up to 500,000 records.</p> <p>All reports are sent as a zip file and can be extracted with most common file compression utilities.</p>
maxSlices	This is the maximum number of distinct slices (6 by default, maximum is 10) 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).

Parameter	Description
reportBatchCount	When sendMulti is enabled, this option is used to determine the number of records (by default, 10,000) included in a scheduled report. Any value over 500,000 is ignored and 500,000 is used.
PageSize	Investigative report results can be output to Portable Document Format (PDF) for easy distribution or printing. The page size (by default, Letter) can be: <ul style="list-style-type: none"> • A4 (8.27 X 11.69 inches) • Letter (8.5 X 11 inches)

Self-reporting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring reporting preferences, page 385](#)
- ◆ [Accessing self-reporting, page 156](#)
- ◆ [Investigative reports, page 134](#)

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 investigative reports with their network credentials. This may have been done previously to enable policy application by user and group names. See [Connecting web protection software to a directory service, page 59](#).
2. Go to the **Settings > Reporting > Preferences**, and mark the **Allow self-reporting** check box. See [Configuring reporting preferences, page 385](#).

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:

```
https://<IP_address>:9443/mng/login/pages/
selfReportingLogin.jsf
```

Replace <IP_address> with the IP address of the TRITON management server.

Remind users that they can save the URL as a favorite or bookmark for future use.

- ◆ What user name and password to use during logon.

Self-reporting users must enter their network user name and password during logon.

18

Configure Network Agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Network Agent global settings, page 418](#)
- ◆ [Configuring Network Agent local settings, page 419](#)

Network Agent serves different purposes in TRITON AP-WEB and Web Filter & Security deployments.

- ◆ In TRITON AP-WEB deployments, and when Web Filter & Security is integrated with a third-party proxy or firewall product, Network Agent is an **optional** component that may be used to:
 - Manage non-HTTP requests
 - Provide enhanced logging
 - Manage Internet access based on bandwidth

In TRITON AP-WEB deployments, Content Gateway can offer some of the same features as Network Agent, by detecting protocols that tunnel over HTTP (see [Configuring tunneled protocol detection, page 92](#)) and offering some bandwidth management capabilities ([Using Bandwidth Optimizer to manage bandwidth, page 261](#)).

- ◆ When Web Filter & Security is installed as a standalone product, Network Agent is a **required** component that enables:
 - Policy enforcement
 - Network protocol and Internet application management
 - Bandwidth management
 - Logging of bytes transferred

Network Agent works by monitoring overall network usage, including bytes transferred. It logs usage summaries—which include start time and end time, overall bytes used, and bytes used per protocol—at predefined intervals.

When used, Network Agent is typically configured to see all traffic in your network, and can distinguish between requests sent from internal machines to internal machines (hits to an intranet server, for example) and requests sent from internal machines to external machines such as web servers (user Internet requests, for example).

Configuring Network Agent global settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Network Agent local settings, page 419](#)
- ◆ [Configuring Network Agent NIC settings, page 421](#)
- ◆ [Adding or editing IP addresses during Network Agent configuration, page 423](#)

Use the **Settings > Network Agent > Global** page to define basic monitoring and logging behavior for all instances of Network Agent connected to the current Policy Server (the Policy Server whose IP address appears in the Web module toolbar).

The **Describe Your Network** list identifies the IP addresses that are part of your network in IPv4 or IPv6 format. 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 Network Agent NIC settings, page 421](#)). 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 internal IPv4 or IPv6 addresses (encompassed by the “Describe Your Network” list) 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 IPv4 or IPv6 address or range to the appropriate list. See [Adding or editing IP addresses during Network Agent configuration, page 423](#), for more information.
- ◆ To edit an entry in the list, click the IP address or range. See [Adding or editing IP addresses during Network Agent configuration, page 423](#), 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 and Deploy**.

Configuring Network Agent local settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Network Agent global settings, page 418](#)
- ◆ [Configuring Network Agent NIC settings, page 421](#)

Use the **Settings > Network Agent > Local Settings** page to configure Internet traffic management, proxy information, and other settings for the selected instance of Network Agent.

- ◆ To get to the Local Settings page for a Network Agent instance, navigate to the **Settings > Network Agent** menu and position the mouse over the **Global** option. A list of IP addresses appears. Select the IP address of the instance you want to configure.
- ◆ 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.

To ensure that user requests are monitored, managed, 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 IPv4 or IPv6 address or range to the list (see [Adding or editing IP addresses during Network Agent configuration](#), page 423).
- ◆ 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 Network Agent NIC settings](#), page 421, 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** (with TRITON AP-WEB, or when Web Filter & Security is integrated with a third-party product) 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.
In TRITON AP-WEB deployments, 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 and Deploy**.

Configuring Network Agent NIC settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Network Agent local settings, page 419](#)
- ◆ [Configuring Network Agent monitoring settings for a NIC, page 422](#)
- ◆ [Adding or editing IP addresses during Network Agent configuration, page 423](#)

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 Network Agent monitoring settings for a NIC](#), page 422, 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 Web Filter & Security in **Stand-Alone** mode, **Filter and log HTTP requests** is selected, and cannot be changed.
3. If you have TRITON AP-WEB, or if Web Filter & Security is integrated with a third-party product, use the **Integrations** options to indicate how this Network Agent should treat HTTP requests. Options that do not apply to your environment are disabled.
 - Select **Log HTTP requests** to make Network Agent responsible for generating all HTTP-related Internet activity log records.
When this option is selected, only the log records that Network Agent creates are forwarded to Log Server for inclusion in the Log Database.
 - Select **Filter all requests not sent over HTTP ports** to use Network Agent to manage only those requests not sent through the integration product.
4. Under Protocol Management, indicate whether Network Agent should use this NIC to manage non-HTTP protocols:
 - Check **Filter non-HTTP protocol requests** to activate Network Agent protocol management. See [Managing access to categories and protocols](#), page 32, and [Protocol-based policy enforcement](#), page 255, 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 261, for more information.

Configuring Network Agent monitoring settings for a NIC

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the **Local Settings > NIC Configuration > Monitor List** page to specify which IP addresses 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 IP addresses 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 that Network Agent should monitor in IPv4 or IPv6 format. See [Adding or editing IP addresses during Network Agent configuration, page 423](#), for more information.

**Note**

You cannot enter overlapping IP address ranges. If ranges overlap, network bandwidth measurements may not be accurate, and bandwidth-based restrictions may be applied incorrectly.

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.
 - a. To identify a machine, click **Add**, and then enter its IP address in IPv4 or IPv6 format.
 - 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 and Deploy**.

Adding or editing IP addresses during Network Agent configuration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Related topics:

- ◆ [Configuring Network Agent global settings, page 418](#)
- ◆ [Configuring Network Agent local settings, page 419](#)
- ◆ [Configuring Network Agent NIC settings, page 421](#)

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, Proxies and Caches, Monitor List, or Monitor List Exceptions.

- ◆ Both IPv4 and IPv6 addresses and ranges are supported.
- ◆ 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 and Deploy**.

19

Troubleshooting

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use this section to find solutions to common issues before contacting Technical Support.

The Websense Support website features an extensive Knowledge Base and customer forums, available at support.websense.com. Search for topics by keyword or phrase, or browse content by product and version.

Troubleshooting instructions are grouped into the following sections:

- ◆ *[Web protection installation and subscription issues](#)*
- ◆ *[Web protection Master Database issues](#), page 426*
- ◆ *[Filtering Service alert messages](#), page 430*
- ◆ *[Network Agent issues](#), page 435*
- ◆ *[User configuration and identification issues](#), page 438*
- ◆ *[Health alerts and Usage Monitor issues](#), page 447*
- ◆ *[Policy Server and Policy Broker issues](#), page 451*
- ◆ *[Log Server and Log Database issues](#), page 453*
- ◆ *[Investigative report and presentation report issues](#), page 465*
- ◆ *[Other reporting issues for web protection solutions](#), page 470*
- ◆ *[TRITON AP-WEB interoperability issues](#), page 475*

Web protection installation and subscription issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

There is a subscription problem

A valid subscription key is needed to download the Websense Master Database and perform Internet policy enforcement. When your subscription expires or is invalid, and when the Master Database has not been downloaded for more than 2 weeks, a warning appears on the **Status > Alerts** page.

- ◆ Verify that you have entered your subscription key exactly as you received it on the **Web > Settings > General > Account** page in the TRITON Manager. The key is case sensitive.
- ◆ Make sure that your subscription has not expired on the **Web > Settings > General > Account** page in the TRITON Manager.
- ◆ Ensure that the Master Database has been downloaded successfully within the last 2 weeks. You can check download status in the Web module of the TRITON Manager: click **Database Download** on the **Status > Dashboard** page.
See [The Master Database does not download](#), page 427, 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.

Unable to verify the subscription key

After you enter your subscription key, Filtering Service attempts to connect to the Websense database download servers to both verify the key and download the Master Database.

If Filtering Service is unable to connect to a download server, both subscription errors and database download errors appear in the Web module of the TRITON Manager.

- ◆ 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 [Verify Filtering Service Internet access](#), page 428, and [Verify Filtering Service firewall or proxy server settings](#), page 429, to make sure that Filtering Service and the network environment are properly configured to enable the connection.

Web protection Master Database issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to help troubleshoot Master Database problems:

- ◆ [The initial database is being used](#), page 426
- ◆ [The Master Database is more than 1 week old](#), page 427
- ◆ [The Master Database does not download](#), page 427
- ◆ [Contacting Technical Support for database download issues](#), page 430

The initial database is being used

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Websense Master Database houses the category and protocol definitions that provide the basis for managing 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 functionality from the time you enter your subscription key.

You must download the full database for full policy enforcement to occur. See [The Websense Master Database](#), page 11, 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Websense Master Database houses the category and protocol definitions that provide the basis for managing Internet content. Websense software downloads changes to the Master Database according to the schedule defined on the **Settings > General > Database Download** page. By default, download is scheduled to occur once a day.

To manually initiate a database download:

1. Go to the **Status > Dashboard** 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 categories, protocols, and filters while a database is being updated.

The Master Database does not download

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you are unable to download the Websense Master Database successfully:

- ◆ Go to the **Web > Settings > General > Account** page in the TRITON Manager and make sure that:
 1. You have entered your subscription key correctly.
 2. The key has not expired. If the expiration date has passed, contact your reseller or Websense, Inc., to renew your subscription.

If you make changes to the key, click **OK** and **Save and Deploy** to activate the changes and start a database download.
- ◆ Verify that the Filtering Service machine is able to access the Internet (*Verify Filtering Service Internet access*, page 428).
- ◆ Check firewall or proxy server settings to make sure that Filtering Service can connect to the Websense download servers (*Verify Filtering Service firewall or proxy server settings*, page 429).
- ◆ Make sure that there is enough disk space (*Insufficient disk space on the Filtering Service machine*, page 432) and memory (*Insufficient memory on the Filtering Service machine*, page 433) on the download machine.
- ◆ Look for any application or appliance in the network, such as anti-virus software, size-limiting applications, or intrusion detection systems that might prevent the download connection. Ideally, configure Filtering Service to go straight to the last gateway so that it does not connect to these applications or appliances.

Verify Filtering Service Internet access

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 make sure the Filtering Service machine has the Internet access necessary to communicate with the download server, you can:

- ◆ For non-appliance installations, open a browser on the Filtering Service machine and 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.
- ◆ From the command prompt or shell, 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.

If the Filtering Service machine cannot connect to the download server:

- ◆ Enable communication on port 80, or the port designated in your network for HTTP traffic, for the network interface used by Filtering Service. On Websense V-Series appliances, this is usually the C interface.
- ◆ Verify that the Filtering Service network interface is using the correct DNS settings.
- ◆ Make sure that Filtering Service is configured to use any necessary proxy servers to connect to the Internet (see [Verify Filtering Service firewall or proxy server settings](#), page 429)

Also make sure that your gateway or firewall does not include any rules that block HTTP traffic from the Filtering Service machine.

Verify Filtering Service firewall or proxy server settings

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If the Master Database is downloaded through a firewall or proxy that requires authentication, use the following steps to check your proxy authentication settings:

1. Go to the **Web > Settings > General > Database Download** page in the TRITON Manager.
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 Filtering Service 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 support.websense.com.

If a firewall restricts Internet access at the time Filtering Service normally downloads the database, or restricts the size of a file that can be transferred via HTTP, database updates cannot occur. 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 time ([Configuring database downloads](#), page 12), if necessary.

If Filtering Service is not running on a Websense appliance, you can check your Filtering Service proxy settings against browser proxy settings on the machine. First verify 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 browser.

- ◆ Microsoft Internet Explorer:
 1. Display the **Menu** bar, then navigate to **Tools > Internet Options** and select the **Connections** tab.
 2. Click **LAN Settings**, then make a note of the settings that appear under **Proxy server**.
- ◆ Mozilla Firefox:

1. Navigate to **Tools > Options >**, then select the **Advanced** tab.
2. On the **Network** tab (usually selected by default), 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.

Contacting Technical Support for database download issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 Filtering Service machine 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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default): **websense.ini**, **eimserver.ini**, and **config.xml**.

Go to support.websense.com for Technical Support contact information.

Filtering Service alert messages

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to respond to health alerts from Filtering Service:

- ◆ [Filtering Service is not running, page 431](#)
- ◆ [High CPU usage on the Filtering Service machine, page 431](#)
- ◆ [Insufficient disk space on the Filtering Service machine, page 432](#)
- ◆ [Insufficient memory on the Filtering Service machine, page 433](#)

- ◆ [Filtering Service can't communicate with a transparent identification agent](#), page 434

Filtering Service is not running

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When Filtering Service is not running, policy enforcement and logging cannot occur.

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 432).
- ◆ A Master Database download failed due to lack of disk space (see [The Master Database does not download](#), page 427).
- ◆ 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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default), 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 (in the Websense **bin** directory) for error messages from Filtering Service.
- ◆ Log off of the TRITON console, restart Policy Server, and then restart Filtering Service (see [Stopping and starting Websense services](#), page 370).

Wait 1 minute before logging on to the TRITON console again.

High CPU usage on the Filtering Service machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 policy enforcement.

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.

Insufficient disk space on the Filtering Service machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Filtering Service needs adequate space to download compressed Master Database updates to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default). 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 the drive and select **Properties**.
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
```

If Filtering Service is not installed in the default directory, use the custom path instead.

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 370).
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 **Web > Status > Dashboard** page in the TRITON Manager, and then click **Database Download**).

Insufficient memory on the Filtering Service machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 4 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 hardware requirements, 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 use the Windows Performance monitor 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.

Filtering Service can't communicate with a transparent identification agent

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When you use DC Agent, Logon Agent, eDirectory Agent, or RADIUS Agent for transparent user identification, Filtering Service must be able to communicate with the agent to correctly apply user-based policies. If this communication fails, the user may be filtered by an IP-address-based policy or the Default policy.

To address this problem:

1. Verify that the agent service or daemon is running.
 - **Windows:** Use the Windows Services tool to make sure that Websense DC Agent, Websense Logon Agent, Websense eDirectory Agent, or Websense RADIUS Agent is running.
 - **Linux:** Navigate to the `/opt/Websense/` directory and use the following command to verify that Logon Agent, eDirectory Agent, or RADIUS Agent is running:

```
./WebsenseAdmin -status
```

2. You can **ping** the transparent identification agent machine from the Filtering Service machine. Try both the IP address and the hostname of the transparent identification agent machine, to make sure that DNS is properly configured. For example:

```
ping 10.55.127.22
ping transid-host
```

3. The transparent identification agent communication port is open between the Filtering Service machine and the agent machine. The default ports are:
 - DC Agent: 30600
 - Logon Agent: 30602
 - eDirectory Agent: 30700
 - RADIUS Agent: 30800
4. The correct agent IP address or hostname and port appear on the **Web > Settings > General > User Identification** page in the TRITON Manager.

If the service appears to be running normally, and there does not appear to be a network communication problem between the Filtering Service and agent machines:

- ◆ Use the Windows Services tool or the `/opt/Websense/WebsenseDaemonControl` command to restart the agent.

Check the Windows Event Viewer or **websense.log** file (in the Websense **bin** directory) on the agent machine for error messages from the transparent identification agent.

Network Agent issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to respond to health alerts about Network Agent.

- ◆ [Network Agent is not installed, page 435](#)
- ◆ [Network Agent is not running, page 435](#)
- ◆ [Network Agent is not monitoring any NICs, page 436](#)
- ◆ [Network Agent can't communicate with Filtering Service, page 436](#)
- ◆ [Insufficient memory on the Network Agent machine, page 437](#)
- ◆ [High CPU usage on the Network Agent machine, page 438](#)

Network Agent is not installed

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

With Web Filter & Security, Network Agent is required to enable policy enforcement for Internet protocols other than HTTP, HTTPS, and FTP. With some integrations, Network Agent is also used to provide more accurate logging.

If you are using TRITON AP-WEB, or if you have integrated Web Filter & Security with a third-party product, and do not require Network Agent protocol management or logging, you can hide the “No Network Agent is installed” status message. See [Reviewing current system status, page 381](#), for instructions.

For standalone Web Filter & Security installations, Network Agent must be installed for network monitoring and policy enforcement to occur. See the [Network Agent Quick Start](#) for installation instructions, and then see [Configuring Network Agent global settings, page 418](#).

Network Agent is not running

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

For TRITON AP-WEB and integrated Web Filter & Security deployments, Network Agent offers enable full management of protocols other than HTTP, HTTPS, and FTP, plus full bandwidth-based policy enforcement.

For standalone Web Filter & Security installations, Network Agent must be running to monitor and manage network traffic.

To troubleshoot this problem:

1. Look for low memory problems on the Network Agent machine that may be preventing the service or daemon from starting.
2. Check the status of the Network Agent service or daemon:
 - **Windows:** Use the Windows Services tool to see if the **Websense Network Agent** service has started.

- *Linux*: Use the **/opt/Websense/WebsenseDaemonControl** command to check the status of **Network Agent**.
 - *V-Series Appliance*: Use the Appliance manager to check the status of the **Network Agent** module.
3. Make sure all administrators are logged off of the TRITON console, then restart the **Websense Policy Broker** and **Websense Policy Server** services (see [Stopping and starting Websense services](#), page 370).
 4. Start or restart the **Websense Network Agent** service.
 5. Wait 1 minute, and then log on to the TRITON Manager again.

If that does not fix the problem:

- ◆ Check the **Windows Event Viewer** for error messages from Network Agent.
- ◆ Check the **Websense.log** file (in the Websense **bin** directory) for error messages from Network Agent.

Network Agent is not monitoring any NICs

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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. Open the Web module of the TRITON Manager, and select the **Settings** tab
2. In the left navigation pane, under Network Agent, place the mouse over the **Global** menu entry until one or more IP addresses is displayed in a submenu.
3. Select the IP address of the Network Agent machine.
4. Verify that all NICs for the selected machine are listed.
5. Verify that at least one NIC is set to monitor network traffic.

See [Configuring Network Agent local settings](#), page 419, for more information.

Network Agent can't communicate with Filtering Service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

With TRITON AP-WEB, Network Agent must be able to communicate with Filtering Service to record protocol management, bandwidth, and other logging information. In standalone Web Filter & Security environments, no policy enforcement can occur when Network Agent and Filtering Service cannot communicate.

- ◆ Did you change the IP address of Filtering Service machine or reinstall Filtering Service?

If so, the new Filtering Service internal unique identifier (UID) or IP address is not automatically registered with Network Agent and the TRITON Manager.

To re-establish connection to the Filtering Service:

1. Log on to the Web module of the TRITON Manager.
A status alert 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, position the mouse over the **Global** menu item until a sub-menu is displayed, then 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 and Deploy**.
- ◆ Do you have more than 2 network interface cards (NICs) on the Network Agent machine?
If so, see [Configure Network Agent, page 417](#), to verify your Websense software settings.
 - ◆ Have you reconfigured the switch connected to the Network Agent?
If so, refer to the [Network Agent Quick Start](#) to verify your hardware setup.

If none of these apply, see [Configuring Network Agent local settings, page 419](#), for information about associating Network Agent and Filtering Service.

Insufficient memory on the Network Agent machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, policy enforcement 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 standalone environment, this can mean that all user requests for websites 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

- ◆ [User and group-based policies are not applied, page 438](#)
- ◆ [Unusually high directory server connection latency, page 439](#)
- ◆ [User Service is not available, page 440](#)
- ◆ [DC Agent has insufficient permissions, page 440](#)
- ◆ [DC Agent unable to access required file, page 441](#)
- ◆ [I cannot add users and groups to the Web module of the TRITON Manager, page 443](#)
- ◆ [User Service on a Websense appliance or Linux server, page 446](#)

User and group-based policies are not applied

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If Filtering Service is applying computer or network policies, or the **Default** policy, to 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 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 the Web module of the TRITON Manager, 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 446.

- ◆ If User Service resides on a Websense appliance or Linux server, and you are using DC Agent (with Active Directory in any mode) or Logon Agent (with Active Directory in native mode) to identify users, verify your WINS server configuration. See [User Service on a Websense appliance or Linux server](#), page 446.

If none of these steps addresses your issue, check the following topics, or search support.websense.com for additional information.

- ◆ [Directory service connectivity and configuration](#), page 443
- ◆ [Directory service configuration](#), page 444
- ◆ [User identification and Windows Server](#), page 444

Unusually high directory server connection latency

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Websense User Service communicates with user directories in your network to:

- ◆ Populate the Clients page and other TRITON Manager pages with user, group, and OU information.
- ◆ Find group information for users so that Websense Filtering Service can enforce the correct policy.
- ◆ Provide user and group information to other Websense components to ensure consistency in policy enforcement, reporting, and alerting.
- ◆ Offer manual authentication via browser-based logon prompts.

When User Service experiences unusually high connection latency to the directories that it queries, users may:

- ◆ Experience slow browsing
- ◆ Receive an IP address-based policy or the Default policy instead of the appropriate user or group policy

Administrators may experience delays when trying to work with clients in the TRITON Manager.

To address this issue, look for:

- ◆ Network problems between the specified User Service machine and each of the directory server machines noted in the health alert message
- ◆ Problems on the domain controller that might slow down directory connections or searches

User Service is not available

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

User Service must be running, and Policy Server must be able to communicate with User Service, in order for user-based policies to be applied correctly.

User Service may appear to have stopped if you restarted Policy Server after restarting other Websense services. To correct this issue:

1. Log off of the TRITON Manager.
2. Restart the Websense Policy Server service (see [Stopping and starting Websense services](#), page 370).
3. Start or restart Websense User Service.
4. Wait 1 minute before logging on to the TRITON Manager again.

If the previous steps do not fix the problem:

- ◆ Check the Windows Event Viewer or **websense.log** file (in the Websense **bin** directory) for error messages from User Service.
- ◆ Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default), 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.

DC Agent has insufficient permissions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

DC Agent may have been installed as a service using the Guest account, equivalent to an anonymous user to the domain controller.

In order to perform domain discovery (required for DC Agent to create and maintain the dc_config.txt file), or to perform computer polling, the Websense DC Agent service requires **domain admin** permissions. In some environments (typically very large enterprise networks), DC Agent requires **enterprise admin** permissions.

If you have disabled domain discovery and computer polling, and are just using domain controller polling while maintaining the dc_config.txt file manually, DC Agent can run as any network user with read access to the domain controller.

To grant DC Agent domain admin privileges:

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 tool:
 - *Windows Server 2012*: Server Manager > Tools > Services
 - *Windows Server 2008*: Start > Administrative Tools > Services
3. Scroll to the **Websense DC Agent** service, right-click the service name, and then select Stop.
4. Right-click the service name again, select **Properties**, and then click 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 tool.
7. Right-click the service name again, and then select **Start**.
8. Close the Services tool.

You may also need to assign User Service the same administrative privileges as DC Agent.

DC Agent unable to access required file

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

DC Agent works by identifying domain controllers in the network, then querying them 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 the domain 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 been corrupted.
- ◆ DC Agent is unable to create the file, because it cannot find any domain controller information.
 - If User Service is installed on a Websense appliance or Linux server, make sure that you have performed required WINS setup steps. For complete instructions, see [User Service on a Websense appliance or Linux server](#), page 446.
 - 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 446.
 - 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 445.
- ◆ 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
```

DC Agent Domains and Controllers page is blank

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

By default, DC Agent performs automatic **domain discovery**, identifying domain controllers in the network. Domain and controller information is stored in a file called **dc_config.txt**. The information from the dc_config.txt file is collected and displayed in the Web module of the TRITON Manager on the **Settings > User Identification > DC Agent Domains and Controllers** page.

This page may display only error text if:

- ◆ DC Agent was recently installed, and domain discovery is still underway.
- ◆ An administrator has modified the **dc_config.txt** file to turn off polling for all domain controllers in the network.
- ◆ Something is preventing DC Agent from performing domain discovery.

Make sure that:

- ◆ DC Agent domain discovery is enabled on the **Settings > User Identification > DC Agent** page for each DC Agent instance in your network.
- ◆ DC Agent has had enough time to complete its domain discovery process.
- ◆ No DC Agent alerts appear on **Status > Alerts** page.

If a DC Agent alert appears, see [DC Agent has insufficient permissions](#), page 440, and [DC Agent unable to access required file](#), page 441. These articles provide instructions for ensuring that DC Agent has the required permissions and network access to complete the domain discovery process and create the dc_config.txt file.

I cannot add users and groups to the Web module of the TRITON Manager

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

A number of problems can prevent the list of users and groups from appearing when you attempt to add clients in the TRITON Manager. Check the following topics, and check the [Knowledge Base](#) for additional information.

- ◆ [Directory service connectivity and configuration, page 443](#)
- ◆ [Directory service configuration, page 444](#)
- ◆ [User identification and Windows Server, page 444](#)

Directory service connectivity and configuration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 machine, the account used to run User Service may require administrative privileges on the directory.
To check or change the User Service account, see [Changing DC Agent, Logon Agent, and User Service permissions, page 446](#).
- ◆ 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 **Web > Settings > General > Directory Services > Add Global Catalog Server** page in the TRITON Manager.
- ◆ If you are running User Service on a Linux machine and communicating with a Windows-based directory service, make sure that you have set up a WINS server and performed all necessary configuration steps (see [User Service on a Websense appliance or Linux server, page 446](#)).
- ◆ Determine whether a firewall is blocking communication between the TRITON Manager and User Service on port 55815. If so, open the blocked port.

Directory service configuration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you encounter problems adding users and groups in the Web module of the TRITON Manager, make sure that you have provided complete and accurate configuration for your directory service.

1. Go to the **Settings > General > Directory Services** page.
2. Select the directory service used by your organization.
3. Verify the configuration. See [Connecting web protection software to a directory service, page 59](#), 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 a Websense appliance or Linux server, page 446](#), for additional configuration requirements.

User identification and Windows Server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

You may encounter problems adding users and groups in the Web module of the TRITON Manager if you install one or more of the following components on a supported Windows Server version:

- ◆ 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. This service was turned on by default in earlier versions of Windows, but it is disabled by default on Windows Server 2008 and 2012.

In addition, when User Service is installed on Windows Server, 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, additional configuration is required. See [User Service on a Websense appliance or Linux server, page 446](#).

To enable the Computer Browser service on a relevant machine, see [Turning on the Computer Browser service, page 445](#).

To configure User Service with rights to access directory information, see [Changing DC Agent, Logon Agent, and User Service permissions, page 446](#).

Turning on the Computer Browser service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Websense installer offers the option to turn on the Computer Browser service during installation of the following components on Windows Server.

- ◆ 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.
 - Windows Server 2012:
 - a. On the desktop, point the mouse to the top, right corner of the screen, then go to **Settings > Control Panel**.
 - b. In the Control Panel, click **Network and Internet**, then **Network and Sharing Center**.
 - c. Click **Change advanced sharing settings** in the left navigation pane, then select **Turn on file and printer sharing**.
 - d. Click **Save Changes** to save and exit.
 - Windows Server 2008 R2:
 - 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. Open the Windows Services tool:
 - Windows Server 2012: **Server Manager > Tools > Services**
 - Windows Server 2008 R2: **Start > 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 tool.
8. Repeat these steps on each Windows Server machine that hosts an affected component.

Changing DC Agent, Logon Agent, and User Service permissions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, open the Windows Services tool.
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 tool, 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 a Websense appliance or Linux server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you plan to apply policies to individual users and groups in your network, and User Service runs on a Websense appliance or Linux server, special configuration steps are required when you:

- ◆ Use Active Directory in mixed mode
- ◆ Want to use Websense Logon Agent to transparently identify users via Active Directory in native mode
- ◆ Want to 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 the TRITON Manager, go to the **Web > 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 and Deploy**.

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 and Deploy** 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 [Connecting to Windows Active Directory \(Native Mode\)](#), page 61, for assistance.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

Health alerts and Usage Monitor issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to learn more about health alerting, and to find Usage Monitor troubleshooting information.

- ◆ [Where do I find error messages for Websense components?](#), page 448
- ◆ [Websense Health alerts](#), page 448

- ◆ [Usage Monitor is not available](#), page 450
- ◆ [Usage Monitor is not running](#), page 450

Where do I find error messages for Websense components?

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When there are errors or warnings related to core Websense components, alert messages are listed on the **Status > Alerts** page in the Web module of the TRITON Manager. In addition, by default, short alert messages are displayed in the **Health Alert Summary** list at the top of the System tab of the **Status > Dashboard** page (see [Websense Health alerts](#), page 448).

- ◆ Click an alert summary in the dashboard to see more detailed information on the **Status > Alerts** page.
- ◆ Click **Solutions** next to the detailed health alert message 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).

For Websense software components installed on Windows machines, you can also check the Windows Event viewer.

Websense Health alerts

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

By default, the System tab of the **Status > Dashboard** page includes a **Health Alert Summary** that lists potential concerns encountered by monitored components of your Websense software. These include:

- | | |
|--|---|
| ◆ The initial database is in use | ◆ The Master Database is being updated |
| ◆ The Master Database is downloading for the first time | ◆ The Master Database did not download successfully |
| ◆ The Master Database is more than 1 week old | ◆ Low disk space on the TRITON management server machine |
| ◆ WebCatcher is not available | ◆ The primary Policy Broker is now available |
| ◆ 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 |
| ◆ The Log Database ETL job has not completed successfully after 4 hours | ◆ There is no Log Server configured for a Policy Server |
| ◆ Low disk space on the Log Server machine | ◆ A Log Server cache directory contains more than 100 cache files |
| ◆ Log Server has not received data from Filtering Service for over an hour | ◆ No Filtering Service has been configured for a Network Agent |

- No monitoring NIC has been configured for a Network Agent
- Low memory on the Network Agent machine
- Filtering Service is not running
- Low disk space on the Filtering Service machine
- High CPU usage on the Filtering Service machine
- DC Agent has insufficient permissions
- Filtering Service is unable to communicate with Logon Agent
- Filtering Service is unable to communicate with eDirectory Agent
- Usage Monitor is not running
- The forensics repository location could not be reached
- A configuration problem is interfering with Threats forensics data collection
- High CPU usage on the Network Agent machine
- There is no Network Agent configured for a Policy Server
- Network Agent is not running
- Low memory on the Filtering Service machine
- A DC Agent instance is unable to access a required file
- Filtering Service is unable to communicate with DC Agent
- Filtering Service is unable to communicate with RADIUS Agent
- A Policy Broker replica has not synchronized with the primary in more than 24 hours
- Usage Monitor is not available
- The forensics repository has reached 90% of its maximum size
- Some records in the forensics repository are scheduled to be deleted within 1 week

TRITON AP-WEB, alerts are also provided for the following Content Gateway issues:

- Content Gateway is not running
- Content Gateway is not available

If you have the Web Hybrid module, health alerts may also appear for 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 the hybrid service.
- 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 policy enforcement or reporting (for example, the Master Database is more than 1 week old, or the subscription key is about to expire).



A component is not functioning (has not been configured or is not running), which may impair policy enforcement or reporting, or your subscription has expired.

Click an alert message in the Health Alert Summary to go to the **Status > 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 Service Alerts table 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 381](#), for more information.

Usage Monitor is not available

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 tool, select **Websense Usage Monitor**, right-click the service, and select **Start**.
- ◆ Linux: Use the `/opt/Websense/WebsenseDaemonControl` command.

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 Broker issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to help troubleshoot problems with Policy Server and Policy Broker:

- ◆ [I forgot my password, page 451](#)
- ◆ [The Websense Policy Database service fails to start, page 452](#)
- ◆ [Policy Server stops unexpectedly, page 452](#)
- ◆ [A Policy Broker replica cannot synchronize data, page 453](#)

I forgot my password

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you are a Super Administrator or delegated administrator using a local account to log on to the TRITON Manager, 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 login page.

- ◆ A temporary password is 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 are prompted to enter a new password after you have logged on using the temporary password.

The Websense Policy Database service fails to start

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.

To do this on a Windows:

1. Log on to the Policy Database machine as a local administrator.
2. Open the **Computer Management** tool:
 - Windows Server 2012: Use the **Server Manager > Tools** menu.
 - Windows Server 2008: Use the **Start > Administrative Tools** menu.
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. Open the Windows **Services** tool:
 - Windows Server 2012: Use the **Server Manager > Tools** menu.
 - Windows Server 2008: Use the **Start > Administrative Tools** menu.
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 tool.

Policy Server stops unexpectedly

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 warns that disk space is getting low.
- ◆ When Policy Server stops, a health alert message is displayed.

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.

A Policy Broker replica cannot synchronize data

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

In a replicated Policy Broker deployment, each replica synchronizes its policy and configuration data with the primary Policy Broker on a regular basis to ensure that current information is available to all components in the deployment.

If a replica is unable to connect to the primary Policy Broker for more than 24 hours, a health alert is displayed. To resolve this issue:

- ◆ Make sure that 2-way network communication is possible between the primary and replica host machines on port **6432**. (The firewall must allow both inbound and outbound connections on this port.)
- ◆ Make sure that the primary Policy Broker machine is up, and the Policy Broker service or daemon is running.
- ◆ Make sure that the Policy Broker replica machine is up, and the Policy Broker service or daemon is running.
- ◆ The replica Policy Broker must use the synchronization password configured during primary Policy Broker configuration. If you have recently replaced the primary Policy Broker, make sure that the correct synchronization password was used.

Log Server and Log Database issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to troubleshoot reporting issues involving Log Server or the Log Database:

- ◆ [*Log Server is not running*, page 454](#)
- ◆ [*Log Server has not received log files from Filtering Service*, page 454](#)
- ◆ [*Low disk space on the Log Server machine*, page 456](#)
- ◆ [*No Log Server is installed for a Policy Server*, page 457](#)
- ◆ [*Log Database was not created*, page 459](#)
- ◆ [*Log Database is not available*, page 460](#)
- ◆ [*Log Database size causes reporting delays*, page 461](#)
- ◆ [*More than 100 files in the Log Server cache directory*, page 461](#)
- ◆ [*Last successful ETL job ran more than 4 hours ago*, page 463](#)
- ◆ [*Log Server is not recording data in the Log Database*, page 464](#)
- ◆ [*Log Server cannot connect to the directory service*, page 464](#)
- ◆ [*Wrong reporting page displayed*, page 465](#)

Log Server is not running

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If Log Server is not running, or if other Websense components are unable to communicate with Log Server, Internet usage information is not stored. Charts on the **Status > Dashboard** page stop updating, and you may not be able to generate reports.

Log Server may be unavailable if:

- ◆ It is unable to contact the Log Database after 20 attempts.
Make sure that the Log Database machine is running, that Microsoft SQL Server is operating properly, and that network communication has not been interrupted between the Log Server and Log Database machines.
- ◆ There is insufficient disk space on the Log Server machine.
Verify the amount of free disk space on the Log Server machine, and remove extraneous files, as needed.
- ◆ You changed the Microsoft SQL Server password without updating the ODBC or Log Server connection.
- ◆ It has been more than 14 days since the Master Database was downloaded successfully.
See [The Master Database is more than 1 week old, page 427](#), and [The Master Database does not download, page 427](#), for information about addressing this issue.
- ◆ The **logserver.ini** file is missing or corrupted.
Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin, by default) 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.
- ◆ You stopped Log Server to avoid logging Internet usage information.
Check the Windows Services tool to verify that Log Server has started, and restart the service if necessary (see [Stopping and starting Websense services, page 370](#)).

If none of these address the issue, check the Windows Event Viewer and **websense.log** file (in the Websense **bin** directory) for error messages from Log Server in order to better understand the problem.

Log Server has not received log files from Filtering Service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Log Server receives Internet usage information from Filtering Service and stores it in the Log Database. If Log Server is not receiving files from Filtering Service, no data is being logged, recent data is not displayed on the **Status > Dashboard** page, and you cannot generate Internet usage reports that include recent data.

Log Server may not be receiving files from Filtering Service if:

- ◆ Filtering Service is not running.

See [Filtering Service is not running](#), page 431, for information about addressing this issue.

- ◆ The two services are not able to communicate across the network.
 - Verify that there have been no recent changes to firewall rules that might affect traffic between the machines on port 55805 (default), or the custom port your organization uses.
 - Use a utility like **telnet** or **ping** to verify that the machines can communicate.
 - Verify that the Log Server IP address and port (55805, by default) is correct on the **Web > Settings > General > Logging** page in the TRITON Manager. If the loopback address (127.0.0.1) or “localhost” is shown, enter the actual IP address of the Log Server machine.
 - Use the **Check Status** button on the **Settings > General > Logging** page to verify that it is possible to connect to Log Server.
If the status check fails:
 - a. Verify there is no firewall blocking the port.
 - b. Run the following command on the Log Server machine to verify that Log Server is listening on the port:


```
netstat -ban > port.txt
```
- ◆ Content Gateway, Network Agent, or a third-party integration product is not configured properly and not receiving Internet traffic.
 - See [Network Agent issues](#), page 435, and [Configure Network Agent](#), page 417, for information about addressing Network Agent configuration issues.
 - See the [Content Gateway Online Help](#) for information about addressing Content Gateway configuration issues.
 - See the [Deployment and Installation Center](#) and your vendor’s documentation for information about other supported integrations.
- ◆ There is not sufficient disk space for Log Server to create new cache files.
See [Low disk space on the Log Server machine](#), page 456, for more information about addressing this issue.
- ◆ Filtering Service is associated with a Policy Server that is not configured for logging or is sending logs to TestLogServer.
See [No Log Server is installed for a Policy Server](#), page 457, and [Configuring how requests are logged](#), page 386, for more information about addressing this issue.
- ◆ Files cannot be written to the cache or BCP folders.
Verify that the path defined for ODBC cache files or BCP files on the **Settings > Reporting > Log Server** page is correct, and that the account used to run Log Server has permissions to write to the path.
- ◆ Log Server did not install properly.
Use the following steps to verify that the Log Server service is registered properly with the Windows operating system:
 1. Use the Windows Services tool to stop the Websense Log Server service.

2. Open a command prompt (**Run > cmd**) and navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin, by default).
3. Enter the following command.

```
LogServer -c
```

 - If no errors display, the service is registered correctly.
 - If errors display, continue with the next step.
4. To remove the Log Server service, enter:

```
LogServer -u
```
5. To register the executable, enter:

```
LogServer -i
```
6. Once again, enter the following command. Verify that no errors appear.

```
LogServer -c
```

If none of the items above addresses your issue:

- ◆ Verify that the Log Server executable version matches the installed product version. To find the Log Server version:
 1. Open a Windows command prompt on the Log Server machine.
 2. Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin, by default).
 3. Enter the command:

```
LogServer -v
```

This should match the version shown on the **Help > About the TRITON Console** page in the TRITON Manager for your web protection software.
 - ◆ Occasionally, the Filtering Service on a Websense V-Series appliance does not restart as expected after a settings change. If the Filtering Service on an appliance stops running, go to the **Status > Modules** page and restart the entire Web module.
 - ◆ If Log Server stops running immediately after restarting and the runtime error “C error (Visual C Runtime Error)” displays, delete the **LogServer.state** file located in the Log Server **Cache** folder (C:\Program Files\Websense\Web Security\bin\Cache, by default) and restart the **Websense Log Server** service.
 - ◆ If you are using TestLogServer, verify that the tool is set up to forward log data to Log Server.
- See support.websense.com for detailed information about TestLogServer.

Low disk space on the Log Server machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Websense Log Server stores Internet activity 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. 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 **Web > Settings > Reporting > Log Server** page in the TRITON Manager. See [Configuring Log Server](#), page 388.



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 System tab of the Status > Dashboard 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 Dashboard page in the Web module of the TRITON Manager.

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 the TRITON Manager.
- ◆ 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 the management console:

1. Log on to the TRITON Manager, open the **Web** module, and click the **Settings** tab of the left navigation pane.
2. Go the **General > Logging** page.
3. Enter the IP address of the Log Server machine in the **Log Server IP address or name** field.
4. Click **OK** to cache your change, and then click **Save and Deploy**.

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.

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.

More than one Log Server is installed for a Policy Server

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Each Policy Server instance can connect to only one instance of Log Server. When multiple instances of Log Server attempt to connect to the same Policy Server, log data is not recorded properly, causing problems with multiple reporting tools.

To resolve this issue:

- ◆ If multiple, active Log Server instances are running, uninstall all but one of the Log Server instances connecting to the Policy Server that is reporting the error.
If you would like to configure multiple Log Server instances to communicate with a central Log Server that is responsible for recording data in the Log Database, see [Extending your web protection deployment](#) in the Deployment and Installation Center.
- ◆ If this error appears, but only one instance of Log Server is active, it is likely that:
 - Policy Server was not running when a Log Server instance was uninstalled.
 - The Policy Server IP address was changed after Log Server was installed.
 - During installation, Log Server connected to a Policy Server instance on another machine. Later, a Policy Server instance was installed on the Log Server machine.

In all of these cases, the safest way to address the problem is:

1. Uninstall the Log Server instance or instances currently connected to the Policy Server instance displaying the error.
2. Stop Websense Policy Server (via the Windows Services tool or the `/opt/Websense/WebsenseDaemonControl` command).
3. Navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin) and make a backup copy of **config.xml** in another location. **Do not skip this step.**

4. Open the original **config.xml** file in a basic text editor (not an XML or HTML editor).
5. Near the top of the file, locate the WebsenseLogServer container. This contains the ID for the “ghost” Log Server instance.

```
<container name="WebsenseLogServer">
```
6. Delete the entire container, including the close tag. For example:

```
<container name="WebsenseLogServer">
  <data name="0c65012f-93af-11e1-8616-f215ee9c7d9d">10.201.136.34</data>
</container>
```
7. Save and close the **config.xml** file.
8. Delete the **config.xml.bak** file from the Websense **bin** directory.
9. Use the Windows Services tool or `/opt/Websense/WebsenseDaemonControl` command to start Websense Policy Server.

Log Database was not created

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 Standard or Enterprise requires **dbcreator** server role membership, **db_datareader** role membership, and membership in one of the following roles:
 - SSQAgentUserRole
 - SQLAgentReader Role
 - SQLAgentOperator Role
 - SQL Server Express: **sysadmin** permissions required

Update the logon account or log on with an account that already has the required permissions, 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.
 To address this issue:
 - If you don't want to upgrade the existing database files, remove or rename them, and then run the installer again.
 - If the existing database files are from a version that can be upgraded, and you want to continue using them, use the SQL Server Management Studio to attach the files to the database engine, 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The Websense Log Database stores Internet usage information for use in presentation reports, investigative reports, and the charts and summaries on the Dashboard page in the Web module of the TRITON Manager.

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 tool 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 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 tool 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 464).
- ◆ 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.

- ◆ Make sure that there are no network interruptions that are preventing the TRITON Manager 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 tool 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 457).

Log Database size causes reporting delays

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 the Web module of the TRITON Manager, go to the **Settings > Reporting > Log Database** page.
2. Locate the **Available Partitions** section of the page.
3. Mark the check box next to any partitions that are not required for current reporting operations, then click **Disable**.
4. Click **OK**, then **Save and Deploy** to implement the change.

See [Log Database sizing guidance](#), page 406, for more information about estimating database size.

More than 100 files in the Log Server cache directory

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Normally, Log Server ODBC cache files or BCP files are moved to the Log Database at a steady rate. If temporary files are accumulating on the Log Server machine, current Internet usage information is not being sent to the Log Database.

Log Server may be unable to process temporary files if:

- ◆ The Log Database is not running, the connection to the Microsoft SQL Server machine is down, or the database is busy. See [Log Database is not available](#), page 460.
- ◆ The Log Database is not installed properly or is the wrong version. See [Log Database was not created](#), page 459.
- ◆ The ETL job has stopped running and the incoming buffer is full.
- ◆ The Log Database is out of allocated disk space. See [Log Server is not recording data in the Log Database](#), page 464.
- ◆ The database creation path is invalid.
- ◆ There is no current active partition.
- ◆ There is a problem with BCP insertion.

- ◆ There is a problem with the size of **tempdb**.

To troubleshoot the problem:

- ◆ Make sure Microsoft SQL Server is running (see [Log Database is not available](#), page 460), and that no other processes that use significant resources, such as a full backup or antivirus scan, are running.
Also check the disk IO to verify that the machine is able to handle a fast insertion rate into the database.
- ◆ Verify that you are using a certified version of Microsoft SQL Server:
 - SQL Server 2008 SP4 (or the latest service pack from Microsoft)
 - SQL Server 2008 R2 SP3 (or the latest service pack from Microsoft)
 - SQL Server 2012 SP2 (or the latest service pack from Microsoft)
 - SQL Server 2014 (or the latest service pack from Microsoft)
 - For very small networks and evaluation environments, SQL Server 2008 R2 Express SP 2 (packaged in the TRITON Unified Installer).
- ◆ Use SQL Server Management Studio to verify that the ETL job is running.
If you are using SQL Server Enterprise or Standard, and the ETL job is not running, make sure the SQL Server Agent service is running on the machine.
If SQL Server Agent is running:
 - Expand the catalog database (wslogdb70) and verify that there are records in the INCOMINGBUFFER. If the INCOMINGBUFFER is full, Log Server cannot add additional records.
 - If records exist in the INCOMINGBUFFER table:
 - a. Locate the **wse_etl_config** table.
 - b. Right-click, then select **Open Table**.
 - c. Change the value for **max_buffer_size** to **40000**.
- ◆ Use SQL Server Management Studio to verify that the **Auto Growth** option is enabled for the catalog database.
- ◆ Go to the **Web > Settings > Reporting > Log Database** page in the TRITON Manager and verify that:
 - The **File Path** entries under **Partition Management** are valid.
 - There is at least one active partition listed under **Available Partitions**.
- ◆ If Log Server has been configured to use BCP insertion, but BCP files are not being processed, change the insertion method to ODBC and see if new cache files are processed:
 1. Go to the **Web > Settings > Reporting > Log Server** page in the TRITON Manager.
 2. Expand the **Log Record Creation** section.
 3. Select the **ODBC (Open Database Connectivity)** radio button.
 4. Click **OK** to cache your changes, then click **Save and Deploy** to implement them.

By default, ODBC cache files are created in the **C:\Program Files\WebSense\Web Security\bin\Cache** directory.

- ◆ The log (ldf) file for the database **tempdb** may be full. Restart the Microsoft SQL Server (MSSQLSERVER) services to clear the tempdb database.

Last successful ETL job ran more than 4 hours ago

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

The ETL (Extract, Transform, and Load) job is responsible for processing data into the partition database. If the job does not run regularly, data is delayed in being written to the Log Database, resulting in reports and Dashboard charts that are out of date.

Typically, the ETL job runs quickly and is scheduled to start again 10 seconds after it completed its last process. If no records are being passed to the database, however, (for example, because there's no traffic due to a network problem, or because Filtering Service or Log Server is not running), the job does not run until it starts receiving data again.

If the job has not run recently:

- ◆ Make sure Microsoft SQL Server is running (see [Log Database is not available, page 460](#)), and that no other processes that use significant resources, such as a full backup or antivirus scan, are running.
 - Also check the disk IO to verify that the machine is able to handle a fast insertion rate into the database.
 - Use the linked procedure to [check for Log Database problems](#).
- ◆ Verify that you are using one of the following certified versions of Microsoft SQL Server:
 - SQL Server 2008 SP4 (or the latest service pack from Microsoft)
 - SQL Server 2008 R2 SP3 (or the latest service pack from Microsoft)
 - SQL Server 2012 SP2 (or the latest service pack from Microsoft)
 - SQL Server 2014 (or the latest service pack from Microsoft)
 - For very small networks and evaluation environments, SQL Server 2008 R2 Express SP 2 (packaged in the TRITON Unified Installer).
- ◆ (*Microsoft SQL Server Standard and Enterprise*) Use the Windows Services tool on the SQL Server machine to verify that the SQL Server Agent service is running.
- ◆ Use SQL Server Management Studio to make sure the ETL job is running. If it isn't, check for errors in the job history and restart or manually run the job.
- ◆ Use the following procedures:
 - [Make sure Filtering Service is sending data](#)
 - [Look for problems on the Log Server machine](#)

You can also use the TestLogServer utility to verify logging behavior. See [Using TestLogServer for Troubleshooting](#).

Log Server is not recording data in the Log Database

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 the **Settings > Reporting > Log Database** page to configure the Log Database to roll over when it reaches approximate 90% of the maximum size. See [Configuring database partition options](#), page 397.

If your IT department has established a maximum amount of disk space for SQL Server operations, contact them for assistance.

Log Server cannot connect to the directory service

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If either of the errors 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.

- ◆ 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If you have deployed a V-Series appliance, the time zone settings on the TRITON management server 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 in the Web module of the TRITON Manager. A logon page or a “logon failed” message is displayed instead of the expected functionality.

To resolve this issue, update the time zone on the TRITON management server and Log Server machines to match the time zone on the appliance, then restart the off-box services.

Investigative report and presentation report issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to troubleshoot issues with investigative and presentation reports:

- ◆ [Presentation Reports Scheduler not connected to Log Database, page 465](#)
- ◆ [Inadequate disk space to generate presentation reports, page 466](#)
- ◆ [Scheduled jobs in presentation reports failed, page 466](#)
- ◆ [Wrong reporting page displayed, page 465](#)
- ◆ [All reports are empty, page 468](#)
- ◆ [Error generating presentation report, or report does not display, page 469](#)
- ◆ [Investigative reports search issues, page 470](#)
- ◆ [General investigative reports issues, page 470](#)

Presentation Reports Scheduler not connected to Log Database

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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, use the Windows Services tool to restart the **Websense TRITON - Web Security** service.

Inadequate disk space to generate presentation reports

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

By default, to generate presentation reports, Websense software uses space in the following folder on the TRITON management server machine:

`C:\Program Files (x86)\Websense\Web Security\ReportingOutput`

If the space available at this location falls below 1 GB, a warning message appears in the Health Alert Summary on the System tab of the **Status > Dashboard** page.

When this message appears, clear disk space on the appropriate disk of the TRITON management server to avoid problems generating presentation reports or other system performance problems.

Scheduled jobs in presentation reports failed

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If one or more scheduled jobs cannot run successfully in presentation reports, the Health Alert Summary on the System tab **Status > Dashboard** 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 385](#), for instructions.
- ◆ There is insufficient disk space on the TRITON management server machine to generate presentation reports. See [Inadequate disk space to generate presentation reports, page 466](#), for more information.
- ◆ Connectivity with the Log Database has been lost. See [No Log Server is installed for a Policy Server, page 457](#), 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.

Trend data is missing from the Log Database

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Trend data is inserted into the Log Database first by the ETL job (which generates daily trend data) and then by the trend job (which generates weekly, monthly, and yearly tables). This data is then used in presentation trend reports.

If there is no trend data in your database, or some trend data is missing:

- ◆ Verify that you have enabled trend data retention on the **Settings > Reporting > Log Database** page in the Web module of the TRITON Manager. The **Store trend data** check box (under Trend Data Retention) must be marked in order for any trend data to be generated and stored.
- ◆ If you are using Microsoft SQL Server **Standard** or **Enterprise**, verify that the SQL Server Agent service is running, and that it is running as the correct user. Use the Windows Services tool to verify that the **SQL Server Agent** service is running.
- ◆ Verify that the **ETL** and **trend** database jobs are running.
The ETL job generates daily trend data, and the trend job runs on a nightly basis to generate weekly, monthly and yearly trend values. Use SQL Server Management Studio to verify that both jobs are running. If they're not, check for errors in the job history and restart or manually run the jobs. See [Last successful ETL job ran more than 4 hours ago](#), page 463, for additional information.

Trend reports are not displaying data

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When you use presentation reports, you can generate trend reports to provide trend information by day, week, month or year. There are separate tables in the Log Database that maintain trend data for each of these time periods.

If the trend reports you generate contain no data, first see:

- ◆ [All reports are empty](#), page 468
- ◆ [Error generating presentation report, or report does not display](#), page 469
- ◆ [Trend data is missing from the Log Database](#), page 467

If these topics don't help to determine the problem, verify that:

- ◆ The report you are running is defined for a trend period that has valid data.
There are 4 different time periods for which trend data can be stored, for which reports can be defined: daily, weekly, monthly, or yearly. Make sure there is trend data for the time option selected for the trend report you are running.
- ◆ Presentation reports is connected to the Log Database.
If the connection to the Log Database has been lost, the presentation reports tool cannot create the reports. See [Log Database is not available](#), page 460.
- ◆ There is disk space available for the report to be created and stored.

The presentation reports tool writes to the disk when it generates a report. See [Inadequate disk space to generate presentation reports](#), page 466.

All reports are empty

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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.
To make sure the appropriate database partitions are active:
 1. Go to the **Web > Settings > Reporting > Log Database** page in the TRITON Manager.
 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.
Click **Save Now** to implement the change.
- ◆ The SQL Server Agent job is active on the Microsoft SQL Server machine. (This service is not used with SQL Server Express.) With Standard or Enterprise additions of Microsoft SQL Server, this job must be running for the log records to be processed into the database by the ETL database job.
 1. Open the Windows Services tool.
 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 is correctly set up to receive log information from Filtering Service.
See [Verify your Log Server configuration](#), page 468.

Verify your Log Server configuration

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Configuration settings must be correct in the Web module of the TRITON Manager 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 the TRITON Manager is connecting to the Log Server successfully.

1. Log on to the TRITON Manager with unconditional Super Administrator permissions and select the **Web** module.
2. Go to the **Settings > General > Logging** page.
3. Enter the IP address or hostname 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 the TRITON Manager 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 and Deploy**.

Next, verify your Log Server settings.

1. Go to the **Settings > Reporting > Log Server** page.
2. Under **Location**, verify that the **Port** matches the value on the **Settings > General > Logging** page.
3. Click **OK** to validate and cache any change, then click **Save and Deploy** to implement it.
4. If you changed the Log Server port setting, use the Windows Services tool to restart the **Websense Log Server** and **Websense TRITON - Web Security** services.

Error generating presentation report, or report does not display

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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) and select HTML format, the report is displayed in the content pane. If you select PDF or XLS format, you are given the option to open the report or save it. 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.

If you encounter this issue, navigate away from the Presentation Reports page in the TRITON Manager, and then run the report again. If that does not work, log off of the TRITON Manager 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 or Chrome, rather than Internet Explorer, when generating reports.

Investigative reports search issues

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 the TRITON Manager, 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 the TRITON Manager 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

- ◆ 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 for web protection solutions

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Use the following articles to troubleshoot problems with Real-Time Monitor, the Status > Dashboard page, and Websense Multiplexer:

- ◆ [Low memory on the Real-Time Monitor machine, page 471](#)
- ◆ [Real-Time Monitor is not running, page 471](#)
- ◆ [Real-Time Monitor is not responding, page 472](#)
- ◆ [No charts appear on the Status > Dashboard page, page 472](#)

- ◆ [There is a forensics data configuration problem, page 473](#)
- ◆ [The forensics repository location could not be reached, page 473](#)
- ◆ [Forensics data will soon exceed a size or age limit, page 473](#)
- ◆ [Websense Multiplexer is not running or not available, page 474](#)

Low memory on the Real-Time Monitor machine

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

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 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

This alert is displayed when the Websense RTM Server service is stopped.

Use the Windows Services tool 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

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

If Real-Time Monitor is installed on a different machine from the TRITON Manager, 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.

No charts appear on the Status > Dashboard page

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

Typically, the **Status > Dashboard** page displays charts and other elements showing the status of your deployment.

- ◆ If you have just deployed a web protection solution, there may not be any reporting data to display. You can use a tool like TestLogServer to see whether traffic is currently being logged. See the [Using TestLogServer for Troubleshooting](#) technical article for instructions.

- ◆ In organizations that use delegated administration, review the reporting permissions for the delegated administrator's role. If **View dashboard charts** is not selected, these chart do not appear for delegated administrators in that role.
- ◆ If the TRITON console loses its connection to the Log Database (for example, because of a network problem, or because the Microsoft SQL Server instance hosting the database is down), no data can be displayed. Check the **Status > Alerts** page for alerts relating to the Log Database.

There is a forensics data configuration problem

Administrator Help | TRITON AP-WEB | Version 8.1.x

When forensics data collection is enabled on the **Settings > Reporting > Dashboard** page in the Web module of the TRITON Manager, transaction details related to attempts to send data out of your network, as well as the actual data files involved, are recorded in a forensics repository.

In rare circumstances, files used to enable collection and storage of this forensic data may be damaged or corrupted. Assistance from Websense Technical Support is required to resolve such issues.

The forensics repository location could not be reached

Administrator Help | TRITON AP-WEB | Version 8.1.x

When forensics data collection is enabled on the **Settings > Reporting > Dashboard** page in the Web module of the TRITON Manager, the administrator provides a location (local directory or UNC path) for storing the file, and credentials for an account with read, write, and delete permissions to the specified directory.

If a health alert indicates problems reaching the forensics repository location, verify that:

- ◆ The path and credential information on the **Settings > Reporting > Dashboard** page is correct.
- ◆ The specified account does have read, write, and delete permissions to the directory.
- ◆ No network problem impedes communication between the TRITON Management Server and the remote machine.

Forensics data will soon exceed a size or age limit

Administrator Help | TRITON AP-WEB | Version 8.1.x

When forensics data collection is enabled on the **Settings > Reporting > Dashboard** page in the Web module of the TRITON Manager, the administrator sets both a maximum size (in GB) for the repository, and a maximum length of time (in days) for storing forensic data.

When the size of the forensics repository approaches the limit, a health alert is displayed. When the limit is reached, the oldest records are deleted, one day's worth of time, until there is room for new, incoming records to be stored.

If the size limit has not been reached, but records are approaching the maximum record age, the health alert is also displayed. When the age limit is reached, records exceeding that limit are deleted.

There is no mechanism available for retrieving deleted forensic data.

Websense Multiplexer is not running or not available

Administrator Help | TRITON AP-WEB and Web Filter & Security | Version 8.1.x

When SIEM integration is enabled, Websense Multiplexer passes Internet activity data (logs) from Filtering Service to both Log Server and the configured SIEM product.

If Multiplexer is not running or not available, a failover feature ensures that Filtering Service passes log data to Log Server. No data, however, is sent to the SIEM product.

To resolve this issue in a Websense appliance-based deployment:

1. If you have not enabled the Multiplexer service, go to the **Administration > Toolbox > Command Line Utility** in the Appliance manager.
Select **TRITON AP-WEB** from the **Component** drop-down list and **multiplexer** from the **Command** drop-down list. Then use the **enable** command.
2. If Multiplexer has already been enabled, but is not running, go to the **Status > General** page in the Appliance manager and restart the Web module.

To resolve this issue in a software-only deployment:

1. Start or restart the Multiplexer service or daemon:
 - Windows: Use the Services tool to start (or restart) the **Websense Multiplexer** service.
 - Linux: Use the **/opt/Websense/WebsenseDaemonControl** command to start (or restart) **Multiplexer**.
2. If the service is unable to restart, the Multiplexer Controller executable may have become unresponsive.
 - Windows: Use Task Manager to stop the **MuxCtrl.exe** process, then use the Services tool to start **Websense Multiplexer**.
 - Linux: Kill the **MuxCtrl** process, then use the **/opt/Websense/WebsenseDaemonControl** command to start **Multiplexer**.

TRITON AP-WEB interoperability issues

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the following articles to troubleshoot issues with Content Gateway and Web Hybrid module components:

- ◆ [Content Gateway is not running](#), page 475
- ◆ [Content Gateway is not available](#), page 476
- ◆ [Content Gateway non-critical alerts](#), page 476
- ◆ [Administrator unable to access other TRITON modules](#), page 479
- ◆ [Sync Service is not available](#), page 479
- ◆ [Sync Service has been unable to download log files](#), page 480
- ◆ [Sync Service has been unable to send data to Log Server](#), page 481
- ◆ [Hybrid policy enforcement data does not appear in reports](#), page 481
- ◆ [Disk space is low on the Sync Service machine](#), page 481
- ◆ [The Sync Service configuration file](#), page 482
- ◆ [Directory Agent is not running](#), page 483
- ◆ [Directory Agent cannot connect to the domain controller](#), page 484
- ◆ [Directory Agent does not support this directory service](#), page 485
- ◆ [The Directory Agent configuration file](#), page 486
- ◆ [Directory Agent command-line parameters](#), page 487
- ◆ [Alerts were received from the hybrid service](#), page 488
- ◆ [Unable to connect to the hybrid service](#), page 489
- ◆ [Hybrid service unable to authenticate connection](#), page 489
- ◆ [Missing key hybrid configuration information](#), page 490

Content Gateway is not running

Administrator Help | TRITON AP-WEB | Version 8.1.x

When a Content Gateway instance registers with Policy Server, that connection is tracked in the Web module of the TRITON Manager.

Information about the Content Gateway instance appears on the **Settings > General > Content Gateway Access** page, and in the Filtering Service Summary on the System tab of the **Status > Dashboard** page. In addition, if the registered instance stops, or is removed, a health alert message is displayed.

- ◆ 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

When a Content Gateway instance registers with Policy Server, that connection is tracked in the Web module of the TRITON Manager. Information about the Content Gateway instance appears on the **Settings > General > Content Gateway Access** page, and in the Filtering Service Summary on the System tab of the **Status > Dashboard** page.

If Policy Server can no longer communicate with the registered instance of Content Gateway, a health alert message is displayed.

- ◆ 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

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	A problem that caused Content Gateway to restart. See the Content Gateway syslog file for information about what caused the reset.
Cache configuration issue	Content Gateway was unable to configure a cache. See “Configuring the Cache” in the Content Gateway Manager Help for more information.
Unable to create cache partition	An error occurred during cache configuration. See “Configuring the Cache” in the Content Gateway Manager Help.

Alert	Description
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 the 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
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>

Alert	Description
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 hostnames 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>
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>

Alert	Description
Content scanning skipped	Content Gateway did not scan content for a requested site that would have ordinarily be scanned. This may occur when Content Gateway is experiencing too many connections, or inadequate system resources (CPU and memory).
WCCP configuration error	See the “WCCP Configuration” section of the Content Gateway Manager Help for configuration parameter details.

Administrator unable to access other TRITON modules

Administrator Help | TRITON AP-WEB | Version 8.1.x

If you receive an error when you click the **Data** or **Email** tab in the TRITON Manager, the local or network account that you use to log on to the TRITON Manager may not have been granted Data or Email module 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 Manager administrator account, **admin**, has full access to all installed modules.

See the TRITON Settings Help (which can be opened from the Help menu on any TRITON Settings page) for more information.

Sync Service is not available

Administrator Help | TRITON AP-WEB | Version 8.1.x

With the Web Hybrid module, 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 your hybrid service account, or if you have attempted to activate the hybrid service, 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 **Web > Settings > Hybrid Configuration > Shared User Data** page in the TRITON Manager 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 management server 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

Sync Service attempts to connect to the hybrid service to download reporting log files at an interval that you configure (see [Schedule communication with the hybrid service, page 216](#)). 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 policy enforcement 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 > Hybrid Service** page to verify that Sync Service is able to connect to the hybrid service. See [Unable to connect to the hybrid service, page 489](#), 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 service account.

Sync Service has been unable to send data to Log Server

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 policy enforcement data does not appear in reports

Administrator Help | TRITON AP-WEB | Version 8.1.x

If Internet activity information for users managed 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 requests are logged](#), page 386.
- ◆ 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 the hybrid service](#), page 216.
- ◆ The **Status > Hybrid Service** page shows that Sync Service has successfully connected to the hybrid service, and retrieved log records. See [Monitor communication with the hybrid service](#), page 222.
- ◆ No health alerts appear on the System tab of the **Status > Dashboard** page indicating Sync Service communication problems or Log Server errors. See [Sync Service has been unable to send data to Log Server](#), page 481.

Disk space is low on the Sync Service machine

Administrator Help | TRITON AP-WEB | Version 8.1.x

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

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **syncservice.ini** file to configure aspects of Sync Service behavior that cannot be configured in the TRITON Manager.

The **syncservice.ini** file is located in the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin, by default).

- ◆ 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 the TRITON Manager. It must match the Sync Service IP address on the **Settings > Hybrid Configuration > Shared User Data** page.
- ◆ **SyncServiceHTTPPort**: The port that Sync Service listens on for communication from Directory Agent and the TRITON Manager (default 55832). It must match the Sync Service port displayed on the **Settings > Hybrid Configuration > Shared User Data** page.
- ◆ **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

Administrator Help | TRITON AP-WEB | Version 8.1.x

With the Web Hybrid module, Websense Directory Agent gathers user information from your directory service and sends it to the hybrid service for use in applying policies.

When Directory Agent is not available, the hybrid service's user data may become outdated.

Make sure that Directory Agent is installed (software) or enabled (appliance), and that the service or daemon is running.

- ◆ Appliance: Go to the **Status > General** page in the Appliance manager and verify that Websense Directory Agent appears as a running service.
 - If Directory Agent is listed as a disabled service (in pale gray), go to the **Administration > Toolbox > Command Line Utility** page and select **directory-agent-service**, then **enable**.
 - If Directory Agent is enabled but not running, restart the Web module.
- ◆ Windows: Use the Windows Services tool to start the service or verify that it is running.
- ◆ Linux: Use the **/opt/Websense/WebsenseDaemonControl** command to start the daemon or verify that it is running.

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 **Administration > Logs** page (Appliance), Event Viewer (Windows), or websense.log file (Linux) for errors.
- ◆ For software installations, navigate to the Websense **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 policy enforcement.

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 **Web > Settings > General > Directory Services** page in the TRITON Manager 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

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

Administrator Help | TRITON AP-WEB | Version 8.1.x

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)
- ◆ Oracle (formerly Sun Java System) Directory
- ◆ Novell eDirectory
- ◆ If you are not using a supported directory service, the hybrid service can still manage filtered locations. User and group-based policy enforcement, however, cannot be performed.

The Directory Agent configuration file

Administrator Help | TRITON AP-WEB | Version 8.1.x

Use the **das.ini** file to configure aspects of Directory Agent behavior that cannot be configured in the TRITON Manager. 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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default).

- ◆ Use a text editor to edit the file.
- ◆ For parameters that can take multiple values, use the pipe symbol (“|”) to separate entries.
- ◆ For parameters that are either enabled or disabled, the only valid values are **0** (for disable) and **1** (for enable). In this file, the values “on” and “off” cannot be used.
- ◆ 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\Web Security\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 the hybrid service.

```
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`
- ◆ Configuration settings if you are using Sun Java System Directory or Oracle Directory Server to send user and group information to the hybrid service. Enable these parameters by removing the # symbol from the beginning of the lines.
`# GroupMembershipAttribute=uniqueMember`
`# MemberOfAttribute=memberOf`
- ◆ Whether or not Directory Agent follows LDAP referrals. Takes a value of 1 (enabled) or 0 (disabled).
`EnableLDAPReferrals=1`

Directory Agent command-line parameters

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

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 the hybrid service 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 System dashboard in the Web module of the TRITON Manager.
- ◆ A more specific alert is shown on the **Status > Alerts** page under Hybrid Service 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 and Deploy 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 the hybrid service 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 System dashboard summary alert and the alerts on the Status > Alerts page are cleared.

Unable to connect to the hybrid service

Administrator Help | TRITON AP-WEB | Version 8.1.x

With the Web Hybrid module, on-premises and hybrid components must communicate regularly to ensure consistent policy enforcement 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 Service 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

In environments that use the hybrid service, 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 service account. The security token is sent to the **contact email address** specified on the **Settings > General > Account** page in the Web module of the TRITON Manager.

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.
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 service 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 **Web > Settings > General > Account** page in the TRITON Manager.
5. Scroll down to the Hybrid Service 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

Administrator Help | TRITON AP-WEB | Version 8.1.x

In environments that use hybrid the hybrid service, 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 service 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 the hybrid service 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 service account.

Copyrights

Websense® TRITON® AP-WEB and Web Filter & Security

©1996–2015, Websense, Inc.

All rights reserved.

10900 Stonelake Blvd, 3rd Floor, Austin, TX 78759, USA

Published 2015

Printed in the United States of America

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.

For other copyright information, refer to:

- ◆ *Trademarks*
- ◆ *Other Acknowledgements*

Trademarks

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

Microsoft, Windows Server 2008, Windows Server 2012, Windows XP, Internet Explorer, and Active Directory are trademarks or registered trademarks of Microsoft Corporation.

The following is a registered trademark of Novell, Inc., in the United States and other countries: Novell Directory Services.

Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates.

Firefox is a registered trademark of the Mozilla Foundation.

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

Other Acknowledgements

Portions of this Websense Software may utilize the following copyrighted material, the use of which is hereby acknowledged.

7-Zip

License for use and distribution

7-Zip Copyright (C) 1999-2012 Igor Pavlov.

Licenses for files are:

1) 7z.dll: GNU LGPL + unRAR restriction

2) All other files: GNU LGPL

The GNU LGPL + unRAR restriction means that you must follow both GNU LGPL rules and unRAR restriction rules.

Note:

You can use 7-Zip on any computer, including a computer in a commercial organization.

You don't need to register or pay for 7-Zip.

Apache HTTP Server

httpd-2.4.3

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache Ant

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache APR

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache APR Utilities

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache Axis2

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache log4cxx

Version 0.9.8

Copyright 2004-2007 The Apache Software Foundation

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

Licensed under the Apache License, Version 2.0.

Apache Log4j

Version 1.2.14

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

Licensed under the Apache License, Version 2.0.

Apache Maven

Versions 2.0.8 and 3.0.4

Apache Maven

Copyright 2001-2011 The Apache Software Foundation

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

This product includes software (Plexus and Classworlds) developed by The Codehaus Foundation (<http://www.codehaus.org/>).

This product includes software (Aether, Sisu, Spice, Plexus Ciper and Sec Dispatcher) developed by Sonatype Inc. (<http://www.sonatype.org/>).

This product includes software (NekoHTML) developed by SourceForge (<http://www.sourceforge.net/>).

This product includes software (Guice) developed by Google Code (<http://code.google.com/>).

Licensed under the Apache License, Version 2.0.

Apache POI

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>). Licensed under the Apache License, Version 2.0.

Apache Tomcat

Copyright 1999-2010 The Apache Software Foundation

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

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

The Windows Installer is built with the Nullsoft Scriptable Install System (NSIS), which is open source software. The original software and related information is available at <http://nsis.sourceforge.net>.

Java compilation software for JSP pages is provided by Eclipse, which is open source software. The original software and related information is available at <http://www.eclipse.org>.

APACHE TOMCAT SUBCOMPONENTS:

Apache Tomcat includes a number of subcomponents with separate copyright notices and license terms. Your use of these subcomponents is subject to the terms and conditions of the following licenses.

For the jasper-jdt.jar component:

Eclipse Public License - v 1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS ECLIPSE PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and
- b) in the case of each subsequent Contributor:
 - i) changes to the Program, and
 - ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents" mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This

patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

- a) it complies with the terms and conditions of this Agreement; and
- b) its license agreement:
 - i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;
 - ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;
 - iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and
 - iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

- a) it must be made available under this Agreement; and
- b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and co-operate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of use

ing and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. The Eclipse Foundation is the initial Agreement Steward. The Eclipse Foundation may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

For the Windows Installer component:

- * All NSIS source code, plug-ins, documentation, examples, header files and graphics, with the exception of the compression modules and where otherwise noted, are licensed under the zlib/libpng license.
- * The zlib compression module for NSIS is licensed under the zlib/libpng license.
- * The bzip2 compression module for NSIS is licensed under the bzip2 license.
- * The lzma compression module for NSIS is licensed under the Common Public License version 1.0.

zlib/libpng license

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

3. This notice may not be removed or altered from any source distribution.

bzip2 license

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
3. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Julian Seward, Cambridge, UK.

jseward@acm.org

Common Public License version 1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and b) in the case of each subsequent Contributor:

- i) changes to the Program, and
- ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

- a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

- b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include

the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

a) it complies with the terms and conditions of this Agreement; and

b) its license agreement:

i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

a) it must be made available under this Agreement; and

b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and co-operate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

Special exception for LZMA compression module

Igor Pavlov and Amir Szekely, the authors of the LZMA compression module for NSIS, expressly permit you to statically or dynamically link your code (or bind by name) to the files from the LZMA compression module for NSIS without subjecting your linked code to the terms of the Common Public license version 1.0. Any modifications or additions to files from the LZMA compression module for NSIS, however, are subject to the terms of the Common Public License version 1.0.

Beecrypt

Beecrypt is distributed under terms of the LGPL. Websense makes no modifications to any LGPL copyrighted source code. Information about the LGPL and code covered under this license can be found at <http://www.gnu.org>. Information about Beecrypt can be found here - <http://beecrypt.sourceforge.net/>.

Boost Regex Libraries

Version 1.54.0

Copyright (c) 1998-2000

Dr John Maddock

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is

hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Dr John Maddock makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

cURL

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2012, Daniel Stenberg, <daniel@haxx.se>.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Electric Fence

Version 2.1.13

Distributed under a GNU General Public License.

FusionCharts Enterprise

Version 3.2

Copyright 2010 FusionCharts Technologies LLC.

gperftools

Version 2.0

Copyright (c) 2005, Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

GNU IDN Library

Version 1.19

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

gSOAP

Version 2.8.10

Part of the software embedded in this product is gSOAP software.

Portions created by gSOAP are Copyright (C) 2001-2009 Robert A. van Engelen, Genivia inc. All Rights Reserved.

THE SOFTWARE IN THIS PRODUCT WAS IN PART PROVIDED BY GENIVIA INC AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Gzip

Version 1.0.4

Gzip is covered under the LGPL. Websense makes no modifications to any LGPL copyrighted source code. Information about the LGPL and code covered under this license can be found at <http://www.gnu.org>.

iTextPDF

GNU AFFERO GENERAL PUBLIC LICENSE

Version 3, 19 November 2007

Copyright (C) 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

TERMS AND CONDITIONS

0. Definitions.

"This License" refers to version 3 of the GNU Affero General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact

all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

(a) The work must carry prominent notices stating that you modified it, and giving a relevant date.

(b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".

(c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.

(d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

(a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.

(b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

(c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

(d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

(e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other

information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- (a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- (b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- (c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- (d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- (e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- (f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder

notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) pri-

marily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Remote Network Interaction; Use with the GNU General Public License.

Notwithstanding any other provision of this License, if you modify the Program, your modified version must prominently offer all users interacting with it remotely through a computer network (if your version supports such interaction) an opportunity to receive the Corresponding Source of your version by providing access to the Corresponding Source from a network server at no charge, through some standard or customary means of facilitating copying of software. This Corresponding Source shall include the Corresponding Source for any work covered by version 3 of the GNU General Public License that is incorporated pursuant to the following paragraph.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the work with which it is combined will remain governed by version 3 of the GNU General Public License.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU Affero General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU Affero General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU Affero General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU Affero General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

Jasper Reports

Version 4.0.1

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed. [This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Jasper is distributed under the terms of the GPL. Websense makes no modifications to any GPL copyrighted source code. Information about the GPL and code covered under this license can be found at <http://www.gnu.org>.

Java and JVM

Version 1.7.0_76

Welcome to the Oracle Web sites (the "Site"). Through the Site, you have access to a variety of resources and content. These include: (a) software and software as a service offerings ("Software"); (b) Web pages, data, messages, text, images, photographs, graphics, audio and video such as podcasts and Webcasts, and documents such as press releases, white papers and product data sheets ("Materials"); and (c) forums, discussion groups, chat areas, bulletin boards, blogs, wikis, e-mail functions, and other services in connection with which you can upload, download, share, email, post, publish, transmit or otherwise access or make available Content (as defined below) ("Community Services"). Software, Materials, Community Services, and other information, content and services are collectively referred to as "Content." The following are terms of a legal agreement between you ("You" or "Your") and Oracle Corporation and its affiliated companies ("We" or "Oracle"). By accessing or using the Site or the Content provided on or through the Site, you agree to follow and be bound by the following terms and conditions concerning your access to and use of the Site and the Content provided on or through the Site ("Terms of Use") and our Privacy Policy. Oracle may revise the Terms of Use and Privacy Policy at any time without notice to you. The revised Terms of Use and Privacy Policy will be effective when posted. You can review the most current Terms of Use at <http://www.oracle.com/us/legal/terms/index.html> and Privacy Policy at <http://www.oracle.com/us/legal/privacy/overview/index.html>.

JDOM

Version 1.0

This product includes software developed by the JDOM Project (<http://www.jdom.org/>).

Copyright (C) 2000-2004 Jason Hunter & Brett McLaughlin. All rights reserved.

JNI

Oracle Binary Code License Agreement for the Java SE Platform Products and JavaFX

ORACLE AMERICA, INC. ("ORACLE"), FOR AND ON BEHALF OF ITSELF AND ITS SUBSIDIARIES AND AFFILIATES UNDER COMMON CONTROL, IS WILLING TO LICENSE THE SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS (COLLECTIVELY "AGREEMENT"). PLEASE READ THE AGREEMENT CAREFULLY. BY SELECTING THE "ACCEPT LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND/OR BY USING THE SOFTWARE YOU ACKNOWLEDGE THAT YOU HAVE READ THE TERMS AND AGREE TO THEM. IF YOU ARE AGREEING TO THESE TERMS ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO BE BOUND BY THE TERMS, THEN SELECT THE "DECLINE LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND YOU MUST NOT USE THE SOFTWARE ON THIS SITE OR ANY OTHER MEDIA ON WHICH THE SOFTWARE IS CONTAINED.

JSON

Copyright (c) 2002 JSON.org

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

The Software shall be used for Good, not Evil.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Jsvc

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>).
Licensed under the Apache License, Version 2.0.

JUnit

Version 3.8

Copyright (c) 2000-2006, www.hamcrest.org

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of Hamcrest nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

libarchive

Version 3.1.2

Copyright (c) 2003-2009 by Tim Kientzle

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer in this position and unchanged.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR(S) "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.

IN NO EVENT SHALL THE AUTHOR(S) BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

libcurl

Version 7.30.0

Copyright (c) 1996 - 2013, Daniel Stenberg, <daniel@haxx.se>.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN

NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Libnet

libnet 1.1.3

Copyright (c) 1998 - 2002

Mike D. Schiffman <mike@infonexus.com>

<http://www.packetfactory.net/libnet>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

LibPCAP

Version 0.9.7

License: BSD

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Libpcxx

Version 4.0

Copyright (c) 2001-2011 Jeroen T. Vermeulen <jtv@xs4all.nl>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the author, nor the names of other contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Net-SNMP

----- Part 1: CMU/UCD copyright notice: (BSD like) -----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE,

DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara,

California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse

or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2004, Sparta, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Novell Header Files and Novell NLM and Network Libraries for C

Novell Developer License Agreement

1.PURPOSE. Novell makes available the Novell developer software and documentation (the "Software") under the terms of this Novell Developer License Agreement. If You do not agree to these terms, You may not access or use the Software.

2.DEFINITIONS. The following terms have the meanings assigned to the below:

"Derivative Software" means the Binary code Software and/or Binary code that results from your compilation of modified or unmodified source code Software. Derivative Software may not include Early Access Release Materials or Internal Tools.

"Developer Product" means Your own computer product that incorporates Derivative Software and does not substantially duplicate the capabilities or compete with the Derivative Software or the Software.

"Early Access Release Materials" means Software and/or Documentation Novell identifies as "beta", "pre-release," "futures," or as an "early access release."

"Internal Tools" means compatibility criteria, test suites, test tools, end user products, and other programs designed to aid in the development of, but not be incorporated in, Derivative Software Products.

3.LICENSES. Use of the Software and Documentation is subject to the restrictions contained in this Agreement. The Software contains various software programs with different license rights. Novell grants You the non-exclusive, non-transferable right to: a) internally use the Early Access Release Materials and Internal Tools in support of Your efforts to develop Derivative Software hereunder (but You may not include any portion of Early Access Release Materials or Internal Tools in Derivative Software); b) use, modify, and compile source code Software for the purpose of creating Developer Products, if the source code Software is identified as licensed for use in creating Developer Products; and, c) reproduce and distribute Derivative Software as part of a Developer Product.

OpenLDAP

Version 2.4.23

The OpenLDAP Public License

Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain copyright statements and notices,
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City,

California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

OpenSSL

Version 1.0.1m

Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OPENSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF

MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the rouines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

OpenSSL FIPS

Version 2.0.5

Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (ey@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (ey@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (ey@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes cryptographic software written by Eric Young (ey@cryptsoft.com)"

The word 'cryptographic' can be left out if the rouines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,

DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

OSGi

Version 2.0.1

This product includes software developed at The OSGi Alliance (<http://www.osgi.org/>).

Copyright (c) OSGi Alliance (2000, 2007).

Licensed under the Apache License 2.0.

PostgreSQL

Version 9.1.15

PostgreSQL Database Management System (formerly known as Postgres, then as Postgres95)

Portions Copyright (c) 1996-2010, PostgreSQL Global Development Group

Portions Copyright (c) 1994, The Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Rosette

Version 4.2.6

Basis Technology Corporation

Website/Services Terms of Use

This page states the terms of use for the Basis Technology Corporation ("Basis") website (the "Website") and the services available through the Website, such as the web-based demonstration and other services (such services, collectively, the "Services"). Please read these Terms of Use carefully. Basis may occasionally update and amend them. By using this Website and/or the Services, you signify that you have read, understood, and agree to be bound by these Terms of Use. These Terms of Use represent a legal agreement between you and Basis. If you do not accept these Terms of Use, then do not use the Website or the Services.

1. Description of Services. The Website is offered to you, enabling access to a variety of resources, including download areas, product and other information, and the Services. Basis will try to provide uninterrupted access to the Website and Services; however, access may be suspended, restricted or terminated at any time and without notice, and if the Website and/or the Services are unavailable for any time for any reason, Basis shall not be liable.

2. Conduct When Using the Website or the Services and Limitations on Use. You understand that any content you post, upload, display or otherwise make available via the Services or the Website is solely your responsibility and that you solely are responsible for your conduct on the Website. This means that Basis is not responsible for any such content, and Basis does not guarantee the accuracy, integrity or quality of such content or that it does not violate any laws or rights of others.

RSA_MD5

Version 1.02

Translation and modification (c) 1995 by Mordechai T. Abzug

Thanks to Martin Cleaver for making it happy on Windows NT and Solaris.

This translation/ modification is provided "as is," without express or implied warranty of any kind.

The translator/ modifier does not claim (1) that MD5 will do what you think it does; (2) that this translation/ modification is accurate; or (3) that this software is "merchantable." (Language for this disclaimer partially copied from the disclaimer below).

Based on:

MD5.H - header file for MD5C.C

MD5C.C - RSA Data Security, Inc., MD5 message-digest algorithm

MDDRIVER.C - test driver for MD2, MD4 and MD5

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

RTPatch

Version 5.0.1

RTPatch Server is a trademark of POCKET SOFT, Inc.

Samba

Version 3.0.21a

Samba is distributed under terms of the GPL. Websense makes no modifications to any GPL copyrighted source code. Information about the GPL and code covered under this license can be found at <http://www.gnu.org>.

Slony-I

Version 2.1.2

Slony-I - A replication system for the PostgreSQL Database Management System

Copyright (c) 2003-2009, PostgreSQL Global Development Group

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

SQLite

Version 3.2.0

All of the code and documentation in SQLite has been dedicated to the public domain by the authors. All code authors, and representatives of the companies they work for, have signed affidavits dedicating their contributions to the public domain and originals of those signed affidavits are stored in a firesafe at the main offices of Hwaci. Anyone is free to copy, modify, publish, use, compile, sell, or distribute the original SQLite code, either in source code form or as a compiled binary, for any purpose, commercial or non-commercial, and by any means.

Tcl

Version 8.3.5

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARS. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

tclreadline

Version 2.1.0

Copyright (C) 1998 - 2007, Johannes Zellner. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- * Neither the name of Johannes Zellner nor the names of contributors to this software may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,

STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

tinyJSON

Version 1.3.0

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

UUID

Copyright (c) 1990- 1993, 1996 Open Software Foundation, Inc.

Copyright (c) 1989 by Hewlett-Packard Company, Palo Alto, Ca. &

Digital Equipment Corporation, Maynard, Mass.

Copyright (c) 1998 Microsoft.

To anyone who acknowledges that this file is provided "AS IS" without any express or implied warranty: permission to use, copy, modify, and distribute this file for any purpose is hereby granted without fee, provided that the above copyright notices and this notice appears in all source code copies, and that none of the names of Open Software Foundation, Inc., Hewlett-Packard Company, or Digital Equipment Corporation be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Neither Open Software Foundation, Inc., Hewlett-Packard Company, Microsoft, nor Digital Equipment Corporation makes any representations about the suitability of this software for any purpose.

WinPCAP

Version 4.1.3

Copyright (c) 1999 - 2005 NetGroup, Politecnico di Torino (Italy).

Copyright (c) 2005 - 2010 CACE Technologies, Davis (California).

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the Politecnico di Torino, CACE Technologies nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE

LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software developed by the University of California, Lawrence Berkeley Laboratory and its contributors.

This product includes software developed by the Kungliga Tekniska Högskolan and its contributors.

This product includes software developed by Yen Yen Lim and North Dakota State University.

Portions Copyright (c) 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes software developed by the University of California, Berkeley and its contributors."
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INSTITUTE AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright (c) 1983 Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Portions Copyright (c) 1995, 1996, 1997 Kungliga Tekniska Högskolan (Royal Institute of Technology, Stockholm, Sweden). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes software developed by the Kungliga Tekniska Högskolan and its contributors."
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INSTITUTE AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE

DISCLAIMED. IN NO EVENT SHALL THE INSTITUTE OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright (c) 1997 Yen Yen Lim and North Dakota State University. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes software developed by Yen Yen Lim and North Dakota State University"
4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright (c) 1993 by Digital Equipment Corporation.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Digital Equipment Corporation not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED "AS IS" AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL DIGITAL EQUIPMENT CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Portions Copyright (C) 1995, 1996, 1997, 1998, and 1999 WIDE Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,

TY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright (c) 1996 Juniper Networks, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that: (1) source code distributions retain the above copyright notice and this paragraph in its entirety, (2) distributions including binary code include the above copyright notice and this paragraph in its entirety in the documentation or other materials provided with the distribution. The name of Juniper Networks may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Portions Copyright (c) 2001 Daniel Hartmeier All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright 1989 by Carnegie Mellon.

Permission to use, copy, modify, and distribute this program for any purpose and without fee is hereby granted, provided that this copyright and permission notice appear on all copies and supporting documentation, the name of Carnegie Mellon not be used in advertising or publicity pertaining to distribution of the program without specific prior permission, and notice be given in supporting documentation that copying and distribution is by permission of Carnegie Mellon and Stanford University. Carnegie Mellon makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

winpthread

This file is covered under the following Copyright:

Copyright (C) 2001,2006 Ross P. Johnson

All rights reserved.

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Pthreads-win32 is covered by the GNU Lesser General Public License

Xalan

Version 2.7.2

The Apache Software License, Version 1.1

Copyright (c) 1999-2003 The Apache Software Foundation. All rights reserved.

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

Xerces

Version 3.1.1

Copyright 1999-2004 The Apache Software Foundation.

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

Portions of this software were originally based on the following:

- software copyright (c) 1999, IBM Corporation., <http://www.ibm.com>.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

zlib

Version 1.2.8

(C) 1995-2013 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

